THE BELL SYSTEM TECHNICAL JOURNAL
Vol. 57, No. 3, March 1978
Printed in U.S.A.

Observations of Errors and Error Rates on T1 Digital Repeatered Lines

By M. B. BRILLIANT

(Manuscript received July 29, 1977)

Measurements of errors on T1 repeatered lines were made at five Bell System offices during 1973 and 1974. They included an informal survey of error rates and error-free seconds on lines in service, and detailed recordings, normally of 24 hours duration on selected lines. The detailed recordings show the existence of at least two distinct error mechanisms, differing significantly in diurnal variation of error rate, distribution of intervals between errors, and dependence on the transmitted bit pattern. It was found that certain T1 lines made errors when in service, driven by D1 channel banks, but not when driven by a pseudorandom test signal.

I. INTRODUCTION

The T1 repeatered line is a short-haul digital transmission system using cable pairs to transmit binary information at a rate of 1.544×10^6 bits per second.[1] T1 lines have been in use since 1962, primarily in conjunction with D1 channel banks to provide a carrier system for voice channels. In the design of the T1 line, the distribution of error rates was an essential design parameter.[2] However, since the field trial of an experimental prototype system, in which the error performance of one worst-case repeatered line was briefly reported,* it has become clear that more knowledge is needed on the subject of the error performance of working T1 lines.

Measurements of timing jitter and errors on T1 lines were made at five Bell System offices during 1973–74. The error measurements were of two types: a survey of error rates and error-free seconds; and detailed recordings of the error process, normally for 24 hours duration, on selected T1 lines. The jitter measurements will be reported elsewhere.[4]

The survey consisted of error rate measurements (determined by

* See pp. 95–96, Ref. 3.

counting bipolar violations on lines in service) on 2594 T1 lines, and error-free seconds measurements on 1640 of these lines. The primary purpose of the survey was to select particular lines, known to be making errors, for detailed recordings of the error process. Such recordings were made on 89 lines. However, many of these recordings showed few or no errors. Recordings of 37 of these lines were analyzed to characterize the following properties: diurnal variation of error rate, distribution of the intervals between successive errors, and sensitivity to the pattern of bits transmitted on the T1 line. Not all these properties could be analyzed for every line; there are 22 lines for which all three properties have been characterized.

Most of the lines in the survey were terminated with D1 channel banks; some had D2 banks. All of the detailed recordings were made on lines with D1 banks. Since the pulse stream format generated by a D1 channel bank is different from that generated by other equipment (for example, in regard to density of ones), some specific results of these measurements are valid only for the D1 environment.

II. ERRORS AND THEIR MEASUREMENT

The digital format on the T1 line is bipolar: that is, the absence of a pulse in any position represents a binary "zero," while a pulse of either polarity represents a binary "one," and consecutive pulses, regardless of the number of intervening zeros, have opposite polarity. The pulses become attenuated and distorted in transmission along the cable pair and are reconstructed by repeaters located at intervals of nominally 6000 feet along the line. An "error" is an incorrect reconstruction in any one pulse position: a pulse where originally no pulse was sent, or a blank where a pulse should be.

If one looks only at the presence or absence of pulses—the binary ones and zeros—errors cannot be detected unless one knows what was sent. But by looking at pulse polarity, one can detect violations of the bipolar format: the occurrence of two consecutive pulses (with or without intervening blank spaces) having the same polarity. Any single isolated error—either the omission or insertion of one pulse—always results in exactly one "bipolar violation" (BPV). On the other hand, multiple errors can combine so that a bipolar violation does not occur, while the reversal of the polarity of a pulse would create two bipolar violations without any errors, and, in any case, the occurrence of a bipolar violation does not define precisely where an error occurred or whether it was an insertion or deletion. In practice, however, the rate of occurrence of bipolar violations is generally close to the error rate. Both are quoted as numbers representing the ratio of the number of events (errors or bipolar violations) to the number of bits transmitted.

Another measure of performance, "percent error-free seconds," is used to specify objectives for the Digital Data System.[5] To define percent error-free seconds, the measurement period is divided into 1-second intervals from an arbitrary starting point, and the 1-second intervals in which errors occur—"error seconds"—are counted. In practice, this count is estimated by counting the intervals in which bipolar violations occur. The remaining intervals are "error-free seconds."

The surveys reported here were based on measurement of lines in service, carrying pulse streams that were unknown to us (except that they obeyed the constraints imposed by the D1 channel bank format). Therefore, although the terms "error rate" and "error-free seconds" are loosely used in referring to the results of the survey, all the survey results and measurements are actually in terms of bipolar violations rather than errors.

On the other hand, the detailed recordings of selected lines were originally planned to record true errors on T1 lines removed from service and carrying a known pseudorandom bit stream. As the program developed we also included recordings of bipolar violations on T1 lines carrying unknown pulse streams from D1 channel banks. Therefore, in referring to the detailed recordings, "errors" and "bipolar violations" are not the same; these terms will be used more strictly in that context than in referring to the survey results.

III. THE SURVEY OF ERROR RATE AND ERROR-FREE SECONDS

3.1 Survey procedure

The scope of the survey is summarized in Table I. The activity at the first office visited was a pilot run for the rest of the program. The survey activity at this office consisted of 1-minute error counts using simple bipolar violation counters. After the pilot project a device was built to count bipolar violations for 54-minute intervals on 16 lines at a time, in conjunction with a PDP-11/20 computer (which was used primarily for the jitter measurements and the detailed error recordings). This device was used in the remaining four offices. After the survey at the second office, the software was modified to obtain counts of error seconds concurrently with the error counts. Error-second counts were thus made at the last three offices surveyed.

Since the original purpose of the survey at each office was to select T1 lines for extended error tests, general information about the population of lines surveyed was not recorded. Some information, however, is available about each office as a whole. Thus, for example, it is known that the lines were all two-cable at office 3, all one-cable at office 5, and mixed at the other offices. (A two-cable T1 system is one in which opposite directions of transmission are segregated in separate cables.)

Table I — Scope of the survey

Office	Survey dates	Number of lines	Data recorded
1	3/1/73 to 4/1/73	461	1-min BPV
2	9/20/73 to 9/27/73	493	54-min BPV
3	10/31/73 to 11/9/73	561	54-min BPV and error-seconds
4	11/29/73 to 1/9/74	792	54-min BPV and error-seconds
5	1/21/74 to 1/24/74	287	54-min BPV and error-seconds
Total		2594	

The survey was not based on random sampling. Measurements began at one end of the office repeater bay lineup and continued either until the other end of the lineup was reached (at offices 3, 4, and 5) or until the expiration of the time allotted for activity at that office (at offices 1 and 2). Lines that were out of service for any reason were skipped over and excluded from the survey. Measurements for the survey were made only on business days between the hours of 8:00 a.m. and 6:00 p.m.

At each office repeater included in the survey, a bipolar violation detector was plugged into the receive monitor jack to detect errors in the pulse stream from the distant office. Each office repeater was counted as a separate T1 line. A through system from one distant office to another via the office surveyed was thus effectively split into two parts and measured as though it consisted of two systems, each connecting the survey office with one of the distant offices. One consequence is that there are more short T1 lines (especially single-span lines) in the survey than in the plant.

Error rate measurements at office 1 were made for 1-minute intervals on one line at a time using one of two different instruments. One was a Philco-Sierra 314A T1 error detecting set, which has a built-in counter with a maximum counting rate of 10 per second. The other was a Western Electric J98710G-2 L3 error detecting set connected to a General Radio 1192 counter, which counted bipolar violations as fast as they could occur on the T1 line.

At the other offices, errors were counted by a 16-line bipolar violation detector using a PDP-11/20 computer. The counting program had two versions, as indicated previously. Version 1, used only at office 2, counted only errors. Version 2, used at the last three offices surveyed, counted error seconds as well as errors. The counting rate of this equipment was subject to the following limitations.

(*i*) (Version 2 only)—If the count in any one second exceeded 1544 errors on all 16 lines combined (corresponding to 10^{-3} error rate on any one line) the counting program was terminated, so that its use of computer time would not interfere with concurrent detailed recording or jitter measurement on another line.

(*ii*) Maximum short-term (fractional-second) counting rate on all lines combined was about 10,000 per second, because the computer took about 95 μs to process each error counted (about 60 μs in Version 1).

(*iii*) Resolution varied from normally about 30 μs to several hundred microseconds depending on error activity. Errors closer together than this on the same T1 line would be counted as a single error.

3.2 Survey data processing

The survey data was recorded manually at the test site and later transcribed to punched cards. The survey data consisted of groups of up to 16 lines, monitored for 54 minutes per group, except at office 1 where each line was monitored for 1 minute. Twenty lines that lost signal at some time during the 54-minute test were eliminated. Error counts were converted directly to error rates by dividing by the number of T1 line bits expected in the monitoring interval at the nominal line rate. Computer programs were used to obtain statistical summaries of three quantities for each line (where available): the percent of seconds with error, the error rate, and a clustering factor. These are all represented graphically by cumulative distribution functions in Figs. 1 to 6.

The clustering factor is not simply the ratio of the number of errors to the number of error seconds. This would give large overestimates of clustering at high error rates (above about 10^{-6}) because the number of error seconds in 54 minutes is bounded. The calculated clustering factor therefore compares the actual count of error seconds with the number of error seconds that would be expected if the errors counted had occurred randomly (Poisson process model). This factor gives the same result at low error rates, but approaches unity at high error rates.

For a 54-minute test, the error rate (strictly, the bipolar violation rate) is

$$\mathrm{ER} = \frac{\text{bipolar violation count}}{1.544 \times 10^6 \times 60 \times 54}$$

The fraction of seconds that contain errors is

$$\mathrm{ESF} = \frac{\text{error second count}}{60 \times 54}$$

The percent error-free seconds is simply

$$\mathrm{PEFS} = 100 \times (1 - \mathrm{ESF})$$

Given an error rate ER, a Poisson process model could be used to predict an error second fraction

$$\text{ESF}' = 1 - \exp(-1.544 \times 10^6 \times \text{ER})$$

The clustering factor is defined as

$$\text{CF} = \frac{\text{ESF}'}{\text{ESF}}$$

The graphs were arranged so that the horizontal axis consisted of a logarithmic scale and the vertical axis consisted of a normal probability scale; thus, a lognormal distribution would be plotted as a straight line. Percent of seconds with error and BPV rate graphs were based on total population, i.e., lines with and without errors. The lowest point plotted corresponds to the percent of lines that had either no errors or one error. Clustering factor graphs were based solely on T1 lines with errors. To avoid step functions appearing on the graphs, a continuous curve was created by plotting only the upper "corners" of the step function.

3.3 Results of the survey

Figure 1 shows the distribution of error rate at each office. Each point on a given distribution can be interpreted as a statement that some percent of the lines had error rates equal to or less than (better than) a certain error rate. For example, about 98.8 percent of lines at office 2 had error rates equal to or better than 10^{-6}. Because of the counting rate limitations described previously, the true error rate may be somewhat higher than the measured error rate for some lines. On the average, it is probably not more than twice the measured error rate, based on estimates of clustering from an earlier survey[6] as well as results from our detailed error recordings. Thus, for the same example, we would estimate that at office 2 less than 98.8 percent of lines had better than 10^{-6} error rate, but probably more than 98.8 percent had better than 2×10^{-6} error rate. That is, the true error rate curve might be farther to the right, by less than a factor of 2 in error rate.

