This paper describes an experimental line quality monitoring (LQM) system, which has been in regular operation in several IBM locations. The LQM system uses a System/7 to gather information about the analog phenomena that occur on telephone data transmission lines. It gathers the information without taking the lines out of service. Described are the ways in which the information from the LQM system is used in daily operations; for example, in one case, use of the LQM system contributed to an increase in circuit availability from 82 percent to 99 percent over a one-month period.

Experiments in line quality monitoring

by P. Bryant, F. W. Giesin, Jr., and R. M. Hayes

Data being carried by transmission links in computer networks are often subject to transmission errors resulting from various phenomena occurring on the transmission lines. This paper describes an experimental system for monitoring the quality of data transmission over common carrier leased lines and gives examples of how the information obtained from the system may be used in the daily operations of a computer center. The experiments have been carried out at several locations within IBM.

In the following section, we describe the environment and background of our experiments. Interested readers may pursue the general principles of data transmission in, for example, Reference 1. In succeeding sections, we describe our experiments, illustrate the ways in which we use the information, and give some concluding remarks.

Background of line quality monitoring

environment

A computer network generally includes one or more transmission links of the form portrayed in Figure 1. The boxes at the top and bottom of the figure are deliberately labeled with the general title "Communications Subsystem" to indicate the variety of physical forms they might take. At least one of them usually includes a processor of some sort, and they would both generally include some sort of transmission control unit. They might both include processors, terminal devices, multiplexors, and so forth, in many different combinations. Each subsystem generally sends and receives information. The transmitting subsystem presents information in digital form to a *modem* (modulator/demodulator), which modulates (encodes) this information into an analog waveform. This waveform is transmitted over the transmission line. At the other end of the transmission line, another modem demodulates (decodes) the analog waveform and presents digital data to the receiving subsystem.

Figure 1 Basic components of a data transmission system

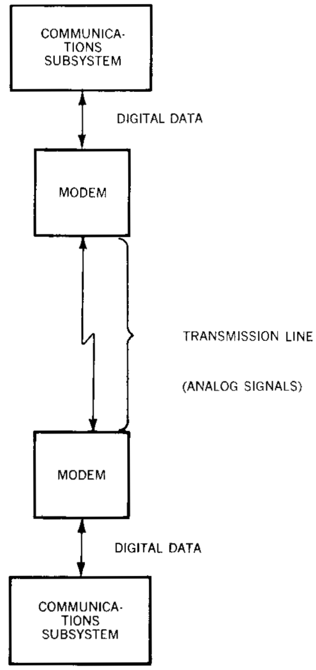

Operating a computer network using such transmission links can be complicated. Some of the reasons for this include:

Multiple equipment types. The various pieces of digital and analog equipment must operate together, and the sources of problems that occur may be difficult to pinpoint.

Errors are more common. Even without hardware failures, transmission errors are much more common than computer or local installation errors. An error rate of one bit per million transmitted is perhaps typical for C2-conditioned leased lines,[2] for example. Because of this, more elaborate coding schemes and error control software must be included in the system.

Complex analog phenomena. Transmission errors result from complex combinations of many analog phenomena (see Reference 2 for a discussion of many of them). The computer operations staff is usually concerned with the *digital* performance of the system and the transmission links—with numbers of bit errors, etc., many of which can (in principle) be determined by software. In contrast, the performance of the transmission line is often specified in terms of limits on the *analog* distortions that will be present. In cases of suspected failure, the operations staff may have to investigate analog problems as well.

Thus, running a computer network can be complicated and may involve considerable manpower. A general approach to such problems that has proven successful in other areas of computer operations (see, e.g., Reference 3) is to have the system measure itself and gather information about what is occurring in it and to develop ways of portraying and using the information to reduce operations problems. In the next section, we describe such an approach to measuring and monitoring the analog portion of a transmission line and discuss ways of using the information in the operation of a computer network.

current approaches

In what follows, the phrase "line quality monitoring" or the abbreviation "LQM" will mean the process of measuring and monitoring the analog phenomena occurring on the transmission line while the line is transmitting. The term does not include

Figure 2 Current approach to line quality measurements

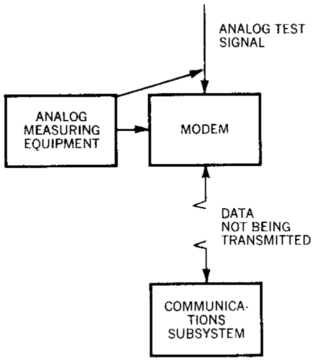

monitoring the protocols, control characters, etc., which may be inserted by transmission control units or the like, or measuring the number of bit or block errors. Currently, measurements are taken as depicted schematically in Figure 2, which shows the receiving portion of a link. When the analog signal arrives at the modem, it is decoded into digital data, which is presented to a communications subsystem (a processor, transmission control unit, or similar system). Measuring is done by removing the line from service and attaching analog measuring equipment to the line, the modem, or various interfaces associated with them. The analog equipment may include meters to measure phase jitter (see below), noise, and so forth. It is common to put a special test signal on the line to make these measurements.

For the operations staff, there are difficulties with this approach:

Knowing when to measure. Operations personnel have many duties and cannot be expected to spend large amounts of time observing the meters of the analog measuring equipment. Often, the only clue they have that there are line problems is when the operating system software indicates an excessive number of retransmissions over a given circuit. Problems may go unnoticed until the number of telephone calls from dissatisfied users at the other end of the transmission line exceeds the operator's personal threshold. Until then, the line is generally left alone.

Intermittent errors. Line disturbances are often transient and intermittent. It may be that when the operator finds time to look at the measuring equipment, the problem is no longer there. Yet, the problem may recur at short enough intervals to degrade service or reduce availability significantly.

Taking the line out of service. When an operator decides that there is a possible line problem, he must normally take the line out of service to do the initial diagnostic checking. It may be necessary to notify the common carrier repair personnel. They may then take the line out of service to perform tests, and it is not uncommon for them to return the line, saying the problem is no longer apparent. The time thus wasted may be anywhere from a few minutes to hours or even days, and during this time, the line is unavailable to the computer network.

Lack of information. Because of the intermittent nature of line and modem problems and the lack of any permanent record of their nature or when they occurred, it often becomes difficult to either diagnose the problem or to be sure that there *is* a *line* problem. Gradual degradation of service may not be noticed until the situation has deteriorated to the point where users complain. If such a trend had been noticed, corrective action might have been taken earlier and at a more convenient time.

Special skills required. The analog measuring equipment involves meters of various kinds. The modems and transmission units have their own error diagnosis lights, switches, and meters. This setup means operators or maintenance personnel must be acquainted with many more concepts, displays, and measurements than they would need to know if they ran a simpler, local installation. It may also mean having someone with engineering training on the premises to handle situations where there really is a line problem. In such cases, it may also require a trained person at the other end of the transmission line to perform the required tests.

a new approach

To attack such problem areas as these, we have experimented with the approach to line quality monitoring described in the following discussion. In general, it works as depicted in Figure 3. The new approach uses extra information about the line's performance—namely, the parameters of the analog signal as that signal is received. These parameters are passed to a line quality monitoring program, which analyzes them and creates appropriate records. The processor containing the LQM program might well be the same as a processor in the communications subsystem, though it would not have to be. The physical source of the information about the analog signal is not very important to our discussion. For a certain class of modems (e.g., References 4–6), though, such information can be obtained from the modem itself, and we took this approach in the experiments described below.

