THE BELL SYSTEM TECHNICAL JOURNAL
Vol. 57, No. 6, July–August 1978
Printed in U.S.A.

Atlanta Fiber System Experiment:

Results of the Atlanta Experiment

By R. S. KERDOCK and D. H. WOLAVER

(Manuscript received October 20, 1977)

In the Atlanta System Experiment, a digital lightwave transmission system operating at 44.736 Mb/s was tested in an environment approximating field conditions. The system included a cable with 144 fibers pulled into a duct, lightguide cable splices, single-fiber connectors, a fiber distribution frame, optical regenerators employing avalanche photodiodes and GaAlAs lasers, and terminal equipment for interfacing to the standard DS3 signal format. A number of experiments were conducted to evaluate and gain experience with this digital lightwave transmission system. These included fiber crosstalk, fiber dispersion, timing jitter, system recovery, and dc powering. In this paper the configuration of the Atlanta lightwave system is briefly described, and experimental results are reported.

I. INTRODUCTION

Tests of an experimental lightwave transmission system operating at 44.736 Mb/s (DS3), the third level of the Bell System digital hierarchy, were begun by Bell Laboratories in January, 1976.[1-3] The system, operating in an environment simulating field conditions, included all the elements necessary for a practical transmission system at the DS3 rate: a small, ruggedized cable containing 144 fibers, lightguide cable splices, single-fiber connectors, a fiber distribution frame, regenerators* employing avalanche photodiodes and GaAlAs lasers, and terminal circuits to interface with the DS3 signal format.

The main goals of the system experiment were to gain experience with a digital lightwave transmission system and to obtain the data necessary to turn this new technology into practical telecommunication equipment.

* A regenerator is a 1-way device, a repeater comprises two 1-way regenerators.

Before the system experiment began, the individual elements comprised by the system were thoroughly tested and characterized. The system experiment brought all these elements together in a simulated field environment. Experiments were performed to measure important system parameters and to evaluate system performance. These included experiments on fiber crosstalk, fiber dispersion, timing jitter, system recovery, and dc powering.

II. SYSTEM DESCRIPTION

A 648-meter length of cable containing 144 fibers[4,14] was looped through an underground duct and two manholes so that both ends of the cable were accessible in a basement laboratory (see Fig. 1). Each end of the cable was spliced[5] to a short section of cable that fanned out to individual fibers terminating in a 12-by-12 optical connector array. The array, called the fiber distribution frame (FDF), provided convenient access to individual fibers. Optical patch cords with single-fiber connectors[6] were used to interconnect the fibers and to connect the fibers to the regenerators.

The regenerators[7] accept at their input an average optical power ranging from −20 to −55 dBm. The data format is an unrestricted, binary, nonreturn-to-zero signal that has been scrambled. The data rate is 44.736 Mb/s, the DS3 rate of the Bell System digital hierarchy. The output is a regenerated lightwave signal of 825-nm wavelength and an average power of at least −3 dBm. (Descriptions of the avalanche pho-

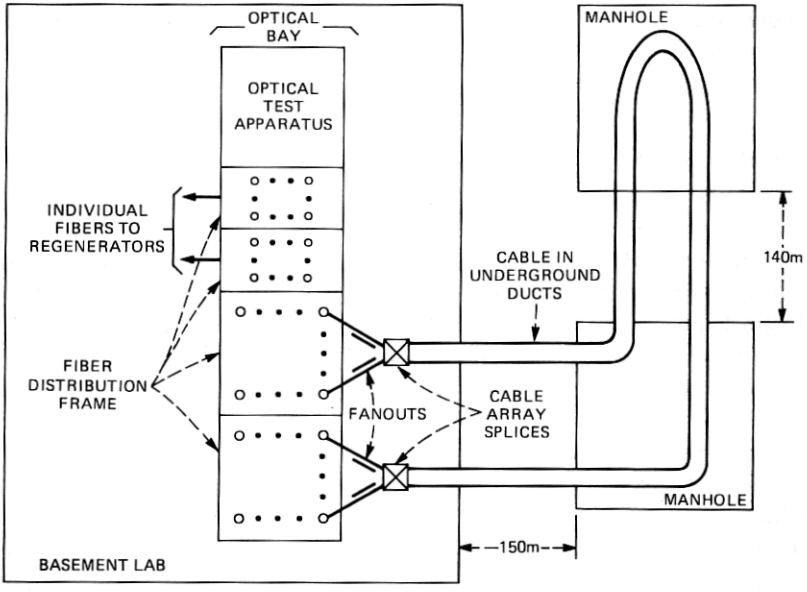

Fig. 1—Lightguide cable configuration showing access to individual fibers.

todiode[8] and GaAlAs laser[9] used in the regenerator are given elsewhere.) Optical patch cords are used to bring the input and output to the FDF.

The terminal circuits[2] provide the interface between the standard bipolar DS3 signal format and the unipolar format required by the regenerators. They also provide parity violation monitoring and removal. Figure 2 is a functional block diagram of the terminal circuits. Three pairs of terminal circuits were available.

These system elements were connected in various arrangements to form experimental setups. A typical experimental setup is shown in Fig. 3. As many as seven loops of the 648-m fibers were included between regenerators, and as many as 14 regenerators were included in a maintenance span.* Up to three maintenance spans could be connected in tandem. A signal source provided pseudo-random data with appropriate framing and parity bits.

