

# On-air Measurements of CLOVER P38 Throughput

Ken Wickwire (KB1JY), Mike Bernock (KB1PZ) and Bob Levreault (W1IMM)

## 1. Introduction

This paper is one of a series treating on-air measurement of throughput of various HF data-transmission protocols available to amateurs (see the references to our other reports at the end of the paper). Here we describe an extensive set of measurements of throughput for text and other files sent over near-vertical-incidence-skywave (NVIS) and one-hop skywave (OHS) paths using the file transfer protocols implemented in the HAL P38-CLOVER terminal package. (NVIS paths, which often produce difficult channel conditions, are used to communicate over 20- to 300-mile ground distances using antennas that can launch energy at high takeoff angles.) The measured throughput data in our experiments were analyzed using software specially written to compute throughput statistics from our CLOVER data.

The HAL CLOVER modem and data-transmission protocols have now been in use for about five years and the system has established itself as one of the most reliable, inexpensive and efficient means of HF data transmission available to hams. Several versions of CLOVER are available (the PCI4000, DSP4100, P38, and others). The versions are distinguished mainly by their maximum speed of data transmission (and corresponding modulation complexity), which is a function of the processing power of the associated hardware (computer or modem or both).

The P38 modem is a low-cost (less than \$400) PC-card version of CLOVER that has a slower TI microprocessor than the Motorola processor used by its more expensive (around \$1000) relatives, the PCI4000 and DSP4100. The P38 does not have the computing power to process received signals in the 8P2A (8-phase, 2-amplitude) and 16P4A (16-phase, 4-amplitude) signaling modes that the other versions can use. Otherwise, the P38 can do everything the PCI4000 and DSP4 100 can do, including uncompressed and compressed file transfers using the CLOVER automatic repeat request (ARQ) protocol. Using the compressed mode, one can send graphics and other “binary” (eight-bit-character) file data, in addition to text files.

CLOVER changes its modulation mode (varying the number of bits per symbol while keeping the symbol *rate* fixed at 3 1.25 symbols per second) in response to changing channel conditions. Channel conditions are assessed (at receivers) by measuring the amount of error-correcting capacity [ECC] used, the number of erroneous data blocks and the phase dispersion [PHS], among other things. Receivers send this information to transmitters at turnarounds, and in Version 23 of the firmware, the transmitter makes decisions about changing the mode using the three items above. This adaptability, combined with the underlying Reed-Solomon forward error correction (FEC) coding, and, when necessary, retransmission of erroneous frames, is the key to CLOVER’s reputation as one of the two fastest and most reliable HF data communication protocols available to hams. (The other is the PacTOR II system, which we have not yet tested.)

For further details on how CLOVER works, see the HAL documentation supplied with the various versions and the descriptions published by Ray Petit and Bill Henry in the *RTTY Journal*, *QST*, *QEX* and elsewhere between 1990 and 1994.

The P38 has been said to come close to the throughput of its more powerful CLOVER relatives on the basis of the claim that only the very best (high signal-to-noise ratio and little or no multipath) HF channels allow use of the 8P2A and 16P4A signalling constellations. To the extent that this claim is valid for our channels (it might apply on occasion to our NVIS paths), our results may describe CLOVER throughput performance in general. On very good paths, however, you might expect to see average throughput up to 25% higher than what we have measured if you use the PC14000 or DSP4100.

The NVIS paths we have studied, which are thirty to sixty miles long, frequently display strong multipath, high local and propagated noise, D-layer absorption at mid-day and occasionally strong interference from other stations operating in both voice and digital modes. (Horizontal antenna polarization at all our NVIS stations allowed us to be fairly certain we were using NVIS rather than surfacewave propagation, and this was confirmed by the fading we observed most of the time.)

We measured one-hop skywave throughput on a 600-mile path. This path produced less multipath than our three NVIS links and therefore somewhat higher throughput (see below).