Figure 2 shows the aggregate distributions of error rate compared with previous surveys. All the surveys refer to lines terminated predominantly or entirely at D1 channel banks. Two curves are shown for our survey because measurements of error rates below 10^{-8} are not available for office 1. However, the two curves are very close together.

The previous surveys had suggested a possibility that T1 error performance was changing. The latter of these surveys,[6] in 1971, showed error rates about 10 times as high as the earlier survey,[7] in 1966. A survey at one site in 1968 showed a distribution (not shown in Fig. 2) lying between the other two.[8] However, the results of our survey show such large differences between offices that the difference between the 1966 and 1971

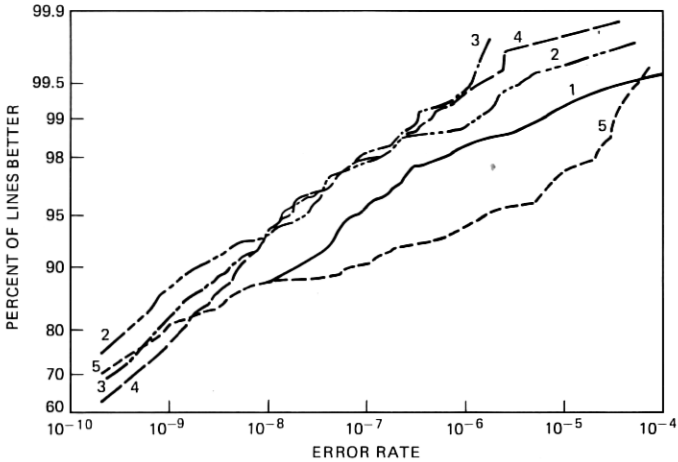

Fig. 1—Error rate distributions for individual offices.

surveys cannot be said to show a significant trend over time. Qualitative observations suggest that the important difference between offices is in maintenance organization and effort, which tends for practical reasons to be correlated with (although not fully determined by) both office size and T-carrier network size.

Offices 2, 3, and 4 in our survey had error rate distributions very close to the 1966 survey. Even allowing for a possible underestimate of a factor

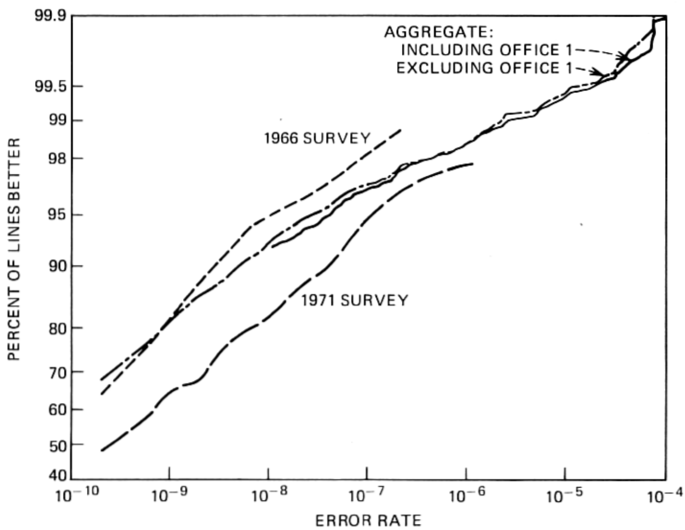

Fig. 2—Aggregate error rate distributions. Results of previous surveys are shown for comparison.

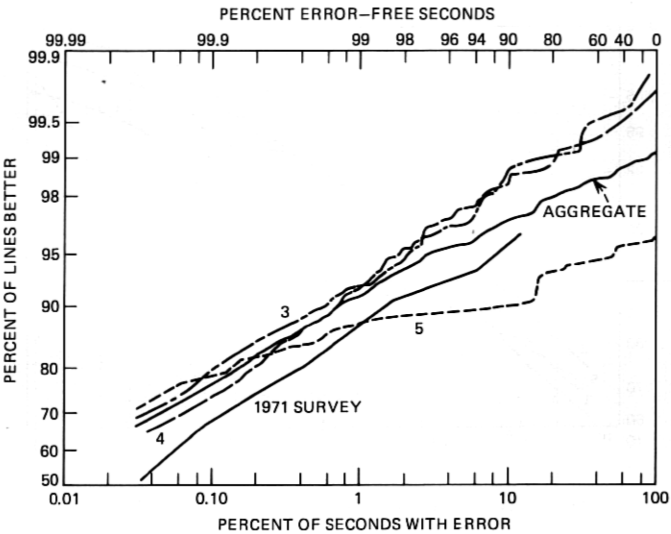

Fig. 3—Distributions of percent seconds with error, compared with the aggregate of the previous survey.

of 2, they are still much better than the 1971 error rate distribution, and much closer to the 1966 curve than to the 1971 curve. These three offices are all large offices in large metropolitan networks, as were the offices in the 1966 survey. In the 1971 survey, on the other hand, two of the three offices were suburban offices.

The office showing the poorest distribution in our survey was also the smallest. This building housed two step-by-step switching machines, and a combined maintenance force was responsible for both switches and T-carrier equipment. This distribution was noticeably affected by the existence of high error rates (between 10^{-6} and 10^{-4}) on 12 lines in the same cable (4 percent of the sample), apparently due to a single cause; omitting these lines would have moved the curve closer to the others.

Error-free seconds results are plotted in Fig. 3 in terms of percent seconds with error (which is 100 minus percent error-free seconds), so that the distribution curves are similar to the error rate distributions. These results are not affected by the counting-rate limitations of the measurement apparatus. The best error-seconds distribution in our survey (office 3) is better than the aggregate of the 1971 survey for about a factor of 4 in percent seconds with error (see Fig. 3); our aggregate is better by a factor of 2. The scatter plot in Fig. 4 indicates that most lines had about the same number of errors as seconds with error (except at high error rates, where the number of seconds with error approaches the number of seconds in 54 minutes). However, clustering can occur at all levels. Figs. 5 and 6 show the distribution of clustering factor (calculated

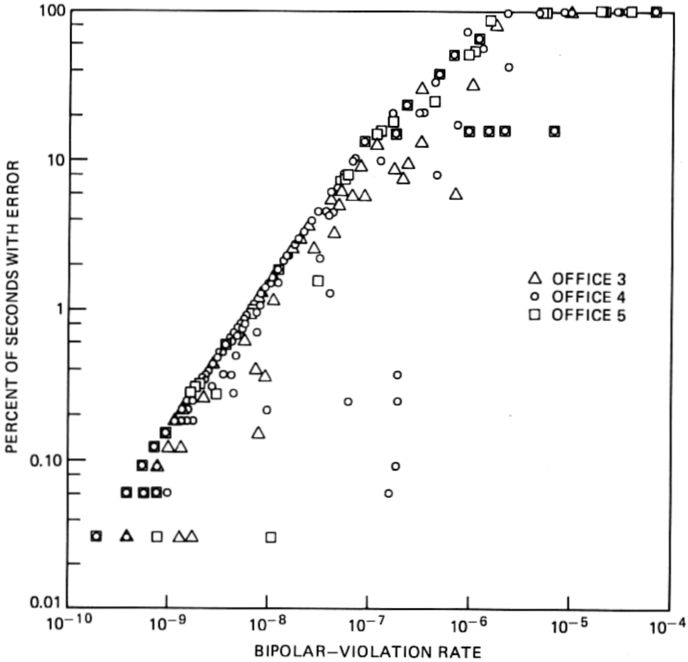

Fig. 4—Scatter plot of percent seconds with error vs. error rate.

so as to avoid false indications of clustering at high error rates, as described in the preceding section). Table II shows the mean and standard deviation of the clustering factor for each office and for the aggregate. In general the standard deviation is much larger than the mean, except at office 3. Overall, the mean is about two events per second-with-error.

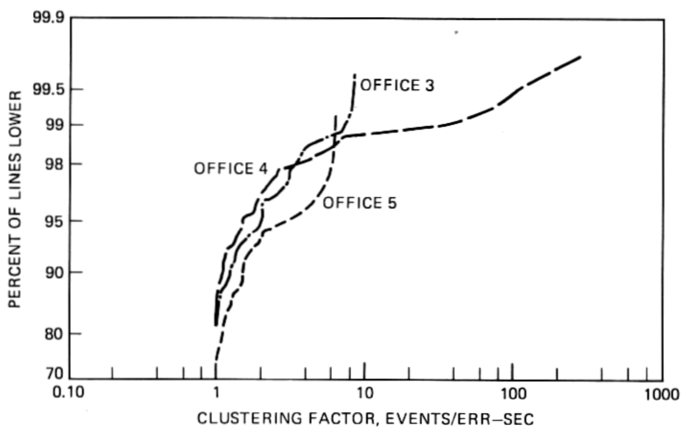

Fig. 5—Distributions of clustering factor for individual offices.

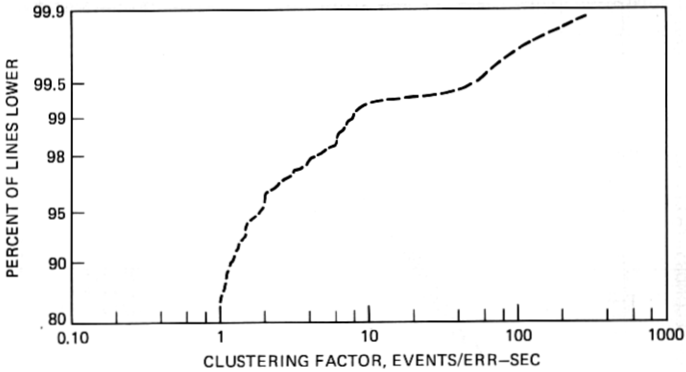

Fig. 6—Aggregate distribution of clustering factor.

Table II — Statistics of clustering factor

Office	Mean	Standard deviation
3	1.20	0.98
4	3.13	22.9
5	1.70	4.93
Aggregate	2.26	16.5

The difference between clustering factors at different offices should probably not be regarded as significant. The unusual statistics for office 4 are due to just four lines with unusually high clustering factors (1 percent of the lines with errors). At office 5, the sample was small, and it is known that several lines could be disturbed by the same condition (although the 12 lines with high error rate in the same cable, mentioned earlier, did not have high clustering factors).

IV. DETAILED RECORDINGS OF ERRORS AND BIPOLAR VIOLATIONS

4.1 Experimental procedure

This phase of the measurement program was designed to obtain complete and continuous recordings of errors on T1 lines for periods of 24 hours (or over a weekend), recording at what time, and on which bit, each individual error occurred. We made no effort to find or fix the cause of the errors.

