$1.75

QEX

ARRL Experimenter's Exchange

October 1997

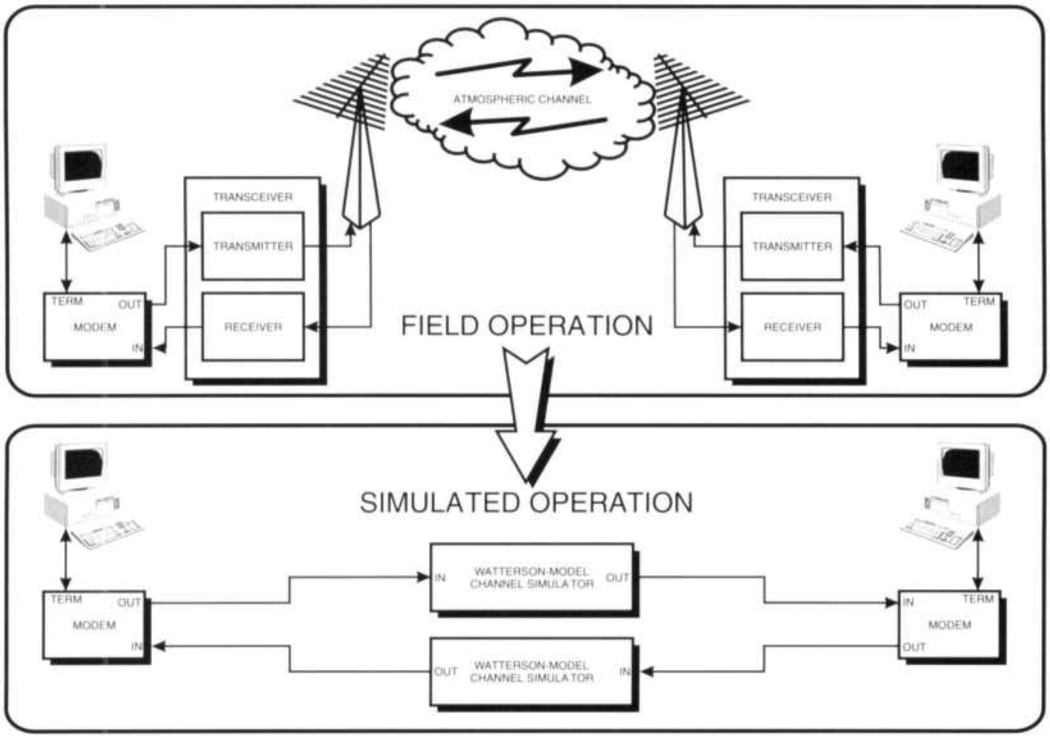

The Institute for Telecommunication Science HF Modem and Protocol Tests

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

QEX

QEX (ISSN: 0886-8093) is published monthly in July 97, August 97, September 97, October 97, November 97 and December 97 by the American Radio Relay League, 225 Main Street, Newington CT 06111-1494. Subscription rate for 12 issues to ARRL members is $15; nonmembers $27. Other rates are listed below. Periodicals postage paid at Hartford CT and at additional mailing offices.

POSTMASTER: Form 3579 requested. Send address changes to: QEX, 225 Main St, Newington CT, 06111-1494

Issue No. 183

Advertising Information Contact:

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

Circulation Department

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

Offices

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

Subscription rate for 12 issues:

In the US: ARRL Member $15, nonmember $27;

US, Canada and Mexico by First Class Mail: ARRL Member $28, nonmember $40;

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

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

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

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

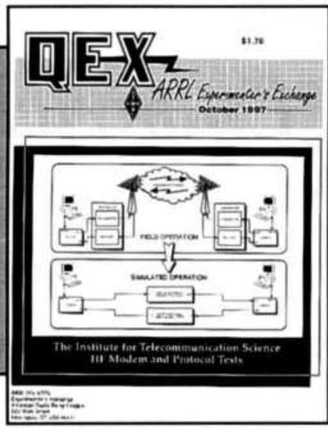

About the Cover

The results of the ITS tests appear in the July 1996 *QST*. Tim Riley and Dennis Bodson, W4PWF, explain how that testing was done.

Features

Columns

October 1997 QEX Advertising Index

Purpose of *QEX*:

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

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

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

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

Empirically Speaking

Inputs to the Editor

In the past few weeks I have been contacting many well known hams for their thoughts and suggestions on *QEX*. The answers have been thoughtful and very helpful, but I have been a bit disappointed with the response from our general readers. Last month in this column I told you that we were looking for letters to the editor and would be publishing some of them. Your letters are a very important input to us, and they can provide valuable and interesting technical dialog for all our readers. If you have some neat idea that is too brief for a fully fledged article, by all means send us a letter. The length and subject matter you see in *QST's* "Technical Correspondence" column are what we have in mind.

For letters challenging the technical content of an article, our standard procedure is to send a copy of the letter to the author for comment. We then publish both the original letter and the reply. We only ask that you stick to technical arguments. We make an effort to see that each article is technically correct, but some things do get past us—or the subject may be genuinely controversial. Letters to the Editor are a way to keep us honest.

I have been very happy to see an increase in article submissions: in particular, antenna articles. This is great, but we need more. In the past, many *QEX* articles have been at quite a high level, and there may be a perception that we only take engineering-level stuff. *Not so!* We will continue to publish engineering-level material in *QEX*, but we also want more balance. That means straightforward practical articles on how to build it, etc. These kinds of articles are our greatest need, and I encourage you to submit them. Remember, we pay $50 per published page. You won't get rich, but that slick idea you have sitting on your bench right now might just pay for itself. Also, keep in mind that we can get your ideas into print faster than other publications. Even if you don't need the money, do it for the glory (well a little bit maybe).

This Month in *QEX*

This month we have two articles on the practical side of amplifier design. Ian White, G3SEK, takes a hard look at screen-grid circuits for tetrode amplifiers. He points out some common problems, along with their solutions. Dave Kirkby, G8WRB, gives us very useful advice on selecting fans for forced-air cooling in amplifiers. His article takes the by-guess and by-golly process and replaces it with some sound but practical science.

For those of you interested in the details of HF modems and their protocols, Dennis Bodson, W4PWF and Tim Riley have a lengthy exposition of just how they tested the HF modems and protocols, the results of which were reported in July 1996 *QST*. Besides the test methods, I found this article gives a good review of the different systems and their strengths and weaknesses, a subject on which I had little knowledge. —*73, Rudy Severns, N6LF,* **rseverns@arrl.org**

HF Modems and Protocols: An Approach to Testing

Which protocol is best? Here's how some top-level agencies pursue that question.

By Tim Riley* and Dennis Bodson, W4PWF**

Certain commercial equipment and programs are identified in this report to adequately explain the operation of the test procedure and the equipment used for testing. In no case does such identification imply recommendation or endorsement by the National Telecommunications and Information Administration, nor does it imply that the program or equipment identified is necessarily the best available for this application.

1. INTRODUCTION AND BRIEF TEST SUMMARY

In 1995, the National Telecommunications and Information Administration's (NTIA) Institute for Telecommunication Science (ITS) was sponsored by the National Communication System (NCS) to support the Federal Emergency Management Agency's (FEMA) mission of reestablishing and maintaining communications during and after emergency situations. To that end, ITS undertook a comparative study of various HF modem protocols, performed under laboratory (simulated) conditions to assure consistency of results. One of the end products of this testing was an article published in the July 1996 issue of *QST*.[1] That article gave a brief overview of the test setup and procedures and reported the outcome of the testing. This article will further elaborate on ITS's HF test bed, its use in the HF modem/protocol testing and background on the various protocols tested, as well as a description of the HF channel simulator used, the channel model it is based on and the channel conditions used in the tests.

The main goal was to test protocols, independent of their implementation in various modems. Obviously, this goal is unobtainable; the only way to test a pure protocol is through simulation and modeling. We wanted these test results to be applicable to real users in the real world. The manner in which a protocol was implemented in a manufacturer's product significantly affects the outcome of the tests. We tried to minimize implementation-specific effects by using the modems in their rawest possible state; no vendor-supplied software was used. The modems were controlled directly by ITS-developed software, common to all modems. The modems were configured to their optimum settings; when these were not specified by the manufacturer, default settings were used. Since data compression in some modems could not be turned off, data compression in all modems was turned on, even though not all modems used the same type of compression and compression is not part of the protocols. In addition to these caveats, others will be covered in the remainder of this article.

There were five protocols covered by the tests: AMTOR (Amateur Teletype-Over-Radio), CLOVER, G-TOR(Golay-Teletype-Over-Radio), PACTOR

[1]Notes appear on page 13.

*US Department of Commerce
NTIA/ITS
325 Broadway
Boulder, CO 80303

**National Communications System
701 S Court House Rd
Arlington, VA 22204-2198

(Packet-Teletype-Over-Radio) and PACTOR2. Four pairs of modems, implementing one or more of these protocols, were subjected to tests that involved throughput measurement under various simulated HF channel conditions, while in a symmetrical, back- to-back configuration (Fig 1). Testing involved transferring a 15-kB plain text, mixed case (upper case for the AMTOR test) file, supplied by FEMA, between the two modems. The modems were configured to operate in an automatic repeat request (ARQ) mode with optimized settings, where allowed. This would support FEMA's requirements for maximum, reliable throughput under adverse conditions, since ARQ is intended to be a method of transferring data error-free. Testing continued until either the file transfer was competed or the link was lost (usually due to degraded channel conditions). While ARQ should be able to transfer a file under extremely adverse conditions, it could require an unlimited amount of time to do so. Time constraints kept us from allowing all tests to run to completion or to run successfully. In the laboratory, as in the real world, there is a limit to how long we can wait for a task to finish, successfully or not.

The channel conditions during testing corresponded to International Telecommunications Union Radiocommunications Sector (ITU-R) recommended simulated channel conditions of good and poor (see section 3.2). The signal-to-noise ratio (SNR) varied from 20 dB to 0 dB (in 1 dB increments) over the range of tests for all but one modem/protocol, but remained fixed for each test. All channel conditions remained static during each individual test to simplify data analysis and reporting. In any case, the ITU-R does not offer any recommended procedures or parameters for testing under dynamic channel conditions. Since this was intended to be a comparison test, the main objective was to establish consistent and uniform test conditions for each modem/protocol pair.

Once a suite of tests was competed on a modem/protocol pair, the data was culled. Tests that were aborted due to link loss, and tests that resulted in the receipt of erroneous data (for whatever reason), were dropped. The remaining results were plotted and presented in the *QST* article. It was not ITS intention to draw any conclusions on the results of the tests, nor did we wish to make any recommendations based on the results. The information was offered to FEMA (as well as existing and future users of HF radio modems and protocols) to help them reach their own conclusions.

Given the complexity of the various modem/protocol pairs, the varying implementations of a protocol in different modems, as well as the number of different protocols implemented in the same modem, an impossibly large number of configurations could have been tested. As a result of the relatively low throughput of these protocols, time was a major factor in the decision to limit the configurations that would be tested. The testing was designed primarily to meet FEMA's main requirements of reliable and fast throughput. Neither ease-of-use, compatibility nor cost was considered. Not surprisingly, those modem/protocols with the highest throughput were the most expensive and complex. Like many users, ITS has to watch its budget (both time and money) as well; consequently, we couldn't test all implementations of a particular proto-

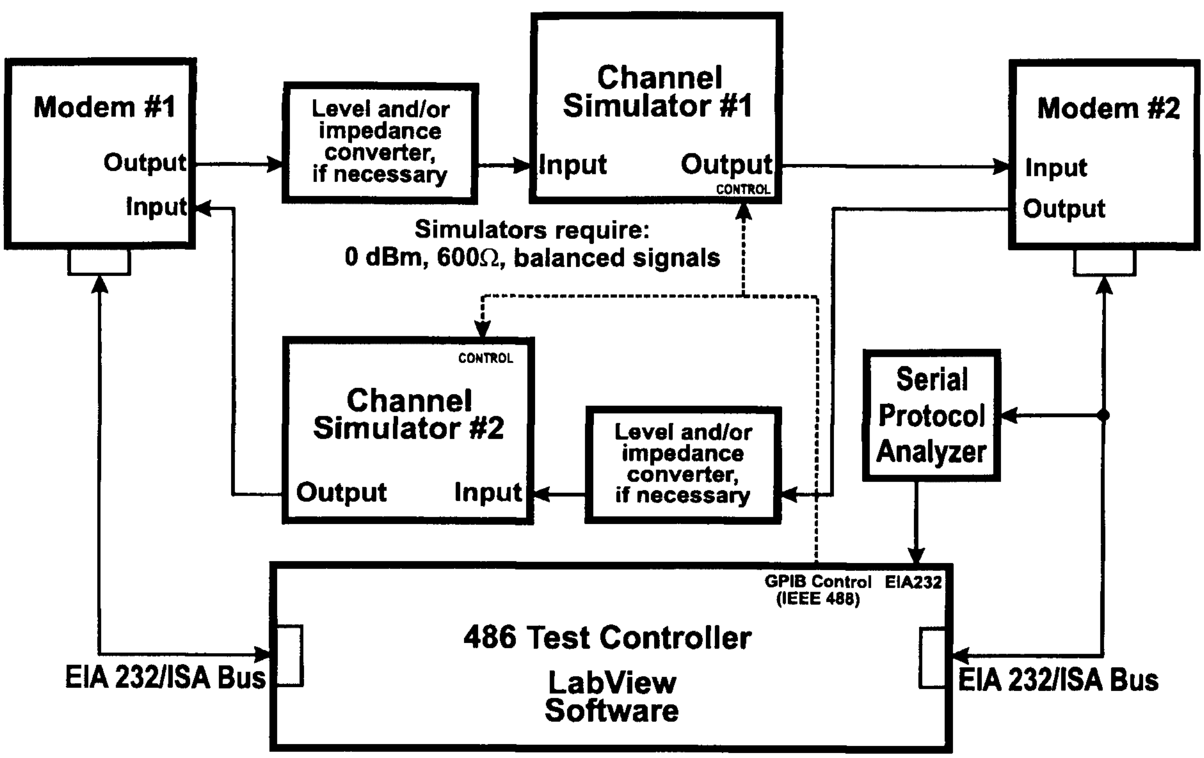

Fig 1—A block diagram of two modems in a symmetrical, back-to-back configuration with channel simulators and interface hardware.

col. It was also impractical to test the most recent and up-to-date modem/ protocol version. The tests were intended merely as a baseline for comparison purposes. We hope to conduct future tests to help answer the question of compatibility between different modems implementing the same protocol and expand the baseline to include additional protocol implementations and more up-to-date versions of those protocols.

2. HF MODEM/PROTOCOL TECHNIQUES

When digital communication was adapted to HF radio, it used technology borrowed from the computer networking arena. Due to the difference in media characteristics, the technology was not an ideal fit. Over the years, newer methods have been developed by Amateur Radio operators, and several have shown enough promise to be accepted by users and manufacturers.

2.1 Protocol Characteristics

When packet protocol was implemented in HF radio communication, it was a version of the X.25 communication protocol developed by the Defense Advanced Research Projects Agency (DARPA). It was initially used over hard-wired and telephone-circuit-based computer networks.[2] Designed with the quality (SNR, interference and error rate) and bandwidth of a telephone circuit in mind, it was effective, but not very efficient. The signal contained a large amount of overhead (synchronization, addressing and a robust error-detection scheme), necessary for transmitting data error-free. When used over an HF transmission channel, with its bandwidth limitations, high noise, fading and frequency shifts and spreads, X.25 is not very effective. It is a half-duplex system; transmitted blocks must be received, verified and acknowledged before subsequent blocks are sent. Consequently, the channel is idle much of the time. The only way the protocol can adapt to worsening conditions is through limited adjustment of the packet length; shorter packets are sent when conditions are poor. If a packet is irreparably damaged during transmission, less time was wasted sending the original packet, and less time is needed to resend the packet. However, the overhead remains the same, resulting in decreased efficiency; more time is spent sending less data. X.25's benefits were its standardization and its proven operation.

Coming from the area of radio teletype (RTTY), AMTOR was another protocol employed to transmit data digitally. As with RTTY, AMTOR was designed to transmit text messages in a broadcast mode; there was no specific destination or receiving site, no handshaking or acknowledgment involved, and minimal error correction. AMTOR uses the 5-bit BAUDOT character set (capital letters, numbers and a basic set of punctuation), limiting the type of information that can be sent. Due to the low level of error correction, degraded channel conditions result in corrupted messages; it is up to the receiving party to interpret the erroneous text. AMTOR uses a binary frequency-shift keying (BFSK) scheme. Due to its high overhead, AMTOR's throughput is limited to a maximum of 6 or 7 characters per second. Finally, AMTOR's architecture is fixed; there was no way to adapt to changing channel conditions.

Subsequent protocols were developed, better suited to an HF radio environment. PACTOR[3] was developed by German Amateur-Radio operators to overcome the limitations imposed by the HF medium. The improvements include dual packet lengths, longer transmission-acknowledgment cycle times (1.25 seconds as opposed to 0.45 second for AMTOR), an enhanced acknowledgment packet, implementation of a 16-bit cyclical redundancy check (CRC), on-line data compression and the introduction of memory ARQ. The 16-bit CRC lessens the possibility of erroneous data being received undetected, while the on-line data compression (Huffman coding, using variable length characters) increases the efficiency of the data transmission; fewer bits are required to send the same amount of data. The packet length can be chosen manually or requested by the receiving site through the acknowledgment packet. Like AMTOR, PACTOR uses a BFSK modulation scheme.

Memory ARQ is a method of recreating an error-free packet from multiple receptions of erroneous packets. Erroneous packets occur when a degraded channel corrupts the signal to the point that the thresholds between the states of modulation are indistinct. In this case, analog representations of successive retransmissions are summed together (using the synchronization information contained in the header to aid in alignment), resulting in a signal that can be demodulated successfully.

Since the introduction of PACTOR, other protocols have been developed to improve on it. They share common goals of increased throughput, improved error detection and correction and enhanced functionality, without increasing the bandwidth usage. ITS tested three of these next-generation protocols. Two of these, CLOVER and G-TOR, are proprietary designs, developed and implemented by their respective manufacturers. The third is an improved version of PACTOR, appropriately named PACTOR2.

G-TOR[4] improves on the PACTOR protocol with minimal increases in processing power necessary to implement it. It adds Golay forward error-correction coding to improve error detection and allow for some level of error correction without resorting to data retransmission. It adds a second type of Huffman coding, optimized for upper-case-only text transmission, as well as run-length coding, which is effective when repeated characters are sent. It implements full-frame data interleaving (transmitting data and parity frames alternately) to improve the likelihood of error recovery. The cycle time was lengthened to 2.4 seconds. Instead of varying the size of the packet to contend with variable channel conditions, the baud rate varies between one of three rates: 100, 200 or 300 baud. Like PACTOR, the selection may be made manually or by request from the receiving site. Like AMTOR and PACTOR, G-TOR uses BFSK as its modulation scheme.

Taking advantage of the improvement in affordable processing power available, both CLOVER and PACTOR2 have implemented more complex techniques to improve throughput and contend with degraded channels. This includes adaptive modulation techniques, employing multiple modulation schemes; more efficient compression schemes; adaptive cycle lengths; and adjustable header size.

CLOVER[5] began as Cloverleaf and got its name from its observed signal pattern; it is shaped like a four-leaf clover due to its use of *m*-ary phase shift keying. CLOVER uses 10 modulations, varying between simple bipolar phase-shift keying (BPSK) to a combination 16-level phase, 4-level amplitude shift keying (16P4A). In addition, CLOVER transmits four simultaneous pulses, 125 Hz apart, each independently modulated. This allows for 256 different possible states; the 8 bits of a byte can be transmitted simultaneously. CLOVER also implements Reed-Solomon block coding and allows for four different block lengths and four coding rates (the ratio of redundant, error-correcting bits to total bits). The modulation selection is chosen adaptively by the receiving site, based on analysis of the received signal. The block lengths and coding rate are chosen manually at the transmitter site, based on the perceived conditions and the priority of the data being sent. CLOVER does not implement any on-line data compression but depends on software-implemented compression algorithms when transmitting large data files.

PACTOR2[6] implements a more efficient compression scheme known as Markov coding, in addition to both versions of Huffman coding (for mixed- and upper-case text). The protocol automatically applies the most efficient technique depending on the entire content of the packet. Four steps of differential phase-shift keying are implemented (from 2-level to 16-level). In addition, a convolutional coding scheme is implemented as opposed to a block coding scheme such as Golay, for better error correction. The coding rate can be varied; fewer redundant bits are used for more robust modulation levels. The modulation levels (and their related coding rates) are chosen automatically, depending on the number of successive erroneous or error-free packets received on the other end of the link.

One benefit of these improved protocols is that, as throughput has increased, the occupied bandwidth of their signal has decreased. Starting with 2 kHz for X.25 (taking advantage of the 3kHz or greater bandwidth of voice quality telephone lines), the bandwidths have decreased to 1 kHz for AMTOR and PACTOR, 500 Hz for G-TOR and CLOVER and 450 Hz for PACTOR2.[7] In addition to occupying less space in the HF spectrum, the signal becomes less susceptible to noise and adjacent-channel interference.