The majority of NVIS measurements were at 3.6155 MHz LSB, with some at 1.8 15 MHz LSB. The OHS measurements were done at 10.141 MHz LSB. Output power (about 100 watts) and antennas were typical of those used by hams. Measurements were made over all daylight hours and a few were made in the evening while the maximum usable frequency (MUF) was still above the operating frequency. The NVIS tests covered the twelve-month period from May 1996 to May 1997. The OHS tests covered the eight-month period from October 1996 to June 1997. The average sunspot number during these periods was about ten, so MUFs were near their cyclical minima.

The rest of the paper describes the paths between stations and layout of antennas, the HAL P38 interface and file-transfer operation, the file types sent, the recorded data format, the statistical analysis software, a statistical summary of the data, a discussion of the statistical results and concluding remarks.

## **2. Layout of Paths and Discussion of Antennas**

The stations used for the NVIS tests are in Bedford, Mass. (KB 1 JY), Norfolk, Mass. (WI IMM) and Derry, N.H. (KB1PZ). Bedford used an 80m dipole up 30 feet for most tests, and occasionally a terminated, bottom-fed 125-ft longwire pointing southwest (for NVIS tests on 160m). Norfolk used an 80m dipole up 40 feet or a longwire for 160m. Derry used an off-center-fed 80m dipole up 30 feet. The stations in the OHS tests are in Bedford, Mass. (KB 1 JY) and Raleigh, N.C. (KF4KL). In the OHS tests Bedford used the resistively terminated sloping

longwire running southwest. Raleigh used a 135-foot dipole running NW-SE and fed with ladder-line. The links (followed by lengths and rough estimates of the percentage of data collected over each link) are

#### NVIS

Bedford-Norfolk (35 miles, 45%)

Bedford-Derry (25 miles, 45%)

Derry-Norfolk (60 miles, 10%)

#### OHS

Bedford-Raleigh (600 miles, 100%)

(All of these links run more or less north-south.)

### 3. The HAL P38 User Interface and CLOVER File transfer

Although there are graphical user interfaces (GUIs) for running CLOVER that may be more sophisticated than the DOS-based one provided by HAL (Express and XPWIN, for example), the HAL GUI is presently the best one to use for making throughput measurements because it gives the transfer time of compressed file transfers. During the course of our tests we moved through versions 20 to 23 of the P38 firmware. (Some of the firmware upgrades made improvements in CLOVER's ARQ scheme.) The majority of our tests used Versions 21 and 22. The improvements appear to have been in the nature of fine-tuning since we have not noticed large increases in performance.

\*

To send a file with the HAL GUI, one first establishes a CLOVER ARQ link with the receiving station by selecting (or entering) the callsign of the station and sending a carriage return.

Once the link is established (which is usually confirmed in a few seconds by a LINKED message on the GUI and the appearance of "HIS" channel statistics on the display), one can enter the name of a file to be transferred. (The file needs to be in the P38 directory.) Since distribution of Revision 22 of the P38 firmware, channel statistics can be recorded during file transfers; these often provide fascinating insight into the workings of a transfer (see the references at the end for our paper on CLOVER channel statistics).

After the filename has been entered, one has a choice of sending a text file from the TX Buffer in uncompressed mode or sending any file from the P38 directory ("Send from Disk") in compressed mode. Only generic text can be sent in uncompressed mode; files with control characters in them will abort a transfer attempt. Files sent in compressed mode can have any format, although some files may not benefit from compression, and may even be expanded slightly by the compression process (more on this below).

Some people view the throughput of uncompressed files as the inherent, or "true" performance of the CLOVER (or any other) waveform, error-control scheme and file-transfer system. Others (such as ourselves) consider compression to be part of a protocol if it is provided as a constituent of the "common" interface provided with a modem's hardware. To satisfy both camps, we have measured throughput of both compressed and uncompressed files.