As originally conceived, the basic procedure was to remove a line from service and apply a test signal with the line looped back to the test site, so that the received signal could be directly compared with the transmitted signal. The test apparatus is shown schematically in Fig. 7. The test signal was a pseudorandom linear shift-register sequence[9] with a period of 1,048,575 bits, generated by a 20-bit shift register, as shown

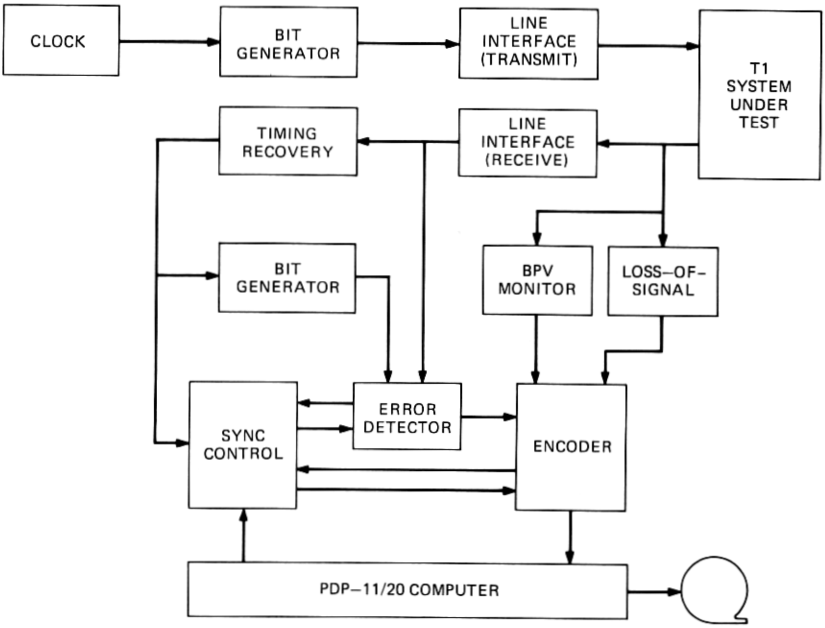

Fig. 7—Error measurement system block diagram.

in Fig. 8. The bit stream returned via the looped T1 line was compared, bit by bit, with the output of a similar sequence generator to detect errors on the line. The occurrence of each error was recorded on magnetic tape. A Digital Equipment Corporation PDP-11 computer controlled test sequencing and formatted the output to the tape.

In most cases the entire T1 line was not removed from service. A T1 line from one office to another generally passes through several intermediate offices, which divide it into spans of the order of 5 miles long. Spare spans are normally provided. The least troublesome procedure

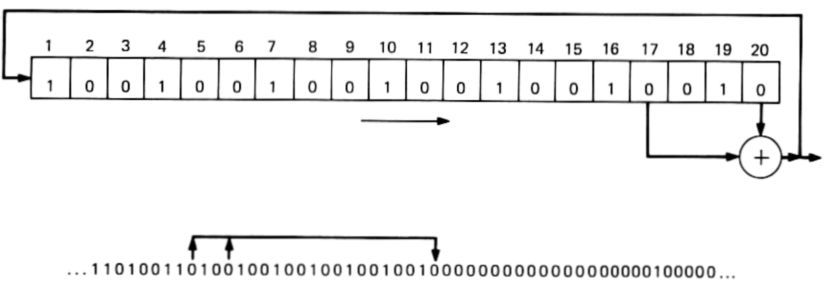

Fig. 8—Pseudorandom sequence generator. The digits in the shift register move from left to right, and hence appear in the output as read from right to left. Generation of the first of the string of 19 zeros is illustrated.

to obtain a length of T1 line for testing is to patch service on one span onto a spare. This procedure requires patching only at the test site and one adjacent office, and does not involve any other offices.

In the earliest measurements at office 1, lines were chosen at random to record both jitter and errors. Since most of the lines so chosen did not make any errors, the bipolar violation survey was introduced as a means of selecting lines for error test that were known to make errors, continuing to use random selection for the jitter measurements. We reasoned that if bipolar violations were observed in the pulse stream coming from the D1 channel bank at the distant terminal office, errors must have been occurring somewhere on the line, and there was a reasonable probability that these errors would be in the span adjacent to the test site.

This procedure resulted in a reasonable yield of error observations at office 1 and office 2. However, it was not successful at office 3. In an effort to improve the yield at this office, lines were chosen that terminated at the next office, so that looping one span would necessarily include the part of the line where the bipolar violations originated. But when these lines were looped and the pseudorandom test signal was applied, errors either did not occur at all, or occurred only occasionally, at rates far lower than the BPV rates observed in service in the survey. Such cases had been observed occasionally at offices 1 and 2. In addition, at office 4, we were able to loop T1 lines back to the test site at distant terminal offices, and again we observed that most lines that showed bipolar violations when in service did not make errors when carrying the pseudorandom test stream. In most cases we verified that the bipolar violations did not originate in the channel bank.

The error recording equipment was therefore modified, late in the measurement program, so as to be able to record bipolar violations in the same manner as errors. The procedure was modified to include pretests that would indicate systematically whether the line should be tested using bit generators and the error detector to record errors ("error run"), as shown in Fig. 9a, or restored to service and tested as in Fig. 9b, to record bipolar violations in service ("BPV run").

The modified procedure included two pretests designed to investigate the effect of the signal on the error performance, using the configuration shown in Fig. 9c. The T1 span was out of service and looped at the next office as in Fig. 9a, but use was made of the signal from the D1 channel bank, which was now routed to its destination via the spare. The dashed lines show alternative connections. One pretest recorded BPVs on the spare. In the next pretest, a signal tapped from the spare was fed into the looped span, and BPVs were recorded on the signal coming back from the looped span. The significant result was that in many cases, BPVs were not seen on the spare, but were seen on the signal that came from the spare via the looped span, although errors did not occur when the signal

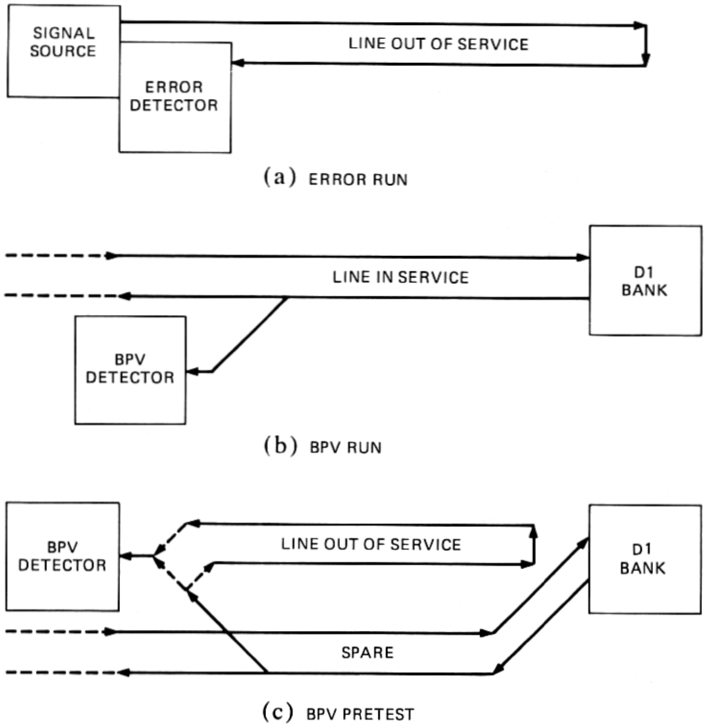

Fig. 9—Error and BPV test configurations: (a) error run, (b) BPV run, (c) BPV pretest.

from the bit generator was transmitted on the looped span (as in Fig. 9a). That is, errors would, or would not, occur on the looped span, depending on the signal fed into it.

Each continuous recording of errors or bipolar violations on a particular line or span is referred to as a "run." Normally the duration of a run is about 24 hours. Longer runs occurred when a run was successfully started on a Friday, since in that case the run was allowed to continue over the weekend. Shorter runs occurred when a line under test turned out to have few or no errors in a period of a few hours, or had an abnormally high error rate; the operator would then terminate the run and select another line for test. Each run was recorded as a distinct file on magnetic tape.

A plot of error rate versus time was derived for each run as a whole. For further analysis, "samples" of up to 1600 consecutive errors were extracted. At low error rates, a "sample" could include a whole 24-hour run. Where possible, separate samples were taken when the plot of error rate versus time showed different conditions occurring at different times. These samples were analyzed to study the "error-free interval" distri-

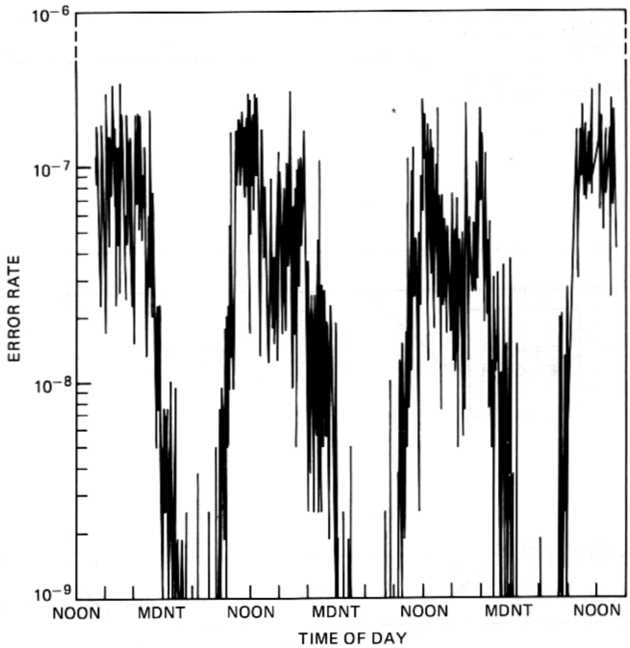

Fig. 10—Normal diurnal variation of error rate, recorded over a weekend, in an error run.

butions (distributions of intervals between successive errors) and the dependence of error probability on the content of the bit stream.

4.2 Diurnal variation of error rate

Since most of the error and BPV runs lasted about 24 hours, diurnal variation of error rate could be observed, but could not be entirely separated from long-term and short-term variation. Even when a run lasted over a weekend, diurnal components cannot be fully separated because a weekend does not consist of identical diurnal cycles. However, some distinct types of diurnal variation were identified.

In roughly one-third of the error and BPV runs of 1 day or longer, the error rate had a broad maximum during the working day (with short-term variations superimposed), falling off gradually during the evening to a minimum error rate at about 4:00 a.m., and rising again to its workday level about 8:00 a.m. Figure 10 shows such a pattern observed over a weekend. We refer to this pattern as "normal" diurnal variation.