2.2 Comparison of ARQ and FEC Modes

All of the later protocols can transfer data using one of two modes: automatic repeat request (ARQ) and forward error correction (FEC) or broadcast mode. ARQ is an addressable mode of data exchange implementing full handshaking, receipt acknowledgment and retransmission-request capabilities. Data is sent to a specific site, identified by a unique identifying address. Once a connection is established, data is transmitted from sender to receiver. The receiver either returns an acknowledgment when data is received error-free, or a request for retransmission when data is received corrupted beyond the ability to correct errors. A retransmission request prompts the transmitting site to resend the erroneous data. Theoretically, ARQ is an error-free method although, under severely degraded conditions, it could take an infinitely long time to transmit error-free data. While the use of error-detection and correction coding is not necessary to ensure error-free transmissions, its use reduces the need for retransmissions. The addressable nature of ARQ minimizes the amount of unnecessary information a receiving site must contend with and assures the transmitting site that the data was received correctly. ARQ is best suited for critical text messages and binary files where errors cannot be tolerated.

FEC, or broadcast mode, operates much like a commercial broadcast station; the transmit site broadcasts data to whoever may be listening. No addressing is involved, except to identify the transmitting site. It depends on error-detection and correction schemes to minimize errors during transmission and will usually transmit the information multiple times; but there is no mechanism to guarantee error-free transmission. FEC is used for broadcasting text messages; any errors hopefully can be resolved manually by the receiving party.

3. HF CHANNEL CHARACTERISTICS

Before laboratory simulation and testing can take place, we need to have a clear understanding of the characteristics of the entire communications system. A number of attempts have been made to characterize and model the atmospheric channel. One model in particular has gained wide enough acceptance to be used as a standard, implemented in commercially available hardware simulators.

The parameters of a model can be varied to extremes, however, and the results will no longer resemble real-world conditions. There must be limits to the parameters so that the model/simulator will correlate to actual observed channel conditions.

3.1 The Watterson Narrowband HF Channel Model

When radio communication uses the HF band (3-30 MHz), the path the radio wave travels usually involves reflections off the ionospheric layer.[8] Ionospheric reflections do not come into play during communication over ground-wave paths, or when ducting (where the wave is guided within a narrow layer of the atmosphere bounded by layers with significantly different refractivity) occurs. Except for these limited cases, users of the HF band must deal with ionospheric reflections and the effects caused by both short- and long-term variations in the ionospheric layer's height, thickness, refractivity and turbulence. These variations result in the effects known as multipath, frequency (Doppler) shifting, component fading, flat fading and frequency spread. All of these effects cause degradation of the signal and system performance.

These degrading effects vary depending on the location of the path, the time of day, the time of year and the time of multiyear periods such as the sunspot cycle. Consequently, testing and comparing radio systems over real paths is subject to these uncontrollable variations. Systems must be tested simultaneously to compare performance, with the hope that the systems do not interfere with one another. Even simultaneous testing does not guarantee that different systems will experience the same conditions. In addition, unless the systems are tested over periods significantly longer than the longest variation cycle, there is little chance that the systems will experience the full range of ionospheric effects. When testing is not performed simultaneously (which is generally the case), there is little likelihood that the results will correlate with any degree of certainty.

Because of these drawbacks to over-the-air testing, there has been significant effort to mathematically model HF-channel characteristics and implement a model in a laboratory-based HF-channel simulator. The use of such a simulator allows for absolute control and repeatability of channel effects and permits testing and comparison of communication systems without the variability and uncertainty induced by real atmospheric channels.

One of the most commonly accepted and implemented channel models is the Gaussian-scatter (or Watterson) model, developed by Clark Watterson, formerly of ITS.[9] A block diagram of the Watterson model is shown in Fig 2. It consists of an ideal delay-line that generates multiple, time-delayed paths. Each path is modulated in amplitude and phase by a complex, random and independent tap-gain function $(G_i(t))$. The modulated paths are summed together, along with additive noise and interference, to form the output signal.

A radio signal can be viewed as the sum of two magneto-ionic components, generally referred to as the I and Q components. These two components are 90° out of phase with one another and tend to be affected differently by the atmospheric channel effects. The Watterson model defines the tap-gain function as:

$$G_i(t) = \hat{G}_{ia}(t)e^{(j2\pi v_{ia}t)} + \hat{G}_{ib}(t)e^{(j2\pi v_{ib}t)} \quad \text{(Eq 1)}$$

where the a and b subscripts refer to the tap-gain components that operate on the I and Q components independently. $\hat{G}_{ia}(t)$ and $\hat{G}_{ib}(t)$ are sample functions of two independent, complex Gaussian random processes with zero-mean values and independent real and imaginary components with equal RMS values that produce Rayleigh fading. v_{ia} and v_{ib} are the Doppler frequency shifts.

Each tap-gain function has a spectrum that, like the function, consists of two components that are Gaussian functions of frequency:

$$v_i(v) = \frac{1}{\tilde{A}_{ia}\sqrt{2\pi}\,\sigma_{ia}} e^{\left[\frac{-(v-v_{ia})^2}{2\sigma_{ia}^2}\right]} + \frac{1}{\tilde{A}_{ib}\sqrt{2\pi}\,\sigma_{ib}} e^{\left[\frac{-(v-v_{ib})^2}{2\sigma_{ib}^2}\right]} \quad \text{(Eq 2)}$$

In this equation, $\tilde{A}_{ia}$ and $\tilde{A}_{ib}$ are the attenuations, $2\sigma_{ia}$ and $2\sigma_{ib}$ are the frequency spreads, and v_{ia} and v_{ib} are the frequency shifts. These values are illustrated in the plot of the tap-gain function spectrum shown in Fig 3. When this channel model is implemented in a hardware simulator, these values are the variables controlled by the simulator.

To validate the Watterson channel model, measurements of an actual 1294 km midlatitude HF ionospheric channel were made at various times of day.[10] The measurements verified the accuracy of the model and uncovered one limitation of the model. The model's accuracy decreases with increasing bandwidth. This is due to the effect of time spread, where the time delay varies with respect to frequency. This causes one end of a signal's bandwidth to be delayed differently than the other end. Consequently, wider bandwidths increase the effect of time spread on the signal. The measurements indicated the model was accurate for:

$$BW \approx \frac{1}{4(2\rho_e)} \quad \text{(Eq 3)}$$

where $2\rho_e$ is the effective time spread on the ionospheric modes. The bandwidth limits are approximately 2.5 kHz for nighttime and 8 to 12 kHz for daytime. This is the reason the Watterson model is also known as the narrowband channel model. Since the bandwidths of the modem/protocols we tested are well under 1 kHz, the Watterson model, and any simulator implementing it, is sufficiently accurate.

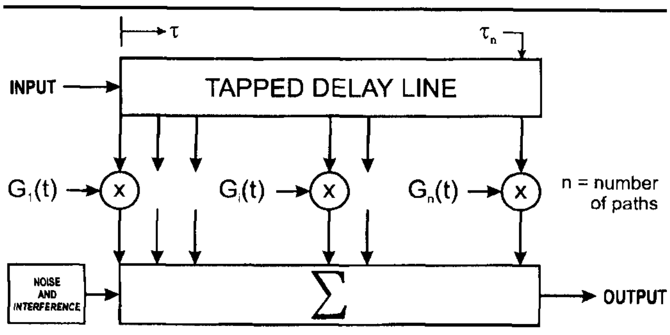

Fig 2—A block diagram of the Watterson HF channel model.

3.2 CCIR-Defined Channels and their Relationship to Real Conditions

The International Telecommunications Union Radiocommunications Sector (ITU-R, formerly known as the International Radio Consultative Committee, or CCIR), an organization that issues recommendations regarding radio communications, recognized that laboratory testing had definite time and cost benefits over field testing. A number of recommendations cover fixed services at frequencies below 30 MHz. Recommendation 520-2,[11] lists a number of transmission-channel parameter combinations for tests of HF components and systems—both to predict expected performance in the field and to compare different systems. For the purposes of qualitative testing, the ITU-R lists three parameter combinations:

Gaussian noise and flat fading (single path)

Gaussian noise, multipath and fading (two paths, no frequency shifting)

Doppler, multipath and fading (two paths with frequency shifting)

For testing the HF modems/protocols, ITS chose the second set of parameter combinations: Gaussian noise, two-path multipath, no frequency shifting and flat fading (attenuation). For this combination, the ITU-R suggests four sets of parameter values, representing different path conditions: good, moderate and poor conditions and flutter fading. The parameter values for each set is shown in Table 1.

In Report 203-1,[12] the ITU-R lists typical interpath time delays caused by multipath on HF channels, based on actual channel measurements. Given that channel conditions vary due to daily and seasonal cycles, the worst-case scenario of nighttime during the winter was chosen. There is an even greater variation due to sunspot activity. Table 2 shows the effective path length for each channel condition differential time delay (DTD), for a worst-case situation, during minimum and maximum sunspot activity.

While noise, fading and delay vary over time, the ITU-R suggested conditions do not cover time fluctuations; consequently, the fading (attenuation) and delay are fixed during testing. The noise is changed in steps during a test set, but remains fixed during a single test (the transfer of a single file).

4. ITS TEST DESCRIPTION

When attempting to compare different systems, the test procedure should be as consistent and repeatable as possible. An automatic test bed was developed by ITS to test all modem/protocol sets under conditions as identical as

Table 1—ITU-R Suggested Channel Parameter Combinations

Condition	*Differential Time Delay (ms)*	*Frequency Spread (Hz)*
Good	0.5	0.1
Moderate	1.0	0.5
Poor	2.0	1.0
Flutter Fading	0.5	10.0

possible, even though the characteristics of the modems/protocols may differ significantly.

4.1 ITS Test Lab Equipment, Computer and Interfaces

Four pairs of modems, implementing one or more of five protocols, were placed in a symmetrical, back-to-back configuration (Fig 4). Throughput was measured under various simulated HF channel conditions. The channel conditions were produced by two HF channel simulators. A serial protocol analyzer monitored the received data and was used to detect a link loss. All equipment, including the modems under test, were controlled by a 33-MHz 80486 computer with 16 Mbyte of memory, an eight-port serial interface and an IEEE-488 (GPIB) interface, running Microsoft *Windows* and National Instruments' *LabView* program development application.

Given the relatively low throughput of these protocols (which are ultimately constrained by bandwidth limitations), their operational demands on the test bed controller were slight. To assure that the test setup did not artificially affect the results of the tests, manual measurements (hand timing) were made for each modem/protocol setup and compared to results produced by the automated setup.

4.1.1 ITS Test Lab Software:

An Intel-processor-based PC, running Microsoft *Windows*, was used as the controller because that system common at ITS. It is also the combination most familiar to those of us involved in the testing and was readily available from existing in-house equipment when the test bed was being developed. *Windows* gave us the ability of multitasking, which was of great use during test-bed development and configuration of the individual protocols.

LabView was chosen as the test-bed controlling software due to its suitability for the task and our prior experience with it. *LabView* is an object- or module-based program-development application.[13] It can be considered a graphical equivalent to text-based programming languages such as *C*. *LabView*'s modules, or virtual instruments (VIs), are the equivalent to subroutines. Low-level VIs can be used to develop more complex VIs. The program code consists of a block diagram (Fig 5) comprised of VIs interconnected by wires. These wires represent different data types (integers, floating-point numbers, strings and arrays) that are passed into and out of various ports of the VIs, much like parameters are passed to and from subroutines. The user interface to a *LabView* program is called a front panel. The front panel for the HF modem test bed is shown in Fig 6. In addition to the front panel for the main program, each VI at any level may have a front panel (Fig 7). Each item on the front panel corresponds to a subVI, displays a data type or indicator, or operates a control/input to the VI.

LabView comes with a wide range of VIs ranging from simple functions such as Boolean and arithmetic operators, file I/O and type conversions, to high-level, complex VIs such as virtual oscilloscopes, complex string and array operators, or instrument interfaces. Most high-level VIs are built from more basic units and can be customized to fit the user's needs. Along with a VI's front panel, *LabView* has built-in debugging tools such as single-step, trace and breakpoint capabilities. One benefit of using *LabView* is its support from most, if not all, of the major instrument manufacturers. Manufacturers will supply, either directly or through National Instruments, the necessary VIs needed for controlling and communicating with their various instruments, interfaces and equipment. This eliminates unnecessarily development work and speeds our development of the test bed.

4.1.2 Hardware

The channel simulators (Fig 8) are based on the Watterson narrow-band HF channel model and have been in use by ITS for some time.[14] Recently, a newer channel simulator has been developed. It is based on a commercially available (DSP) card that can be plugged into a PC's ISA bus. This

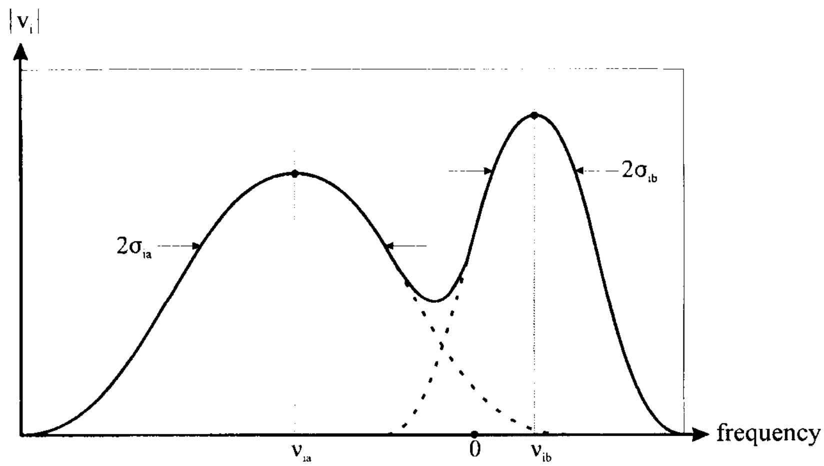

Fig 3—A plot of the tap-gain function spectrum.

Table 2—Effective Path Length for Various ITU-R Suggested Channel Conditions

Condition	*Maximum Sunspot Activity*	*Minimum Sunspot Activity*
(DTD)	*Lower limit ionospheric height = 300 km*	*Lower limit ionospheric height = 220 km*
	Ionospheric layer half-thickness = 120 km	*Ionospheric layer half-thickness = 80 km*
Good (0.5 ms)	>5000 km	>3000 km
Moderate (1.0 ms)	> 5000 km	~2200 km
Poor (2.0 ms)	~2000 km	~1000 km
Flutter Fading (0.5 ms)	>5000 km	>3000 km

simulator has been approved for use by NATO. As a result, future testing by ITS will use this approved simulator. This means that the results of previous testing with the old simulator will have to be correlated with the results produced by the new simulator. That work is currently being done at ITS.

The simulators used in the first set of modem/protocol tests are capable of producing two independent paths with the ranges of parameters shown in Table 3, covering the ITU-R good and poor conditions with which we wish to test. The simulators' settings can be manually set with front-panel controls or through an IEEE-488 (GPIB) interface by the test-bed controller. Two identical simulators were used, one for each direction between the two modems. Both simulators were adjusted to identical settings for each test.

While some modems can report the loss of link to their attendant computer or terminal, the method used varies between implementations, while some modems do not have this capability at all. At the protocol level, this capability is not covered. Therefore, to determine when the link is lost, the controller needs to periodically poll a timer that is reset by any output from the modem. This increases the possibility of biasing the actual throughput of the modem/protocol, unless care is taken in programming this function. We chose to off-load this function to an external instrument. An RS-232 serial interface protocol analyzer was used to detect a lack of data output from the modem (link loss). It was programmed so that a timer was reset every time a character was received from the modem. When no characters were received during the timed interval, the analyzer sent a signal to the test controller through a separate serial port. This interrupted the *LabView* program, aborted the current test and allowed the program to continue on to the next test. The timer was set to five minutes; this allowed the protocols to function under extreme channel conditions, while limiting the test time when throughput dropped below usable levels.

All modems (except one) and the protocol analyzer were connected to the

Table 3—Channel Simulator Specifications

Parameter	*Value/Range*
Minimum Delay (each path)	2.2 msec
Multipath Delay Spread	0 to 9.3 msec in 0.1 msec steps
Fading Bandwidth (each path)	0.1, 0.2, 0.5, 1.0, 2.0, 5.0, 10.0 Hz
Fading Depth	0 to -60 dB
Signal Bandwidth	350 to 3050 Hz
Noise Bandwidth	350 to 5000 Hz
SNR	+50 to -40 dB in 1 dB steps
Differential Group Delay	500 msec

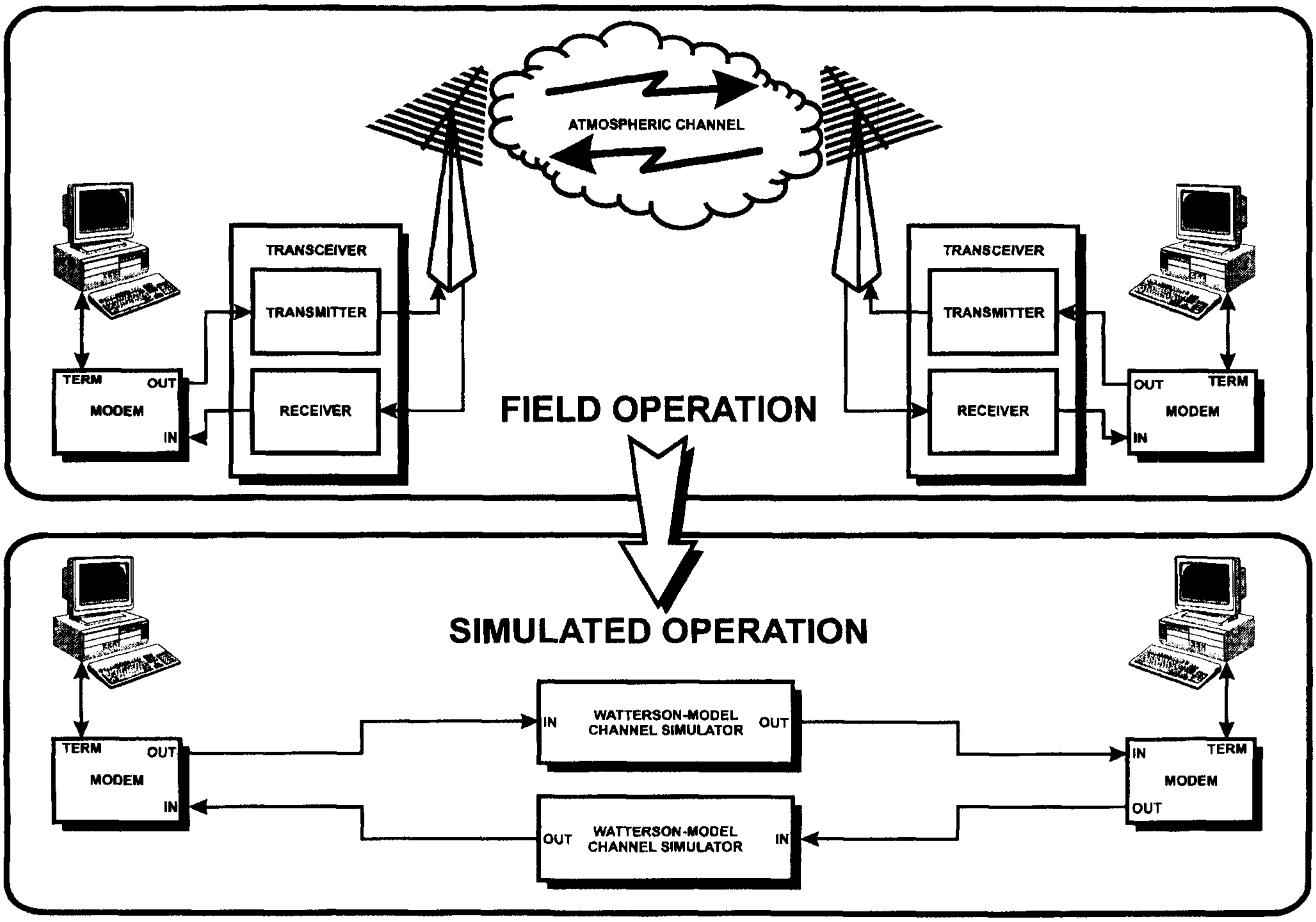

Fig 4—A comparison of the field- and lab-test systems.

controller through an eight-port RS-232 interface. The interface came with its own set of VIs, which eliminated the need to develop a program interface. This setup allowed us to keep all modems connected to the controller, which reduced the development and reconfiguration times between modem/protocol tests. The exception was the modem implementing the CLOVER protocol. It consisted of a card that plugged into the PC's ISA bus. Rather than use the manufacturer's associated software package, we wrote VIs that allow *LabView* to control and communicate with the modem directly and test all modems on an equal basis. Control of and communication with all modems were accomplished through *LabView*; since the goal of the tests was to test the protocol independent of the modems, no vendor's software was used. In any case, none was compatible with our automated testing.