From this general outline, it should be clear that such an approach offers the following advantages:

1. The line stays in service during the monitoring. Thus, initial diagnosis does not require making the line unavailable. Further, measurements taken correspond to actual use of the line, not to a special test situation.
2. Whatever form of permanent record is desired is, at least in principle, available as one output of the LQM program. Such a record can be labeled so as to be readily comprehensible to the operator, providing him with as much or as little information as he needs.
3. The monitoring can be continuous, reducing the need to continually observe analog measuring equipment. Intermittent disturbances, even of short duration, can be detected and recorded.
4. In principle, such a monitoring program could be linked to the transmission control software. When a line was discovered to be malfunctioning, the software might choose, in addition to notifying the operator, to reroute messages, using a different line, or at least invoke orderly shutdown procedures, rather than letting a link fail suddenly.

Figure 3 Experimental new approach to line quality monitoring

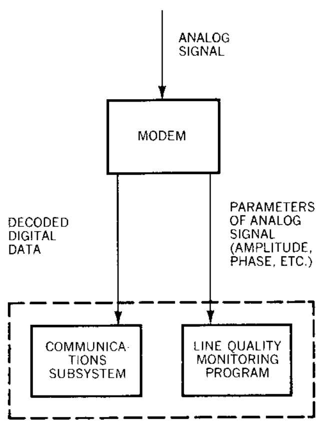

Figure 4 Signal constellation—transmitted signals

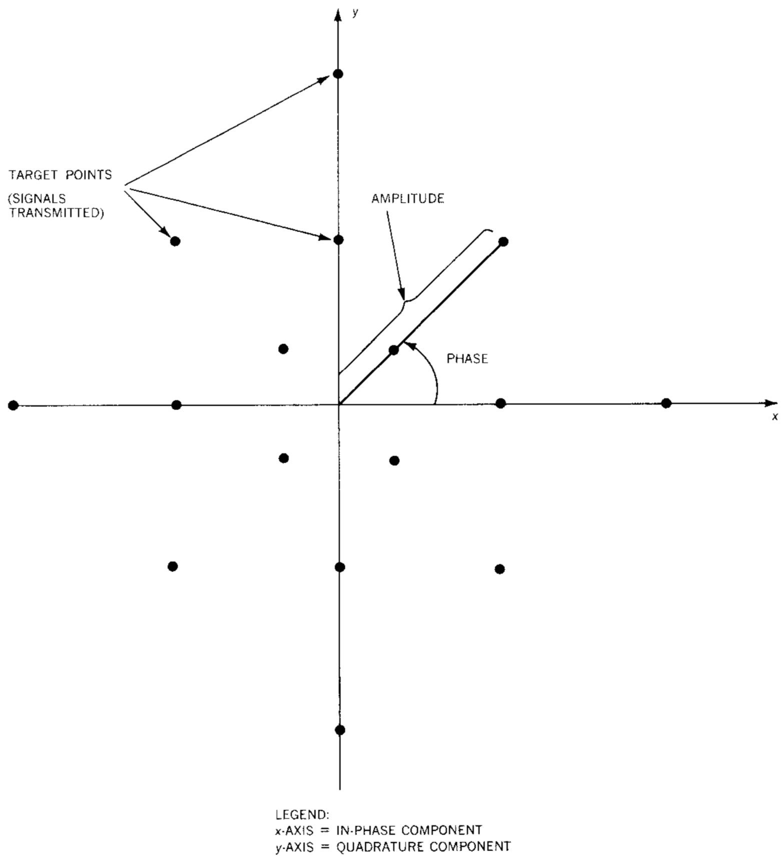

Experimental setup

We begin the description of the LQM experiments with a simplified discussion of the modulation schemes used by the modems we have used in our experiments and of how various analog distortions are reflected in the data we gather. Following that, we discuss the means we use to gather the data and describe the programs and their output.

modulation schemes

Modems of the class we are concerned with encode different bit patterns as analog signals with different *amplitudes* and *phases*.[7] Though many different methods are available, the following general scheme is used by these modems, and it illustrates the basic principles involved. Let us suppose that we are transmitting 2400 signals per second and wish to transmit 9600 bits per second. The modem may do this by breaking up the incoming bit stream into four-bit chunks. There are 16 possible four-bit combinations (0000, 0001, . . . , 1111), and each such combination corresponds to one of 16 different analog signals, each with different

Figure 5 Signal constellation—received signals

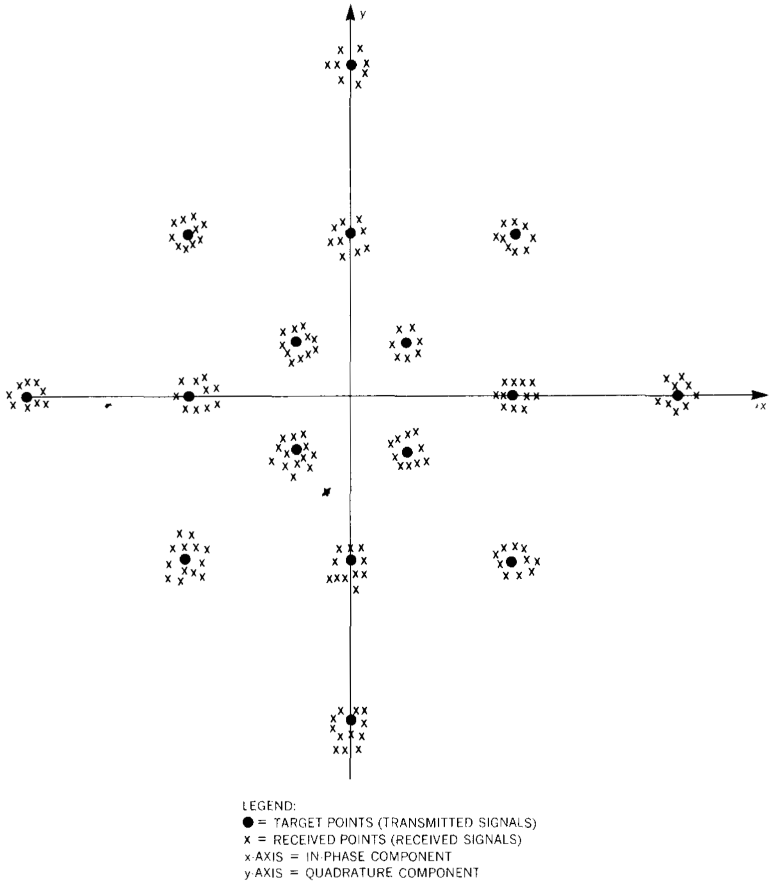

amplitude and phase characteristics. At the receiving end, the modem decides which of the 16 transmitted signals the received signal is closest to, in some sense, and decodes the signal into the bit pattern corresponding to that signal.