Most of the remaining error runs, and a few of the BPV runs, showed essentially steady error rates over the day, as in Fig. 11. Irregular variations, on the other hand, were observed in most of the remaining BPV

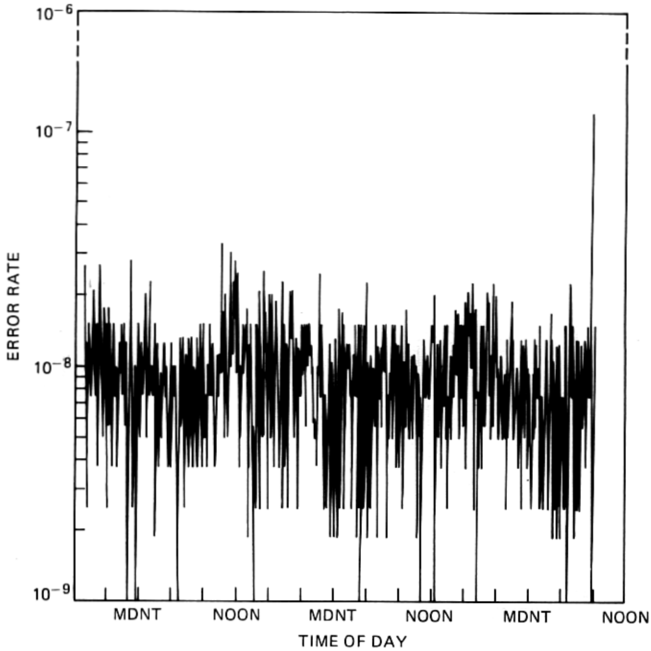

Fig. 11—Steady error rate, recorded over a weekend, in a BPV run. Each plotted point is computed from the number of errors in a cell about 5 minutes long, containing an average of about 5 errors; hence the point-to-point variation is entirely accounted for by assuming errors occurring independently at random (Poisson process), except at the very end of the run.

runs, and a few error runs. Fig. 12 shows distinctly irregular variation, with large nonperiodic changes. However, any variation that did not fit a recognizable pattern was classified as irregular.

The predominance of irregular variation in the BPV runs (as contrasted with steady error rates predominantly in the error runs) may be due to the variability in the content of the bit stream from the channel bank. As described in a later section, the lines with steady or irregularly varying error rates tended to be sensitive to the bit pattern on the lines. The invariably repeating pseudorandom sequence in the error runs would tend in such cases to give error rates that were constant with time; a pulse stream with variable content, such as the channel bank output, would result in varying error rates.

Two T1 lines (one BPV run, one error run) showed two distinct error rates during the diurnal cycle, with abrupt transitions from one error rate to the other, as in Fig. 13. This pattern has been classified as "exaggerated" diurnal variation.

Table III lists the total number of error and BPV runs with each type of diurnal variation. Only runs that actually extended overnight are

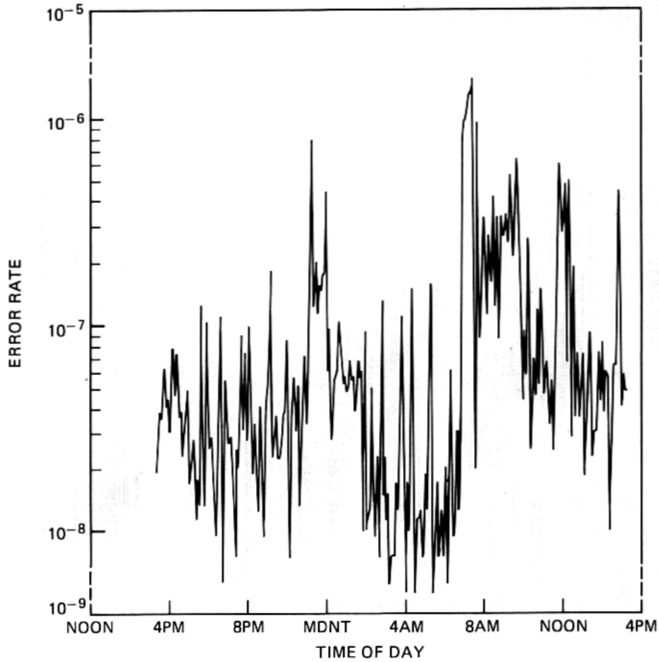

Fig. 12.—Irregular variation of error rate, recorded for 24 hours, in a BPV run.

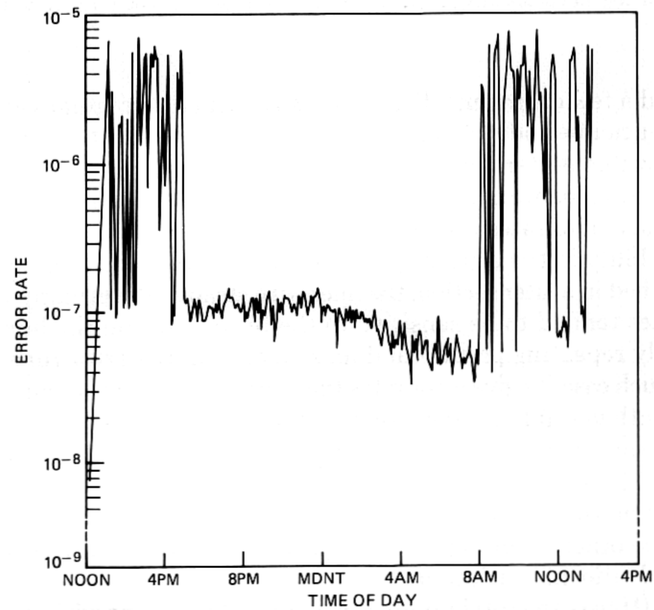

Fig. 13—Exaggerated diurnal variation, recorded for 24 hours, in a BPV run.

Table III — Frequency of occurrence of different types of diurnal variation

Variation type	Error-run lines	BPV-run lines	Total lines
Normal	6	4	10
Steady	6	1	7
Irregular	3	5	8
Exaggerated	1	1	2
Total	16	11	27

shown; shorter runs are excluded because the error rate minimum near 4:00 a.m. was an essential feature for identification of the "normal" type.

In the diurnal variation curves in Figs. 10 through 13, each point plotted represents the average error rate in a 5-minute cell. Since the cell boundaries were defined by the starting times of the data records on the tape, which occurred irregularly, the cells are not exactly 5 minutes long. However, the error rate for each cell is correct, being determined by dividing the number of errors by the actual number of bits in the cell.

4.3 Distribution of intervals between errors

Distributions of "error-free intervals" (the intervals between successive errors, in time units equal to the reciprocal of the bit rate) were plotted for every error sample. Three principal types were observed, identified as unimodal, bimodal, and trimodal. The same types were observed in both BPV runs and error runs. Table IV shows the number of lines observed with each type of distribution.

A typical unimodal distribution is shown in Fig. 14. This curve may be interpreted (except for the numbers on the vertical axis) as the probability density function of the logarithm of the interval, considering the interval as a continuous random variable. The curve was actually derived by setting up a logarithmic interval axis, dividing it uniformly into 40 cells per decade, and plotting the number of intervals that fell

Table IV — Frequency of occurrence of different types of distribution of the intervals between successive errors or BPVs

Distribution type	Error-run lines	BPV-run lines	Total lines
Unimodal	2	4	6
Nearly unimodal	6	5	11
Bimodal	5	2	7
Trimodal	6	3	9
Total	19	14	33

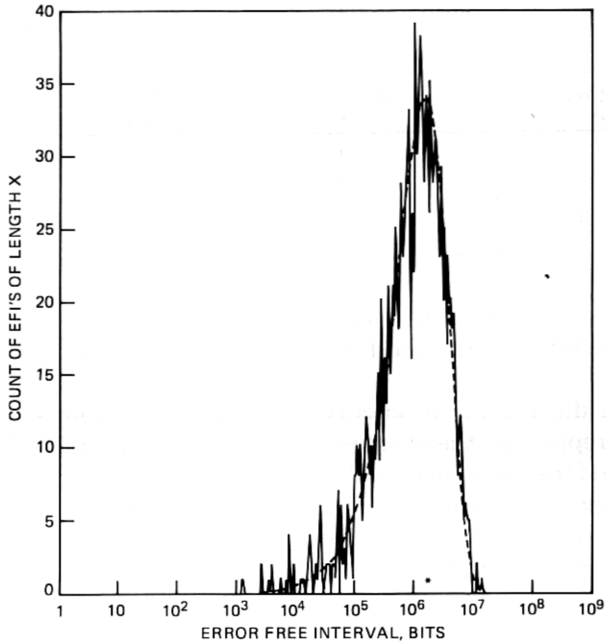

Fig. 14—Unimodal distribution of "error-free intervals," observed in a BPV run.

in each cell. The presence of only one mode indicates the absence of clustering. The asterisk just above the horizontal axis, in the vicinity of the mode, is plotted at the reciprocal of the mean error rate, that is, at the mean interval between errors.

Also shown in Fig. 14, as a dashed curve, is the probability density function of an exponential distribution with the same mean, scaled to fit the same logarithmic horizontal axis and to represent the same total number of errors on the vertical axis. The apparent fit suggests that the errors occurred independently (that is, as a Poisson process). As will be shown later, errors on this line were not independent, but depended strongly on the immediate pattern of bits on the line. However, there is a clear indication of large-scale independence: error-sensitive patterns occurred regularly, but the occurrence of an actual error in each pattern was independent of previous occurrences.

Fig. 15 shows a phenomenon that (fortunately for service) did not happen very often, or last very long: a very high error rate. The distribution shows a very low mean interval between errors. For the shortest intervals, the number of occurrences of each discrete interval are resolved in the plot. Clearly the errors are not completely independent (as in Bernoulli trials), because the count would then be highest at 1 and would decrease with increasing interval length. But the independent-error

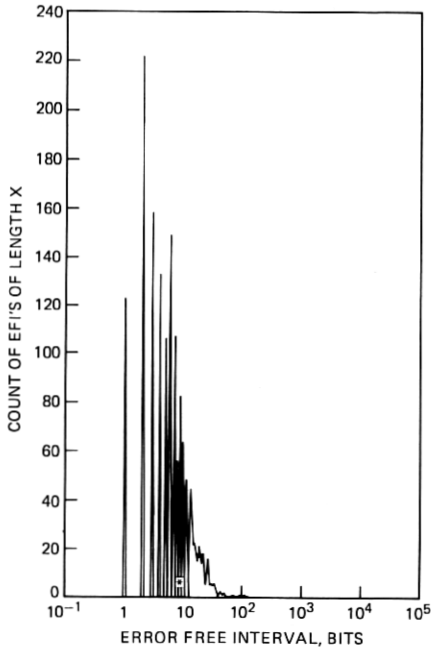

Fig. 15—Distribution of "error-free intervals" at very high error rate, observed in an error run.

model is conceptually useful as a point of departure. These distributions are not included in Table IV.

The bimodal distribution shown in Fig. 16 indicates error clustering. The mode at the right represents the intervals between clusters, and its shape again suggests large-scale independence (Poisson process) as in Fig. 14, but in this case also the occurrence of errors actually depended on local bit patterns. (The small amplitude of this mode may be deceptive; it actually includes about one-fifth of the total number of intervals.) The mode at the left represents the intervals between errors within a cluster, and, as in Fig. 15, errors are not quite independent within clusters. However, the general appearance is not very different from what one would expect from a combination of two exponential waiting-time distributions.