III. CROSSTALK EXPERIMENT

In the configuration of the lightguide cable,[4] 12 fibers are placed side by side to form a ribbon, and 12 ribbons are stacked to form the cable

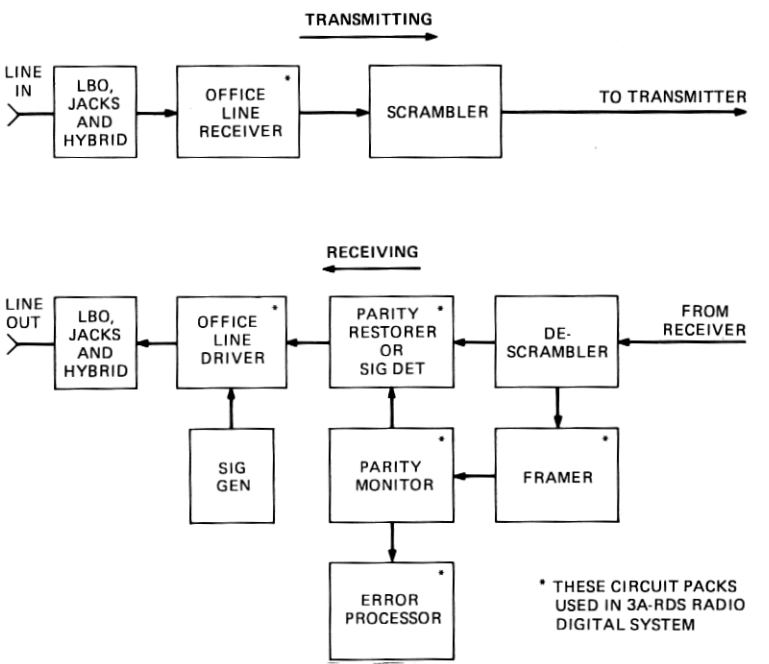

Fig. 2—Functional block diagram of the terminal circuits.

* For one experiment, fibers in a second lightguide cable 644 meters long were joined with low-loss splices to provide extra-long fiber lengths with up to 17 loops between regenerators.

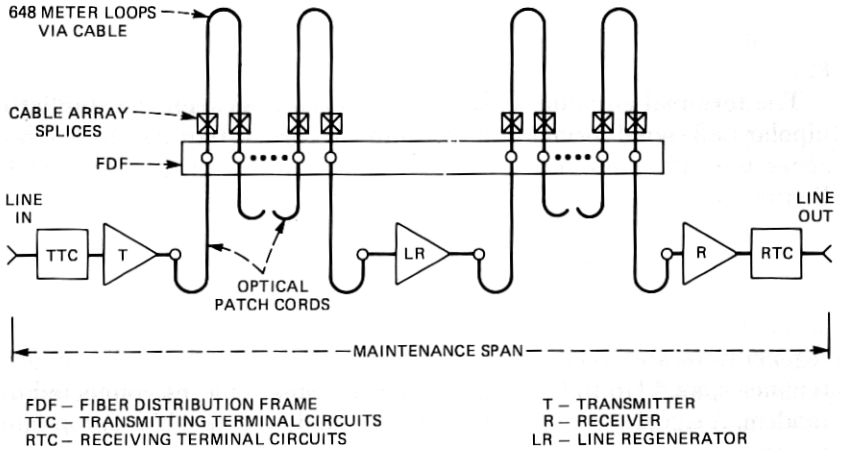

Fig. 3—Typical interconnection of system elements.

core. Thus most of the fibers have four neighbors. If stray light coupling is a problem, it would most likely be between neighboring fibers either along the length of the fibers or at cable splices.

The primary aim of this experiment was *not* to measure the amount of crosstalk coupling between fibers. Such data, obtained by a more sensitive technique, are reported elsewhere.[10] The purpose here was to measure the effect of crosstalk on system performance in terms of increased error rate. We also develop a model for the effect of crosstalk and compare measured results with a computer simulation.

3.1 Measured effect of crosstalk

Figure 4 shows the test arrangement for measuring the effect of crosstalk. Fiber A carries a lightwave signal with pseudorandom data, and the received signal is monitored for error rate. The level of the lightwave signal is adjusted by an optical attenuator so that the error rate is about 10^{-6}. (With a baud of 44.7 Mb/s and a count time of 10 seconds, there are about 447 error counts.) Then a lightwave signal 38 dB more powerful* is introduced into fiber B adjacent to fiber A. If this causes the measured error rate to increase by 14 percent or more (3σ of the count variation), there is defined to be measurable crosstalk.

Each of the 138 transmitting fibers in the cable was taken in turn as fiber A, and in each case the adjacent fibers (as many as four) were taken in turn as fiber B. In all, 456 fiber pairs were tested for crosstalk. If the measured error rate increased by a factor less than 1.14, the effect was attributed to statistical variation, and the pair was not investigated

* This is intended to be near the maximum signal level difference between two adjacent fibers in an actual application.

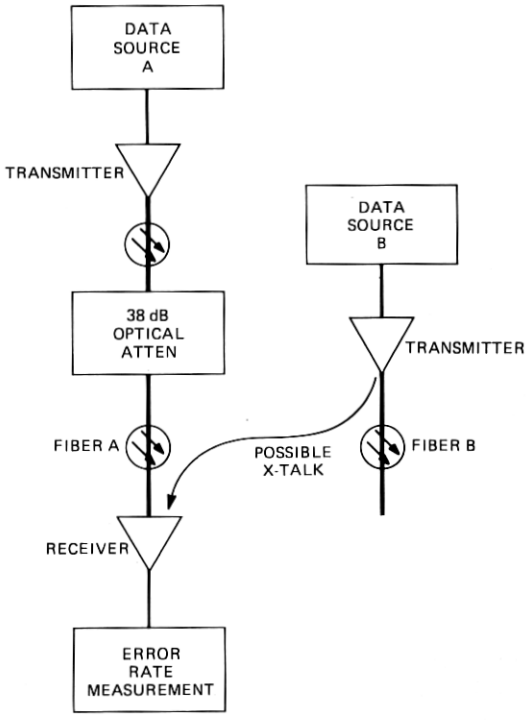

Fig. 4—Arrangement for crosstalk measurements.

further. There is some probability that actual crosstalk effects were overlooked. For example, an actual error rate increased factor of 1.30 would be overlooked with a probability of 3 percent.

Those fiber pairs with a measured error rate increase factor α more than 1.14 were tested again, measuring the error rates for 10 seconds. If α was again less than 1.14, the effect was attributed to statistical variation.

Two cases of fiber pairs with measurable crosstalk remained. The results are illustrated in Fig. 5. The arrows indicate the direction of the crosstalk, and α is the measured error rate increase factor. In determining α here, the error rates were measured for 100 seconds, giving a standard deviation of about 0.03 for α. The case with the greatest amount of crosstalk had $\alpha = 2.0$. This corresponds to a decrease of 0.25 dB in system sensitivity (see Fig. 5 of Ref. 7).