4.2 Test Procedures

An attempt was made to create a universal test bed, capable of interfacing with and testing any modem/protocol pair with a minimum of customization work (hardware or software). Given the differences between the various modem interfaces, operations and protocol specifics, preparing the test bed for a different modem turned out to be the most effort consuming portion of the testing. Interfaces, commands and responses were different among all the modems. In one instance, two modems interpreted the specifications for the AMTOR protocol differently (one implemented the basic BAUDOT character set, while the other implemented an enhanced BAUDOT character set).

This is where the modularity of the software was beneficial. Basic modules for file access, reading, writing, error detection, throughput calculation and protocol analyzer initialization and interrupt handling could be used for all modems. Separate modules for modem and protocol initialization and operation were written for each modem/protocol.

4.2.1 LabView Modules:

Fig 9 displays the overall structure of the controlling program as implemented in *LabView*. Each block indicates a separate module, or VI, while the lines indicate dependencies between VIs and subVIs. Most of the VIs shown are either standard functions (open/read/write a file, build/search/modify an array) that come with *LabView* or are supplied by various hardware manufacturers. (Eg, the manufacturer of the eight-port serial interface supplied the serial port read/write/data-present VIs.) The higher-level VIs were built from these basic blocks and perform the actual testing

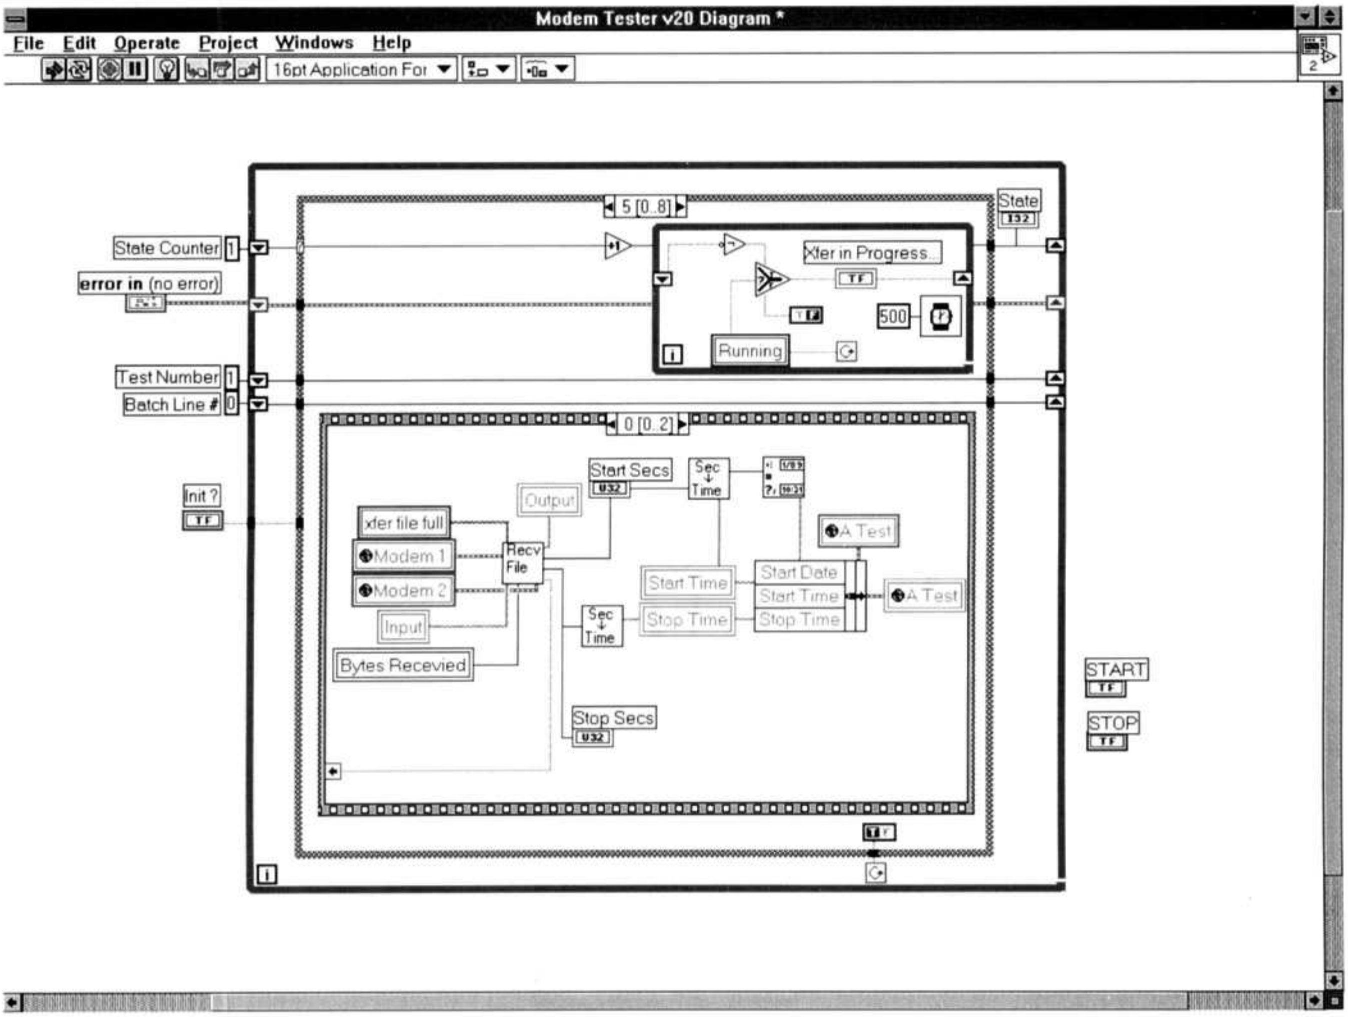

Fig 5—*LabView*'s program code consists of a block diagram comprised of Virtual Instuments (VIs) interconnected by wires that represent different data types (integers, floating-point numbers, strings and arrays) that are passed into and out of various ports of the VIs.

functions. They are:

- Get next test information from test batch file
- Initialize channel simulators
- Initialize protocol analyzer
- Set modem parameters from parameter array
- Connect modems
- Write file to modem
- Receive file from modem
- Check file for errors
- Write test results to master log file

Information in a test batch file included the type of modem/protocol to be tested, the modem/protocol settings to be used during testing, and a series of channel conditions, one for each test (file transfer). For each modem/protocol tested, an array containing the various parameters and settings unique to that particular combination had to be constructed. Based on the test batch file's contents, the appropriate array is chosen and passed to the modem initialization VI, while the channel conditions are passed to the channel simulator initialization VI. Following the completion of the file transfer, the received file is checked for errors, the throughput is calculated and the results are appended to the end of the master log file.

Because one of the modems plugged into the PC's ISA bus instead of connecting through a serial port, the software turned out to require more customization than we had planned. Another set of VIs had to be written specifically for this modem, including a unique initialization routine, get/put word and in/out port routines. As with the other modem/protocols, test runs were performed and compared to manual timings to assure that the program was operating properly and was reporting correct throughput values.

4.2.2 Error and Exception Handling:

By measuring throughput only, the test procedure was greatly simplified. The hardest part in error detection regarding complex data is determining what constitutes a single error. Error detection not only involves identifying the location of the start of the error incident, but the end of the error as well. This involves resynching the incoming data stream with known good data. The process increases in complexity as the data stream contents become more complex. A pseudorandom text file will contain short duplicated patterns: ie, certain words are used many times during typical communications. Resynching after an error of unknown length creates the possibility of synching to the wrong point in the data, resulting in false errors. A truly random, nonrepeating text message could be generated and used, but this defeats the objective of using a typical, or actual, message during testing. Error detection under these conditions is complex and time consuming, although the conditions of

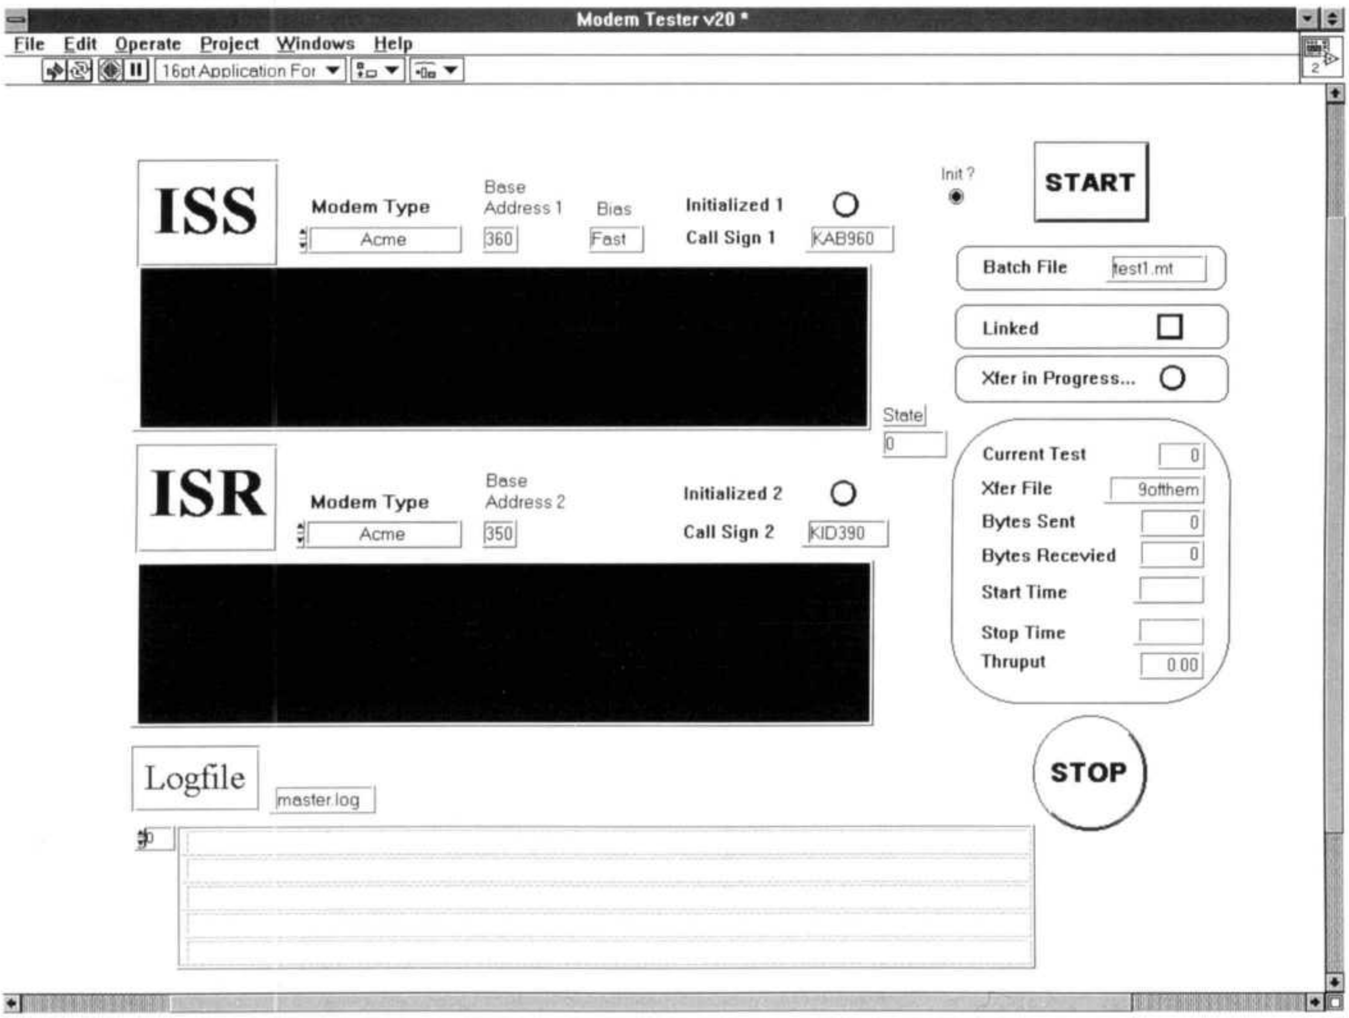

Fig 6—The front panel for the HF modem test bed.

our tests would allow detection and measuring to be done off-line, following the complete transmittal of a file.

Without the overhead of error detection and measurement, the test program needs to account for four possible problems/errors, as shown in Table 4.

Given the amount of testing to be done, tests generally were not rerun when aborted. Three tests, at the most (except during development), were performed at each simulator setting. The tests were scheduled such that the channel conditions proceeded from higher to lower SNRs. Overall, a modem's ability to establish a connection decreased as the SNR decreased. While the inability to establish a connection could be blamed on excessively degraded conditions, we could not end testing based on a modem's inability to establish connection at a particular SNR. Some modems had difficulty establishing connection at one SNR, but could connect at a lower SNR.

4.3 Test Results Analysis and Reporting

Each test had three possible outcomes:

- Successful—entire file was transferred error-free.
- Erroneous—entire file was transferred but contained errors.
- Aborted—link was lost before the entire file was transferred and could not be re-established.

Only the results of the first instance

Table 4—Error conditions handled by the test controlling program

Problem	*Symptoms*	*Resultant action(s)*
Unable to initialize modem under test	Return of unexpected response string	• Retry initialization routine • Abort test
Unable to establish connection between modems under test	Modems do not return the proper connect response string	• Retry connection routine • Retry initialization routine • Go to next test
Loss of connection during data transfer	Throughput timeout (detected by the protocol analyzer)	• End test
		• Go to next test
Receipt of erroneous data	Received data does not match original	• Mark test as erroneous • Go to next test

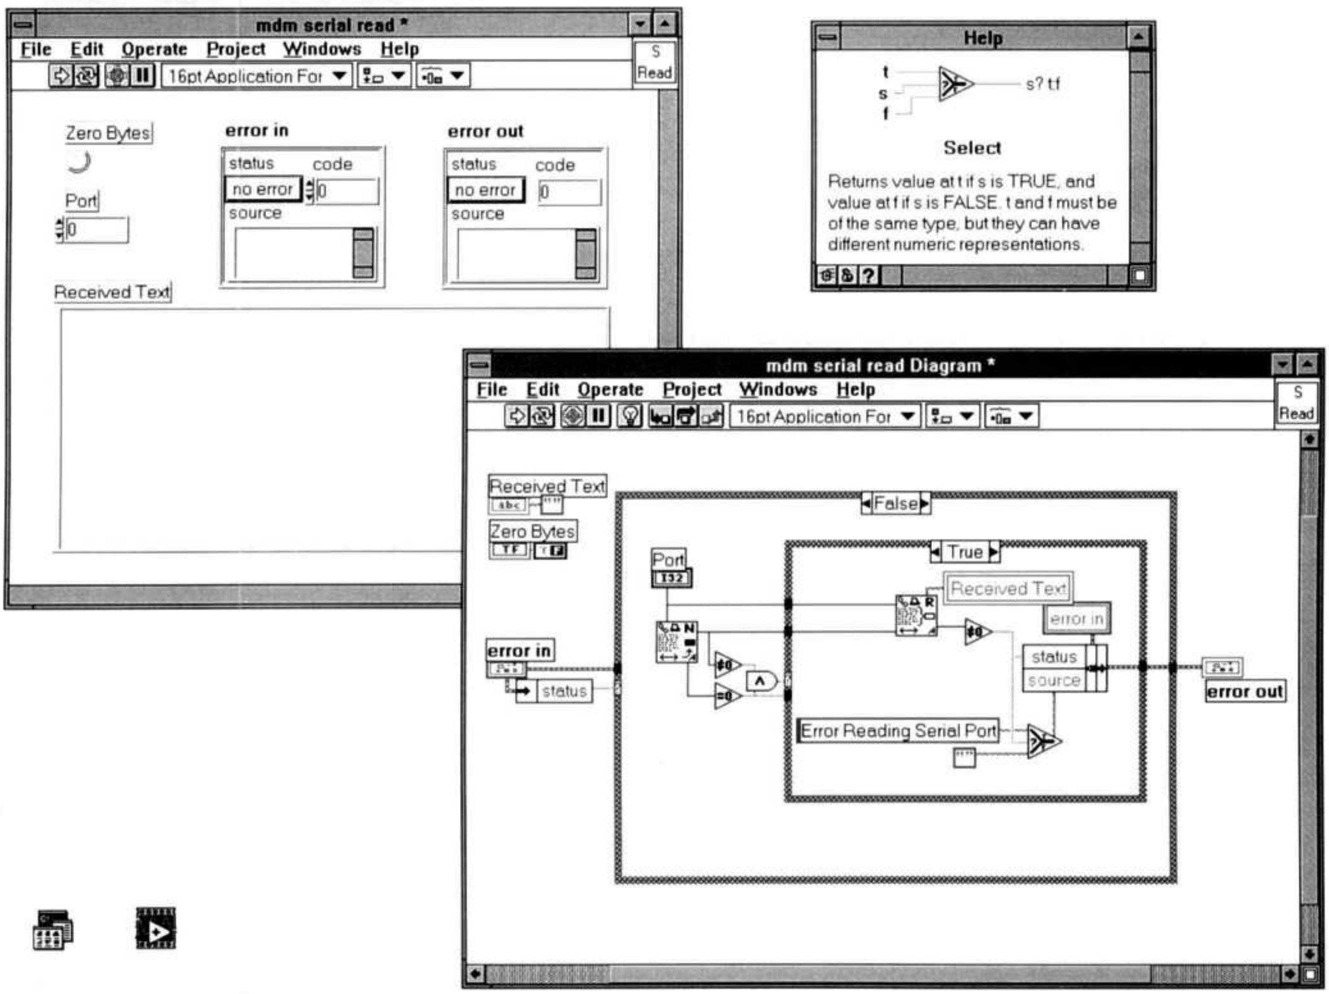

Fig 7—In addition to the front panel for the main program, each VI may have a front panel.

were included in the final report. Due to time limitations, no attempt was made to rerun erroneous or aborted tests. There was no analysis done on the type or cause of errors contained in erroneous files. While the cause of aborted tests is understandable, the existence of errors in files transmitted in ARQ mode is disturbing. Errors could corrupt the meaning of critical text files, while binary files would become useless at best, and potentially damaging at worst. Rather than needing to retransmit a single packet, the entire file would have to be resent. This situation needs to be studied further.

The data was plotted as SNR versus throughput for each modem and each CCIR condition (good or poor). Because the basic premise of testing had been simplified, the data presentation was relatively straightforward.

5. CONCLUSIONS

As with any project, we discovered better and more efficient ways to perform the tests as time progressed. Since the first tests were conducted, more advanced channel simulators have entered the market; software has been upgraded and, hopefully, improved. The test bed now uses a channel simulator based on advanced DSP boards. *LabView* version 3.1 has been replaced with version 4.0. Some of the VIs have been or will be rewritten as our experience with the product grows. *Windows* 3.1 has been replaced with *Windows NT*. As testing continues, we must confirm that these changes do not affect the basic operation of the tests and that the results for a given modem/protocol remain the same regardless of test-bed changes. This verification will be performed before any further testing is conducted.

Plans for the future include testing the same protocols implemented in different modems, newer versions of modems/protocols, as well as investigating the interoperability of different modems implementing the same protocol. This will allow FEMA to recommend the use of a protocol without requiring a specific manufacturer's hardware implementation.

6. REFERENCES

[1]Riley, Tim, Stephen Rieman, Teresa G. Sparkman and Dennis Bodson, "Comparison of HF Digital Protocols," *QST*, Jul 1996, pp 32-39.

[2]Dijak, J. T., "AMTOR, AX.25, and HERMES: a performance analysis of three systems," *Ham Radio*, Dec 1985.

[3]Helfert, Hans-Peter and Strate, Ulrich, "PACTOR—Radioteletype with Memory ARQ and Data Compression," *QEX*, Oct 1991.

[4]Prescott, Glenn, Phil Anderson, Mike Huslig and Karl Medcalf, "G-TOR: A Hybrid ARQ Protocol for Narrow Bandwidth HF Data Communications," *QEX*, May 1994.

[5]HAL Communications Corp, "PCI-4000 CLOVER-II HF Radio Modem Operator and Reference Manual," Nov 1992.

[6]Special Communications Systems GmbH, *SCS PTC-II, Manual for Version 1.1*. 1995.

[7]Anderson, Phil, "HF ARQ Protocols," *RTTY Digital Journal*, Apr 1994.

[8]CCIR Report 549-3, *HF Ionospheric Channel Simulators*, ITU, Geneva (1990).

[9]Watterson, C. C., G. G. Ax, L. J. Demmer and C. H. Johnson, "An Ionospheric Channel Simulator," *US Dept of Commerce, Technical Memorandum ERLTM-ITS 198*, Sep 1969.

[10]Watterson, C. C., J. R. Juroshek and W. D. Bensema, "Experimental Confirmation of an HF Channel Model," *IEEE Transactions on Communication Technology*, Vol COM-18, No. 6, Dec 1970, pp 792-803.