This opened a small box of bother, because it turned out that how much a file gets compressed by the HAL compression software (a dictionary-based algorithm contained in the PKWARE Library) depends on what kind of a file it is. “Data” files (defined as text files with regular structure) are compressed more than arbitrary text files, and certain graphics files with large monochromatic parts can be compressed to much less than 50% of their original size. (See below.)

To measure the throughput of uncompressed (text) files, we had to record transfer time in our transfer-data files by hand, since the HAL software does not display transfer times of data sent from the text buffer. (This was not as onerous as it sounds since once we decided on a recording format we were able copy, paste and edit with a text editor to speed up data recording. We could, as an alternative, have written a program using HAL’s control-code specification and developer guidelines to automate throughput testing, but we didn’t have the leisure for that.)

To measure the throughput of compressed (text or other) files, we took advantage of the HAL software’s calculation and display of file transfer time on the GUI screen. Through experimentation we deduced that the optimal (highest throughput with smallest test time) file size for throughput assessment of CLOVER was between 10 and 30k bytes, and most of our files had sizes in that range.

Precompressed files (for example, GZIP-ed ones) may be sent by CLOVER in even smaller form than those sent compressed only by CLOVER’s PKWARE algorithm, but this does not correspond to use of the “common” implementation. Nevertheless, precompression in HF data communication is a very useful technique that deserves further study.

#### 4. File Types Sent

We have sent uncompressed and compressed text files, uncompressed and compressed “data” files, graphics files and “hybrid” files. Data files are defined to be fairly rigidly formatted text files that have a relatively large number of repeated strings and relatively short lines. An example of such a file is a CLOVER channel statistics file (recorded on command by the HAL interface). A snippet of a channel stats file is given below.

```
034601,KB1JY      ,BPSM  ,2,02,000, 000,000,000,000
034601,W1IMM     ,BPSM  ,2,02,031, 003,040,050,000
034604,KB1JY     ,BPSM  ,2,02,000, 000,000,000,000
034604,W1IMM     ,BPSM  ,2,02,028, 005,045,100,000
034606,KB1JY     ,BPSM  ,2,02,036,-6  ,054,000,000
034606,W1IMM     ,BPSM  ,2,02,026, 003,042,050,000
034609,KB1JY     ,BPSM  ,2,02,033,-3   ,049,000,000
034609,W1IMM     ,BPSM  ,2,02,026, 002,044,000,000
```

Graphics files come in many formats (\*.bmp, \*.gif, pict, etc.). Some formats (JPEG, MOV) usually come already compressed and get bigger when sent by CLOVER. Other graphics

formats (\*.bmp, for example) can be compressed by CLOVER's PKWARE library and sometimes get much smaller (with resulting very high throughput) if they have large one-color sections with no detail in them. Generally speaking, the more detailed in layout and color a graphics file is, the less it can be compressed.

"Hybrid" files are those we define as having combinations of text and unprintable control characters. Examples of such files are Microsoft Word and Microsoft Excel files. Because the effect of compression on a hybrid file depends on both its layout and character content, the percent of compression of a hybrid file can be hard to predict. In our discussion of statistical results we will comment on the effect of file type on compression ratio and throughput.

Some files, like zipped files and executables, apparently look to most compression schemes like random sequences of characters and cannot be compressed much by the algorithms of PKWARE. These files may even be expanded somewhat by addition of useless compression-implementing strings. We did not include such files among those we sent. Their throughput would probably be about the same as that of uncompressed text files of similar size (see below).

## 5. Recorded Throughput-Data Format

The data archive file into which the results of each transfer are entered contains the file type ("g" for graphics, "d" for data, "h" for hybrid; see above), date-time group, callsign of receiving station, callsign of sender, CLOVER interface (P38 = HAL), uncompressed file size in bytes, compressed file size in bytes, transfer time in seconds, predominant waveform used by CLOVER during the transfer, HF frequency, observed channel condition (Q = quiet, etc.) hand-calculated throughput in characters per second (cps, which doesn't have to be accurate since it's calculated later by the analysis software) and the CLOVER "bias" chosen by the sender (F[ast], N[ormal] or R[obust]). In the case of text (including "data") files sent from the TX buffer, the transfer time is read from the computer clock and may therefore be off by a second or two. This makes little difference to throughput calculations for files such as ours that take several minutes to send.