A substantial number of lines had interval distributions intermediate between the unimodal and bimodal types. In these lines, most of the intervals were grouped around a single mode at the right, with less than 10 percent of the intervals (sometimes only one) defining another mode at the left. These distributions are referred to as "nearly unimodal."

The trimodal distribution in Fig. 17 indicates two levels of clustering, that is, clusters within superclusters. The mode at the right represents

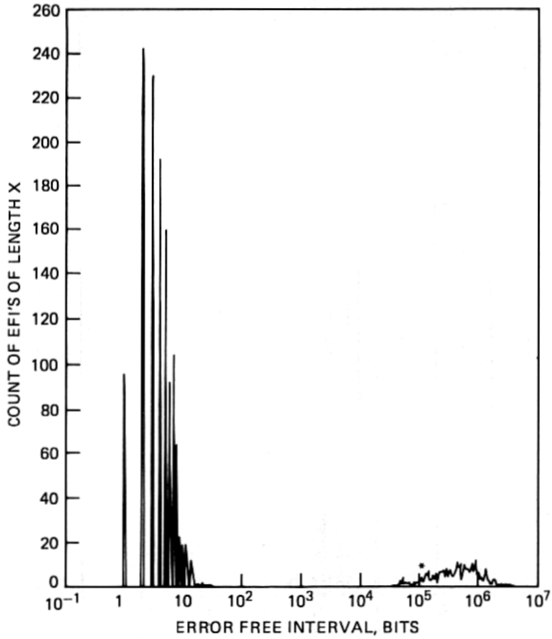

Fig. 16—Bimodal distribution of "error-free intervals," observed in an error run.

intervals between superclusters, the mode in the middle represents intervals between clusters within superclusters, and the mode on the left represents intervals between errors in a cluster. Actually, in this terminology, a "cluster" or "supercluster" could consist of a single error, since the average clustering factor was not large (usually about two to three errors per supercluster) and errors often occurred in isolation. On the other hand, it was not uncommon to find superclusters containing as many as 30 errors, occurring both alone and in clusters within the same supercluster.

Whenever a trimodal distribution was observed in an overnight run (permitting identification of the diurnal variation type), the line showed "normal" diurnal variation. Conversely, most of the "normal" diurnal variation runs showed trimodal interval distributions. In these cases, the interval distribution remained trimodal throughout the diurnal cycle. The middle mode, which indicates structure within the supercluster, remained usually in the same place, extending from about 100 to 500 bits, while the right mode, which described the occurrence of superclusters, moved left or right as the error rate on the line went up or down (respectively) with time of day.

A number of distributions classed as bimodal (or nearly unimodal) differed in appearance from Fig. 16 because of other features. In two

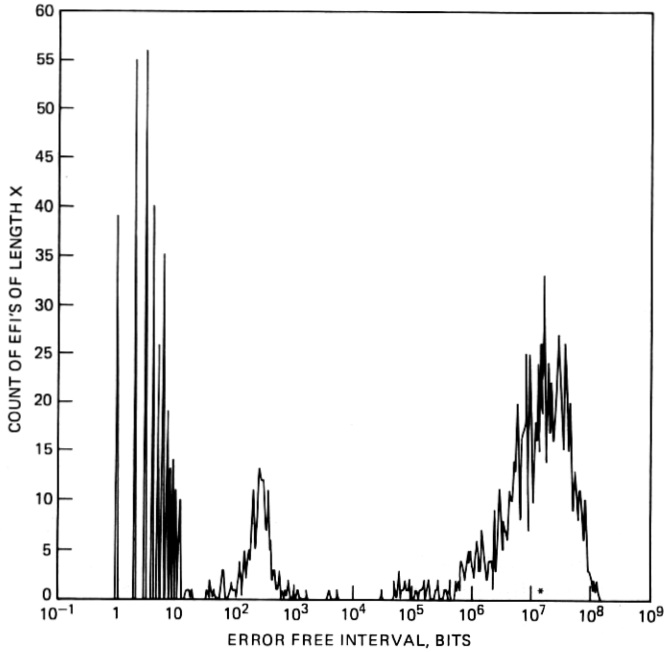

Fig. 17—Trimodal distribution of "error-free intervals," observed in a BPV run.

runs, discussed in Section 4.5, only certain values of short interval occurred. In a BPV pretest for an error run, multiples of about 193 (the D1 frame duration) predominated in the long intervals. In another sample, representing a burst of relatively high error rate in a BPV run, the long intervals are almost all multiples of about 20,000 bits (1/75 second). In most of these cases, longer and shorter intervals are interspersed, indicating error clustering. On the other hand, we obtained a sample from one error run that looks much like Fig. 11, but actually represents a transition from a low error rate, as in Fig. 14, to a very high error rate, as in Fig. 15; all the long intervals came first, followed by the short ones.

4.4 Effects of local bit patterns

On most lines, the probability of error depended on the local bit pattern being transmitted. The probability of error on any given line was usually not the same for ones as for zeros (bias effect), and also depended on the values of preceding and following bits (intersymbol interference effect). The pattern sensitivity varied greatly from line to line in both error and BPV runs. Pattern sensitivity was categorized as high, medium, or low for each line; Table V shows the number of lines observed in each category.

Table V — Frequency of occurrence of different levels of pattern sensitivity of probability of error or BPV

Pattern sensitivity	Error-run lines	BPV-run lines	Total lines
High ($\geq$50.0)	7	3	10
Medium	0	8	8
Low ($\leq$4.0)	7	3	10
Totals	14	14	28

Pattern sensitivity was directly observable in the error runs (except at office 2) because the bit stream on the line was a known pseudorandom sequence. (At office 2, a broken wire forced the second bit of every 16-bit data word on every tape to zero; the evaluation of error rate, diurnal variation, and the interval distribution was not seriously affected by the consequent loss of data, but pattern dependence could not be determined because identification of bit positions in the sequence was ambiguous.) In the error runs, pattern sensitivity fell clearly into two groups: very high and very low (except for short periods of very high error rate, which showed no discernible pattern sensitivity at all). In the BPV runs, evidence of pattern sensitivity appeared as a consequence of the periodic structure of the D1 channel bank frame. The bit stream is organized into frames of 193 bits at 8000 frames per second, in which one bit (the framing bit) is alternately one and zero in successive frames to enable the receiving channel bank to determine the frame phase. The remaining 192 bits consist of 24 words of 8 bits each, one word for each channel, in which one bit carries signaling and the other seven represent the voice by pulse code modulation (PCM). Pattern dependence could be inferred by relating the occurrence of bipolar violations to the periodicities of the D1 channel bank frame, and indications of pattern dependence presumably depended on the content of the channel bank output.

In the error runs, bias effects were characterized by identifying each error as an insertion (error in a zero) or a deletion (error in a one) and computing the percent of errors that were deletions. The percent deletions observed on 14 T1 lines varied from 0 to 100 percent, in a manner consistent with a uniform distribution.

Preliminary analysis suggested that intersymbol interference effects were usually confined to the bit following the error and the 6 bits preceding it. These bits, together with the bit in error, comprise an 8-bit string with the error in the seventh bit. For each error sample, a tally was then made of the number of times each of the 256 possible 8-bit strings occurred with an error in the seventh bit. Figs. 18 and 19 show the results of two such tallies, representing examples of low and high sensitivity respectively. (In these figures, the seventh bit is shown as transmitted; the slash through it indicates that this was not the value received.)

As a numerical measure of pattern sensitivity, the parameter

$$D = \frac{256}{N^2} \sum_{i=1}^{256} N_i^2 - \frac{255}{N} - 1$$

was computed for each error sample, where N_i is the number of times the ith 8-bit pattern occurred with an error in its seventh bit, and N is the total number of errors. This parameter is related to chi-square as it would be evaluated to test the null hypothesis of equal probabilities for all 256 patterns. However, while chi-square is useful primarily as a measure of statistical significance, D is designed to be nearly unaffected by the sample size for strongly pattern dependent errors, so that it measures pattern sensitivity as a property of the T1 line. The largest possible value of D, which would be attained if all errors occurred in the same pattern, is nearly 255 (actually, 255 − 255/N). If errors were independent of pattern, so that all 256 patterns were equally likely, D would have zero mean (based on a binomial distribution, $p = N/256$, for N_i) and a standard deviation of approximately $22.6/N$ (based on the related chi-square distribution, valid for large N, and verified by enumeration as roughly correct for $N = 2$ and $N = 3$).

The values of D fell into two groups: high (50 or above) and low (2.2 or below). The only exception is one sample that has an intermediate value because it spans a transition from one type to the other. However, in almost all cases the value of D is much larger than $22.6/N$, and hence significantly different from zero. The exceptions are of two types: small samples, and very high error rates (above 10^{-3}). Pattern sensitivity showed no consistent relation to percent deletions.

In the BPV runs, the bit patterns on the line were unknown. However, it could be presumed that some bit patterns might tend to recur periodically either at the 193-bit frame period or at the 8-bit word period within the frame. Accordingly, in each BPV sample each bit was assigned a frame position from 1 to 193, starting arbitrarily at the beginning of the sample, and a tally was made of the number of bipolar violations found in each of the 193 positions. These tallies were analyzed both numerically and graphically.

As a general numerical measure of pattern sensitivity for BPV runs, the parameter

$$D' = \frac{193}{N^2} \sum_{i=1}^{193} N_i^2 - \frac{192}{N} - 1$$

was computed for each BPV sample, where N_i is the number of BPVs that occurred in the ith position in the 193-bit frame. This parameter has similar properties to the parameter D computed for the error samples. Specifically, its largest possible value is $192(1 - 1/N)$, and its standard deviation, for random errors, is $19.6/N$.