[11]CCIR Recommendation 520-2, *Use of High Frequency Ionospheric Channel Simulators*, ITU, Geneva (1994).

[12]CCIR Report 203-1, *Path Time-Delays and Shifts Caused by Multipath Propagation on Radio Circuits*, ITU, Geneva (1986).

[13]National Instruments Corporation, *LabView User Manual*, Jan 1996.

[14]Harris Corporation, *RF-3460 HF Simulator Instruction Manual*, Apr 1990.

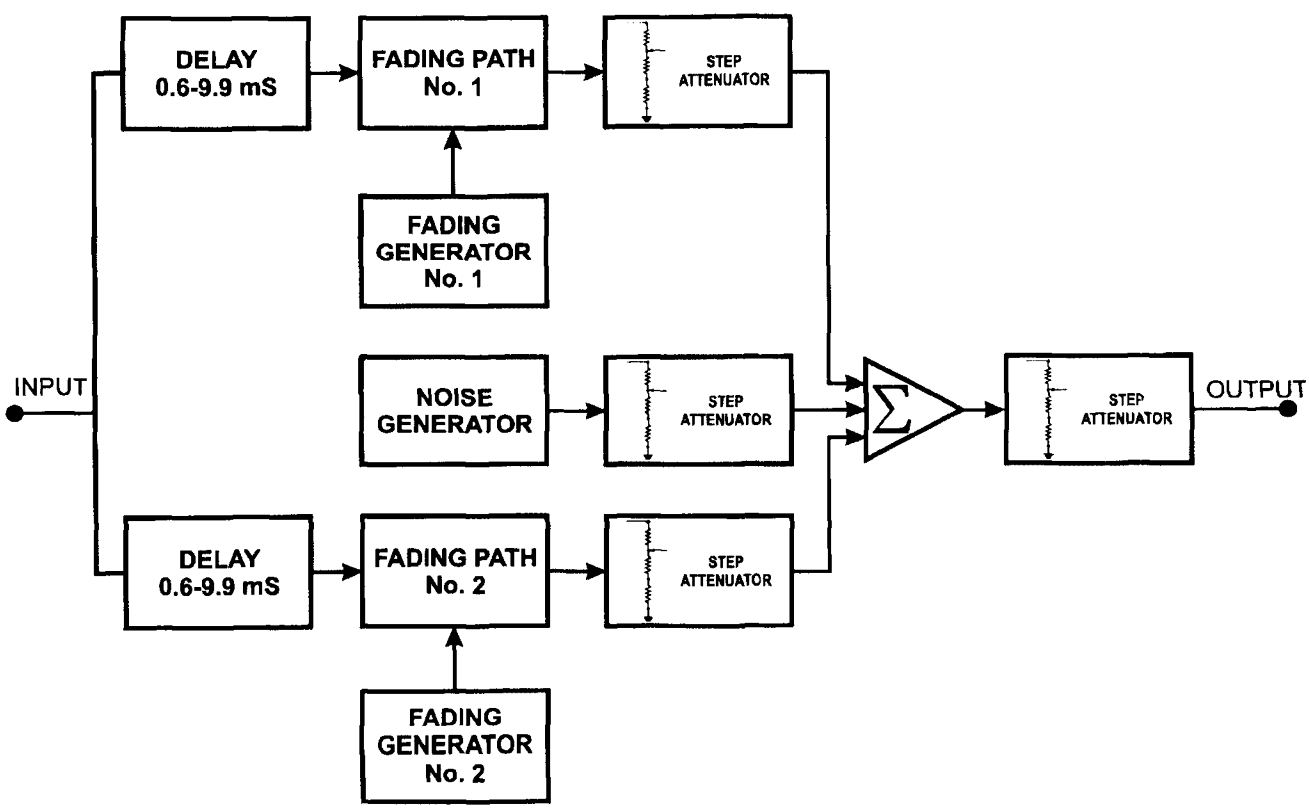

Fig 8—A channel-simulator block diagram.

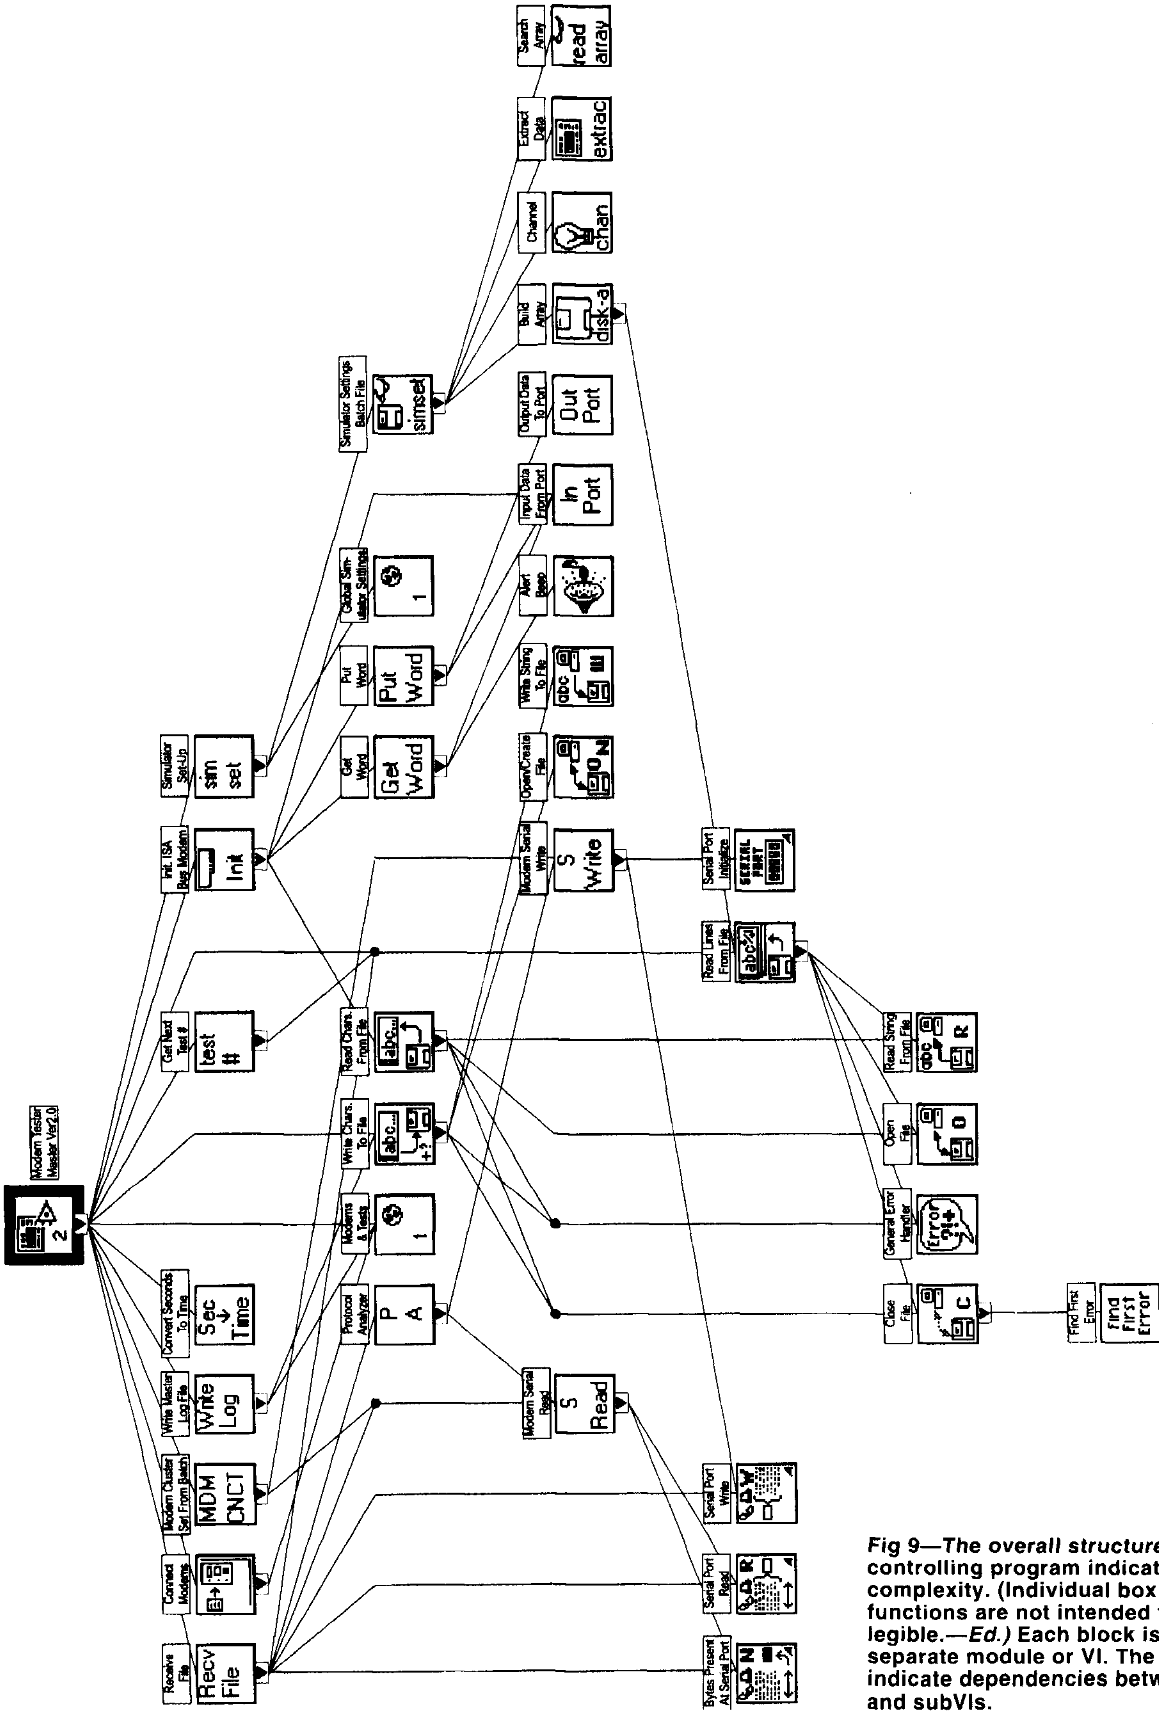

Fig 9—The overall structure of the controlling program indicates its complexity. **(Individual box functions are not intended to be legible.—*Ed.*) Each block is a separate module or VI. The lines indicate dependencies between VIs and subVIs.**

Power and Protection for Modern Tetrodes

A top-notch screen-bias supply can yield excellent linearity, prolong tube life and protect RF and power-supply circuits against mishaps. Here are some ideas from across the pond.

By Ian White, G3SEK

Tetrode power amplifiers are coming back into fashion, after many years in which US amplifier builders have focused almost exclusively on triodes. The situation in Europe has been very different, because tetrodes never went away. In particular, there have been many developments in tetrode power-supply designs that US amateurs have missed—a situation that this article aims to rectify.

There's a great temptation to think of a tetrode as "a triode with an extra grid," and to treat the screen-grid supply as a minimal afterthought. That's a big mistake! The screen grid of a large transmitting tetrode has very specific needs, and if these are met the tube will reward you with excellent linearity on SSB. A screen-current meter will show you whether the tube is tuned and loaded correctly, and a power supply that continuously monitors the screen current can protect the whole amplifier from a wide range of faults.

52 Abingdon Rd, Drayton
Abingdon OX14 4HP
England
e-mail **g3sek@ifwtech.demon.co.uk**

This article describes a modern stabilized screen-grid power supply that provides adjustable voltage and excellent dynamic regulation. The supply also includes very effective circuits to protect the tube and the rest of the amplifier. Although many of these circuit ideas may be new to you, they have been widely used for several years in Britain and the rest of Europe. With some adaptation to meet different requirements for screen voltage and current, these ideas can be used as a 'drop-in' upgrade for almost any existing tetrode PA.

I will begin by explaining why it's a good idea to stabilize the screen voltage to a much higher standard than has been regarded as normal (at least in the USA). It isn't difficult, and there are several good reasons for doing it.

DC Stability

The most basic reason for stabilizing the screen voltage is to achieve dc operating stability for the tube. The screen current in many tetrodes can be either positive or negative, in both normal and fault conditions, and this creates unusual requirements for the screen power supply. In normal operation, some of the electrons flowing from cathode to anode inside the tube are intercepted by the screen grid and flow outward to the screen supply; this is observed as a positive screen current, flowing into the tube (Fig 1a). However, the electron beam striking the screen grid will also result in secondary emission of electrons from the surface of the grid, especially when the anode voltage is swinging very high because the tube is lightly loaded. Electrons leaving the screen grid and joining the main cathode-anode flow will be observed as a negative screen current, coming out of the tube (Fig 1b). Here's where trouble can start, because the reverse current is dumped into the screen supply and tends to drive the voltage upward. The higher screen voltage leads to more secondary electron emission, which in turn leads to even higher voltage, and a runaway situation that can end in serious damage. That is why a screen supply **must always** have the capability to absorb negative screen current without allowing the voltage to increase appreciably. The old-fashioned way to do this was by bleeding a generous standing current to ground[1] through a resistor (Fig 2a), so that the bleed current swamps any changes in voltage caused by the screen current. This circuit can absorb negative screen current but it has no voltage regulation at all. The next step forward was to stabilize the screen voltage using VR tubes, later replaced by Zener diodes (Fig 2b), and that's as far as most tetrode power supplies have ever progressed.

[1]Notes appear on page 26.

At this point it's interesting to take a sideways look at the radically different power supply arrangement of Fig 3. This circuit was used by the Collins company with great success, in the 30S-1 amplifier for the amateur market and also in commercial linear amplifiers. None of the voltages is formally regulated at all, but there are two separate high-current supplies, one for the anode and another for the

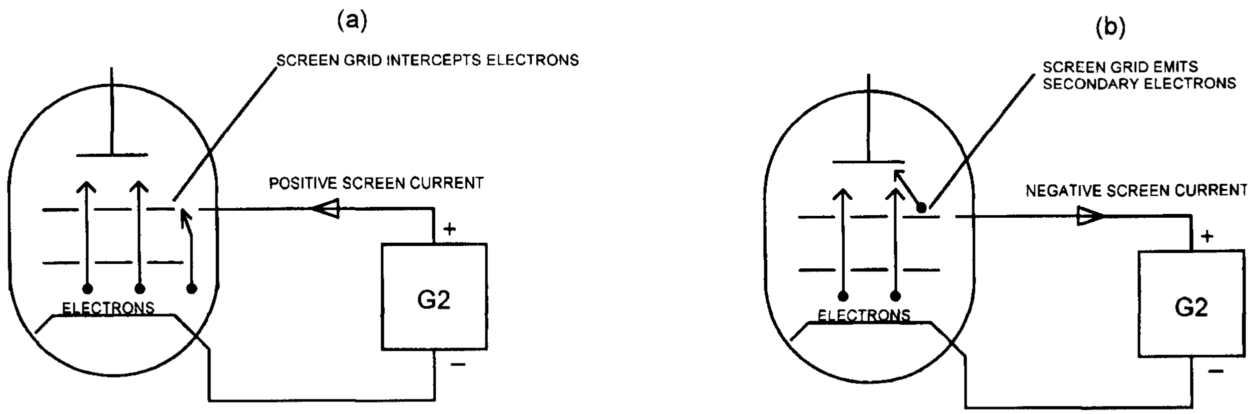

Fig 1—Screen-grid current can flow in either direction. At (a), the screen grid intercepts some electrons, drawing positive screen current from the supply. At (b), the screen grid emits more secondary electrons than it intercepts, driving negative screen current back into the supply.

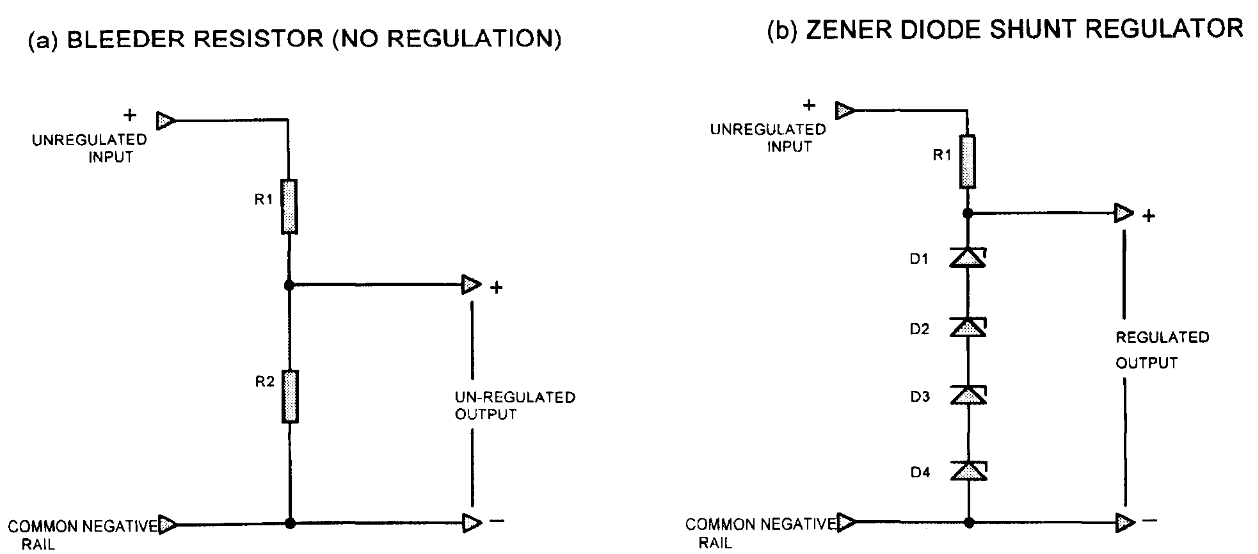

Fig 2—Historical screen supplies: (a) Bleeder resistor gives no voltage regulation. (b) Zener diode or VR-tube regulator (Zeners shown).

cathode/screen. This means that any variations in screen current are swamped by the much larger variations in cathode current. In addition, both the anode and the cathode/screen supplies in the 30S-1 used choke-input filtering, which gives better voltage regulation than today's capacitor-input supplies, and this helped to prevent the operating point from wandering around under dynamic modulation conditions. As authors from the Collins company explained,[2] if no voltages are stabilized, any variations in the mains voltage will change all the supply voltages in the same proportion, so the zero-drive current of the tube hardly changes. Unfortunately this approach is not as simple as it seems, for the Collins authors also made it very clear that if the control-grid bias is stabilized, the screen-grid voltage needs to be stabilized too—a point that later designers missed when they tried to borrow selected features from the 30S-1 without realizing that it's an all-or-nothing deal. Today, there are better ways to achieve dc stability in tetrode amplifiers, involving a little more electronics but much less heavy iron.

A final point in favor of improving the dc stability is that secondary electron emission from the screen of many tetrodes tends to increase with time. Older tubes may not be usable in amplifiers that have poor screen regulation, because of the runaway effect. With a power supply that takes a very tight grip on the screen voltage, you can often continue to use these tubes for hundreds of hours more.

Reduced IMD

As the *ARRL Handbook* points out,[3] "The power output from a tetrode is very sensitive to screen voltage, and any dynamic change in the screen potential can cause distorted output. In a linear amplifier, the screen voltage should be well regulated for all values of screen current." How well must we regulate the screen voltage? The answer will depend partly on the type of tetrode that you're using, but mostly on the standards you're setting for low intermodulation distortion (IMD). The screen supplies described in this article were designed to meet the exacting standards of European VHF DXing and contesting. Compared with HF, background noise levels at VHF are much lower, yet worst-case signal strengths between local stations using stacked arrays of long Yagis can be very much higher. In an IARU Region 1 2-meter contest, "kilowatt alley" covers most of Western Europe! As well as testing the dynamic range of receivers to the limit, these operating conditions place extreme demands on the IMD suppression of transmitters—demands that are reinforced by tough contest rules against persistent poor-quality transmissions.

Traditionally, amateurs measure IMD by on-air tests, listening to each other's signals and by two-tone testing. Informal on-air tests seem less respectable than two-tone tests using laboratory equipment, yet in many ways they are more meaningful because they exercise the whole amplifier—including the power supply—under realistic modulation conditions. A normal two-tone test hardly exercises the power supply at all. The meters never move, so even an amplifier with totally unregulated power supplies can produce good-looking IMD performance in this essentially static test. Real-life speech modulation tells a very different story. If you have access to a modern digital spectrum analyzer, you can perform a very revealing test by setting the analyzer to peak-hold mode and simply talking into the microphone. In the course of a few minutes, a very broad IMD spectrum will build up, as the analyzer records even the transient peaks of splatter. Unlike the static two-tone test, a peak-hold test is likely to reveal high-order IMD extending far out on either side of your main signal—as your neighbors on the band may already know! This is different from a static two-tone test in that real speech exercises the dynamic regulation of your power supplies over a wide range of audio frequencies, from about 3 kHz all the way down to powerful syllabic pulses at a few Hertz. It's very simple to improve the regulation at 500 Hz to 3 kHz by connecting a large reservoir capacitor across the output of the screen supply; that's an easy way to make a two-tone test look good, but the capacitor has no effect at syllabic frequencies.