The bias is a parameter that describes the error-correcting overhead (and thus the number of data-bytes per ARQ block) chosen by the transmitting station's operator for a communications session. Robust bias is usually chosen when the channel appears (sounds) bad and fast bias when it appears to be good.

Here's an excerpt from the NVIS transfer-data file for tests run in May 1997:

```
h16.05.97 20:54:00 W1IMM KBLJY P38 37376 10063 332 8PSM 3.6155 Q [113_cps]F
h16.05.97 20:59:00 W1IMM KBLJY P38 18944 6476 229 8PSM 3.6155 Q [83_cps]F
h16.05.97 22:05:00 W1IMM KBLJY P38 37376 10063 333 8PSM 3.6155 Q [112_cps]F
h16.05.97 22:09:00 W1IMM KBLJY P38 18944 6476 246 8PSM 3.6155 Q [77_cps]F
h21.05.97 12:29:00 KBLPZ KBLJY P38 37376 10063 343 8PSM 3.6155 Q [109_cps]F
h21.05.97 12:33:00 KBLPZ KBLJY P38 18944 6476 215 8PSM 3.6155 Q [88_cps]F
h21.05.97 12:40:00 KBLPZ KBLJY P38 37376 10063 341 8PSM 3.6155 Q [110_cps]F
h21.05.97 12:44:00 KBLPZ KBLJY P38 18944 6476 217 8PSM 3.6155 Q [87_cps]F
d21.05.97 12:59:00 KBLPZ KBLJY P38 26148 26148 915 8PSM 3.6155 Q [29_cps]F
d21.05.97 13:04:00 KBLPZ KBLJY P38 26148 6502 228 8PSM 3.6155 Q [115_cps]F
```

This file is opened and analyzed by a data-analysis program described in the next section.

## 6. The Data-analysis Software

The results in the data archive are analyzed off-line by a program called `summary clo.c`. This program reads the archive file line-by-line looking for various strings. As it moves through the file to the end-of-file character, the program keeps running totals of throughput and other data corresponding to the strings, from which it calculates statistics such as the average and standard deviation of the throughput. The statistics are written to a summary file after the pass through the archive file. Switches in the summary code are set before each run to pick out specific data (corresponding to various string combinations) for analysis. For example, we select lines starting with the letter “d” to calculate the throughput for “data” files (see below). Since the summary program was written to analyze archive files of fixed format but arbitrary length, summaries of the data collected so far can be made at any time.

Shown below is the output of the summary program for all P38 NVIS tests run from May 1996 to May 1997. For this output we set the software switches to compute throughput statistics for *uncompressed text files*.

```
Statistical Summary of P38-CLOVER Throughput Tests:
21.06.97 21:34:00 UNCOMPRESSED TEXT

NUMBER OF P38 TRANSFERS IN SAMPLE = 193
E(FILE_SIZE) = 13917.9 bytes, E(COMPRESSED_SIZE) = 13917.9 bytes
E(TRANSFER TIME) = 652.3 s, sd(TRANSFER TIME) = 385.5 s
E(THRUPUT) = 24.40 cps, sd(THRUPUT) = 6.80 cps, sd(mean_THRUPUT) = 0.490 cps
MAXIMUM THRUPUT = 35.71 cps, E(THRUPUT/Hz) = 0.049 cps/Hz
E(COMPRESSON RATIO) = 100.00%, sd(COMPRESSON RATIO) = 0.00%
Lowest compression ratio = 100.00%; Compressed-size:File_size = 0:0
Highest compression ratio = 100.00%; Compressed-size:File_size = 10000:10000
NUMBER OF UNCOMPRESSED TRANSFERS IN SAMPLE = 193
```

The output shows that the average throughput for 193 uncompressed-text-file transfers was about 24 characters per second (cps) and that the largest observed throughput in this mode was about 36 cps. The `sd(THRUPUT)` reflects the spread of the throughput measurements about their average. Roughly speaking, about two-thirds of a set of measurements will be within one standard deviation (here 6.8 cps) of their mean and over 90% will be within two standard deviations of their mean.