Both crosstalk situations involved fiber number 1-11. Subsequent investigation by Buckler and Miller[10] determined that crosstalk did not take place along the length of the fibers. An anomaly apparent in the splice of fiber 1-11 at the end labeled "W" produced the measurable crosstalk.

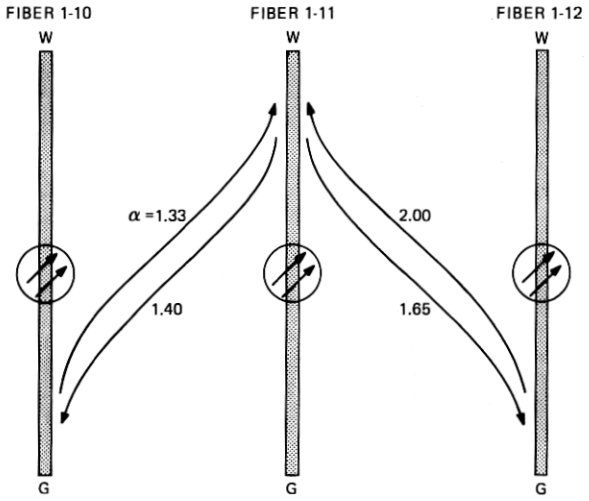

Fig. 5—Only case of measurable crosstalk. Arrows show direction of crosstalk. α is error rate increase factor.

The conclusion is that crosstalk measurably affecting the lightwave transmission system is indeed rare (two cases out of 456). Even the worst crosstalk case had a negligible effect on the system performance (0.25 dB decrease in sensitivity).

3.2 Model for effect of crosstalk

In this section, we develop a relation between strength of crosstalk interference and the error rate increase factor. The work is based in part on a computer simulation. There is one experimental data point.

As the interference from crosstalk is high (logical one) and low (logical zero), it increases the error rate of the main signal in two ways. Because of ac coupling in the receiver, it shifts the main signal both down and up by an amount Δ proportional to the interference power. This is equivalent to shifting the decision threshold up and down by Δ. When the interference is high, it also increases the error rate by increasing the shot noise and excess multiplication noise.[11]

In our case, the interference is high 50 percent of the time. We assume that the effect of low threshold and added noise is the product of the effect of each separately. Then the error rate increase factor is

$$\alpha = \tfrac{1}{2}(\alpha_t(\Delta) + \alpha_t(-\Delta)\alpha_n), \tag{1}$$

where α_t and α_n are the error rate increase factors due to threshold offset and light-associated noise, respectively.

The function $\alpha_t(\Delta)$ has been measured (see Fig. 12, "Error Rate Versus

Decision Threshold," of Ref. 7), and we can relate the offset to the interference power by

$$\Delta = 500(P_I/P_S)\text{mV}, \tag{2}$$

where P_I is the interference power and P_S is the signal power. To relate α to P_I/P_S we must also find α_n as a function of P_I/P_S.

If the crosstalk interference is constantly high (not PCM modulated), it will not shift the signal relative to the threshold because of ac coupling. Its only effect on the error rate will be through light-associated noise. This technique was used to measure α_n for one of the crosstalk cases in Fig. 5. For the case $\alpha = 1.65$, we measured $\alpha_n = 1.76$. Then from eqs. (1) and (2), P_I/P_S was calculated to be -17.3 dB for the modulated case. Since the signal in fiber B was 38 dB greater than the signal in fiber A, the crosstalk coupling was -55.3 dB for this case. This is in agreement with measurements by a different technique of the same crosstalk coupling.[10]

Only that one experimental point of α versus P_I/P_S was determined. To generate a complete curve, a computer simulation incorporating a Chernoff bound[12] was used to relate α_n to P_I/P_S. This, together with eqs. (1) and (2), gave the curve of error rate increase factor α versus interference-to-signal ratio in Fig. 6. The simulated results and the experimental point agree within a decibel.

IV. DISPERSION EXPERIMENT

When a pulse of light is launched into a fiber, the components of the pulse will arrive at the end of the fiber with different delays, depending upon the transit times for the modes. Therefore a transmitted pulse becomes dispersed in time as it travels through the fiber. Fibers with a graded index[13] are designed to reduce this effect, but some pulse spreading persists.

A spreading of the impulse response is equivalent to a reduction of the bandwidth. The bandwidths of the fibers in the cable are given in Ref. 14. The average 3-dB bandwidth (half-light power) of the 648-m cabled fibers is 690 MHz with a standard deviation of 286 MHz.

At a bit rate of 44.7 Mb/s, the effect of dispersion on system performance was expected to be small, so the six fibers with smallest bandwidth (high dispersion) were selected for a dispersion study. The average 3-dB bandwidth for the six is 177 MHz with a standard deviation of 27 MHz.

4.1 Measured effect of dispersion

The effect of dispersion was determined by comparing the system performance using high-dispersion fiber with the system performance using dispersionless attenuation (a neutral-density filter) in place of the

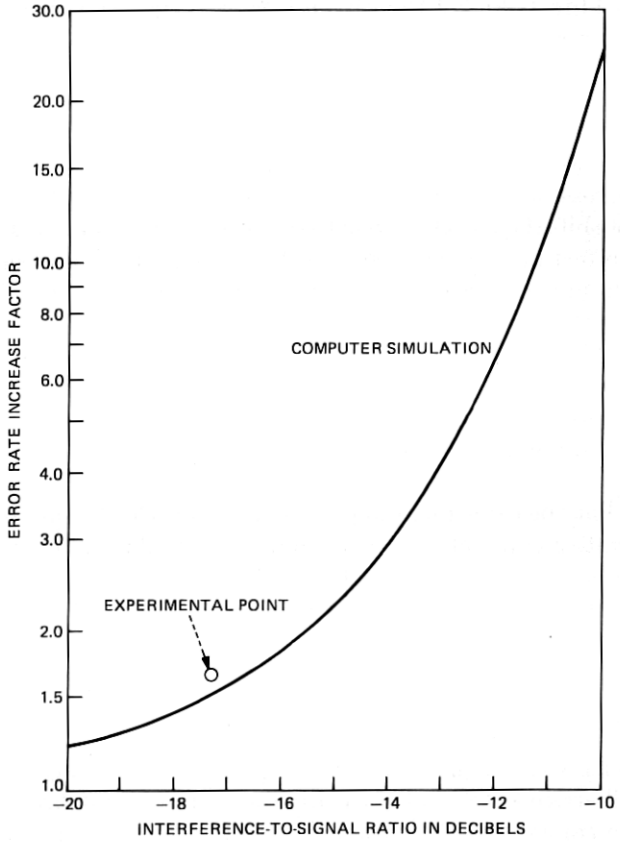

Fig. 6—Effect of crosstalk interference.

fiber. The measure of performance used here is the "system sensitivity"—the detected light power* for a 10^{-6} error rate. The decrease in sensitivity due to dispersion is called the "dispersion penalty."