John Nelson, GW4FRX, has been a constant campaigner for cleaner signals, and has been responsible for many key developments in tetrode power supplies.[4, 5, 6] In particular he has carried out many series of two-tone IMD tests on 4CX250 and 4CX350 amplifiers, and more recently peak-hold tests, as well. The first major conclusion is that better screen-voltage regulation gives significantly better IMD performance, especially for the higher-order products that make your signal wide. Fig 4 shows the IMD performance of a pair of 4CX250Rs in class AB1 at 500 W PEP output, with three different types of screen supply. The best IMD performance comes from GW4FRX's own extremely well regulated supply. Second best is that same supply with a 150 Ω series resistor added to artificially increase the output impedance. A poor third—especially for higher-order IMD—is the traditional "chain of Zeners" stabilizer with a parallel reservoir capacitor.

The second major conclusion is that improved screen stabilization can give

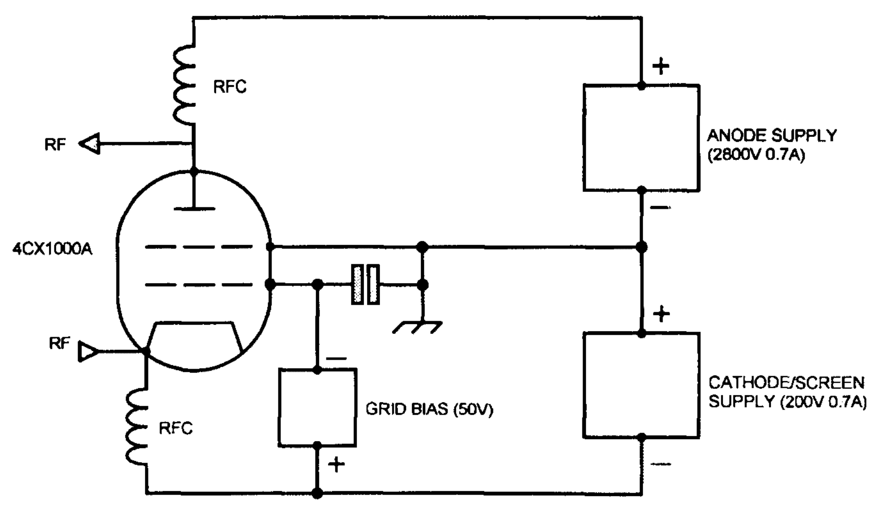

Fig 3—The Collins 30S-1 used two separate high-current supplies for the cathode and screen, with choke-input filters but no other voltage regulation.

IMD performance that is notably better than stated in the Eimac data sheets. On the air, these conclusions have been verified by many British and European stations. Tight screen voltage regulation really does make a difference to your reputation!

Effective Protection

The next reason for paying close attention to the screen supply is that it can protect the whole amplifier. Almost anything wrong in a tetrode power amplifier will result in incorrect screen current. The range of faults that can be detected by monitoring the screen current includes: incorrect plate-circuit tuning; loading too light, or too heavy; excessive drive; loss or major change in anode, screen or control grid voltage; high-voltage RF and dc arcs, flashovers and other glitches—even overheating. All of these faults will result in too much screen current, either positive or negative. Using the protection circuit described later in this article, I've been able to keep the same pair of 4CX250Bs delivering 1 kW output on 432 MHz moonbounce for more than 10 years. If the circuit hadn't worked well and reliably, the tubes would have been destroyed several times over.

That's enough philosophy. I hope I've convinced you that improved screen supplies can bring major benefits, so now let's get down to some circuits.

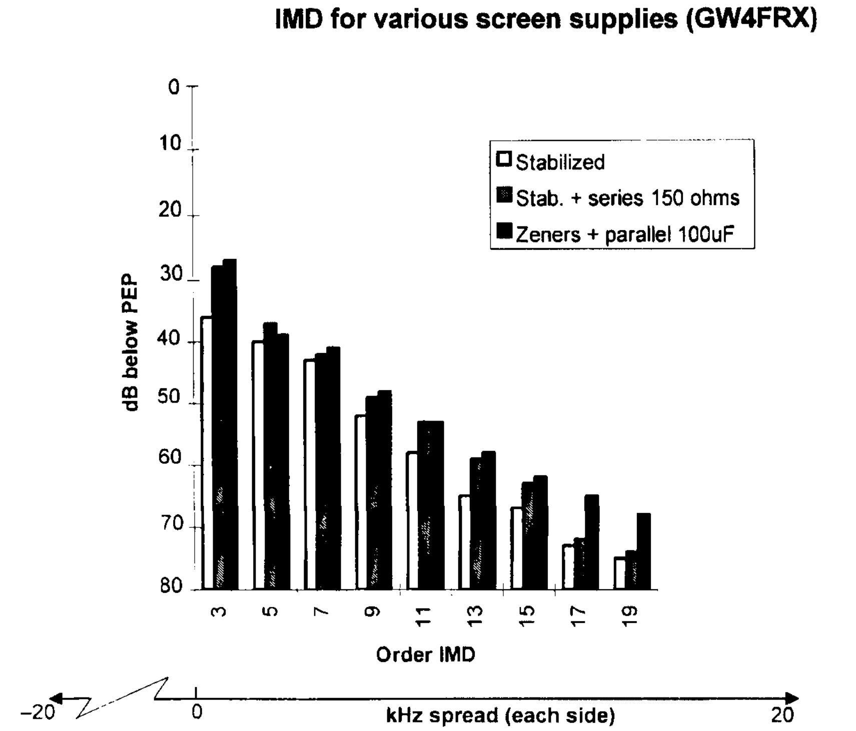

Fig 4—Better-regulated screen supplies give lower IMD: two-tone performance of a pair of 4CX250Rs with three different screen regulators.

Shunt Regulator Basics

Because of the need to handle negative screen current, all screen voltage supplies **must** have a dc path to ground.[1] In other words they need to be shunt regulators rather than series regulators.[7] Fig 5 recalls the basic shunt regulator configuration, a potential divider with a resistor (R1) from the unregulated supply to the screen and some kind of constant-voltage circuit from screen to ground.

Let's look at the current-flow budget in a shunt stabilizer. The current that flows through R1 is almost the same under all conditions; what varies is the fraction of the total current that is either delivered to the screen grid or shunted to ground through the voltage stabilizer. The unregulated supply and R1 must be capable of delivering the maximum positive screen current required by the tube, but the current also needs to be limited to protect the screen from excessive dissipation. The current that bleeds to ground through the voltage stabilizer must always be greater than the maximum positive screen current that the tube ever requires; otherwise the screen voltage will sink with excessive current demand. Also the stabilizer element must be capable of sinking the maximum negative screen current that the tube ever generates, plus the bleeder current provided to handle with positive current demands. If the stabilizer can't handle all this extra current, it will allow the screen voltage to rise, which can lead to the runaway effect.

In an SSB amplifier the peak positive and negative screen currents can occur at unexpected points in the speech modulation waveform. Screen current will be close to zero with no drive, but in some tetrodes the current may peak negatively at moderate drive levels, then pass through zero again with increasing drive and finally reach a positive peak. In other tetrodes, screen current may peak almost exclusively in the negative region, except at very high drive levels. These negative and positive variations can occur within a single audio cycle, and the screen current meter cannot follow them. Large positive and negative current swings may average out to almost zero on the meter and lead you to assume that there are no real demands on the power supply—but if the screen voltage drops out of regulation for even an instant, your neighbors will know it!

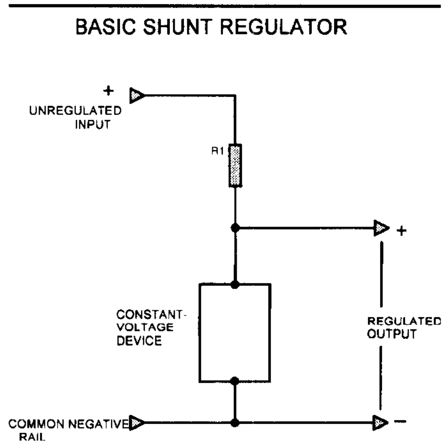

Fig 5—Basic shunt regulator configuration.

What's wrong with conventional chains of VR tubes or Zener diodes in regulated screen supplies (Fig 2b)? VR tubes have a very significant dynamic impedance—the ratio of (voltage variation)/(current variation)—and this effect is cumulative when devices are connected in series to obtain the required total voltage. A typical series string of two VR105s and a VR150 has a dynamic impedance of about 500 Ω. Unfortunately high-voltage Zener diodes are not much better than VR tubes, so either type of stabilizer may allow the screen supply to swing by several volts when the screen current changes by about ±10 mA. Passive devices for screen-voltage regulation are hardly adequate, as the IMD spectra in Fig 4 clearly show.

For excellent voltage regulation that will allow the tube to develop its optimum IMD performance, the solution is an active-feedback regulator. The next question is to find the right level of circuit complexity, namely the simplest circuit that will achieve all three of the following:

- Give excellent voltage regulation.
- Protect the tube and amplifier against faults.
- Survive major faults such as arcs and flashovers without damage to the regulator itself.

Two or three transistors in a simple feedback loop can make quite an effective shunt regulator,[8] but my view is that once you've made the decision to "go the feedback route," you might as well go all the way. With a bipolar power transistor or MOSFET as the active shunt element, controlled by an op-amp, the improvement is dramatic. Voltage fluctuations, hum and noise levels all drop to a few tens of millivolts (on a 350 to 400 V rail!), which is better than any tetrode could possibly need. You simply need not worry about voltage regulation any more.

Active Shunt Regulator

The basic circuit shown in Fig 6 was originated by G4JZQ[4,5,6] The shunt regulator element is power MOSFET Q1, which is fed from the unregulated power supply by resistor R1. The gate of Q1 is driven by op-amp U1. A divided-down sample of the output voltage is fed to the noninverting input of U1, and the inverting input is held at a stable reference voltage. To analyze how this circuit works, think what would happen if the output voltage tried to rise. Current through the voltage divider R2-R3 would cause the voltage at the noninverting input of U1 to rise, and therefore the output voltage from U1 to the gate of Q1 would rise by a much greater amount. This would make Q1 conduct more heavily, pulling down the output voltage and compensating for its original tendency to rise. Exactly the reverse would happen if the output voltage tried to fall; Q1 would conduct less, and allow the output voltage to rise again by exactly the correct amount.

This is a very high-gain feedback loop, so it requires stabilization over a wide range of frequencies. An ordinary internally compensated op-amp is not suitable—in fact it will oscillate. The simple trick, courtesy of G4JZQ, is to use an uncompensated op-amp such as the 748 with heavy external compensation from the network R4-C1. (If you're not familiar with the 748, it's the good old 741 without its built-in compensation capacitor.) C2 also affects the loop's stability and HF response, as does C3 to a lesser extent. This basic circuit has shown reliable margins of stability in several variants, using both bipolar and MOSFET power transistors, and also in configurations involving much higher loop gain than shown here.

To conclude the description of the feedback loop, Q1 always has to operate in its turn-on threshold region, which requires a gate voltage of about +2 V. Since the output of U1 will not swing reliably down to this voltage when used with a single supply rail, the potential divider R5-R6 allows U1 to operate at a more comfortable output voltage of about +4 V.

There are two reasons for using a power MOSFET at Q1 rather than the more familiar bipolar power transistor. One is the high gate impedance—MOSFETs are easy to drive at these low frequencies. The other reason is that screen-regulator usage involves an unpleasant combination of high voltage, relatively high current and high heat dissipation that can cause bipolar transistors to fail unexpectedly by a phenomenon called "second breakdown." Power MOSFETs are immune from second breakdown and are therefore the best choice for Q1. With a little care to avoid electrostatic damage, they are very easy to use, and are very rugged once installed in circuit. You'll like the prices too—1000 V devices rated at more than 100 W dissipation at 25°C cost less than $5 each.

The value of the HV (high voltage) feed resistor R1 is important. Together with the unregulated power supply it controls the maximum current available and the resting power dissipation of Q1. When the tetrode draws positive screen current, that current no longer flows to ground through Q1 so its dissipation decreases. The maximum current available without losing voltage stabilization is when Q1 draws no current at all. The worst situation for power dissipation in Q1 is when the tube is continuously producing negative screen current, which Q1 must bleed away to ground in addition to the normal current supplied through R1. The maximum power dissipation is therefore:

$$P_{DISS} = V_R \times (I_{Gmax-} + I_{Gmax+}) \quad \text{(Eq 1)}$$

where,

V_R = regulated voltage

I_{Gmax-} = maximum negative screen current

I_{Gmax+} = maximum positive screen current

Fortunately, not all of this power must be dumped into Q1, because you can add high-power resistor R7 in series with Q1 to share the load. R7 narrows the ranges of both positive and negative screen currents that the supply can handle without losing voltage stabilization, so you need to choose the value carefully.

Flashover Protection

Up to now we've mainly been thinking about normal operation—but what happens when things go wrong? Many amateur amplifier builders seem to ignore this possibility, or resign themselves to extensive damage in the event of a major fault such as a flashover. I find this totally unacceptable. A reasonable design aim is **zero damage** from any kind of minor fault—just push the RESET button and be back on the air immediately. Even a major flashover doesn't have to result in anything worse than a blown fuse. It shouldn't be necessary to switch on the soldering iron.

Flashovers are the main cause of tube and circuit damage. If your amplifier can handle one of those, it can probably handle most other kinds of faults too. They can occur either inside or outside of the tube envelope, and can be caused by incorrect tuning, dust or bugs in the cooling air, a sudden release of gas within the tube (especially in the first few hundred hours), and sometimes there seems to be no reason at all—the amplifier just goes BANG! Whatever the reason, the effect of a flashover is to crowbar the HV supply with a low-resistance arc from the anode, which can be highly destructive. When a tetrode flashes over, an internal arc will hit the screen grid and an external arc will hit the contact ring and the socket. Then the surge current will head back toward the power supply. It is vitally important to protect all these components on a timescale of microseconds, and then to kill the arc as quickly as possible.

Whatever your views about screen supplies, the following precautions are **absolutely essential**. Most of them apply to triode amplifiers too.

• Use a current-limiting resistor in series with the HV+ supply. For a typical 2 to 3 kV power supply, Eimac[9] recommends a resistor that will limit the peak fault current to 40 A; in other words, about 50 Ω. The resistor must be capable of withstanding the full HV for a few milliseconds without internal arcing, so use a long-bodied 50 to 100 W component.[10]

• Connect a surge voltage protector such as a voltage-dependent resistor (VDR) or a Siemens spark-gap from screen to cathode. When the arc hits, this device will conduct heavily and divert damaging current away from the tube, the screen bypass capacitor and the power supply. Surge voltage protectors cost a few dollars; tubes and sockets cost hundreds!

• Interrupt the mains supply to the HV transformer as quickly as possible, to limit the follow-through energy in the arc. Don't wait for a fuse to blow—use a fault-detection circuit to control a high-speed circuit breaker. A solid-state relay can interrupt the mains power in less than 10 milliseconds, at the next zero crossing of the ac cycle.

• Protect the meters and the HV-negative rail from the effects of the current surge (very important,[11] but outside the scope of this article).

• Protect the screen supply—but with-out risk to the tube.

Fig 7 shows some bad screen circuits that are either ineffective or could actually endanger the tube.[12] Many of

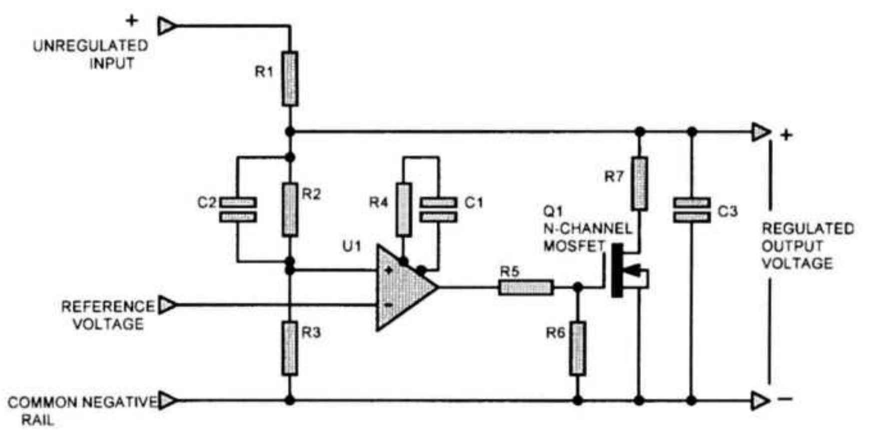

Fig 6—Simplified circuit of G4JZQ's active shunt regulator.

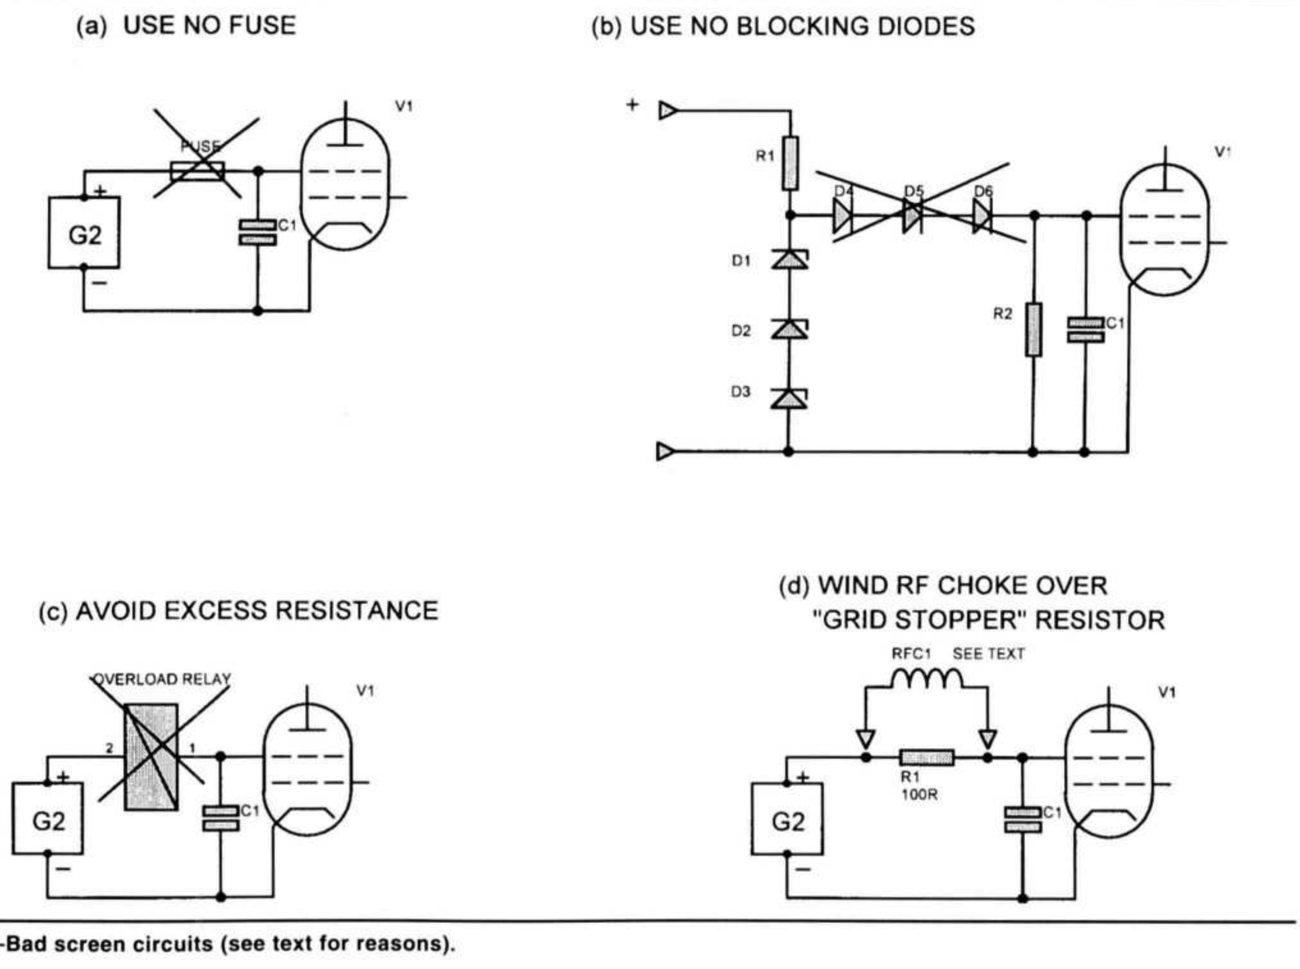

Fig 7—Bad screen circuits (see text for reasons).

them seem to originate from the fuzzy notion that it's more important to protect the power supply than the tube, or that you must be willing to sacrifice either one or the other. *Wrong!* A good circuit will reliably protect everything!