We also calculate the “standard deviation of the mean throughput” [`sd(mean THRUPUT)`] in characters per second and the average throughput per Hertz of signaling bandwidth. The standard deviation of the mean throughput (equal to the standard deviation of the throughput divided by the square root of the sample size) is an assessment of the variability of the mean itself (which has its own statistical variability). The `sd(mean)` above suggests that our sample size in this case is big enough to give us pretty high confidence that if we collected many more throughput measurements under roughly the same conditions, we would not get an average throughput that differed from the one above by more than about half a character per second.

To estimate the average throughput per Hertz [ $E(\text{THRUPUT}/\text{Hz})$ ], we divide the average throughput by the signaling bandwidth. For P38 CLOVER, the signaling bandwidth is 500 Hz (see the Clover documentation).

The compression ratio is defined as the ratio of compressed size in bytes to uncompressed size in bytes. A ratio of 100% therefore means no compression.

Here's the corresponding statistical summary for *compressed text files* sent over NVIS links.

```
Statistical Summary of P38-CLOVER Throughput Tests:  
11.08.97 08:26:44 COMPRESSED TEXT
```

```
NUMBER OF P38 TRANSFERS IN SAMPLE = 266  
E(FILE_SIZE) = 19434.6 bytes, E(COMPRESSED_SIZE) = 9958.6 bytes  
E(TRANSFER TIME) = 468.1 s, sd(TRANSFER TIME) = 284.0 s  
E(THRUPUT) = 43.13 cps, sd(THRUPUT) = 11.30 cps, sd(mean_THRUPUT) = 0.693 cps  
MAXIMUM THRUPUT = 65.57 cps, E(THRUPUT/Hz) = 0.086 cps/Hz  
E(COMPRESSSION RATIO) = 51.41%, sd(COMPRESSSION RATIO) = 2.16%  
Lowest compression ratio = 46.76%; Compressed_size:File_size = 18316:39168  
Highest compression ratio = 69.90%; Compressed_size:File_size = 699:1000  
NUMBER OF UNCOMPRESSED TRANSFERS IN SAMPLE = 0
```

Note that the text files we sent in compressed mode were reduced on average to about half (51.4%) of their original size, a fact reflected in the average throughputs for compressed and uncompressed text files.

## 7. Statistical Summary of Throughput Results

The rest of the results of our NVIS and OHS tests (as of summer 1997) are summarized in Tables 1 and 2 below. The first column gives the average throughput and its standard deviation, the average throughput per Hertz, the standard deviation of the mean throughput and the maximum observed throughput. The second column gives the number of transfers in each case. The third column gives the mean and standard deviation of the compression ratio for compressed transfers. The fourth column gives the mean and standard deviation of the transfer time and the fifth column the average number of bytes in the original, uncompressed files.