From one to six of the high dispersion fibers were concatenated using optical jumpers. These fibers were included in the transmission path along with enough dispersionless attenuation to make the error rate 10^{-6}. Then the fibers were removed from the path, and the dispersionless attenuation was increased until the error rate was again 10^{-6}. The difference in detected light power for the two cases is the dispersion penalty for that length of fiber.

The dispersion penalty is plotted as a function of high-dispersion fiber length in Fig. 7. The highest penalty measured was only about 0.5 dB.

* The detected light power is calculated from the measured APD current. For this measurement, the APD bias is set for a known APD gain, so the primary current is known. The detected light power is 1.5 watts per ampere of primary current (for light with a wavelength of 825 nm).

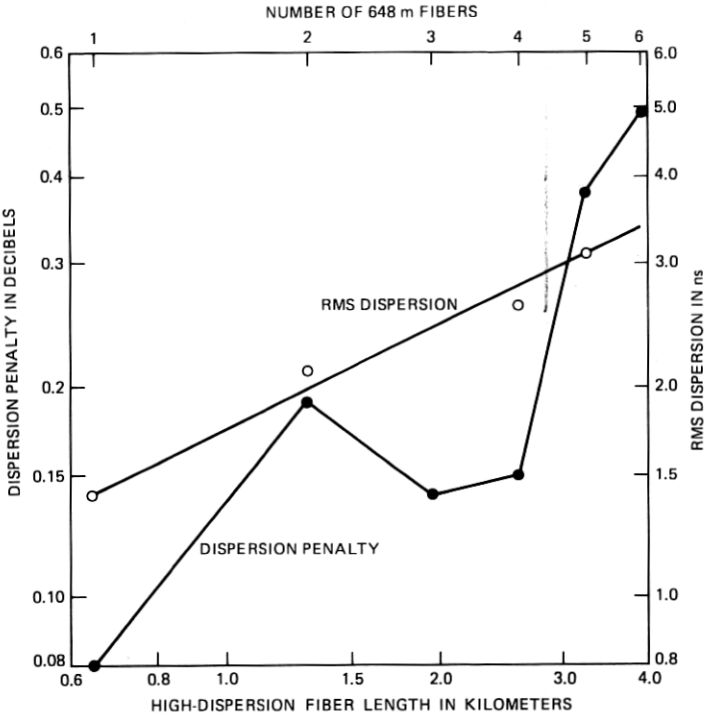

Fig. 7—Effect of high-dispersion fiber.

The curve should be monotonic increasing; the reversals are due to the limits of measurement accuracy.

4.2 Measured RMS dispersion

We have the dispersion penalties for a particular set of concatenated fibers. By measuring the dispersion of those fibers, we can extend the results to any fiber for which the dispersion is known.

The pulse responses for a few lengths of high-dispersion fiber were measured by S. D. Personick using the technique described in Ref. 14. Since the source pulse was only 0.5 ns wide, the pulse response can be considered to be an impulse response. The waveforms for four different lengths of the high-dispersion fiber are shown in Fig. 8. The impulse response broadens with increasing fiber length and approaches what is approximately a Gaussian waveform. Note that the waveform for the shortest fiber length is not smooth, indicating that mode mixing is not complete.

The impulse response is conveniently characterized by one parameter—the RMS dispersion σ defined by[11]

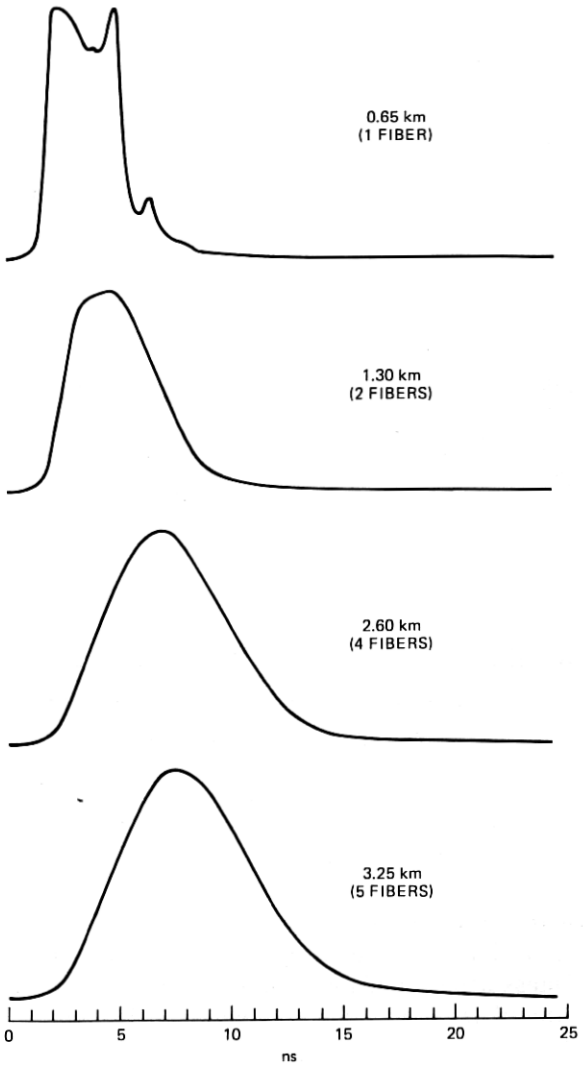

Fig. 8—Impulse responses of high-dispersion fiber.

$$\sigma^2 = \frac{\int t^2 h(t)dt}{\int h(t)dt} - \left[\frac{\int th(t)dt}{\int h(t)dt}\right]^2,$$

where $h(t)$ is the impulse response. The RMS dispersions for the four impulse responses in Fig. 8 are plotted in Fig. 7. The points lie close to a line with a slope of $\frac{1}{2}$, indicating a square-root dependence on length.