27	011010Ø1	11	110101Ø1	8	010100Ø0	6	000100Ø1	5	011000Ø1	4	011001X0	3	110000Ø1	2	110101X1
26	101010Ø0	11	111000Ø0	8	010100Ø1	6	001000Ø1	5	011000X0	4	011110Ø1	3	110000X0	2	110111Ø1
25	011010Ø0	10	000001Ø1	8	010111Ø1	6	001001Ø0	5	011110Ø0	4	100000Ø0	3	110010X0	2	111010X0
25	101010Ø1	10	001100Ø1	8	011001Ø0	6	001001Ø1	5	100011X0	4	100101X0	3	110010X1	2	111010X1
24	111010Ø0	10	011011X0	8	011100Ø1	6	001101Ø0	5	100100Ø1	4	101000Ø1	3	110100X1	2	111101X1
22	010010Ø1	10	100101X1	8	011101Ø0	6	001111Ø0	5	100100X0	4	101101X1	3	110111X1	1	000001X1
20	110010Ø1	10	101011Ø0	8	011111Ø1	6	010000Ø1	5	100110Ø1	4	101110Ø1	3	111000X0	1	001001X0
20	111010Ø1	10	101100Ø1	8	101000X0	6	010101Ø1	5	101000Ø0	4	110000X1	3	111000X1	1	010010X0
19	000010Ø1	10	101101Ø0	8	101011Ø1	6	010101X0	5	101100Ø0	4	110001Ø0	3	111001X0	1	011010X0
19	001010Ø0	10	101110Ø0	8	101011X0	6	010110Ø1	5	101111X1	4	110001X1	2	000000Ø0	1	011010X1
19	001010Ø1	10	110011X0	8	101111Ø1	6	011100Ø0	5	110000Ø0	4	110100Ø0	2	000010X0	1	011100X0
19	100010Ø1	10	111100Ø0	8	111000Ø1	6	100000Ø1	5	110001X0	4	110111Ø0	2	000010X1	1	100000X1
18	000011X0	9	000000Ø1	8	111011Ø0	6	100011Ø1	5	110011Ø0	4	110111X0	2	000100X0	1	100001X0
17	000010Ø0	9	000110Ø1	8	111101Ø1	6	100101Ø0	5	110011Ø1	4	111110X1	2	001001X1	1	100010X0
17	100001Ø1	9	001000X0	7	000000X1	6	101001X1	5	110100Ø1	3	000011Ø1	2	001010X0	1	100010X1
15	100010Ø0	9	001011Ø0	7	000100Ø0	6	101111X0	5	110100X0	3	000101X1	2	001100X0	1	100100X1
15	100011X1	9	001100Ø0	7	000101Ø0	6	110101Ø0	5	110110X0	3	001000Ø0	2	001101Ø1	1	101010X0
14	010110Ø0	9	001110Ø0	7	000101X0	6	111001Ø0	5	111111X1	3	001100X1	2	001101X0	1	101110X0
14	111011X1	9	010011X1	7	001111X0	6	111011Ø1	4	000001X0	3	010000X1	2	001110X0	1	110110X1
13	110110Ø0	9	010111Ø0	7	010001Ø0	6	111100X1	4	000100X1	3	010001X1	2	010010X1	1	111100X0
13	111110Ø1	9	011001Ø1	7	010001X0	6	111111Ø1	4	000110X1	3	010100X0	2	010100X1	1	111101Ø0
12	001011X0	9	011011Ø0	7	011101Ø1	5	000000X0	4	000111Ø1	3	010110X1	2	010110X0	1	111111Ø0
12	001011X1	9	100110Ø0	7	011111Ø0	5	000011Ø0	4	000111X0	3	010111X1	2	011011Ø1		
12	011011X1	9	100111Ø1	7	011111X1	5	000101Ø1	4	001000X1	3	011001X1	2	011100X1		
12	100111X1	9	101001Ø0	7	100011Ø0	5	000110Ø0	4	001101X1	3	011101X0	2	011101X1		
12	110010Ø0	9	101001Ø1	7	100111X0	5	001110X1	4	001110Ø1	3	100001Ø0	2	011110X0		
12	111011X0	9	101101Ø1	7	101000X1	5	001111X1	4	001111Ø1	3	100100Ø0	2	011110X1		
11	000011X1	9	110110Ø1	7	111001Ø1	5	010000Ø0	4	010000X0	3	100111Ø0	2	100000X0		
11	010010Ø0	8	000111Ø0	7	111100Ø1	5	010011Ø1	4	010011Ø0	3	101100X1	2	100001X1		
11	101011X1	8	000111X1	7	111101X0	5	010011X0	4	010101Ø0	3	101101X0	2	100101X1		
11	110001Ø1	8	001011Ø1	7	111110Ø0	5	010101X1	4	011000Ø0	3	101110X1	2	101010X1		
11	110011X1	8	010001Ø1	6	000001Ø0	5	010111X0	4	011000X1	3	101111Ø0	2	110101X0		

Fig. 18—Computer printout of occurrences of errors within 8-bit patterns showing low pattern sensitivity in an error run ($D = 0.5$).

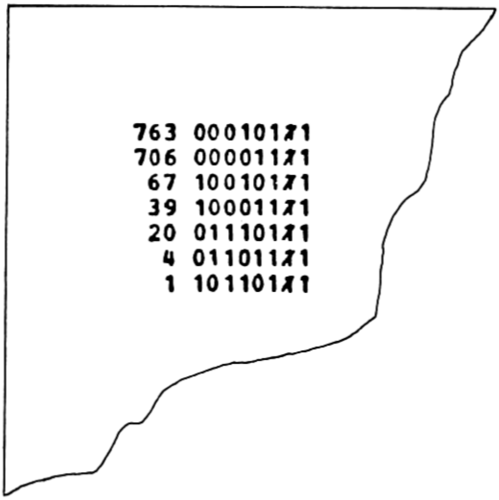

Fig. 19—Like Fig. 18, but showing high pattern sensitivity (D = 107.5).

The values of D' are significantly different from zero in most cases, and vary as widely as the values of D for the error samples. However, they do not fall clearly into two distinct groups; about half fell into the "medium" range, between 2.2 and 50, in which D never fell. This may be explained by the fact that dependence of BPV rate on frame position depends not only on the error mechanism on the T1 line but also on the properties of the transmitted bit stream. The presence of frame periodicity in the bipolar violations actually depends on two factors: the probability of a bipolar violation must depend on the bit pattern, and recurrent patterns must exist in the D1 bank frame.

It should be noted that bipolar violations are necessarily pattern-dependent, even if the errors causing them are not, because a bipolar violation can be detected only when a one is received. In the D1 channel bank format, some frame positions are more likely to contain ones than others; hence, bipolar violations will always be frame-position dependent to some extent. This factor may account for some of the observed frame-position dependence where this dependence is weak. However, it cannot account for the observed cases of strong frame-position dependence, because there are constraints in the channel bank that prevent long strings of zeros (as described in Section 4.5).

Figures 20 through 22 respectively show examples of weak (D' = 0.4), medium (D' = 11.7), and strong (D' = 108.5) dependence of bipolar violations on bit position in the frame. These figures represent 3-dimensional bar graphs, the height of each bar representing the number of bipolar violations at each frame position. The lower left corner represents the first bit in the frame, the first 8 bits are laid out from left to

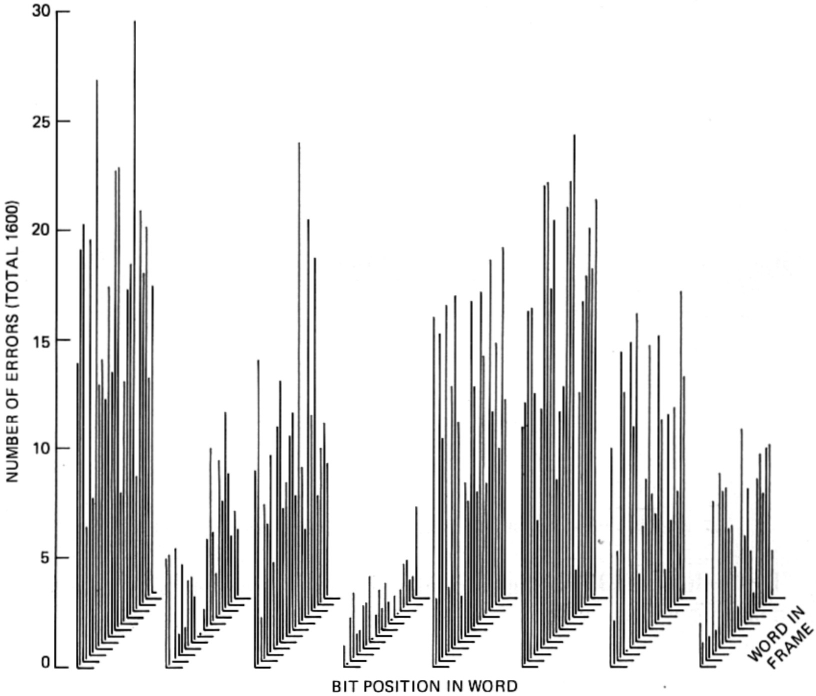

Fig. 20—Three-dimensional bar graph showing weak dependence of bipolar violation probability on position in the D1 bank frame in a BPV run ($D' = 0.4$). Frame alignment is hypothetical.

right, successive 8-bit words are laid out behind the first, and the 193rd bit is in the far left corner of the base of the diagram. Each row running from front to back thus represents a given bit position in each word. In these figures, the hypothetical starting point of the frame was shifted to the position that gave the figure the most regular appearance (by maximizing a parameter similar to D', but based on the number of bipolar violations in each position in the word), on the presumption that such a choice probably approximated alignment with the actual channel bank frame being transmitted.

Figure 20 shows an example of relatively weak position dependence; bipolar violations occur in all positions, but their probability clearly depends on the position of the bit in the word. An example of medium position dependence is shown in Fig. 21. Fig. 22, where most of the bipolar violations occurred in one position in the frame, shows strong position dependence. Fig. 22 is an extreme case; the same line at another time of day showed an intermediate pattern in which no single bit position had a majority of the BPVs. Figure 22 necessarily indicates high sensitivity to some recurrent pattern. Figures 20 and 21, however, could

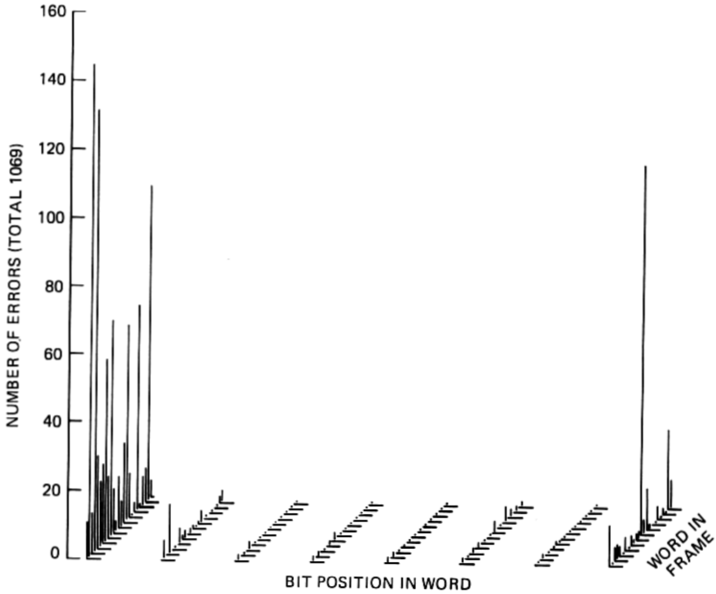

Fig. 21—Like Fig. 20, but showing medium position dependence ($D' = 11.7$).