**Table 1. Statistical Summary of P38-CLOVER NVIS Throughput Data**

<b>File Type</b>	<b>E(thruput) sd(thruput) E(tput/Hz) sd_mn(tput) max_tput</b>	<b>No. Xfers</b>	<b>E(Compr. Ratio) (CR) sd(CR)</b>	<b>E(xfer_tm) sd(xfer_tm)</b>	<b>E(No_char)</b>
<b>Uncompr. Text</b>	24.4 cps 6.8 cps 0.05 cps/Hz 0.49 cps 35.7 cps	193	—	652 s 386 s	13918
<b>Compr. Text</b>	43.1 cps 11.3 cps 0.09 cps/Hz 0.69 cps 65.6 cps	266	51.4% 2.2%	468 s 284 s	19435
<b>Uncompr. Data</b>	26.3 cps 6.1 cps 0.05 cps/Hz 0.53 cps 35.4 cps	130	—	704 s 329 s	17935
<b>Compr. Data</b>	89.8 cps 29.2 cps 0.18 cps/Hz 2.5 cps 149.9 cps	135	26.6% 6.8%	315 s 178 s	24526
<b>Graphics</b>	75.3 cps 62.6 cps 0.15 cps/Hz 5.6 cps 392.3 cps	124	40.4% 14.4%	583 s 418 s	42100
<b>Hybrid</b>	88.9 cps 24.2 cps 0.18 cps/Hz 2.6 cps 234.3 cps	85	29.8% 4.4%	336 s 106 s	29130



**Table 2. Statistical Summary of P38-CLOVER OHS Throughput Data**

<b>File Type</b>	<b>E(thruput) sd(thruput) E(tput/Hz) sd_mn(tput) max_tput</b>	<b>No. Xfers</b>	<b>E(Compr. Ratio) (CR) sd(CR)</b>	<b>E(xfer_tm) sd(xfer_tm)</b>	<b>E(No_char)</b>
<b>Uncompr. Text</b>	29.4 cps 4.0 cps 0.06 cps/Hz 0.57 cps 35.1 cps	49	—	542 s 206 s	15916
<b>Compr. Text</b>	50.4 cps 9.9 cps 0.10 cps/Hz 1.32 cps 62.0 cps	57	51.8% 3.1%	309 s 177s	16053
<b>Uncompr. Data</b>	30.0 cps 4.9 cps 0.06 cps/Hz 0.58 cps 50.6 cps	73	—	746 s 258 s	21640
<b>Compr. Data</b>	101.8 cps 19.8 cps 0.20 cps/Hz 2.3 cps 133.6 cps	71	25.1% 5.7%	236 s 85 s	22929
<b>Graphics</b>	100.6 cps 95.0 cps 0.20 cps/Hz 11.3 cps 396.4 cps	71	39.8% 14.5%	463 s 248 s	49255
<b>Hybrid</b>	93.5 cps 21.8 cps 0.19 cps/Hz 2.8 cps 197.5 cps	62	29.9% 4.8%	302 s 66 s	28449

## 8. Discussion of Results

Tables 1 and 2 show that the inherent (no compression used) over-the-air average throughput of P38 CLOVER on our links is around 24 cps on NVIS paths and 29 cps on the OHS path for daytime operations. (The standard deviations of these mean throughputs are about 0.5 and 0.6

cps, giving us high confidence that additional measurements made under the same conditions would yield nearly the same mean throughputs.)

At night, especially with the sunspot cycle near minimum, we have found that on our paths there is usually so much multipath, noise and interference on the small band of frequencies that lie below the MUF in the ham bands that throughput is generally significantly lower than the daytime values (most of the ones presented here). (An automatic link establishment [ALE] system, such as prescribed in MIL-STD-188-141A, and now widely available, can often find a useful, interference-free frequency even at night if given enough choices and a suitable-usually broadband-antenna.)

The average P38 throughputs for compressed text files are 43 and 50 cps on our NVIS and OHS paths, with standard deviations of mean throughput of about 0.7 and 1.3 cps. These average throughputs, and the average compression ratios of about 52% given in column 4 of the tables, show that CLOVER's PKWARE compression is squeezing our text files down to about half their uncompressed size. This, of course, doubles text-file throughput.