Fig 7a has a low-current fuse. It probably won't blow reliably, and then the arc will drive the screen voltage disastrously high. Even if the tube survives, it will probably blow the screen bypass capacitor and total that expensive base. A surge voltage protector will definitely help, but it's only a band-aid for a fundamentally bad circuit. Fig 7b uses one or more blocking diodes to protect the Zener stabilizer diodes. In normal operation this circuit has absolutely no voltage regulation against negative screen current, leading to potential dc instability and perhaps even provoking a flashover. When a flashover does happen, the circuit relies totally on the surge voltage protector. Fig 7c is rather more sensible; it uses a small relay to detect excessive screen current, but it is slow-acting and also the relay coil adds a significant resistive and inductive component to the dynamic impedance of the screen supply. Fig 7d has a 100 Ω "grid stopper" resistor, and was probably copied from circuits that were published back in the Class-C days. Unfortunately the voltage drop across the resistor degrades the screen-voltage regulation. As described later, it's very simple to convert this resistor into a highly damped RF choke that has a minimal voltage drop.

Screen-Current Trip Circuit

Screen-current metering is essential in any tetrode power amplifier because it's the most reliable tune-up indicator. In addition to monitoring the screen current visually, it's very useful to monitor the current electronically as a basis for fault protection. Electronic circuits can react far faster than you can! Fig 8 shows a screen-current monitor circuit that is optocoupled and can float at any voltage. Bridge rectifier BR1 makes the circuit respond to both positive and negative screen current. The extra resistors and Zener diode D1 protect opto-isolator U1 against flashovers and short-circuits. Like the overload relay in Fig 7c, this current monitor causes a significant voltage drop, but that doesn't matter if the circuit is located inside the feedback loop of an active voltage regulator.

The optocoupler transmits the screen-current signal to the amplifier control circuits. There are several ways to use this signal, for example to trigger a small thyristor as shown in Fig 8. The trigger point is stabilized by the voltage regulator U2 and adjusted by RV1. With the component values shown, the trigger point is adjustable for screen currents in the range ±20 mA to ±40 mA. When the thyristor Q1 triggers, Q4 is biased to cutoff, and takes the amplifier off-line by removing the +24 V dc supply from all relays, including the relay that grounds the screen and the two-pole mains power relay for the HV transformer. The alarm LED lights, and the thyristor Q1 remains latched in this condition until you press the RESET button SW1 (or remove power completely). If it was only a minor fault, you're back on the air as soon as you press the RESET button.

Other fault signals can be linked into the gate or anode of Q1, as shown in inset of Fig 8. The gate of Q1 requires a positive current to trigger the thyristor, and multiple inputs should have steering diodes to prevent interaction. Inputs to the gate of Q1 are "latching," that is once triggered, Q1 will conduct until you press the RESET button. For momentary faults that don't require a latching input, ground the anode of Q1 to hold the amplifier off-line only as long as the ground connection is present. You can decide exactly how much automatic protection you want (at 3 AM in a contest, I want a lot!). The options include: an optocoupled screen current monitor for a second tube; a monitor to detect control-grid current in a Class-AB1 amplifier to prevent overdriving; a warm-up timer to hold the amplifier in standby mode until the cathode has reached oper-ating temperature; and switches to indicate blower failure or excessive exhaust air temperature.

Practical Regulator Circuits

Now we're ready to look at some practical screen-regulator circuits. Fig 9 is the complete working version of Fig 6 with the optocoupled current monitor of Fig 8. With the component values shown, the output voltage is adjustable from 350 to 400 V. The output voltage is controlled by comparing the sample from RV1 with the +12 V reference provided by U2. R14 is the voltage-dropping resistor from the unregulated supply, which should be at least 30 to 40 V higher than the stabilized output voltage. Because the voltage regulator is so effective, the unregulated supply need not be very stiff. The only important consideration is that there is 30 to 40 V of headroom (excess voltage) under all conditions, even on negative peaks of mains ripple at maximum current. The value of R14 depends on the headroom voltage and the maximum positive screen current that the regulator must deliver. Since this is a shunt regulator, the current that flows through R14 is nearly constant; only the proportions delivered to the screen grid and shunted to ground through Q1 vary. Therefore you should adjust R14 to deliver the required current through Q1 when the regulator is disconnected from the tube.

R12 is a power resistor that takes some of the thermal load off Q1, so that the semiconductor can be mounted on a smaller heat sink. To further reduce the long-term power dissipation, relay K1A switches the voltage regulator into a low-power mode on receive. R15 allows about 3 mA to pass through Q1, which is just enough to maintain voltage regulation but allows a quick return to full power on switching to transmit. Relay K1B switches the screen to ground on receive, or whenever the fault circuit is triggered. This also protects the screen grid against failure of the anode HV supply, which results in very high screen current. In the event of any such fault, K1B will quickly change over and ground the screen. In the few milliseconds while K1 is switching, the maximum screen current is limited by R14. Likewise R14, BR1 and the protective components around optocoupler U1 are all rated to survive high currents until K1B opens. R4 is a permanent bleed resistor installed in the amplifier RF deck to maintain ground continuity to the screen grid while the contacts of K1B are changed over, or in case a screen supply in a separate enclosure becomes disconnected from the RF deck. If R4 provides a bleed current of about 10 mA when screen voltage is applied, it will also be plenty low enough to protect the screen grid while K1B is changing over. For a screen voltage of 360 V, R4 should therefore be about 36 kΩ, with a generous power rating of 10 W to ensure cool and reliable operation.

The screen-current meter needs to display both positive and negative currents, and the 10 mA bleed current through R4 means that a conventional left-hand-zero meter will read +10 mA, even when the actual screen current is

Continued on page 24.

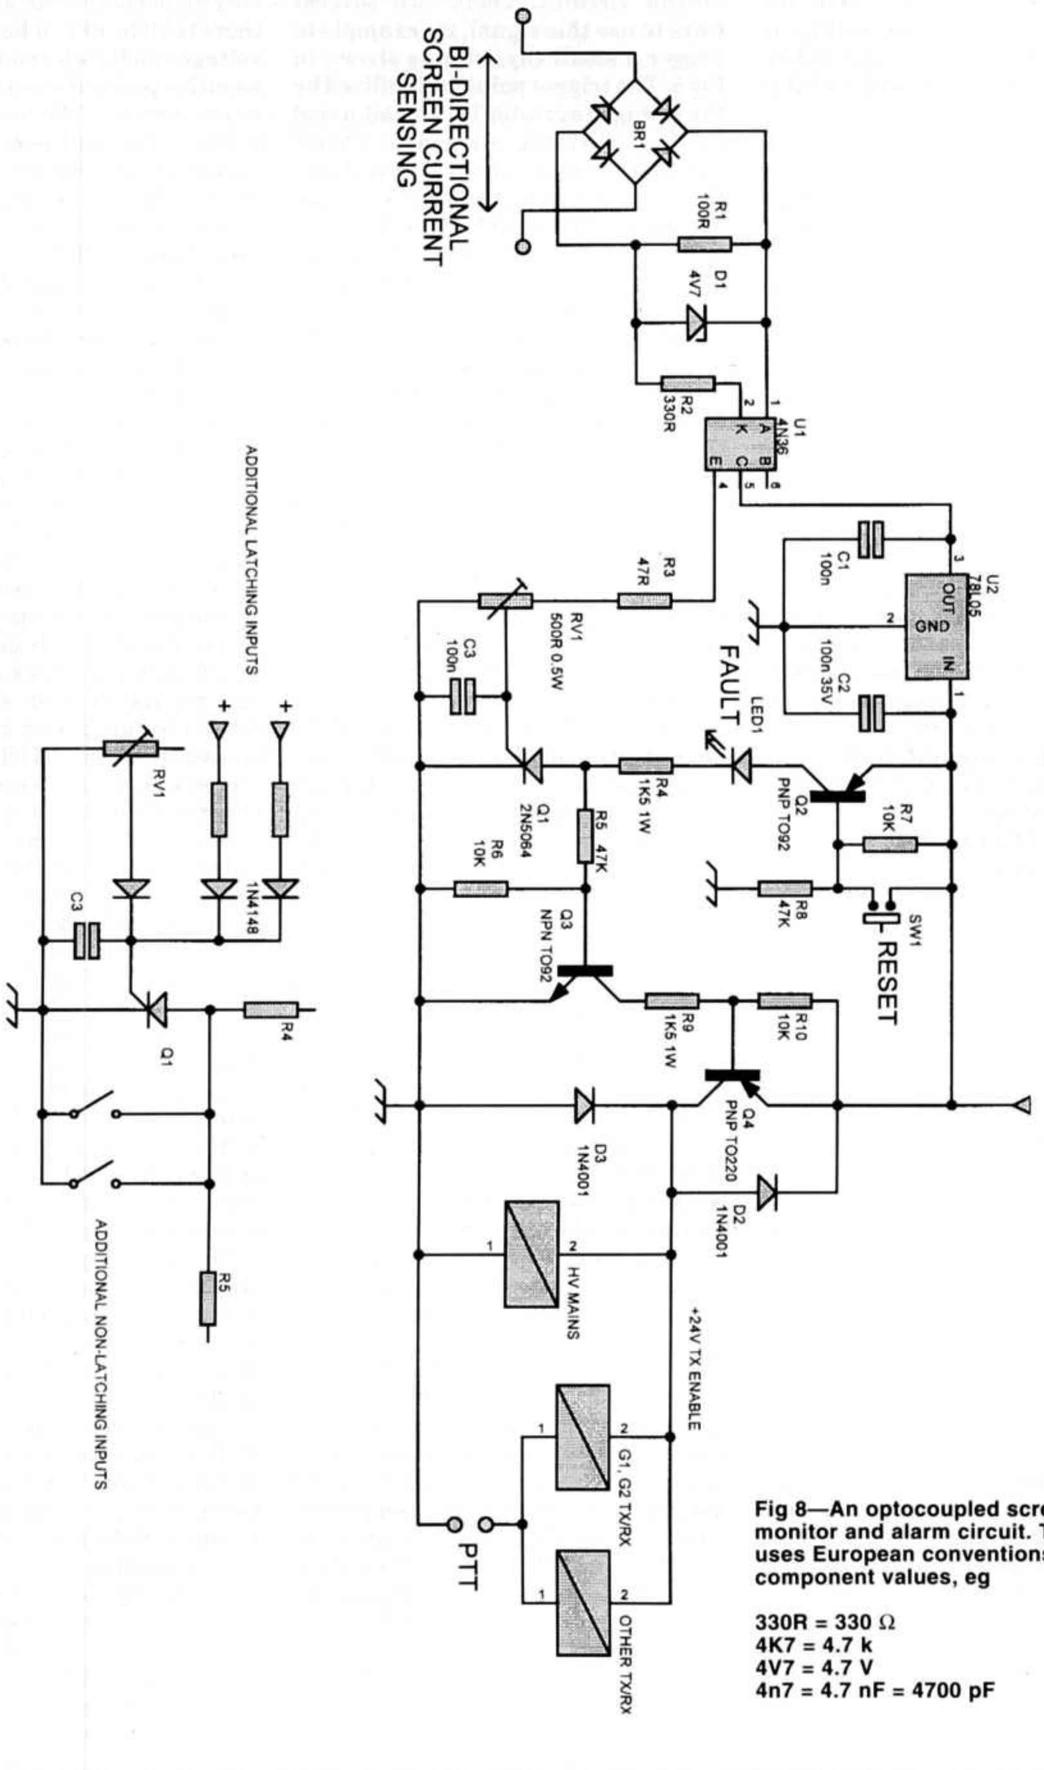

Fig 8—An optocoupled screen-current monitor and alarm circuit. This diagram uses European conventions for component values, eg

330R = 330 Ω
4K7 = 4.7 k
4V7 = 4.7 V
4n7 = 4.7 nF = 4700 pF

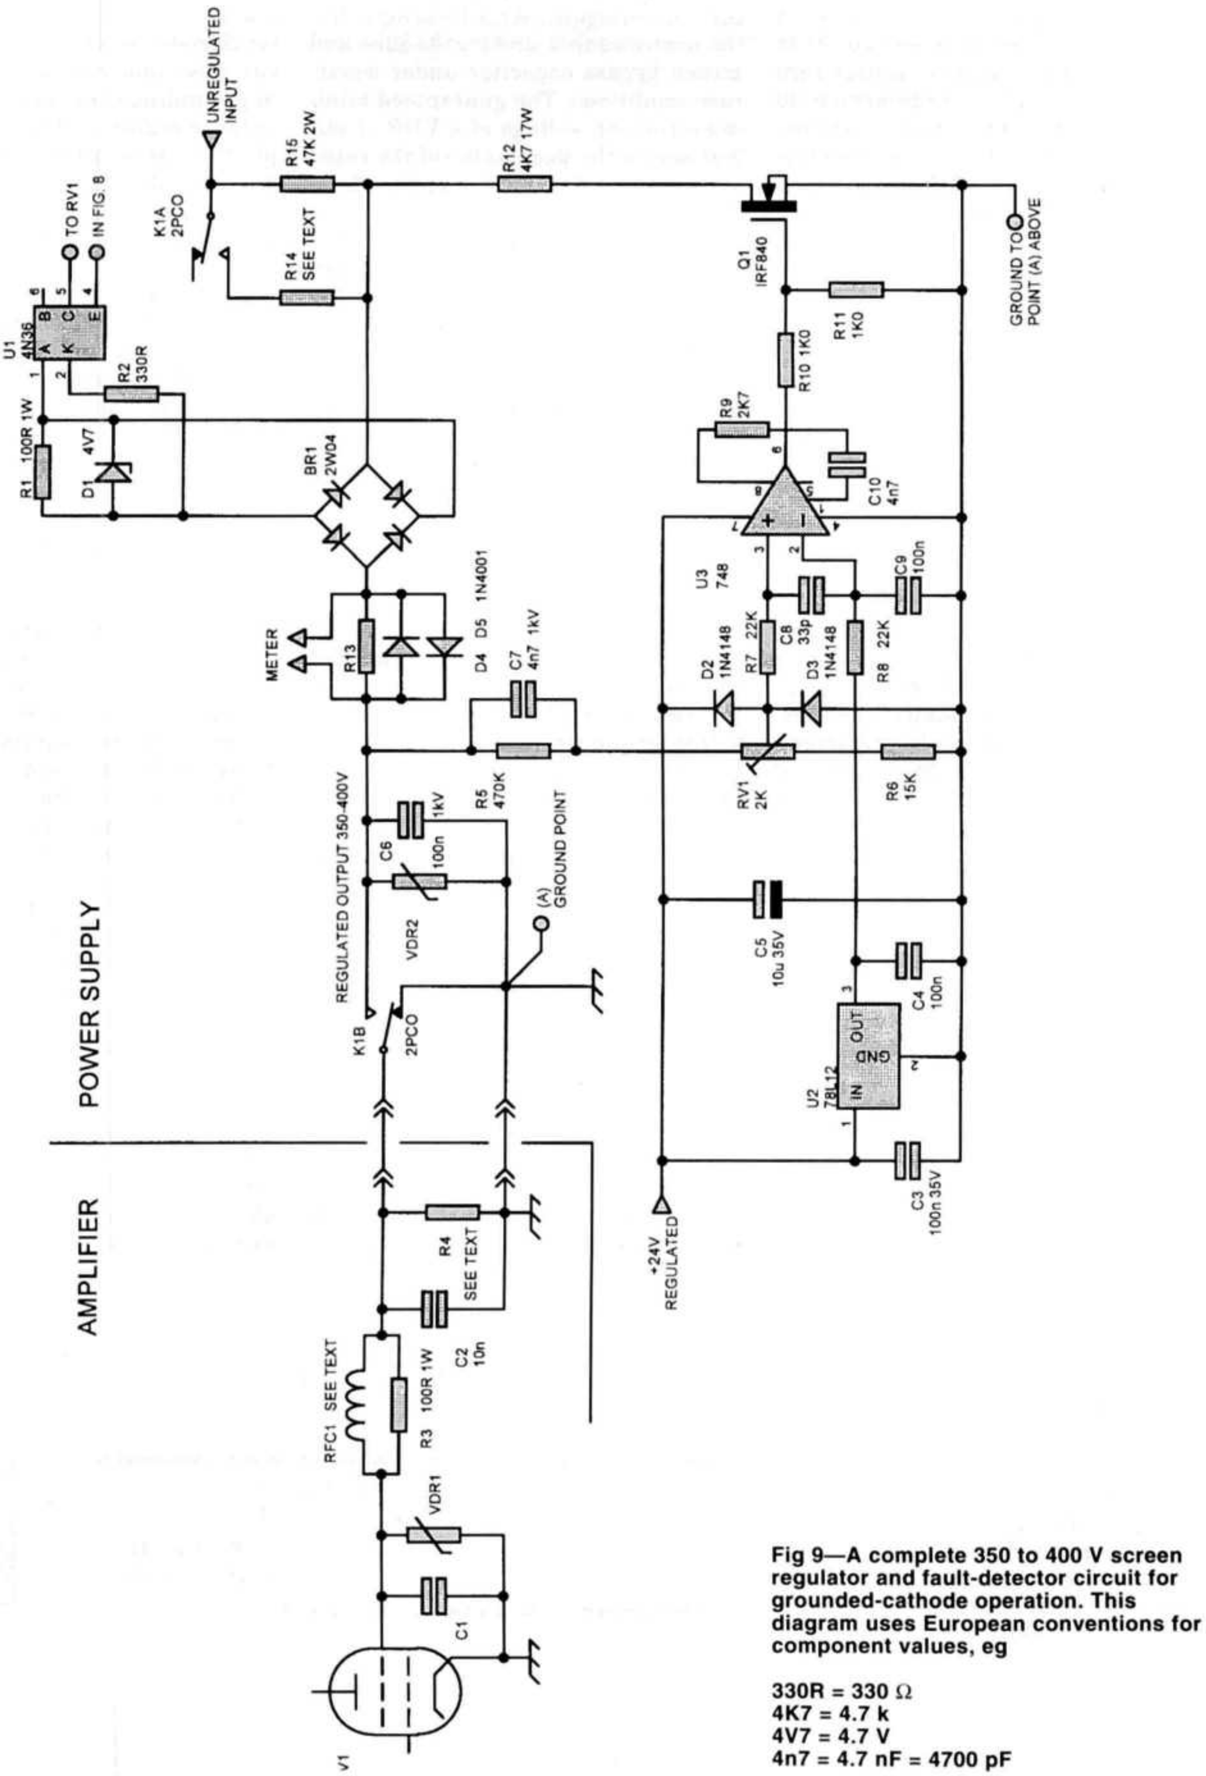

Fig 9—A complete 350 to 400 V screen regulator and fault-detector circuit for grounded-cathode operation. This diagram uses European conventions for component values, eg

330R = 330 Ω
4K7 = 4.7 k
4V7 = 4.7 V
4n7 = 4.7 nF = 4700 pF

zero. This is a very useful feature, because it means that you can see both positive and negative screen currents without using a special center-zero meter. For example, an ordinary 0 to 20 mA meter will display true screen currents of –10 mA to +10 mA, which is exactly what you need for a small tetrode. It doesn't matter if the bleed current through R4 is not quite the value you wanted; simply zero the meter using the adjustment screw when screen voltage is applied but the screen current is zero. A less desirable consequence of the current through R4 is that the screen current trip is asymmetrical. For example, if the current monitor is set to trip at ±25 mA through BR1, a 10 mA bleed current through R4 means that the circuit trips at true screen-grid currents of +15 mA or –35 mA. In practice this is not a problem, because most tetrodes should never approach their screen dissipation limit in normal operation. As Fig 10 shows, you can still set the asymmetrical trip limits to protect the tube.

R3 and RFC1 decouple the screen bypass capacitor C1 from the rest of the circuit. This is important to avoid any parallel resonance between C1 and the self-inductance of capacitors such as C6, which will make the screen "live" in the HF region. For example, using an MFJ-259 I measured a strong parallel resonance at 15 MHz from an Eimac SK-620A socket and the kind of plastic-film capacitor you would typically use for C6. RFC1 is made by winding about 40 turns of thin enameled wire over the body of R3 (100 Ω, 1 W, which must be carbon or metal-film, not wirewound). Inserting this combination between C6 and the tube socket completely kills the unwanted parallel resonance without introducing any significant voltage drop.

Note the two voltage-dependent resistors (or VDRs, also known as varistors, metal-oxide varistors, MOVs, Transzorbs, zinc-oxide nonlinear resistors, ZNRs, etc) to protect the screen voltage from being driven excessively high by an arc or flashover from the HV supply. The VDRs act in nanoseconds, giving front-line protection while the trip circuit catches up with events. VDR1 protects the tube, and VDR2 is a backup to protect the rest of the circuit. VDRs are normally rated for their nominal ac operating voltage and their energy-absorbing capability. For this application you should choose VDRs that have a guaranteed minimum turn-on voltage at least 20 V above the required screen voltage, so that they will not normally conduct at all, but the turn-on voltage must not be so high that the device cannot protect the tube and screen bypass capacitor under worst-case conditions. The guaranteed minimum turn-on voltage of a VDR (1 mA leakage) is the peak value of the rated ac voltage. Taking examples from the GE-MOV product line (Harris Semi-conductors), the 275 V ac-rated V275LA40B is suitable for screen voltages of 350 to 370 V, and the V320LA40B for voltages up to about 440 V. The energy-absorbing capability of these devices is 140 to 160 joules, which proves very adequate. In practice, these devices will give protection against repeated flashovers. Instead of VDRs you could also use similarly rated gas discharge tubes from the Siemens line; the choice is largely a matter of preference and availability.