A useful way of visualizing this process is illustrated in Figure 4. Such a picture is often referred to as the "signal constellation," and the x and y coordinates are often called the "in-phase" and "quadrature" components of the signal. The 16 transmitted signals are indicated as "target points" in Figure 4. The amplitude of a signal is the distance from the target point to the center of the figure, and the phase of the signal is the angle between the x-axis and a line drawn from the center to the target point.

At the receiving end of the transmission line, the signal is, in general, different from the signal sent. This is illustrated in Figure 5, where the received signals are the points surrounding the target points. The demodulating process may be viewed as taking a received point and determining which target point it is closest to. That target point determines which bit pattern was sent.[8]

Figure 6A Gaussian noise

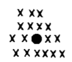

A diagram such as Figure 5 is useful here because different kinds of analog distortions lead to different patterns in the swarms or clusters around the target points at the receiving end. Some of these are illustrated in Figures 6A through 6E, which show patterns of received signals for various distortions and line problems. Gaussian noise causes a spherical swarm around the target point. Phase jitter,[9] however, is a distortion in the phase of the signal rather than the amplitude and is reflected in an arc-like pattern around each target point. A loss of signal is indicated when all points collapse to the center of the diagram. In the modems we have used, harmonic distortion is reflected in the amplitude, and hence, the radius of the received point. Other characteristics may include such things as *retraining*—an automatic re-equalization procedure the modems use when they detect unsatisfactory line performance. This procedure typically involves having all points at two specific training areas instead of having them evenly spread through the 16 possible target areas.

Figure 6B Phase jitter

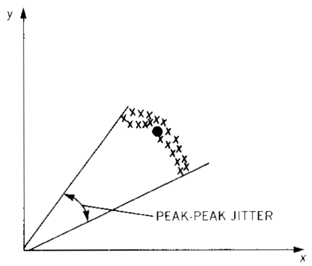

It should be clear from this discussion that if we can gain access to the received points as in Figure 5, a program can be written that can measure, record, and classify the kinds of phenomena outlined above. Checking for loss of signal, for example, consists of determining whether all of the points in a time interval fall near the center. Similarly, the deviation of a received point from its target may be broken into its radial (amplitude) and angular (phase) components to see which is the dominant source of distortion. It *is* possible to gain access to these points, and in the following text, we describe how this has been done.

LQM experimental equipment

A number of modems (e.g., References 4–6) now use the kind of information shown in Figure 5 quite explicitly in their internal logic and make it more or less readily available outside the modem as well. It has been found useful to display pictures like Figure 5 on an oscilloscope at the installation. Problems with the line can often be roughly characterized by observing the oscilloscope patterns as described in Reference 10, for example. The eye-pattern,[11] as this display is sometimes called, does not solve many of the problems outlined in the first section, though, since it is of no use unless someone is looking at it, and there are generally many transmission lines to be considered. It is also difficult to measure things like phase angles very precisely by eye.

Figure 6C Loss of signal ("lineout")

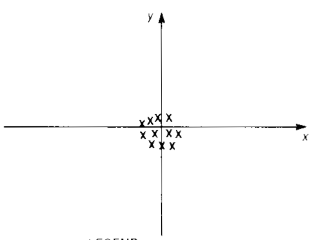

The modem we have used as an experimental tool makes the x and y coordinates of the received signal (see Figure 5) available in digital form outside the modem, from which they are run through an eye-pattern generator that drives an oscilloscope. We have fed these coordinates into the digital input module of an IBM System/7, as well (see Figure 7). The hardware involved consists of a tap box to collect the bits of information and present them in parallel to the System/7, a signal to cause a process interruption at the System/7 when the data has changed, and some System/7 digital input cards to plug into the digital input module.

Figure 6D Retraining

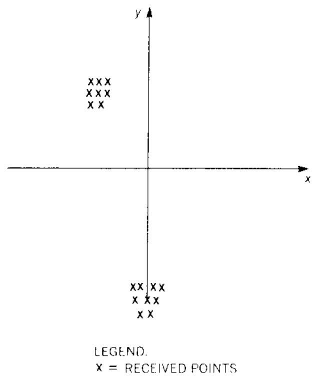

The programs in the System/7 are set up to read each point from the digital input module as it is ready and place it in a buffer. A point consists of the x and y coordinates, eight bits each. When a buffer is filled, a processing program is notified and given the address of the buffer. A buffer size of 120 points (one twentieth of a second) has proven convenient.

data analyses

When a line is either known or suspected to be having problems, the raw (x, y) data may be gathered and shipped to a host computer. There, summary programs that plot the deviations from the target point over time and perform related analyses may quickly reveal the nature of the problem, which may then be reported to the common carrier. Some examples of how disturbances are reflected in such plots are provided by Figures 8 through 11. Figures 8 and 9 are from one of the lines at an IBM location in normal operation. Figure 8 is a recreation of the eye-pattern display for a particular twentieth of a second. In Figure 9, the phase and amplitude deviations are plotted as a function of time for the same twentieth of a second. No particularly striking patterns emerge from these plots. In Figures 10 and 11, the data comes from a line operating with phase jitter present. The amount of phase jitter (which was put on the line by a line simulator) was 20 degrees, peak-to-peak, at a frequency of 120 hertz (Hz). The cyclical pattern in the phase deviations shows up clearly in Figure 10. Figure 11 is a periodogram[12] for the same data, and it shows clearly the frequency of the jitter at 120 hertz.

Figure 6E Harmonic distortion

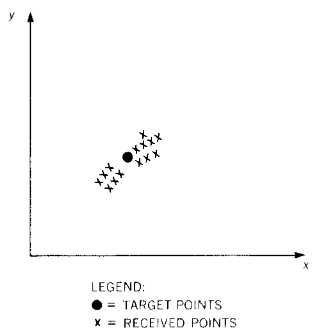

Other analog phenomena are equally clearly reflected in such analyses and plots, but the amount of data involved (38,400 bits per second per line) makes such an approach impractical for the continuous monitoring necessary to detect intermittent problems, or for monitoring more than one line simultaneously. For these reasons, the stand-alone monitor described below was written.

stand-alone monitor

In the specific situations in which they may reasonably be performed, analyses like those above may be all that is needed. For continuous, multiline monitoring, though, we feel that summarizing the data in real time is required. On the one hand, summarizing is needed to extract meaningful insights from the very large amounts of data and to reduce the amount of data which would

Figure 7 LQM hardware configuration

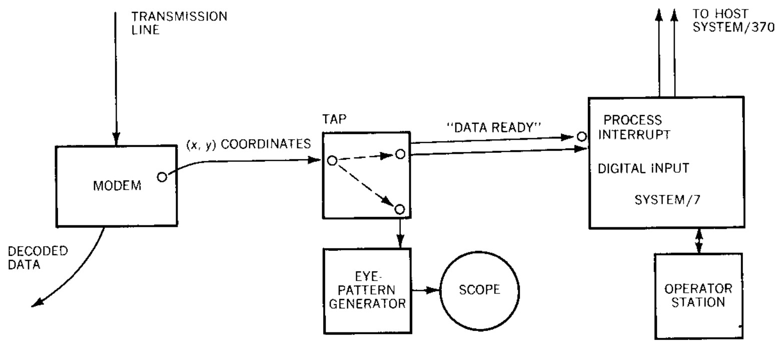

otherwise have to be transmitted and stored. On the other hand, too much summarizing would mean that the variability of line phenomena over time might be inadequately measured, and intermittent disturbances might be lost in the summaries. The stand-alone line quality monitor (LQM3) is an attempt to strike a balance between these extremes.