From the data in Fig. 7, we can determine the dispersion penalty

versus RMS dispersion (see the plot in Fig. 9). The curve through the data points was generated by a computer simulation using a Chernoff bound approximation of noise.[12] This curve holds for both high- and low-dispersion fiber. The corresponding lengths of high-dispersion fiber are indicated on a scale at the top of Fig. 9. For example, 7 km of high-dispersion fiber would have a dispersion penalty of about 1.0 dB. Even for this worst case, the dispersion penalty (for a bit rate of 44.7 Mb/s) is negligible compared to a loss of about 40 dB for 7 km of fiber.

4.3 Long fiber spans

There was some question whether the method of joining fibers at the fiber distribution frame adequately simulated field splicing methods in regard to dispersion characterization. Every time a 648-m fiber was added to a regenerator span, it was necessary also to add two optical fanouts with two splice interfaces and one optical patch cord with two connector interfaces (see Fig. 3). The numerous connectors and splices contribute significantly to the total loss of a fiber span, and they might also have significantly affected the measured dispersion by providing extra mode mixing.

To eliminate most of these connectors, an experiment was performed using a second cable. This 644-m cable was installed in the ducts in the

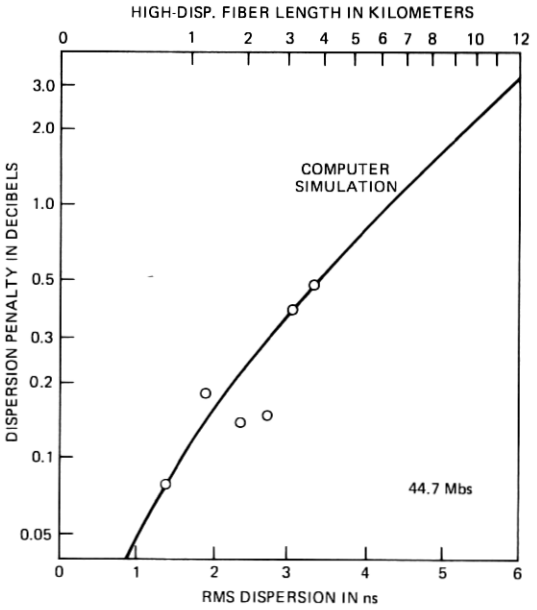

Fig. 9—Effect of RMS dispersion.

same manner as the first cable (see Fig. 1). For these tests, the fiber distribution frame was not used. Instead, fibers were concatenated using low-loss loose tube splices[15] to join individual fibers. One single-fiber connector was used at each end of the long concatenated fiber span to interconnect with the optical regenerators.

Two long fiber spans were constructed by C. M. Miller using these loose tube splices. Assuming that very little mode mixing occurred at the splices, the effect of fiber dispersion could now be accurately measured. Using the technique described above, the dispersion penalty was measured for each long span and compared to the span loss measured by M. J. Buckler.[1] One span was 10.3 km long with a dB loss of 49.2 and a dispersion penalty of 0.5 dB. The other span was 10.9 km in length and had a 47.9 dB loss and a dispersion penalty of 1.1 dB. It is clear that, as in the experiment using numerous connectors, the dispersion penalty is negligible compared to the loss. Therefore it is concluded that fiber-guide dispersion will have only a small effect on regenerator spacing at 44.7 Mb/s.

It should be noted that a length of nearly 11 km is not to be considered a "practical" regenerator spacing. In this experiment, low-loss fibers were selected, splicing losses were minimized, and very little allowance was made for operating margins. It has been estimated that these results translate into a regenerator spacing of approximately 7 km ($\approx$4 miles) for a practical transmission system employing current technology.

V. TIMING JITTER EXPERIMENT

In this experiment, a line of 14 regenerators was set up. The phase of the data at the output of each regenerator was compared with the clock phase of the data source—a maximum length pseudorandom sequence of length $2^{15} - 1$ scrambled by a 17-stage feedback register. In this way, the timing jitter accumulation was measured as a function of the number of regenerators.

Two types of jitter were investigated: random jitter and systematic jitter. In one experiment, as much random jitter as possible was introduced by operating each regenerator near an error rate of 10^{-6} (detected optical power near -55 dBm). The effect on the accumulated jitter was negligible. As will be explained below, random jitter tends to be swamped by systematic jitter as the number of regenerators in the line grows.

Two cases of systematic jitter were investigated. In one, offset current was purposely introduced in the phase-locked loop of each regenerator's timing recovery circuit. This caused a static timing offset μ of -12 degrees in each. In the other case, any offset current was nulled out to reduce μ to less than 1 degree. The measured results are shown in Fig. 10. That the jitter does not always increase monotonically with N is due to measurement error.

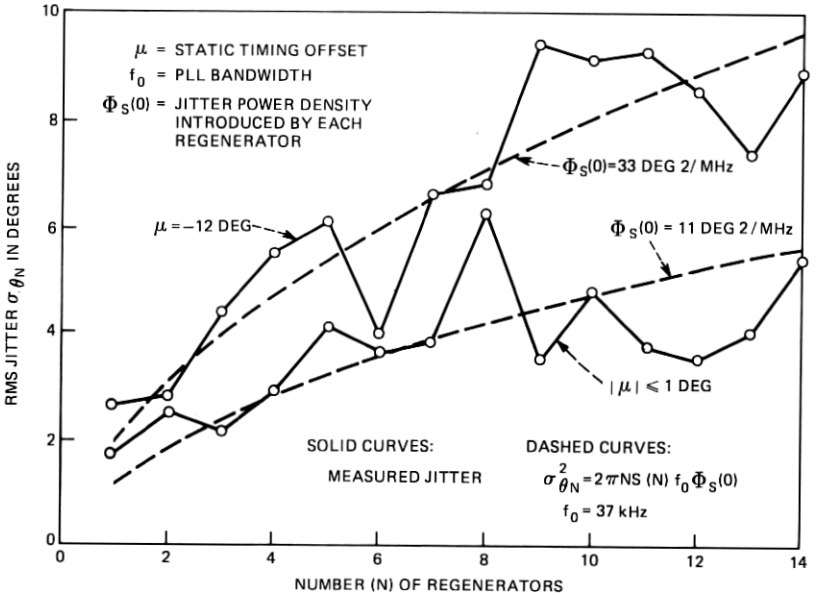

Fig. 10—Cumulative timing jitter. The function $S(N)$ is plotted in Fig. 11.