In a flashover, the current pulse through the ground return to VDR2 could be as high as 30 to 40 A, limited mainly by the resistor in the positive HV rail. If this current passes along a thin ground rail used by sensitive low-level circuitry, the voltage drop could cause component damage due to "ground bounce" (as I discovered when testing an early prototype). Therefore the ground return to VDR2 must also be the main chassis ground for the whole circuit, as shown in Fig 9. With that precaution—and the all-important current limiting resistor in the HV+ rail—this screen supply will survive repeated deliberate flashovers and crowbar short-circuits.

A Floating Regulator

The circuit shown in Fig 9 is mainly suitable for grounded-cathode configurations, because it uses the +24 V dc relay supply for op-amp U3 and voltage reference source U2. There are many other possible dc configurations for tetrode amplifiers, involving cathode drive and various options for RF/dc grounding of the two grids,[13] but all of these require a floating screen supply. The clever circuit of Fig 11 (again due to G4JZQ[6]) solves that problem by borrowing power for U2 and U3 from the screen supply itself. Q2, D6 and D7 form a simple shunt-regulated supply to power the two ICs; and shunt-regulator transistor Q1 sits above this +30 V rail. A level-shifting network (R20 and R21) connects the output of U2 to the gate of Q1. In order to obtain a floating output, this circuit has a common-negative rail rather than the chassis-ground rail of Fig 9. To avoid destructive current surges along the negative rail and into the low-level parts of the circuit, all connections to the common negative must be routed to a single point as shown. As with the circuit of Fig 9, this version has proved highly effective and prototypes have been in use in Europe for several years.

The screen regulators in Figs 9 and 11 are both shown configured for 350 to 400 V output. This voltage range is suitable for the vast majority of modern ceramic-metal tetrodes used by amateurs in Class AB1 or AB2, including all tubes in the 4CX250, 4CX350, 4CX400, 4CX800 and 4CX1600 families. Consult the manufacturers' data sheets for the most appropriate setting. For the 4CX1000 and 4CX1500 family of tubes, which may prefer screen voltages in the range from 300 to 350 V, increase R6 to 16 to 18 kΩ. With appropriate changes, the same circuit should work equally well for screen voltages up to 1000 V, which is about the practical limit for readily available power MOSFETs.

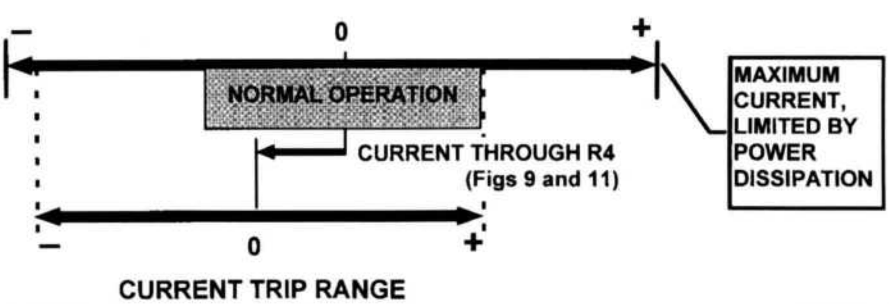

Fig 10—The screen-current trip point can be adjusted to cover the normal range of operation and also protect against excessive dissipation, even through the center-zero is offset from true screen current by the bleed current through R4 (see Figs 9 and 11).

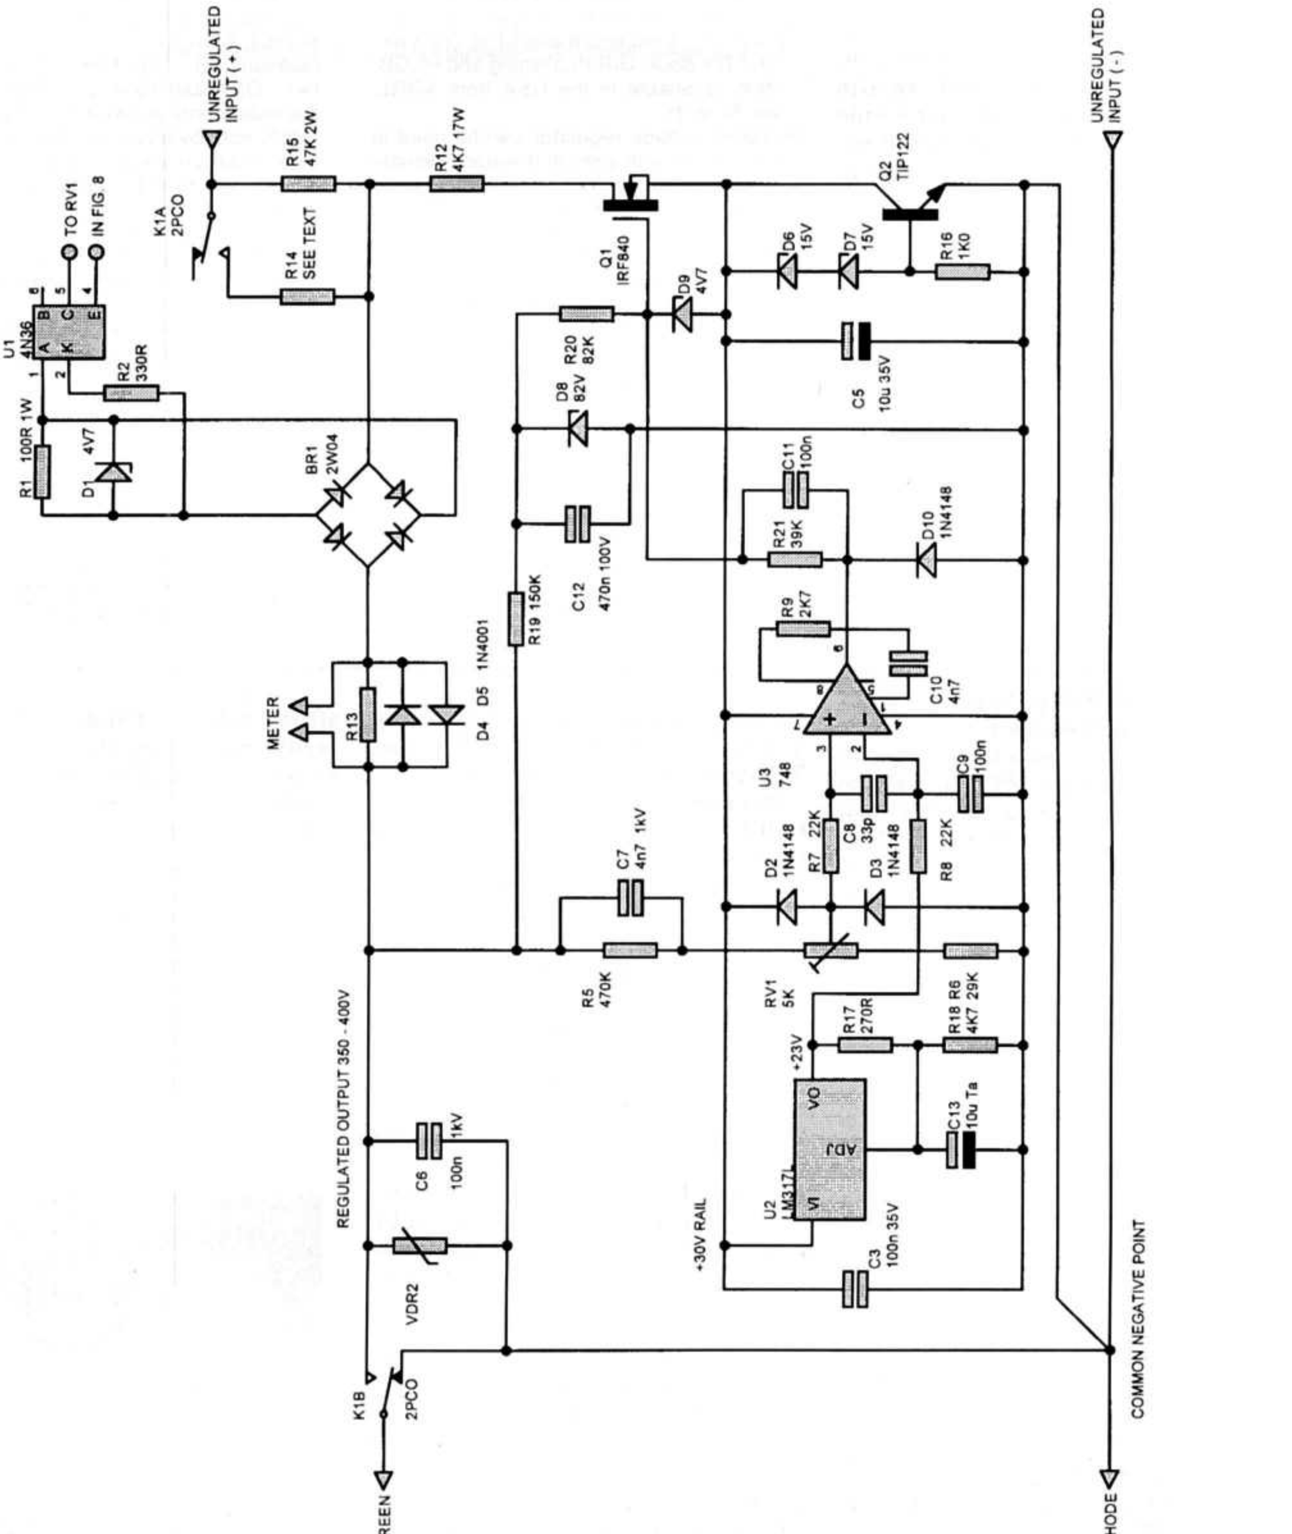

Fig 11—A fully-floating 350 to 400 V screen regulator, generating its own +30 V rail for the op-amp and voltage reference. This diagram uses European conventions for component values, eg

330R = 330 Ω
4K7 = 4.7 k
4V7 = 4.7 V
4n7 = 4.7 nF = 4700 pF

Conclusion

This article is intended as a source of new ideas. The screen regulator and protection circuits described can also be used as drop-in upgrades for a wide variety of existing tetrode amplifiers. They are thoroughly tested and can survive repeated HV flashovers and crowbar short-circuits. Further details and updates can be found at **www.ifwtech.demon.co.uk/g3sek.**

There has always been some resistance to the uncomfortable fact that tetrodes are more complicated than triodes. Maybe the circuits involved are more complex than you'd like them to be, but I have explained the reasoning behind the design decisions so that you can make your own choices and avoid some common mistakes. You only build an amplifier once, and that is the time to build in quality, security and peace of mind for all the hours of successful operating to come.

I am grateful to John Nelson, GW4FRX, and Melvyn Noakes, G4JZQ, for the comprehensive tetrode power supply and control units that started all these developments,[4,5,6] and also to GW4FRX for providing the IMD test results in Fig 4. Many thanks to Mark Mandelkern, K5AM, Bill Sabin, WØIYH and Tom Rauch, W8JI, for their advice about important points to emphasize for readers in the USA.

Notes and References

[1]Through most of this article I will describe circuits for a tetrode in the grid-driven, grounded cathode configuration, where the negative rail of the screen supply is connected to chassis ground. The use of a bleeder resistor is described in Eimac's classic *Care and Feeding of Power Grid Tubes* (1967, out of print).

[2]Bruene, Pappenfus and Schoenike, "Power Supplies for SSB Amplifiers." Chapter 15 of *Single-Sideband Principles and Circuits*, first edition (New York: McGraw Hill, 1964).

[3]*The 1998 ARRL Handbook* and recent editions (ARRL Order No. 1743), the RF Power Amplifiers chapter. ARRL publications are available from your local ARRL dealer or directly from ARRL. Mail orders to Pub Sales Dept, ARRL, 225 Main St, Newington, CT 06111-1494. You can call us toll-free at tel 888-277-5289; fax your order to 860-594-0303; or send e-mail to **pubsales@arrl.org**. Check out the full ARRL publications line on the World Wide Web at **http://www.arrl.org/catalog**.

[4]Nelson, J., "A High Performance Power Supply and Control System for 4CX250/4CX250 Amplifiers," Parts I through VIII, *Short Wave Magazine*, Jul 1981 to Feb 1982.

[5]Nelson, J. and Noakes, M., "A Power Supply and Control System for Tetrode Amplifiers," *Radio Communication*, Dec 1987 and Jan 1988.

[6]Nelson, J., "Transmitters, Power Amplifiers and EMC" and "Power Supplies and Control Units," Chapters 6 and 11 of *The VHF/UHF DX Book*, DIR Publishing and RSGB, 1995 (available in the USA from ARRL, see Note 3).

[7]A series voltage regulator can be used in conjunction with a shunt (bleeder) resistor to ground, but a true shunt regulator is usually more convenient.

[8]Mandelkern, M., "A Luxury Linear," *QEX*, May 1996; "Design Notes for a Luxury Linear," *QEX*, Nov 1996.

[9]"Fault Protection," Eimac Application Bulletin #17, Jan 1987.

[10]G4GCM has successfully wound 50 Ω current-limiting resistors using resistance wire on a long 1-inch-diameter form, spacing adjacent turns by one wire diameter to prevent arcing.

[11]Measures, R., "The Nearly Perfect Amplifier," *QST*, Jan 1994, pp 30-34. Some of the statements in this article have proved highly controversial, but it gives good advice about connecting "glitch protection" diodes to protect meters and hold the negative HV rail near chassis potential in the event of a flashover.

[12]All of these circuits have appeared in published designs. References are omitted to avoid red faces!

[13]White, I., "A Tetrode Isn't a Triode," In Practice, *Radio Communication* (RSGB), Sep 1996, p 76.

Letters to The Editor

Here's a comment on "Synthesizing Vacuum Tubes," by Parker R. Cope, W2GOM/7, in August 1997 *QEX*

Over the years, I have converted a Collins 51S-1 receiver and some Collins VFOs to solid state. I liked W2GOM/7's article. I used 40673 dual-gate MOSFETs for pentodes (The NTE222 is equivalent to the 40673), and had problems with too much G_m giving oscillations.

I recommend, when converting an old tube receiver to solid state, that you get rid of the high voltage entirely, and run the receiver off 12 or 24 V B+. This simplifies substituting FETs for the tubes. The heat production is greatly reduced: The currents stay the same, while the plate voltage is cut by a factor of 10, giving 1/10 the plate dissipation. The heater dissipation is, of course, entirely gone.

The low-level stages will work fine. You may have to modify the B+ decoupling networks to reduce their voltage drop: Put small molded RF chokes across dropping resistors (try 100 µH to start, for RF use).

The audio power stage, of course, was designed for the high B+. Replace the audio power stage with an IC audio power amp, such as the LM380 or LM3875. After this conversion, my 51S-1 draws 500 mA at 12 V dc, half of which is for the dial lamps. The converted receiver is much more stable than when it had tubes.

Of particular note, is the R390 receiver. This radio ran off 24 V dc, using 24 V on both heaters and plates of all the tubes (old-radio buffs, please verify my memory).—*Pete Traneus Anderson, KC1HR, 990 Pine St, Burlington, VT 05401; e-mail* **traneus@emba.uvm.edu**.

QEX invites you to share your ideas and comments with fellow hams. Send them to "QEX Letters to the Editor" at c/o ARRL 225 Main St, Newington, CT 06111-1494; e-mail to **rseverns@arrl .org**. Please include your name, call sign, complete mailing address, daytime telephone number and e-mail address on all correspondence. Whether praising or criticizing an item, please send the author(s) a copy of your comments.

Predicting the Performance of Centrifugal Fans for Valve Cooling

Is that surplus fan adequate to cool your expensive transmitting tube?

By D. R. Kirkby, G8WRB

Most valves with an anode dissipation between 100 W and 20 kW require forced air cooling. The minimum airflow rate is stated on the valve data sheet and the back pressure that must be overcome is also stated. The valve's cooling requirements change with different air inlet temperatures and altitudes[1,2] and, of course, valve dissipation. Fan performance also changes with altitude.[1,2] When selecting a fan for a professional application, the best approach would be to compare the valves air-flow rate requirements to that of data sheets from fan manufacturers, then select an appropriate fan that can provide at least the required flow rate at the required back pressure. For an amateur, the cost of a new fan may be considered prohibitive, so a surplus fan is used, for which full data is often not available when one is seen for sale. Here we look at methods of estimating the performance of a fan from its physical dimensions and motor speed, so allowing one to make an educated guess as to whether a fan is suitable before purchasing.

While the method proposed is not 100% accurate, it is better than trying to estimate the size of a blower from blowers on other equipment, since these are often too small. John Nelson, GW4FRX, mentions this problem in the *VHF/UHF DX Book*[3] in both amateur-built and commercially produced equipment for amateur radio use, but even equipment designed for the professional market suffers this problem. Many dealers selling fans say they were stripped from equipment using large valves, such as the Eimac 4CX1000A (1000-W dissipation) or 3CX1500A7 (1500-W dissipation), when in fact they are totally inadequate to cool such a valve properly. So don't assume a fan stripped from a 2-kW amplifier will be adequate for your 500-W amplifier.

[1]Notes appear on page 31.

Stokes Hall Lodge
Burnham Road
Althorne, Essex CM3 6DT England

Flow Rate and Back Pressure

In a typical amplifier, using a valve such as the ever popular Eimac 4CX250B, the grid compartment is pressurised by a fan. The pressure difference between the grid compartment (above atmospheric pressure) and the anode compartment then forces the air past the grid and cathode, cooling the valve-base seals as it passes through the valve base, up through the chimney, through the anode cooler, and into the anode compartment. The pressure can easily be measured[2,3,4] with a manometer (clear plastic tube filled with water), as shown in Fig 1, by measuring the height of a column of water the air pressure can support. The manometer tube should enter the grid compartment at right angles to the air flow, and should be flush with the surface, so as to not disturb the air flow. As you increase the air flow through a valve (ie, more cubic feet per minute), the back pressure rises with approximately the square of the flow rate. So if a valve requires a flow rate of F_1 cfm at a back pressure of BP_1 inches of water, then at an air-flow rate of F_2 cfm, the back pressure BP_2 will be approximately:

$$BP_2 = BP_1 \left(\frac{F_2}{F_1} \right)^2 \quad \text{Eq 1}$$

So doubling the air flow rate requires approximately four times the pressure. Fig 2 (dashed line) shows a graph of the air-flow rate and back pressure of a pair of 4CX250B valves. Also shown is the characteristic of a hypothetical fan (solid line). This fan produces a maximum flow rate of 191 cfm when its output is unrestricted and zero flow when the back pressure is 1.39 inch of water. These two curves overlap at some point, which is shown at about 15 cfm and 1.34 inch of water. So this fan will produce about 15 cfm through the pair of valves and the back pressure will be 1.34 inch of water.

Strictly speaking, in the SI system of units, air-flow rate should be measured in cubic meters per second and pressure in Newtons per meter squared (Pascals). Eimac, the largest power valve manufacturer, give minimum air-flow requirements in cubic feet per minute (1 cfm = 4.72×10^{-4} m^3/s) and back pressures in inches of water (1 inch of water = 249.07 Pa). Hence, this paper uses units of cfm and inches of water since these are the most useful. It is also easy to measure the back pressure directly in inches of water as shown in Fig 1, but how many people have a pressure gauge calibrated in Pascals?

Fan Characteristics

Fans are incredibly complicated and the performance will depend on many things, such as the pitch of the blades, blade roughness, clearance between blades and casing, etc, so we can not expect to make an accurate judgement. However, a book on fans by Osborne[5] derives a few *Fan Laws* that relate the air-flow rate, back pressure and mechanical drive power of a fan to its dimensions and motor speed. Daly[6] gives an expression for the noise power, which is valid only for centrifugal fans, unlike the other three laws. These laws are intended to allow comparisons of the performance of identically designed fans with different speeds and diameters. These fan laws have been simplified here since it is assumed that the density of air at the outlet of the fan is not significantly different from the density of air at the inlet. This is a valid assumption for the small pressures (relative to atmospheric) developed in cooling even the largest valves amateurs use. (A Standard atmosphere is defined as 760 mm of mercury, which is approximately 29.92 inches. Mercury has a relative density of 13.53, so one atmosphere is $29.92 \times 13.53 = 404.8$ inches of water pressure. The one or two inches of water pressure needed to cool even fairly large tubes is negligible compared to the 404.8 inches of water pressure of the atmosphere.)