LQM3 is set up to process each buffer of data in the System/7 in real time, calculating estimate of phase dispersion, noise, and so forth, continually for each line. The results are used to update an exponentially weighted time average for each quantity. Such an average is a memory whose values are determined from the current reading according to the formula (new value of memory) = (1-constant) (old value of memory) + (constant) (current reading). By adjusting the value of the constant, one may adjust the relative weights placed on old and new readings. When messages are typed out, the value of the memory is displayed. Thus the displayed value is a weighted average of the reading over some recent time interval. The messages may be displayed on the System/7 operator station and/or sent via direct channel connection to the host machine for storage and analysis like that described in the following section. The host machine might forward the information over the network to a network maintenance control machine.

The current LQM stand-alone monitor is designed to handle up to seven lines simultaneously (which is the capacity of one digital input module). The operator specifies which lines are to be monitored at any time. The program issues a message for each line at specified intervals (typically five minutes), whenever a loss of signal ("lineout") or retraining occurs and whenever any of the measured quantities exceed certain limits (except that a minimum interval is imposed between messages to avoid flooding the console with error messages in continuing failure situations). All limits and intervals have default settings but may be adjusted by the operator at the System/7 console.

Figure 8 Eye-pattern display recreation

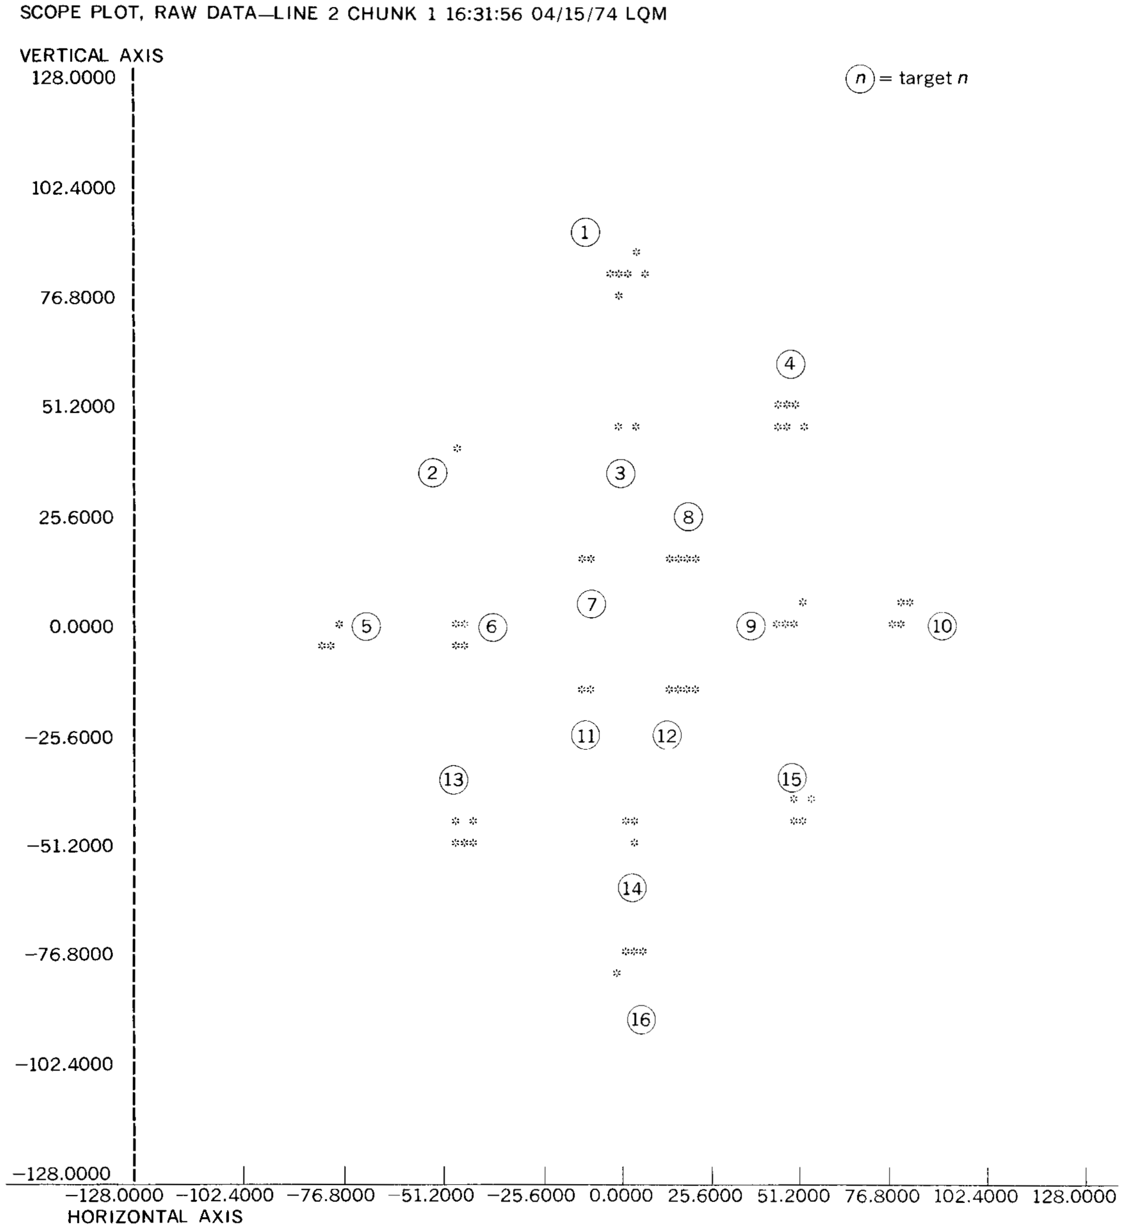

The quantities measured include: average quality (the average distance from a received point to its target), average amplitude deviation, phase dispersion (the angle subtended by a swarm around a target point), radial noise (reported in decibels (dB), indicating the variance of the radial component), the phase jitter (which is an indication of how much of the total phase dispersion is in excess of that attributable to noise alone), the peak-to-