In contrast, PacTOR (with Huffman compression on) and GTOR (with its built-in compression) achieve average throughputs for text files of about 18 and 24 cps on NVIS links and 20 and 32 cps over OHS links (see Ref. 4). Thus, if data compression is viewed as an intrinsic part of what we have called (Ref. 4) the "common" implementation of PacTOR and GTOR, then P38 CLOVER in its common (PKWARE compression on) implementation has about twice the average throughput of PacTOR and GTOR for text-file transfers. Average P38-CLOVER throughput *without* compression is about the same as GTOR (with built-in compression) and a bit higher than PacTOR (with its Huffman compression turned on). P38-CLOVER throughput without compression is almost double that of PacTOR with Huffman off.

The throughput of uncompressed "data" files is about the same as that of ordinary text files, since these are indistinguishable to CLOVER's text-buffer transfer mode. However, because of their characteristic repetitions of strings such as placeholders for null-entries (see the CLOVER channel statistics excerpt given above as an example of a data file), and their generally short lines, data files usually get compressed by more than 50% (i.e., their compression ratios by our definition are less than 50%). The data files we have sent got compressed to about a quarter of their original size with resulting average throughputs of 90 and 102 cps over NVIS and OHS paths.

Our graphics files ranged in size from 9k to 230k bytes. (CLOVER forces one to send graphics files in compressed mode ("Send-from-Disk") since the uncompressed send-from-buffer mode can only handle text files. Some of these graphics files were not very complicated; others were relatively detailed. The average compression ratio for graphics files was about 40% and the average NVIS and OHS throughputs were 75 and 101 cps.

Recall that we define "hybrid" files as those with a combination of text and control (e.g., formatting) characters. Examples are word-processor and spreadsheet files. Our hybrid files ranged in size from 1 Ok to 37k bytes. (CLOVER also forces one to send such hybrid files in

compressed mode.) Because of their often complicated layout and unpredictable distribution of repeated characters, it's hard to guess compression ratios accurately for hybrid files. The average compression ratio of the ones we sent (Microsoft Word and Excel files) was about 30%. The NVIS and OHS throughputs of our hybrid files were 89 and 94 cps over NVIS and OHS paths.

The relatively large standard deviations of uncompressed and compressed file throughput for all file types (the standard deviations range from four to several tens of cps) reflect, in part, CLOVER's ability to adapt itself to changing channel conditions (see Sec. 1). However, these standard deviations are also affected by variability of file size (especially for graphics files), so channel adaptation should not be viewed as the sole source of throughput spread.

## 9. Concluding Remarks

We hope that our data will aid understanding of CLOVER file transfer over HF and perhaps serve as a useful introduction to how CLOVER works. CLOVER is now used all over the world by amateurs (in particular, for BBS mail forwarding). A number of international aid and other communications-providing organizations also use the P38 and other CLOVER versions (some with two-kHz bandwidth) to send medical and other information across parts of Africa and Australia, where alternative means of long-haul communication are not available, or too expensive. CLOVER modems are also used for at least two special-purpose e-mail systems (one used widely by private boating and shipping operators). Our data may shed light on why this inexpensive, amateur-developed modem and its protocols have become so popular.

## Acknowledgments

We are grateful to Doug Hall (KF4KL) for his customary reliability in keeping a P38 on the air in Raleigh, and to Bill Henry (K9GWT) and Drew White (K9CW) of HAL for useful comments on the paper. (The work was not supported financially by HAL.)

## References

1. Ken Wickwire (KB1 JY), "A Software Package for Analyzing CLOVER Channel Statistics." To appear in *QEX*, October 1997.
2. Ken Wickwire (KB1 JY) et al., "On-air Measurements of HF TOR and Packet Throughput, Part I: Near-Vertical-Incidence-Skywave Paths," *Digital Journal*, March 1996.
3. Ken Wickwire (KB 1 JY), "On-air Measurements of HF TOR and Packet Throughput, Part II: One-hop Skywave Paths," *QEX*, June 1996.
4. Ken Wickwire (KB1 JY), "On-air Measurements of HF Data Throughput: Results and Reflections," *Proc. 15th ARRL/TAPR Digital Communications Conference*, ARRL, 1996.