Free Flow Rate

The free air-flow rate (ie, flow with no restriction) is related to the fan characteristics by:

$$F_{FA} = k_1 \, N \, w \, d^3 \quad \text{Eq 2}$$

where k_1 is a constant, N is the fan speed (revs/minute), w is the fan wheel width (inches), d is the outside wheel

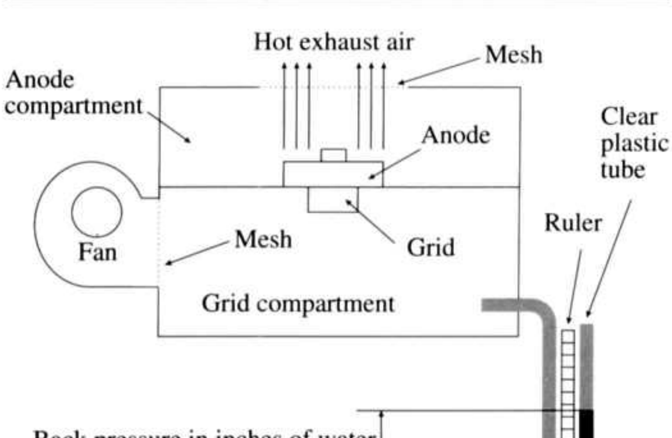

Fig 1—Diagram showing the basic method of cooling a forced air cooled valve and of measuring back-pressure.

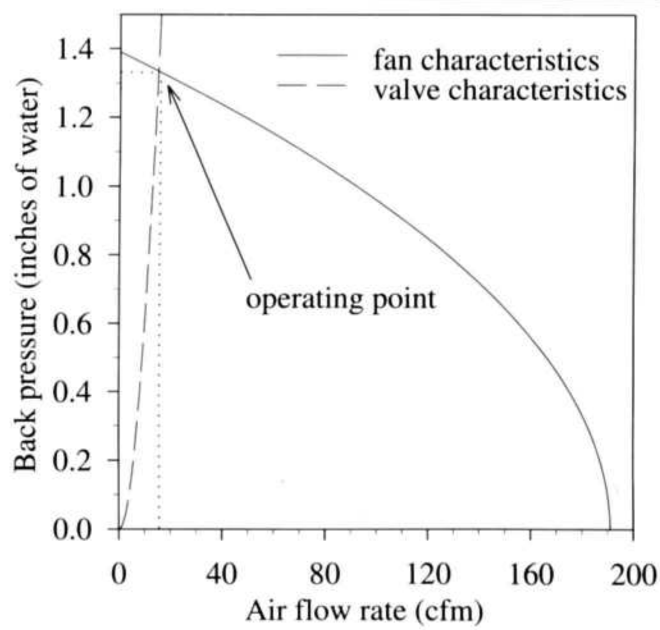

Fig 2—Graph showing a fan performance curve and the air flow into a pair of 4CX250Bs. They cross at the operating point (15 cfm, 1.34 inches of water).

diameter (inches) and F_{FA} is the free flow rate (cfm). ("Wheel" is the "cage" part of a squirrel-cage blower.—*Ed.*)

Back Pressure

The cut-off back pressure BP_{CO} (ie, the back pressure which will stop all flow) will be given by:

$$BP_{CO} = k_2 \ N^2 d^2 \qquad \text{Eq 3}$$

where k_2 is a constant.

Mechanical Drive Power Required

One of the fan laws states that the power required to rotate a fan increases as:

$$P = k_3 \ N^3 \ d^5 \ w \qquad \text{Eq 4}$$

From Eqs 2, 3 and 4, you can easily see that the drive power is related to the free air-flow rate F_{FA} and the cut-off back pressure BP_{CO} by:

$$P = k_3 \frac{BP_{CO}}{k_1} \frac{F_{FA}}{k_2} \qquad \text{Eq 5}$$

therefore,

$$P = k_4 BP_{CO} F_{FA} \qquad \text{Eq 6}$$

where k_4 is another constant. This shows that fans that are capable of either a large pressure or large flow rate, require powerful motors to drive them.

Noise Power

Fans are noisy, since moving a lot of air always creates noise. According to Osborne[5], most of the noise is air, not mechanical. However, the noise power increases rapidly as fan speed increases, or diameter increases, as the following expression shows, which is taken from Daly[6], although I have assumed the noise increases in proportion to the fan width, which was not explicitly stated.

$$\text{Noise power} = k_5 \ N^5 \ d^7 \ w \qquad \text{Eq 7}$$

where k_5 is another constant.

Numeric Values for k_1 and k_2

The equations above are of little use in selecting a fan since we must know the constants k_1 and k_2. The constants k_3, k_4 and k_5 are less important for our purposes. The constants k_1, k_2 were therefore evaluated for 11 single inlet centrifugal fans (ie, the conventional centrifugal fans) and six duplex fans (two wheels on the one motor shaft)(in the case of duplex fans, the wheel width considered is the total wheel width—ie, twice the width of an individual wheel), as well as three fans with double inlets, manufactured by Airflow Developments Ltd of High Wycombe, Buckinghamshire, England. The smallest considered was the model 21 ATXL capable of supplying 12 cfm in free air and capable of a blocked off back pressure of 0.13 inch of water. This runs at 2700 rpm with a wheel width of 1 inch (25.4 mm), a wheel diameter of 2.1 inches (54 mm) and is fitted with a 17-W motor. The largest considered was a model 126 HW capable of supplying 4000 cfm in free air, has a cut-off back pressure of 3.5 inches of water. This runs at 1420 rpm with a wheel width of 5.6 inches (143 mm), wheel diameter 12.7 inches (324 mm) and is fitted with a 4-pole, 3-phase, 4-kW (5.2 hp) motor. In between these, some of the fans used 6-pole (approximately 900 rpm on 50-Hz mains) motors. In effect, the range of fans considered was much larger than those necessary for amateur service in cooling valves. The smallest is incapable of cooling a 2C39A or 3CX100A5 triode (100-W dissipation), whereas the largest would easily cool the 12-kW dissipation 3CX10000H3 industrial triode. The three types of fans were initially considered separately, but it soon became apparent that this was not necessary as the differences between fan types was no greater than the differences between individual fans of the same type. The average (mean) values of k_1 and k_2 were then calculated, along with their standard deviations, so we can see how much variation there is between fans. The mean value of k_1 was 3.11×10^{-4} ft^3 $inch^{-4}$ with a standard deviation (σ) of 6.13×10^{-5} ft^3 $inch^{-4}$ and the mean value of k_2 was 9.54×10^{-9} inch of water $inch^{-2}$ min^2 with a standard deviation of 1.71×10^{-9} inch of water $inch^{-2}$ min^2. These values assume the fan dimensions are measured in inches. If you prefer to measure the fans dimensions in mm rather than inches, then $k_1 = 7.46 \times 10^{-10}$ with a standard deviation (σ) of 1.47×10^{-10} and $k_2 = 1.48$ 10^{-11} with a standard deviation (σ) of 2.65×10^{-12}.

The units of k_1 and k_2 are rather strange, but for those like myself, who usually prefer SI units, (fan dimensions in m, fan speed in s^{-1}, flow in $m^3 s^{-1}$, pressure in Nm^{-2}), then the units are more sensible. Then $k_1 = 21.12$ m^{-1} and $k_2 = 13.22$ Nm^{-4} s^2.

Air-flow rate at a Back Pressure less than Cutoff

A fan will generally produce maximum air-flow rate when in free air, falling to zero when it is blocked completely, as shown in Fig 2. Valve cooling will use the fan somewhere between these two limits, so it begs the obvious question: "How much air will it produce at a given back pressure greater than 0 but less than BP_{CO}?" The answer is impossible to estimate accurately without knowing more about the fan—we really need the full-back pressure vs flow-rate curve from the fan manufacturer. Some fans produce a steadily decreasing flow rate with increased back pressure (like Fig 2), but others can produce more flow at a high back pressure than at a low one! However, let's assume the air flow decreases with increasing back pressure according to:

$$F = F_{FA}\left(1 - \frac{BP^2}{BP_{CO}^2}\right) \qquad (\text{for BP} < \text{BP}_{CO}) \qquad \text{Eq 8}$$

where F is the flow, F_{FA} is the free air-flow rate ($F_{FA} > F$), BP is the back pressure developed. Many fans approximate this characteristic. So by calculating the free air flow F_{FA} and the cut-off back pressure BP_{CO}, we can find the flow for any given back pressure BP using Eq 8.

Air Flow and Back Pressure in an Application

Assuming the square-law relationship of Eq 1, between pressure and flow through the our valve, we would expect the flow rate through the amplifier system (including valve and filters) at any pressure BP to be:

$$F = F_{VA}\sqrt{\frac{BP}{BP_{SY}}} \qquad \text{Eq 9}$$

where BP_{SY} is the system back pressure at the flow rate F_{VA} needed for satisfactory cooling, which is generally considered 20%[1] above the back pressure for the valve BP_{VA}, to allow for some small pressure drop across pieces of RF screening mesh and simple filters.

We have seen that it is possible to calculate the flow rate from the fan at any given back pressure using Eq 8, and also the air flow through the amplifier system at any pressure using Eq 9. If Eqs 8 and 9 were solved as a pair of simultaneous equations, we could find the back pressure that would be developed when the fan was used to blow the system consisting of the valve, RF mesh and any simple filters. This occurs when the system characteristics and the fan performance overlaps as shown in Fig 2.

After my mathematical skills failed me, a computer program *Mathematica*[7] was used in an attempt to find a formula for the flow that will pass through the valve with the fan. *Mathematica* managed to find an expression, but it was so long that it would be impossible to implement without a computer, and then it would require considerable programming as the equation took many pages of printed output. An easier option, given a computer or programmable calculator, would be to evaluate the air flow the fan can provide for various back pressures *BP* starting at zero up to the maximum of BP_{CO}, in steps of 0.01 inch of water, using Eq 8, then calculate the flow through the valve for the same back pressures from Eq 9. (Of course, the brute-force approach of solving Eqs 8 and 9 by trying lots of solutions like this is not very efficient, but this problem does not warrant any more sophisticated numerical technique.) At some particular back pressure, the two flows will be equal, then we know the back pressure the fan will develop cooling the valve and the flow rate. This is better shown by example.

An Example of Cooling Two 4CX250Bs

Let's assume we wish to find a fan to cool a pair of Eimac 4CX250B valves in a two-valve amplifier such as the well-known design by W1QVF and W1HDQ.[8] The 4CX250B valve data sheet specifies we need a flow rate of 6.4 cfm (F_{VA}=6.4) at 0.82 inch of water (BP_{VA}=0.82) for sea level at an incoming air temperature of up to 50°C for a single 4CX250B. Assume we have seen for sale a used centrifugal fan with a 4.65-inch (118-mm) diameter wheel (d=4.65), 2.36-inch (60-mm) wheel width (w=2.36), rotating at 2600 rpm (N=2600), but don't have the full data on the fan. Can this cool our pair of 4CX250Bs at 250-W dissipation per valve at sea level?

Assuming a 20% extra pressure drop in our system, for ducts, mesh and filters, we need 6.4 cfm at 1.2×0.82=0.98 inch of water (BP_{SY}=0.98) for one valve. For two valves the flow must be doubled to 12.8 cfm (F_{VA}=12.8), but the pressure will remain the same at 0.98 inch of water. Using the formula in Eq 2, we estimate the fan will produce a free air flow F_{FA} of:

$$F_{FA} = k_1\, N\, w\, d^3 = 3.11 \times 10^{-4} \times 2600 \times 2.36 \times 4.65^3 = 191 \text{ cfm} \qquad \text{Eq 10}$$

Using Eq 3, our estimate for the cut off back pressure of the fan will be:

$$BP_{CO} = k_2\, N^2\, d^2 = 9.54 \times 10^{-9} \times 2600^2 \times 4.65^2 = 1.39'' \text{ of water} \qquad \text{Eq 11}$$

Before numerically solving Eqs 8 and 9, which is a lengthy procedure without a computer, an estimate of the flow that this fan will produce at a system back pressure of 0.98 inch of water is first found, to check that it is at least F_{VA} (12.8 cfm).

$$F = F_{FA}\left(1 - \frac{BP_{SY}^2}{BP_{CO}^2}\right) = 191\left(1 - \frac{0.98^2}{1.34^2}\right) = 96 \text{ cfm} \qquad \text{Eq 12}$$

This is above 12.8 cfm, so a larger fan is unnecessary. We need not do any more calculations since the fan would appear to be adequate, although we do not yet know how much air it will pass—only that we estimate it will exceed 12.8 cfm. Hence we have determined easily that the fan is adequate for the job.

The flow in our system will not be 96 cfm since any attempt to increase the flow above 12.8 cfm will increase the back pressure above 0.98 inch of water. We can, if we wish, calculate the flow at various back pressures until the flow the valve will pass is equal to the same flow rate the fan can produce. The flows at back pressures of 0, 0.5, 0.98, 1.34, 1.38 and 1.39 inch of water are shown below. At a back pressure of 1.34 inch of water, the fan can provide 14.93 cfm, but the valve and filters will pass 14.93 cfm, so this is the operating point for this fan/amplifier combination. This data is shown in Table 1, and also in Fig 2. This is easy to do in a computer program, but is unnecessary unless you wish to know the flow rate.

Testing a Fan

Once a fan seems to fit the needs, as outlined here, it can be purchased. It can then be tested either in the complete system or a mock-up with nothing more than a cardboard box[3] and RF screening mesh, to ascertain the valve does produce sufficient back pressure. Measuring flow is more difficult (one way to estimate the flow in a working amplifier is to measure the difference between the air inlet and outlet air temperatures at a known valve dissipation. Flow (cfm) = 1.76 × dissipation (*W*) / temperature difference (°C). See the referenced Burle application note for full details), but this is largely unnecessary since if the back pressure is adequate, assuming we are using the correct bases, the flow will be too. If the back pressure seems okay, the system can be assembled properly, then the temperature of the valves checked[1,2], as this is ultimately the best way of checking for sufficient cooling air.

Results from an Actual Measurement

I wished to cool a single Eimac 3CX5000A7 triode. After making allowances for a typographical error in the 3CX5000A7 data sheet, we find that at sea level, with an incoming air temperature of 25°C, the valve requires 181 cfm (F_{VA}=181 cfm) and will cause a pressure drop of 1.7 inch of water (BP_{VA}=1.7 inch of water) for the full 5-kW dissipation. Allowing 20% for pressure drops in the system other than across the valve ($BP_{SY} = 1.2 \times BP_{VA} = 2.0$ inches of water), we need a fan which can provide 181 cfm at a pressure of 2.0 inches of water. A friend, G8WYI, offered

Table 1

Back pressure (inches of water)	*Fan flow (Eq 8)*	*Valve flow (Eq 9)*	*Notes*
0.00	191 cfm	0.0 cfm	Fan runs flat out at 191 cfm.
0.50	166.6 cfm	9.1 cfm	
0.98	96.6 cfm	12.8 cfm	Valve would pass 12.8 cfm, if a pressure of 0.98 inch of water is used.
1.34	14.9 cfm	14.9 cfm	Operating point
1.38	3.6 cfm	15.2 cfm	
1.39	0.0 cfm	15.2 cfm	Fan cannot provide any flow at 1.39 inches of water

me a fan which had a speed of 1300 rpm, a wheel diameter of 10 inches (253 mm) and a wheel width of 9.6 inches (243 mm). This was from a different manufacturer than the fans used in evaluating the constants k_1 and k_2.

Using Eq 2 and 3, we can calculate F_{FA}=3816 cfm and BP_{CO}=1.6 inches of water for the fan. Since BP_{CO} is less than needed, we immediately know the fan is too small. Solving Eqs 8 and 9 iteratively, we calculate a flow of 150 cfm would result at a back pressure of 1.57 inches of water using this fan. However, the fan was tested in an arrangement similar to that in Fig 1 to be certain of this. First, with the box sealed completely, BP_{CO} was measured at 1.85 inches of water, some 16% higher than expected. Next, a hole was made to mount the valve, and the chimney fitted to direct the air into the anode cooler. With air flowing through the valve, the back pressure fell to 1.6 inches of water—some 2% higher than expected.

Having no method to measure flow rate, I was unable to determine how close the flow was to the estimated 150 cfm. However, using Eq 9, the flow at the measured back pressure of 1.6 inches of water would be

$$F = F_{VA}\sqrt{\frac{BP}{BP_{SY}}} = 181\sqrt{\frac{1.6}{2.0}} = 162\ cfm \qquad \text{Eq 13}$$

So the fan is putting approximately 162 cfm through the valve at a back pressure of 1.6 inches of water—not too far from the 150 cfm at 1.57 estimated with the aid of a ruler and computer program. I have now obtained a suitable fan to cool the 3CX5000A7. This has a 2800 RPM motor and a 10-inch diameter, 1-inch-thick fan wheel and weighs 72 pounds (32.7 kg). It is expected to provide around 313 cfm at 6.0 inches of water pressure—some 70% excess flow, which is a very comfortable, but not excessive, safety factor.

Discussion

Some fans, especially larger industrial units, may have regions of the flow/pressure curve where they must not be operated. there may be insufficient airflow to cool the motor, too much load on the motor or unstable airflow, which would damage the fan blades. In particular, many fans should not be discharged straight to the atmosphere—it puts too much load on the motor. This can only be ascertained from the fan data sheet, so make every effort to obtain the data sheet.

The constants k_1 and k_2 are only valid for centrifugal fans. The basic theory should work for any fan, but new values for these constants will have to be ascertained if it is required to use the technique for other types of fans. Since there is a significant spread of actual values for k_1 and k_2 for individual fans, the approximates given here should be treated as such—just approximations, and so some extra should be allowed for the air flow. The fan that was measured had a cut-off back pressure within 16% to that calculated—some fans will almost certainly be predicted less accurately.

Large valves with large diameter anode coolers generate less back pressure for a given air flow than do small valves. Hence, although a fan may have cooled a large 2-kW dissipation valve dissipating 1 kW, it does not necessarily follow that it will be able to cool a smaller 500-W valve dissipating its full 500 W. This is perhaps counter-intuitive, so it pays to calculate and not guess.

The small fan used as an example in cooling 4CX250Bs provided 191 cfm in free air, but when two 4CX250B valves are put in front of it, this falls to 14.9 cfm. That is a reduction by a factor of 12.8:1! The larger fans output fell from 3816 cfm to less than one twentieth of this when a 3CX5000A7 is put in front of it. Hence, free air-flow rates are meaningless on their own, for our purposes.

Putting two identical small fans in parallel will not double the air flow through a valve. It will produce twice the free air flow F_{FA}, but for a given valve, the flow will increase much less. For the example given of the two 4CX250s, using two identical fans will increase the air flow from 14.93 cfm to 15.08 cfm—an increase of only 1%, as can be shown if you try the calculations with a fan of double the width—120 mm (4.72 inches) instead of 60 mm (2.36 inches). Generally speaking, when cooling valves with their high back pressures, putting two fans in parallel is a complete waste of time. For applications where the back pressures are much smaller, this will not be true. Putting two identical fans in series will double the cut-off back pressure, but keep the free air flow rate the same. Using again the example of the two 4CX250Bs, this would increase the air flow to 21.0 cfm, which is a 41% increase compared to one fan.

Remember to put some method of detecting fan failure into your amplifier, such as an air operated switch, temperature sensor, etc. My favorite is a band of plastic material looped around the anode(s) and a microswitch, with the plastic held in tension by a spring fastened to both ends. If the valve anode(s) overheats (for any reason—not just fan failure), the plastic melts, the microswitch opens and the amplifier shuts down. Use an adjustable soldering iron or oven to find a material with a suitable melting point.

Conclusions

A method has been presented for estimating the cooling capabilities of fans that is more accurate than pure guess work, yet simple enough to be used quickly at junk sales when a fan is seen that you consider using for your next project. Finding an approximation of the actual flow rate for a fan/valve combination requires a programmable calculator or computer. However, determining if a fan will provide at least the required air flow can be done with a nonprogrammable calculator in less than a minute. I hope to have proved that guessing the size of a fan to use can be dangerous, since some facts are about as intuitive as quantum mechanics—ie, not very.

Notes

[1]"Application Guide for Forced Air Cooling of BURLE Power Tubes," TP-118 Application note, Burle Technologies Inc, 1000 New Holland Avenue, Lancaster, PA 17601-5688, USA, 1989.

[2]*Care and Feeding of Power Grid Tubes*, Eimac division of Varian, 310 Industrial Way, San Carlos, CA 94070, USA, 1978.

[3]White, I., *The VHF/UHF DX Book*, Chapter 6, DIR Publishing, Buckingham, England, 1992.

[4]*The Radio Amateurs Handbook*, 95th edition, Chapter 6, American Radio Relay League, Newington, CT 06111, USA, 1982.

[5]Osborne, W. C., *Fans*, Pergamon Press, Oxford, UK, 1977.

[6]Daly, B. B., *Woods Practical Guide to Fan Engineering*, Third edition, Published by Woods of Colchester Limited, 1978.

[7]*Mathematica*, Produced by Wolfram Research Inc, 100 Trade Drive, Champaign, IL, 51820-7237, USA.

[8]McMullen, T. F. and Tilton, E. P., "New ideas for the 2-Meter Kilowatt," *QST*, Feb 1971, pp 24-30.