5.1 Jitter power density from fit to data

Each regenerator adds a certain amount of random jitter, caused by noise, and systematic jitter, caused by data dependence. We assume that each regenerator adds the same amount of jitter, and that the jitter power density spectra of the random and systematic components are $\Phi_R(f)$ and $\Phi_S(f)$, respectively. (The unit of the "power" density is degrees^2/Hz.) The spectra can be considered, for all practical purposes, to be flat and equal to $\Phi_R(0)$ and $\Phi_S(0)$. The cumulative RMS jitter $\sigma_{\theta N}$ is given by

$$\sigma_{\theta N}^2 = \sqrt{\pi N}\, R(N) f_0 \Phi_R(0) + 2\pi N S(N) f_0 \Phi_S(0), \qquad (3)$$

where N is the number of regenerators, f_0 is the average PLL bandwidth, and $R(N)$ and $S(N)$ are "correction factors" needed for small N; they go to unity for large N (see Fig. 11). The first term, giving the dependence of N for random jitter, is due to DeLange.[16] The second term (systematic jitter) is due to Byrne, et al.[17]

In both experiments the measured spectra of the jitter power density displayed the notches associated with systematic jitter[17] (see Fig. 12). The depth of the notches indicated that the contribution from random jitter was insignificant. Therefore the RMS jitter $\sigma_{\theta N}$ should grow with $\sqrt{N}$, as seen from the second term in (3). Using this relation, we fitted curves to the measured data by selecting $\Phi_S(0)$ for each case. For $\mu = -12$ degrees, the fitted $\phi_S(0)$ is 33 deg^2/mHz and for $\mu < 1$ degree, the fitted $\Phi_S(0)$ is 11 deg^2/mHz.

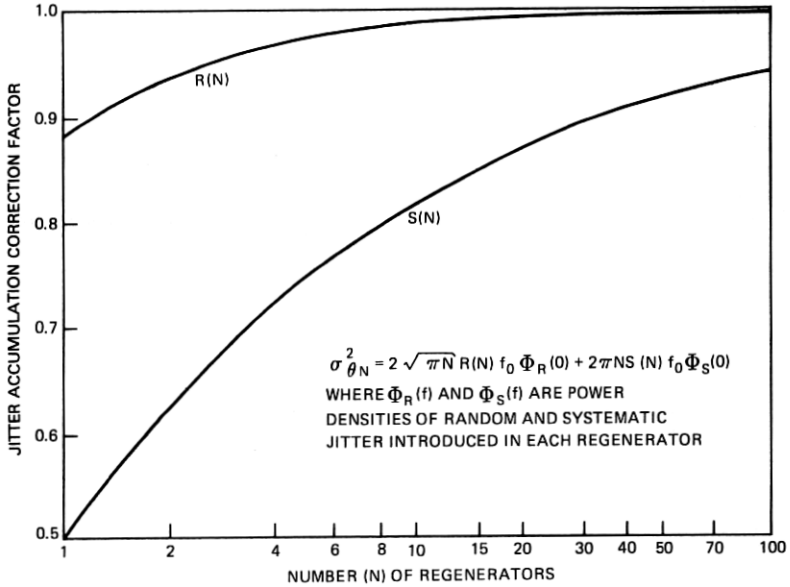

Fig. 11—Jitter accumulation correction factors.

5.2 Jitter power density calculated from model

Duttweiler[18] has done a computation of the jitter power densities $\Phi_R(0)$ and $\Phi_S(0)$ for a PLL that models the one used here fairly well. $\Phi_R(0)$ is calculated from the noise present at the output of the linear channel. $\Phi_S(0)$ is calculated from the static timing offset μ and from the pulse response of the linear channel (intersymbol interference). The analysis assumes a random data signal.

Using the pulse response in Fig. 9 of Ref. 7, the $\phi_S(0)$ for $\mu < 1$ degree was calculated to be 1.5 deg^2/MHz. This would produce only 37 percent of the measured jitter. The $\phi_S(0)$ for $\mu = -12$ degrees was calculated to be 7.7 deg^2/MHz. This would produce only 50 percent of the measured jitter. The discrepancy is not completely understood, but it may be due in part to the difference between random and pseudorandom data.

The measured data can be extrapolated by (3), neglecting random jitter. For $\mu = -12$ degrees and $N = 50$, the RMS jitter would be 19 degrees. This represents a worst case with the offset in the PLL of all regenerators at the maximum and in the same direction. Even this worst case is well within the capability of the terminal circuits, which can handle up to 30 degrees RMS jitter.

VI. SYSTEM RECOVERY EXPERIMENT

The objective of this experiment was to measure the response of the transmission system to signal interruptions of the sort that might be expected in actual application. We also wanted to gain an understanding

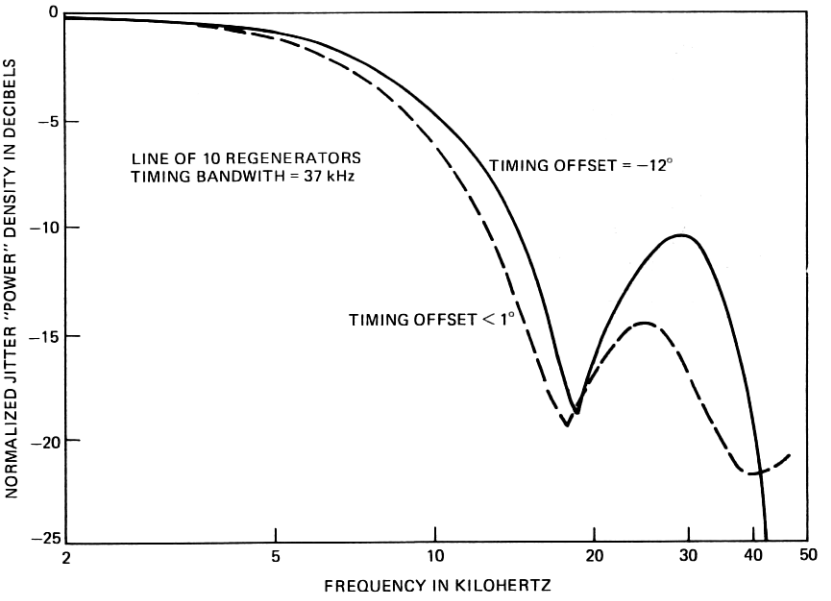

Fig. 12—Spectra of systematic jitter.

of the mechanisms affecting the recovery. To this end, recovery time was measured as a function of the interruption interval and the signal level. First we investigate the response of a single regenerator, then a line of regenerators in tandem, then a maintenance span with terminal circuits, and finally a number of maintenance spans in tandem.