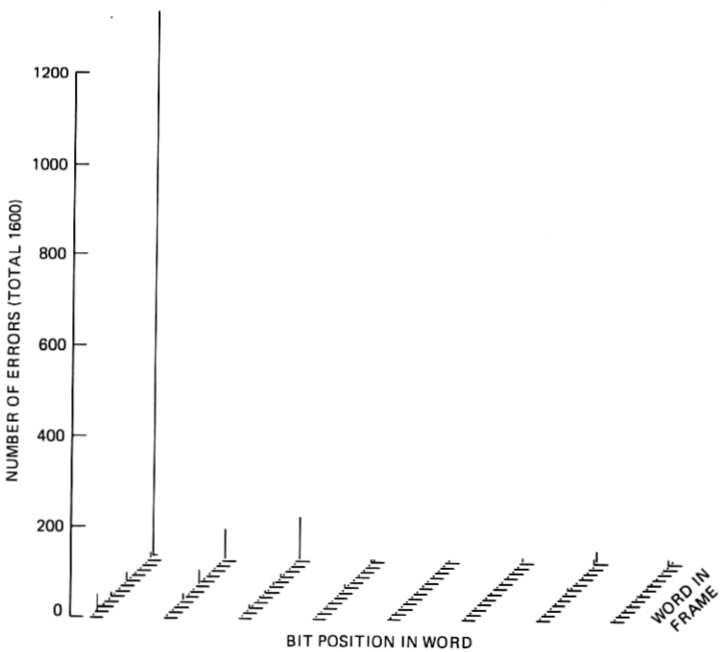

Fig. 22—Like Fig. 20, but showing strong position dependence ($D' = 108.5$).

be explained either by lower pattern sensitivity (possibly none at all), or by high sensitivity to patterns that occur less regularly. Correlation with other properties of the error process, discussed in Section 4.6, suggests that the former explanation is applicable to cases such as Fig. 20, and the latter to cases such as Fig. 21; but such conjectures are speculative.

4.5 Effects of density of ones in the bit stream

The pseudorandom sequence does not meet the constraints on minimum ones density for bit streams transmitted on T1 lines. These constraints are intended to ensure that there is always sufficient energy in the timing tank circuit in each T1 line repeater. This tank circuit gets an input pulse whenever a one is received, and if the density of ones is too low the repeater timing becomes inaccurate and errors are more likely to occur.

It was, therefore, expected that any given T1 line, driven by the pseudorandom sequence generator, would be more prone to make errors than when driven by a D1 channel bank, which does meet these constraints. This expectation was not fulfilled. On different lines, the error rate for the pseudorandom sequence might be greater than, equal to, or less than the error rate for the D1 bank output. In most cases, errors occurred only when the signal source was a channel bank; sometimes, however, errors occurred only when the signal was supplied by the pseudorandom generator. This does not necessarily mean that T1 lines in general have higher error probability when driven by a D1 channel bank, because lines were usually selected for test on the basis of errors observed with a channel bank as the signal source, so that the sample of lines is biased.

In addition, it was expected that errors would be more likely to occur in those parts of the pseudorandom sequence where the constraints were most severely violated. As described below, two of the lines tested showed this tendency to a remarkable degree, but the rest showed no such tendency at all.

The pseudorandom sequence differs from the output of a channel bank both in maximum length of strings of zeros and in average ones density. The D1 channel bank must have a one in each word; the longest string of zeros that it can generate is 15. The pseudorandom sequence, on the other hand, contains eight strings of more than 15 zeros, including one of 19 zeros, in each repeated frame of a million bits (actually 1,048,575 bits, about 0.68 seconds).

The average ones density is nominally one-half for both, but in the channel bank output it varies from one bank to another, while in the pseudorandom sequence it varies over the million-bit frame. In the D1 channel bank, the ones density depends on how the analog-to-digital

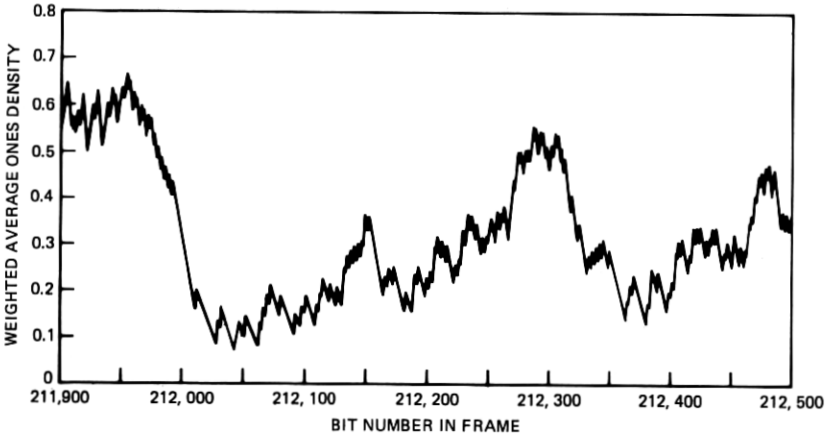

Fig. 23—Effective ones density for a timing tank with $Q = 60$, as a function of bit position in the pseudorandom sequence, in the vicinity of the string of 19 zeros, showing the location of the lowest effective density of ones in the sequence.

encoders (of which there are two in each bank) are adjusted to encode the zero level. The intended zero code is 1000000, which leads to a low ones density. However, with a slight misadjustment, zero might be encoded as the next lower code level, 0111111, which leads to a high ones density. The ones density is also affected by the signaling bit (a one for the on-hook condition) and by the encoding of noise and speech. One-second measurements on D1 channel banks in service have shown ones densities ranging from about one-fourth to three-fourths, with little variation over the day for any given bank.[10] Other observations have shown ones densities as low as about one-eighth on T1 lines in service.[11]

In the pseudorandom sequence, the average ones density over the million-bit frame is almost exactly one-half (actually 524,288/1,048,575). The variation of ones density within the frame can be evaluated as a weighted average of the past bits, using exponential weighting with a time constant of Q/π bits (about 32 bits for the typical timing-tank Q of 100). This average is theoretically proportional to the amplitude of ringing in a repeater timing tank circuit having the specified Q.

Figures 23 through 25 show, for three values of Q, the variations of ones density following the string of 19 zeros. (Since the ones density increases at each one and decreases at each zero, the strings of zeros can be identified by their downward slope. The string of 19 zeros starts at the 211,993rd bit after the framing bit. The framing bit itself was the last in the string of 20 consecutive ones.) The ones density has an absolute minimum of about one-fourth in this segment, but the exact value and location of the minimum depends on the timing tank Q.

Only two of the T1 lines that were tested had a noticeable tendency

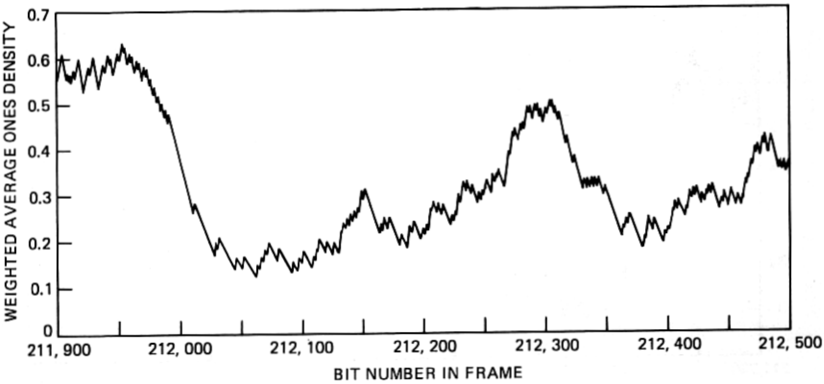

Fig. 24—Like Fig. 23, but for timing tank $Q = 100$. Note the change in the position of the minimum effective ones density as well as the change in minimum value.

to make errors in this part of the frame. One of these lines had an error rate of about 4×10^{-8} for the channel bank pulse stream (55-minute test), but had a steady error rate of only 3×10^{-9} for the pseudorandom sequence. On this line, errors occurred most often on the sixth one after the run of 19 zeros, on the first or second zero after that, or the eleventh one after the 19-zero sequence. These would be the locations of the lowest timing tank amplitudes assuming a Q of about 100, as in Fig. 24.

The other line, recorded at office 1, was a maintenance spare modified (as a stress test) by changing the line buildout (LBO: a circuit installed to equalize the loss of repeater sections of different lengths) in the office repeater. As installed, with an 836A LBO, the line made no errors. Substitutions of 836B through 836E had no apparent effect, an 836F resulted in an error rate of 2.5×10^{-6} (several errors per frame), and with an 936G the errors occurred too fast to be recorded. With the 836F LBO, errors

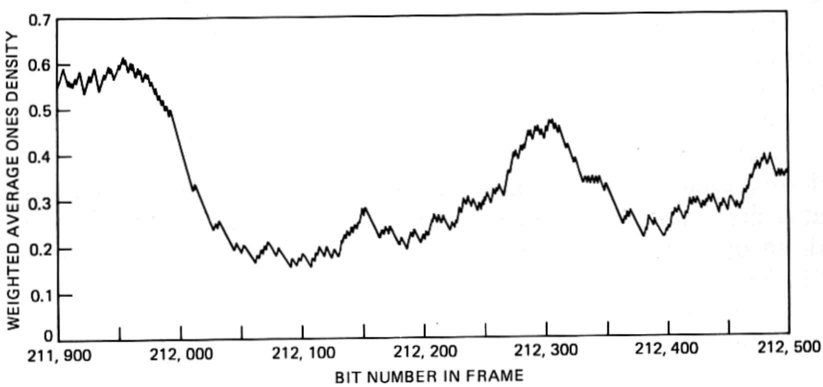

Fig. 25—Like Fig. 23, but for timing tank $Q = 140$.

Table VI — Frequency of occurrence of combinations of characteristics of errors of BPVs on 22 T1 lines*

Distribution type	Diurnal variation type			
	Normal	Steady	Irregular	Exaggerated
Unimodal	*m*	*Hh*	*m*	*h*
Nearly unimodal	*m*	*HH*	*Lmmm*	
Bimodal		*H*	*m*	*L*
Trimodal	*LLLLLll*			

* The letter indicates run type and pattern sensitivity: capital = error run, lower case = BPV run; pattern sensitivity is high (H/h), medium (m), or low (L/l).

occurred most often on the fourth and sixth ones after the string of 19 zeros, and sometimes on the two zeros after the seventh one. The fourth and sixth ones would just follow the two lowest values of effective ones density if the timing tank Q were about 60, as in Fig. 23.

4.6 Association between properties of the error process

Table VI suggests that the different characteristics of the error process tended to occur in certain typical combinations. The Appendix describes how the statistical significance of the apparent tendencies was verified. The most consistent combination was normal diurnal variation with a trimodal interval distribution (remaining trimodal as the error rate varied) and weak pattern dependence. Every line with a trimodal interval distribution had normal diurnal variation. Conversely, nearly every line with normal diurnal variation had a trimodal interval distribution. Every line with a trimodal interval distribution also had weak pattern dependence. This combination occurred in both error and BPV runs.

Another frequent combination was steady error rate with strong pattern dependence and a unimodal or nearly unimodal interval distribution. Every line with a steady error rate had strong pattern dependence. Every line with strong pattern dependence (except the two lines that made errors only at the minimum of ones density in the pseudorandom sequence) had a unimodal or nearly unimodal interval distribution.

Other combinations were less consistent. A substantial number of lines had irregular variation of error rate, most of them in BPV runs; these lines had either unimodal or bimodal interval distributions and usually showed an intermediate degree of pattern dependence. Both the irregular variation and the ambiguous indication of pattern dependence might be explained by the variability of the channel bank output bit stream.