Figure 9 Plot of phase and amplitude deviations

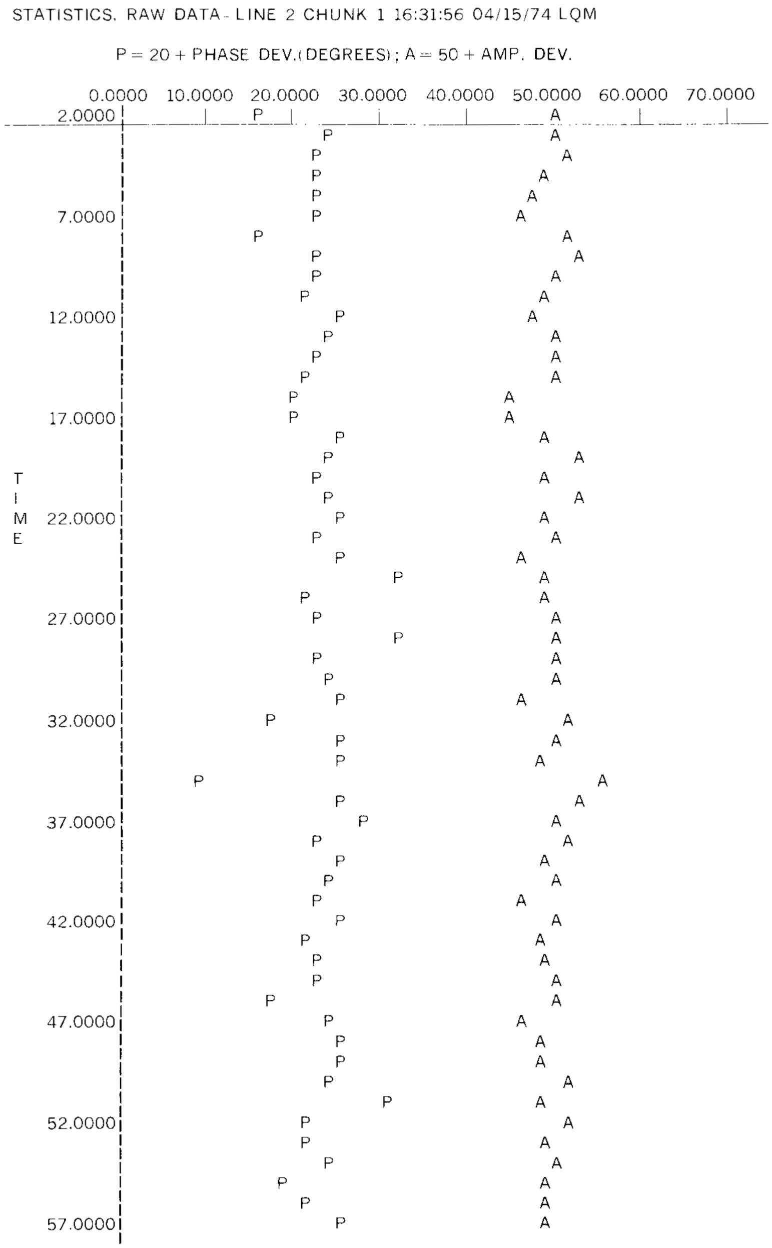

peak average phase deviation (a measure of slow-speed phase variations), and counts of points ("hits") that lie too far away from target points. The most recent reading is displayed as well as the time-averaged value for several quantities. The average value is labeled "A" in Figure 12, which gives a sample printout of the program. The last, or most recent value, is labeled "L."

Figure 10 Phase-jitter plot

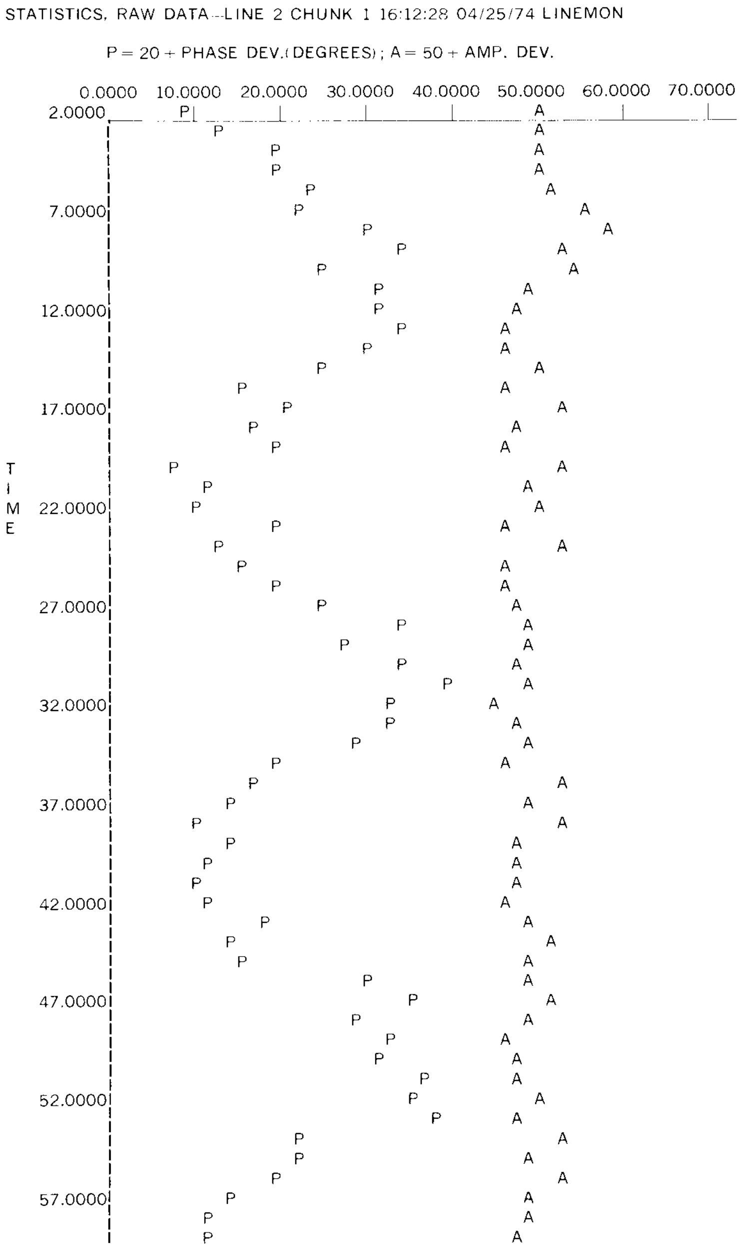

The printout may also be sent to the host for permanent storage in machine-readable form.

The program has been calibrated using special analog test equipment, and over a wide range of phase jitter and noise readings the P-JIT reading in Figure 12 has been shown to be within about