In each test the data input to the system (or subsystem) was interrupted and set to logical "zero" for intervals of 0.2, 1.0, 3.0, and 10.0 ms. Following this, the data resumed.* The 0.2-ms interval is the expected maximum interruption that would be incurred during a protection switch. An interruption of 3.0 ms can occur due to the characteristics of the terminal circuits. A 10-ms interruption is typical of that during manual patching for service restoration or rolling.

6.1 Regenerator recovery

The input to a regenerator was supplied by a pseudorandom data source, and the output of the regenerator was monitored by a bit-error-rate test set (BERTS). The BERTS received a clock signal directly from the source so that it maintained synchronism during the interruption. The recovery time was defined as the time from the end of the interruption to the beginning of the first error-free second.

* Some tests were made in which the resumed data had a baud shifted 1.0 kHz from the original rate. However, this was found to have no measurable effect on the recovery time.

Typical regenerator recovery times are plotted in Fig. 13. The response is a function of both the interruption interval and the input signal level—the detected optical power. For interruptions up to 3.0 ms, the recovery times are 0.1 ms or less. For an interruption of 10 ms, the recovery times are between 0.4 and 4.0 ms. We will examine the mechanisms governing these responses.

Part of the recovery mechanism is illustrated in Fig. 14, which shows the signal at the decision point in the regenerator. After 0.3 ms or more of interruption time, the ac coupling in the linear channel causes the signal to be displaced upward by about 500 mV. When the data signal

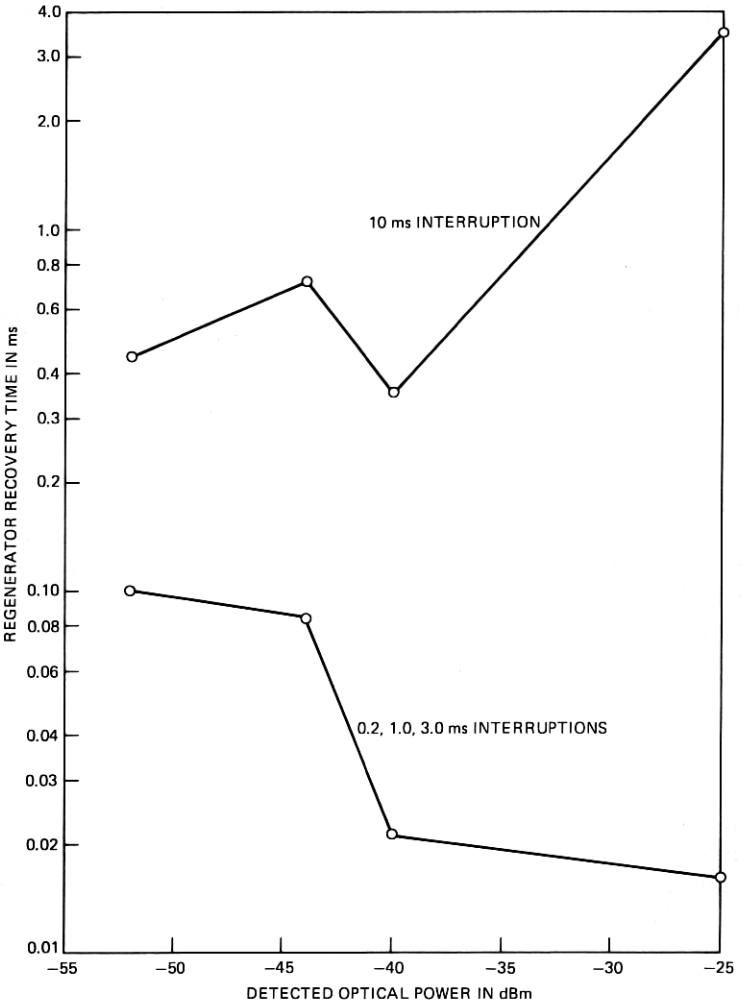

Fig. 13—Regenerator recovery times.

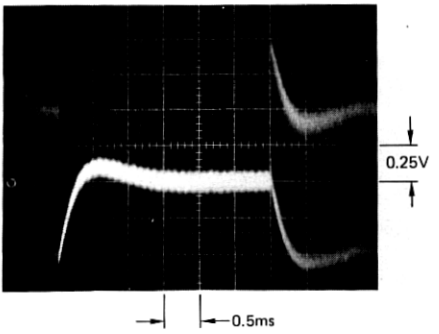

Fig. 14—Signal at decision point for 3.0-ms interruption and −44-dBm detected optical power.

returns, the signal in Fig. 14 drops to its proper position with the time constant of about 0.25 ms. The error rate is high until the "eye opening" moves down to include the decision threshold at −70 mV. Recovery is then complete. With greater detected optical power, the "eye" is more open, and the recovery time is shorter. This explains the characteristic of the lower curve in Fig. 13.

For a 0.2-ms interruption and high (−25 dBm) detected optical power, the "eye opening" does not have time to drift so far that it does not include the decision threshold upon the return of data. For this case the recovery might be expected to be immediate. However, another mechanism becomes important—the 15 μs recovery time of the phase-locked loop (PLL) in the timing circuitry. The result is that the response curve for a 0.2-ms interruption is about the same as that for 1.0-ms and 3.0-ms interruptions.