These observations indicate that at least two different error mechanisms can be observed. In one, the error rate varies with traffic, but depends very little on the bit pattern on the line, and two levels of error clustering occur as indicated by the trimodal interval distribution. In the other, the probability of error is very much dependent on the pattern of bits on the line, the error rate remains steady as long as the same bit stream is repetitively transmitted, and there is little or no error clustering. It might be conjectured that errors are related to switching noise in the former type, and to intersymbol interference in the latter.

Occasional short intervals of very high error rate occurred in several lines. Samples during these intervals usually showed measured error rates above 10^{-3}, and sometimes above 0.1, with little or no pattern dependence, and about 50 percent deletions in the error runs. It is clear that in these events the errors were quite unaffected by the pattern of bits on the line, but the cause is unknown.

The great difference between the error rate in the pseudorandom sequence and the error rate in the channel bank bit stream is unexplained. It probably is not due to short term effects of ones density (on the time scale of variations in the timing tank output), as discussed in Section 4.5. On lines that show strong pattern dependence, BPV runs and error runs might be expected to give different results. But it is not clear why a line can show BPVs in service with very little pattern dependence (evidenced both directly by the distribution of errors over the frame, and indirectly by normal diurnal variation), and yet transmit the pseudo-random sequence without error.

4.7 Practical implications of the detailed recording results

The clearest general implication that can be drawn from the detailed error recordings is that error rate measurements on digital facilities should be interpreted with caution. For many possible reasons, the error rate measured on a T1 line may be quite different from the error rate experienced when the line is used for communication.

An error rate measurement made at night, especially in the early morning hours, can be misleading because many lines consistently have much lower error rates at that time than during the business day. Furthermore, since some lines have irregularly varying error rates, a single measurement during the business day would not necessarily show the typical performance of a line. It would seem that continuous monitoring, or at least frequent sampling, would be required to evaluate the error rate performance of an individual line.

It may also be misleading to test a line with a special test signal, or with any signal that is different from the communication signal that it normally carries. The measurement program showed conclusively that on many lines the error rate for a pseudorandom test sequence is quite different from the error rate for the D1 channel bank output carried by the line in service. Hence the error rates for a digital data signal, or for another channel bank output (either a different type, or a different unit of the same type), might be different from the error rates measured either with the line in service, or driven by a "quasirandom signal source" (a single generator used in telephone central offices, which generates a sequence similar to our pseudorandom sequence, except that ones are inserted to avoid long runs of zeros). However, the properties of the signal that affect the error rate remain unidentified.

Other implications are less clear. The relationships observed among diurnal variation, pattern sensitivity, and error-free internal distribution, suggest the existence of at least two clearly distinct causes of errors on T1 lines. It has been suggested that these are the same as the two major sources of error considered theoretically by Cravis and Crater:[2] The errors characterized by normal diurnal variation and low pattern dependence would be attributed to impulse noise originating in switching machines, which can be strong enough, when it occurs, to cause errors regardless of the bit pattern; the errors characterized by high pattern dependence would be attributed to crosstalk, which would cause errors mainly in those patterns in which intersymbol interference was most severe. This appears plausible as a tentative hypothesis. However, no attempt was made during the measurement program to identify the causes of the errors on individual lines, because of the extensive effort that would have been involved.

V. CONCLUSIONS

The T1 error measurement program of 1973–74 has resulted in both a survey of the population of T1 lines and some detailed observations of the error process.

The survey is in some ways less detailed, less systematic, and less accurate than previous (unpublished) surveys, but is probably overall the best overview we have of the error performance of T1 lines (when terminated predominantly with D1 channel banks). The large differences between offices show that a trend in the distribution of error rates on T1 lines, as a function of either time, or growth over time, cannot be inferred from comparison of surveys taken at different offices at different times. In other respects our results are consistent with previous surveys. Detailed conclusions were discussed in Section 3.3.

The detailed error recordings have shown a few remarkable results. The fact that many T1 lines can make errors in service without making any errors on a pseudorandom signal is both unexpected and unexplained. Another notable result is the observable existence of at least two different patterns of errors: one with a seemingly traffic-related diurnal variation, two levels of error clustering, and less sensitivity to the bit patterns on the line; the other with simpler patterns of variation and clustering, but more sensitivity to bit patterns. Another result is the fact that errors are always dependent to some extent on the bit patterns on the line, except at very high error rates. Detailed results of the intensive error measurements were discussed in Section 4.6.

VI. ACKNOWLEDGMENTS

The author entered the jitter and error measurement experiment after it began, when the measurements at office 2 were being made. The error measurement program was begun by C. A. Richardson, who designed the experiment, built the error measurement system (and the 16-line "BPV box" used for surveys after office 1) and produced the first graphical output and tentative conclusions from the data from office 1. P. C. Lopiparo and T. C. Spang worked with him and continued the experiment after that point. G. Zaccaria was responsible for the real-time monitor software for the PDP-11 computer. B. R. Wittig programmed the analysis of the survey data. A. K. Jain provided suggestions and critical comments on the statistical evaluation of the detailed recording results. The cooperation of the many telephone operating company personnel who worked with us at the test sites, and of the many Bell Laboratories members who operated the equipment, is gratefully appreciated.

APPENDIX

Statistical Significance of Apparent Association of Error Characteristics

The statistical significance of the frequencies of particular combinations in Table VI was verified by a test based on the chi-square test for independence in contingency tables. In order to conclude that Table VI shows an actual tendency toward certain combinations we had to determine that such an apparent tendency could not easily have occurred by chance. Two problems were met in the testing procedure. First, the null hypothesis to be tested must not be the hypothesis that all four dimensions in Table VI are independent, because the apparent tendency for certain characteristics to be associated with BPV runs is explainable as a property of the measurement technique; the test must allow for this and determine whether any further interdependence can be deduced from the data. Second, the sample size is so small, and the number of cells so large, that the chi-square distribution is not applicable.

The first problem was not to be solved by simply combining the error runs with the BPV runs, because this would leave medium pattern sensitivity apparently associated with irregular diurnal variation, actually because both are associated with BPV runs. Instead, the dependence of pattern sensitivity on run type was allowed for by considering the five observed combinations of these variables as a single variable. The dependence of diurnal variation on run type was then removed from the table by combining the steady and irregular types into one. The resulting 3-way contingency table is shown in Table VII.

Table VII — Three-way contingency table derived from Table VI and used in statistical testing of the significance of the results*

Run type and pattern sensitivity	Distribution type: Unimodal	Nearly unimodal	Bimodal	Trimodal	Total	Estimated probability
error, high	$1V$	$2V$	$1V$	0	4	0.182
error, low	0	$1V$	$1X$	$5N$	7	0.318
BPV, high	$1X, 1V$	0	0	0	2	0.091
BPV, medium	$1N, 1V$	$1N, 3V$	$1V$	0	7	0.318
BPV, low	0	0	0	$2N$	2	0.091
Total	5	7	3	7	22	
Estimated probability	0.227	0.318	0.136	0.318		1.0

Total of $N = 9$, $V = 11$, $X = 2$
Estimated probability of $N = .409$, $V = .500$, $X = .091$

* N = normal diurnal variation, V = variable (steady or irregular), X = exaggerated.

The estimated probabilities for each type in Table VII were set equal to the relative frequencies. As in an ordinary chi-square test for independence, the probability of each of the 60 possible combinations of types was derived (by a computer program) by multiplying the corresponding type probabilities. For example, the probability that a line would combine unimodal distribution type ($P = 0.2$), normal diurnal variation ($P = 0.45$), and "error high" run type and pattern sensitivity ($P = 0.2$), is $0.2 \times 0.2 \times 0.45 = 0.018$. The expected number of occurrences of each combination, in 22 lines, is 22 times the corresponding probability. The statistic called chi-square, measuring the deviation of the observed numbers, o_i, from the expected numbers, e_i, is derived as

$$\chi^2 = \sum_{i=1}^{60} (e_i - o_i)^2/e_i$$

For the actual observed numbers, chi-square was 88.04.

In a large sample this value would be compared with critical values based on a chi-square distribution with 50 degrees of freedom (the number of combinations of types, minus one, minus the number of independent probability values—not counting the ones determined by the constraint that probabilities must add up to 1—that were estimated from the observed data). If we did this we would conclude that a value of chi-square as large as 88.04 would occur by chance, if the three properties were independent, with a probability of about 0.01 percent, and that the deviation was therefore certainly significant. However, this reasoning is not valid, because with a sample as small as 22 the chi-square statistic does not have, even approximately, a chi-square distribution.

A Monte Carlo method was therefore used to estimate the relevant distribution. The model that was simulated on a computer considered three independent properties (diurnal variation, waiting time distribution, and run-type combined with pattern sensitivity), each property having 22 perdetermined outcomes distributed according to the type totals shown in Table VII. To implement each trial, a random-number generator assigned one of the 22 outcomes for each property (by drawing without replacement) to each of 22 lines. In 10,000 trials, 62 trials had values of the chi-square statistic greater than 88.04.

We could thus estimate that, under the null hypothesis, the probability that the chi-square statistic would exceed 88.04 is 0.0062, but this estimate has some uncertainty that we cannot easily estimate or allow for. The result, therefore, was interpreted by the following reasoning. If the null hypothesis is true—that is, if the three properties examined are independent—the single field experiment and the 10,000 Monte Carlo trials are replications of the same experiment, and the value of chi-square for the field experiment was among the 63 highest values in 10,001 trials. This would be an event with probability 0.0063, or 0.63 percent. This is small enough to consider that the deviation from independence is statistically significant.

A similar procedure was followed with another statistic, max $|o_i - e_i|$. This had a value of 4.09 for the field trial, which was equaled in 3 out of 10,000 computer trials (and never exceeded). Under the null hypothesis this would be an event with probability 0.0004, or 0.04 percent, showing even more clearly a significant deviation from independence.

Having concluded that independence of the properties does not account satisfactorily for the observations, we are then justified (by Occam's razor) in adopting the simplest hypothesis that does account for them, which is that we have observed two different types of error process, each having different probabilities for the properties.

REFERENCES

1. K. E. Fultz and D. B. Penick, "The T1 Carrier System," B.S.T.J., *44,* No. 7 (September 1965), pp. 1405–1451.
2. H. Cravis and T. V. Crater, "Engineering of T1 Carrier System Repeatered Lines," B.S.T.J., *42,* No. 2 (March 1963), pp. 431–486.
3. J. S. Mayo, "A Bipolar Repeater for Pulse Code Modulation Signals," B.S.T.J., *41,* No. 1 (January 1962), pp. 25–97.
4. P. C. Lopiparo, to be prepared.
5. J. J. Mahoney, Jr., J. J. Mansell, and R. C. Matlack, "Digital Data System: User's View of the Network," B.S.T.J., *54,* No. 5 (May–June, 1975), pp. 833–844.
6. A. K. Reilly, unpublished work.
7. J. E. Kessler, unpublished work.
8. E. W. Geer, Jr., unpublished work.
9. S. W. Golomb, *Shift Register Sequences,* San Francisco: Holden-Day, 1967.
10. H. Hagen, unpublished work.
11. R. F. Ewald, unpublished work.