Figure 11 Periodogram of phase jitter

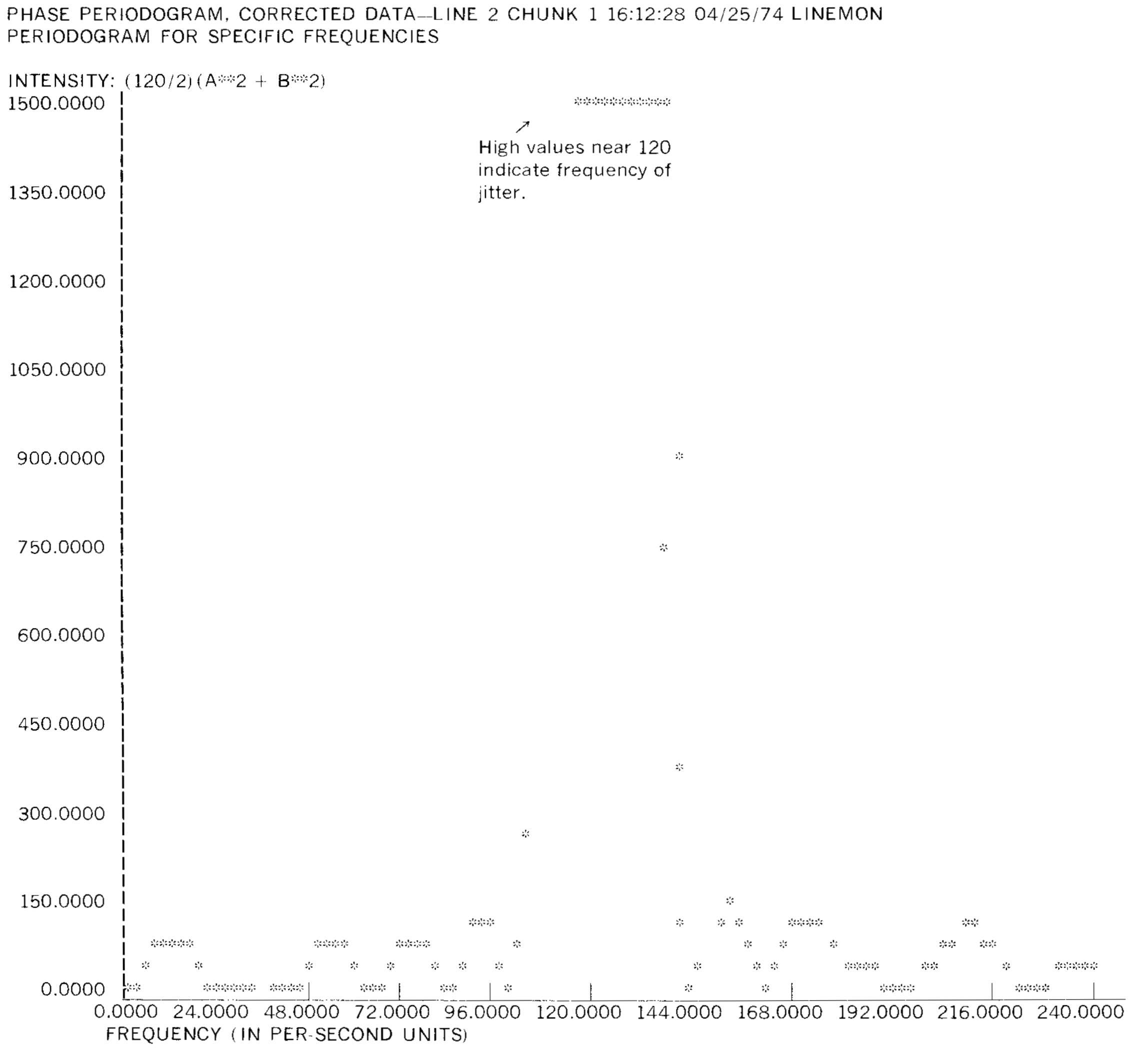

three degrees of the actual phase jitter. The noise is currently reported in decibels with respect to reference noise, corresponding to a specific laboratory test situation. A more realistic view would be to interpret it in terms of a signal-to-noise ratio, as seen by the modem receiving the signal. Work on calibrating this scale of measurement is proceeding. After some experience with the current reading, it becomes clear that adapting to almost any scale is easy. For example, our experience indicates that on the listing in Figure 12, readings of 69 or 70 in the "R-NOS" column represent the point at which we begin to have difficulties with our lines.

Figure 12 Sample output from LQM3 program

```
   05/05/75 / QU/ AMPDEV/PHDEV/P-JIT/R-NOS/APD/ /HITS
TIME     LINE/ A /  A   L/ A  L/ A  L/ A  L/ A / S  M
14:43:42**3 / 9 /- 4 - 4/25 33/ 9 15/63 65/ 2 / 0  2
14:46:23  2 /10 /- 3 - 4/25 24/ 5  1/64 65/ 3 / 0  0
14:48:42  3 /10 /- 3 - 2/25 29/ 9  9/63 65/ 1 / 0  0
14:48:59**3 /10 /- 4 - 2/24 35/ 7 24/63 61/ 2 / 0  5
14:49:00     LINE  3  LEAVING LINEOUT MODE
14:49:00     LINE  3 --RETRAINING
14:51:23  2 /10 /- 3 - 2/25 26/ 5  4/64 65/ 4 / 0  0
14:54:05  3 / 9 /- 4 - 2/24 23/ 7 12/63 61/ 1 / 0  0
14:56:24  2 /10 /- 3 - 4/25 33/ 5 15/64 65/ 4 / 0  0
14:59:06  3 / 9 /- 4 - 2/25 24/ 9 13/63 61/ 2 / 0  0
15:01:24  2 / 9 /- 3 - 2/24 18/ 7 16/63 53/ 3 / 0  0
15:02:34**2 / 9 /- 3 - 6/24 25/ 7  0/63 67/ 3 / 1  1
15:04:06  3 / 9 /- 3 + 0/24 32/ 7 14/63 65/ 1 / 0  0
15:05:58**3 /10 /- 3 + 0/26 76/10  0/63 71/ 1 / 1  3
15:06:00      LINE  3  LEAVING LINEOUT MODE
15:06:00      LINE  3 --RETRAINING
15:07:26**3 /11 /- 3 - 4/31 44/15  8/64 69/ 1 / 0  0
15:07:31**3 /11 /- 3 + 0/31 22/15 10/64 61/ 2 / 0  0
15:07:35  2 / 9 /- 3 - 4/24 19/ 7  0/63 65/ 3 / 0  0
15:07:38**3 /11 /- 4 - 4/31 26/15 16/64 61/ 2 / 0  0
15:07:45**3 /11 /- 3 - 2/31 22/15 10/64 61/ 2 / 0  0
15:07:58**3 /11 /- 3 - 2/31 23/15 12/64 61/ 2 / 0  0
15:10:26**3 /10 /- 4 + 2/25 90/ 9  0/63 75/ 2 / 1  5
15:10:48**3 /10 /- 4 + 0/25 94/ 9  0/63 71/ 2 / 0  4
15:10:49     LINE  3  LEAVING LINEOUT MODE
15:10:49     LINE  3 --RETRAINING

   05/05/75 / QU/ AMPDEV/PHDEV/P-JIT/R-NOS/APD/ HITS
TIME     LINE/ A /  A   L/ A  L/ A  L/ A  L/ A / S  M
15:12:35  2 / 9 /- 3 - 2/23 15/ 6  0/63 61/ 4 / 0  0
15:15:54  3 /10 /- 4 - 4/24 20/ 7  7/63 61/ 2 / 0  0
15:16:43**3 /10 /- 4 + 0/24 61/ 7 24/63 69/ 2 / 0  3
15:17:35  2 /10 /- 3 - 4/24 20/ 7  0/63 67/ 4 / 0  0
15:21:13**3 /10 /- 3 - 2/25 29/ 9 21/63 61/ 2 / 0  2
15:21:14     LINE  3  LEAVING LINEOUT MODE
15:21:14     LINE  3 --RETRAINING
15:22:35  2 /10 /- 3 - 2/25 27/ 5  0/64 67/ 3 / 0  0
15:25:07**3 /10 /- 3 + 0/25 26/ 9  4/63 65/ 2 / 0  4
15:25:09     LINE  3  LEAVING LINEOUT MODE
15:25:09     LINE  3 --RETRAINING
15:27:36  2 /10 /- 3 - 4/24 16/ 7  0/63 61/ 4 / 0  0
15:30:14  3 /10 /- 3 + 0/25 36/ 9 25/63 61/ 1 / 0  0
15:32:36  2 /10 /- 3 - 6/24 31/ 4  0/64 67/ 3 / 0  0
15:33:19**3 /10 /- 3 - 2/25 27/ 9 18/63 61/ 2 / 0  3
15:33:20     LINE  3  LEAVING LINEOUT MODE
15:33:20     LINE  3 --RETRAINING
```

TIME-AVERAGE
LAST
MEDIUM HITS
SEVERE HITS
ASTERISKS INDICATE LIMIT VIOLATIONS
NORMAL MESSAGE INTERVAL
SHORT LINEOUT, THEN RETRAINING

Using the LQM program

one installation's experience

To give a more specific feeling for the use of such a monitoring facility, we summarize here the ways in which the monitor has been used at an IBM facility. The particular values apply to the specific facilities on the dates they were measured and may not be representative of other facilities, but the ways in which the LQM program has been applied are generally applicable.

Prior to August 15, 1974, circuit availability for use during prime shift was 82 percent. On August 15, the LQM program was installed. By September 15, 1974, circuit availability on prime shift was at or over 99 percent. Some of this improvement undoubtedly comes from solving one-time problems of some newly installed lines. But we believe a large portion of it is due to the increased information available through the monitoring and the ways in which the information was used, some of which are described below.

During the first three months of operations, diagnoses of line problems (e.g., "jitter," "noise") were reported to the common

Figure 13 Display of performance of a line

```
LINE : 4                                                          09/17/75

HOURS  8|._________._________._________._________._________._________|
       9|.1_______X._XR______.___1_____._1_______._________.______1__|
      10|._1______XX______1__.______1XXXX_______X&X________._________|
      11|.______XXX._________.______XX_._________._________.______1__|
      12|.LLR______._________._________._________._________.1________|
      13|._________._________._________.___1_____._________.________R|
      14|X_X_______1_____XX__.__1X_R___.___7X_LLLRLR___1___1_________|
      15|.__R______._________.1_1______._________._____1___.1________|
         00        10        20        30        40'       50

KEY : -   = NO INCIDENTS       U   = LINEOUT + LIMIT VIOLATION OR HITS
      L   = LINEOUT            X   = RETRAINING + LIMIT VIOLATION OR HITS
      R   = RETRAINING         S   = MESSAGES SUPPRESSED
      *   = LIMIT VIOLATION    E   = LINE IN ERROR MODE
      & = LINEOUT + RETRAINING + LIMIT VIOLATION OR HITS POSSIBLY
     1-9 = NUMBER OF HITS (0 = 10 OR MORE) + LIMIT VIOLATION POSSIBLY

FOR MORE DETAILS ABOUT ONE OF THESE HOURS, TYPE 'H' AND THE NUMBER (2 DIGITS)
TYPE 'D2' FOR THIS PAGE AGAIN, HIT 'ENTER' FOR THE NEXT PAGE.
h14
HIT 'CLEAR', THEN 'ENTER'
```

carrier in 21 cases. In each case, the carrier took the line out of service to diagnose the problem, and in 17 cases, the System/7 diagnosis of the problem was confirmed to be correct. The knowledge that a particular kind of disturbance seems to be present can be an aid to the carrier in diagnosing the problem and locating that part of the transmission link that is faulty, so that decreased repair time may also be a benefit of this kind of monitoring, provided the information is conveyed to the appropriate carrier personnel.

The lines are monitored daily, and at certain specified times, the readings are recorded in a log. Examination of this log often reveals gradual trends over days or weeks, and when a line appears to be getting gradually worse, preventive maintenance on the line is scheduled with the carrier at a mutually convenient time, such as a second or third shift.

When line problems appear, one of the first actions taken is to monitor the modem involved in local loop-back mode (a situation in which the modem transmits to itself, so that the transmission line is not involved). In such a situation, little or no distortion should be present. If any significant distortions are found, it may indicate a modem problem. Thus the LQM program is used to rule out certain kinds of modem errors before calling in the common carrier's repair personnel.

The monitoring has also proved useful in diagnosing a number of problems in telecommunications software. In one case, the monitoring revealed that the periods of difficulty corresponded exactly to retraining cycles on the modems involved, and that information was sufficient to pinpoint the source of a long-standing problem.

analyses of LQM3 log

We have found it helpful to store the LQM3 output (see Figure 12) on disk files in a host System/370. A collection of experimental interactive analysis programs is then used to examine the

Figure 14 Display of a particular hour

```
LINE : 4                                                                09/17/75

TIME |S  M  -4   0         10        20        30        40        50
14:00|   12 |  A .       J . Q       P         .W        .    N    .         . |
14:05|      | A  .      J  .  Q      . P       . W       .      N  .         . |
14:10| 1    | A  .     J   .   Q     .  P      . W       .       N .         . |
14:15|   15 | A  .    J    . Q      P.         .W        .    N    .         . |
14:20| 1  2 | A  .   J     . Q       P         . W       .     N   .         . |
14:25|      |    .         .         .         .         .         .         . |
14:30|    7 | A  .   J     . Q       P         . W       .      N  .         . |
14:35|    6 | A  .    J    .  Q      .P        . W       .      N  .         . |
14:40|      |    .         .         .         .         .         .         . |
14:45| 1    | A  .    J    . Q      P.         . W       .     N   .         . |
14:50| 1    | A  .    J    .  Q      .P        . W       .      N  .         . |
14:55|      | A  .   J     .  Q       P         . W       .      N  .         . |

S = SEVERE HITS   A = AMPLITUDE DEVIATION  W = AVERAGE PHASE DEVIATION + 30
M = MEDIUM HITS   Q = AVERAGE QUALITY + 5  J = JITTER
N = NOISE - 15    P = PHASE DISPERSION     * = ANY COMBINATION

'H' AND THE NUMBER (2 DIGITS) FOR ANOTHER HOUR, HIT 'ENTER' FOR NEXT PAGE.

HIT 'CLEAR', THEN 'ENTER'
```

data for given lines on given days. The output of the analysis programs may be directed either to the screen of a display terminal (as illustrated by Figures 13–15) or to the off-line printer.

Figure 13 portrays the performance of a line (line 4 in the figure) at an IBM location, prime shift. Each row represents an hour, each column a minute. The various symbols indicate the presence or absence of various phenomena, as indicated in the key. Thus, at a glance, an operator can get a feeling for the line's performance over an extended period.

If the operator wants more information about a particular hour, he may request it, as shown at the bottom of Figure 13. The resulting display is shown in Figure 14. The operator may also ask that the original LQM3 output for specified lines be displayed.

A useful, though conservative, indicator of line availibility for interactive (terminal) use is shown in Figure 15 for the same line, but for the entire day. A "clean" minute is defined to be a clock minute in which no line incident occurs. A measure of interactive availability is then the observed probability that if a user logged on the system at a randomly chosen moment, the line would remain clean for at least five minutes, 10 minutes, etc.

This measure uses a rather restrictive definition of "clean," and would not be appropriate for measuring availability for, say, batch interchange. But for interactive purposes, it seems to be a good indicator of line performance. Figure 16, for example, compares three lines at the same IBM location for the month of September 1975 and shows clearly that line 3 had better availability.

The LQM3 output logs are automatically sent to the host, received, listed, and cataloged. Operations personnel may request

Figure 15 Indicators of line availability for interactive use

```
SUMMARY OF INCIDENTS
TOTAL MINUTES          1440
MINUTES MONITORED         1379
MINUTES MONITORED(FRACTION)       0.96
COUNT OF MONITORED MINUTES(RATE)
   CONTAINING LIMIT VIOLATIONS                   96        0.07
   CONTAINING LINE DROPS                         62        0.04
   IN LINEOUT MODE                                1        0.00
   CONTAINING RETRAININGS                        66        0.05
   CONTAINING ANY INCIDENT                      103        0.07
PROBABILITY OF CLEAN LINE FOR:
AT LEAST   5 MINUTES         0.78
AT LEAST  10 MINUTES         0.64
AT LEAST  15 MINUTES         0.54
AT LEAST  30 MINUTES         0.32
AT LEAST  60 MINUTES         0.12
END OF FILE
R;
```

analyses and plots of any specified line, day, and time. They may thus search for trends as well as current experience. Initial experience has shown that this setup is a valuable one.

Discussion of experiments

The previous sections describe the kinds of information that can be used to attack the problems of line quality in daily operations. We are investigating other kinds of monitoring, too. We would particularly like to combine line quality monitoring with monitoring of the observed error rate to see if insight may be gained into which analog phenomena are the most troublesome.

As mentioned earlier, the fact that we obtain the phase and amplitude information from the modem itself is a convenience, not a logical necessity. It would be possible to look at the analog line directly, rather than tapping into a modem, but this could prove quite complicated. The above discussion of modulation schemes is very simplified, and much more is involved in putting the signal on the line and decoding it than is given above. In essence, by tapping into the modem, the modem is given the hard part of analyzing the analog signal (e.g., carrier recovery) — we see what's left over after the removal of what is (for our purposes) the less interesting information. Since we are not monitoring in the test situation specified by the common carrier for tariff purposes, our figures for phase jitter, noise, etc. are not always exactly the same as those measured by the carrier. On the other hand, the numbers we have obtained have proven close enough to be useful. For example, we have recorded 16 degrees phase jitter, had the carrier test the line and find 17 degrees. The carrier's objective is 10 degrees or less on that kind of line. (See Reference 2). Our measurements have the advantage of being end-to-end measurements, too, including the local loops, whereas carrier measurements are often made between central offices. We doubt that measuring the analog signal directly would enable us to get appreciably more useful figures.

Figure 16 Sample availability plots

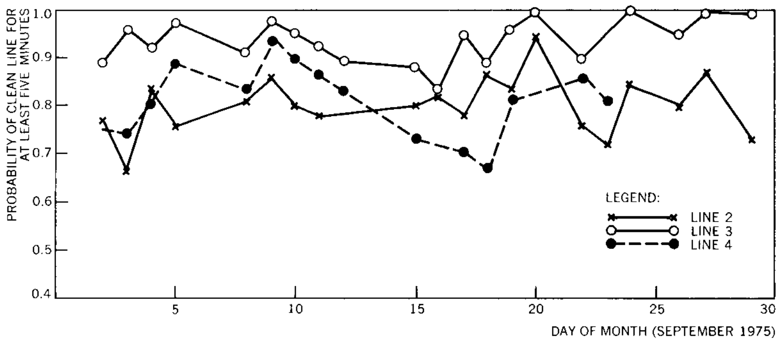

We are continuing our experiments in determining the most useful way in which information such as we are now getting can be displayed, used, and summarized in an operations environment. With such services as all digital transmission (e.g., DDS[13]), we recognize that we will have to measure different specific quantities than we now measure. We suspect, though, that our experience in gathering, displaying, and interpreting information about the lines' performance and in relating it to general system performance and operation may remain applicable.

ACKNOWLEDGMENT

We wish to thank W. P. Timlake for originally suggesting the ideas leading to these experiments and C. I. Johnson for many helpful discussions related to them. We are indebted to C. A. Salisbury for writing the experimental System/7 code to support the digital input module and channel attachment and to Miss F. Vittoz for designing and writing the experimental analysis programs for the System/370. We wish also to thank P. deWitte, H. Harrison, P. O'Neil, A. Springer, T. A. Gilbert, B. Stanton, and the computing operations staff at the Cambridge Scientific Center for their help at various times in this project.

CITED REFERENCES

1. R. W. Lucky, J. Salz, and E. J. Weldon, Jr., *Principles of Data Communication*, McGraw-Hill, New York, New York (1968).
2. *Data Communications Using Voiceband Private Line Channels*, Bell System Technical Reference PUB 41004 (October 1973).
3. W. Freiberger (editor), *Statistical Computer Performance Evaluation* (and many references cited therein), Academic Press, New York, New York (1972).
4. *Codex 9600 Data Modem*, specification sheet, Codex Corporation, Newton, Massachusetts.
5. *MCS9600 Modem*, information brochure, Intertel, Inc., Burlington, Massachusetts.
6. *ICC Modem 96 Multi-Mode*, information brochure, International Communications Corporation, Miami, Florida.
7. Briefly, the amplitude of the signal is the strength of the signal, and the phase is the displacement in time relative to some reference point.

8. Variations on the pattern illustrated in Figure 5 are possible. For example, to transmit 2400 signals per second, but only 4800 bits per second, the modem encodes two bits per signal, and the picture has only four target points instead of 16 (because four two-bit combinations are possible). The exact location of the 16 target points may be varied, too. The problem of locating them optimally, taking into account the various kinds of distortions expected and the power constraints on the line, is of some current research interest as in G. J. Foschini, R. D. Gitlin, and S. B. Weinstein, "On the selection of a two-dimensional signal constellation in the presence of phase jitter and Gaussian noise," *Bell System Technical Journal* **52**, No. 6, 927–965 (July–August 1973). In the discussion that follows, we concentrate on the particular pattern illustrated in Figure 5. The changes necessary to accommodate most variations will be obvious.
9. *An Introduction to Phase Jitter*, TTI Technical Bulletin No. 100, Telecommunications Technology, Inc., Sunnyvale, California (January 1972).
10. S. J. Puchkoff, "A modem eye-pattern looks at telephone lines," *Telecommunications* **6**, No. 8 (August 1972).
11. "Eye-pattern" is the term used by a particular manufacturer.[4] Other manufacturers often use the terms "signal constellation" or "signal space." A plot of amplitude versus time on an oscilloscope is also called an eye pattern, but we shall use the term here in the sense described in the text. A related display is offered by some of the new, off-line test equipment. See, for example, "Does your line measure up?" *Telecommunications* **9**, No. 5, 58–60 (May 1975).
12. A periodogram is a plot showing the relative intensities of the cyclical components of a signal at various frequencies.
13. *Digital Data System Data Service Unit Interface Specifications*, Preliminary Bell System Data Communications Technical Reference PUB 41450 (March 1973).

Reprint Order No. G321-5030.