The upper curve in Fig. 13 indicates significantly greater recovery times for 10-ms interruptions. One reason for this is that the automatic gain control (AGC) has time to drift during the interruption so that the signal can be very large upon return of the data. In high (−25 dBm) detected optical power, this can lead to a saturation effect that incapacitates the linear channel of the regenerator for a number of milliseconds. This effect is shown in Fig. 15. After the data resume and the 0.25-ms transient is over, the signal remains unsymmetric for 2 to 4 ms. This indicates a badly distorted signal that results in errors.

For detected optical powers of −40 dBm and lower, the saturation effect does not occur. In this range, another mechanism dominates in causing longer recovery times for 10-ms interruptions. In some regenerators, there was enough offset in the PLL to cause it to drift out of its seize range during a 10-ms interruption. In that case, the recovery time of the PLL was about 0.5 ms.

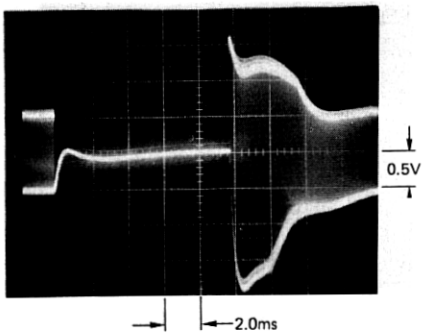

Fig. 15—Signal at decision point for 10-ms interruption and −25-dBm detected optical power.

6.2 Line recovery

Lines of as many as nine regenerators were constructed as in Fig. 3. In this part of the experiment, the terminal circuits were not included. The recovery times for various numbers of regenerators in tandem were measured by the method outlined for single regenerators. Typical results for a 3.0-ms interruption are plotted in Fig. 16. Notice that the recovery time grows more slowly than the number of regenerators.

It was not expected that the recovery time for a line of N regenerators

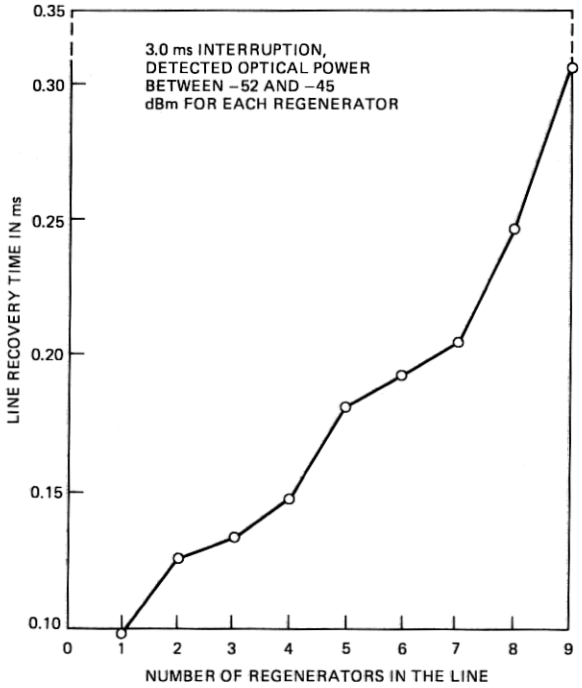

Fig. 16—Line recovery times.

would be N times the recovery time for one regenerator. The recovery mechanisms in the second regenerator can begin while the first regenerator is still recovering and putting out erroneous data.

The recovery of downstream regenerators is also affected by the activity of earlier regenerators *during* the interruption. With no lightwave signal at the input of the first regenerator, its PLL continues to free run. It continues to clock out whatever signal appears at the decision point, including noise. For example, Fig. 14 shows the noise meeting the decision threshold (−70 mV) after about 1.3 ms of interruption. For the rest of the interruption, the regenerator provides noise-like data to the second regenerator. For higher noise at the decision point, the interruption seen by the second regenerator becomes even more filled in.

The filling in of the interruption continues down the line of regenerators until only the first few tenths of a millisecond remain free of noise-like data. Figure 17 shows a typical example of the signal appearing at the decision point of the eighth regenerator in a line during a 10-ms interruption.

One effect of filling in the interruption is to keep the AGC from drifting out of adjustment. Therefore, the saturation effect, which causes long recovery times, can occur only in the first few regenerators in a line.

Another effect of filling in the interruption is to keep the PLL of each regenerator locked to the frequency of the PLL in the previous regenerator. This has little effect on the recovery time unless the PLL offset in one of the first few regenerators is large. In that case, the PLL drifts out of its seize range during a 10-ms interruption. If there is sufficient noise at the decision point, the noise-like data at the output cause all downstream regenerators to be pulled out of their seize range. For such a case, the recovery time for each regenerator is long, and line recovery times as high as 9 ms were measured for a line of nine regenerators.

Extrapolating the measured results to a line of 20 regenerators, we expect line recovery times of less than 1.0 ms for interruptions of 3 ms or less. Recovery times for interruptions longer than 3 ms can be greater.

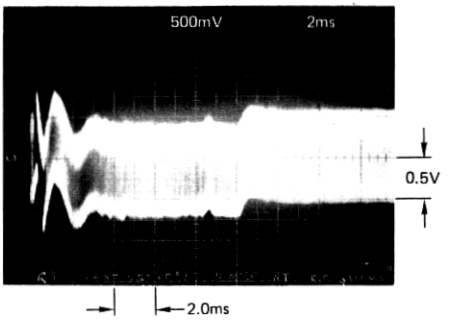

Fig. 17—Signal at the decision point of eighth regenerator in a line, illustrating the filling in of a 10-ms interruption.

It is planned to limit the PLL offset so that the PLL remains within its seize range for interruptions up to 10 ms. In that case, the expected line recovery time would be as much as 4 ms for interruptions from 3 to 10 ms. Such recovery times would be dominated by saturation effects in the first regenerator.

6.3 Maintenance span recovery

In this portion of the recovery experiment, the line of regenerators was extended to include the terminal circuit on each end (see Fig. 3). The recovery time was measured from DS3 interface to DS3 interface. The measurement method was, by necessity, different from that used for the regenerator and line recovery tests. The bit-error-rate test set that operates with the DS3 format requires framing to do bit error measurements, and the test set reframe time prevents accurate determination of recovery time. Therefore we need another measure of recovery time.

For some cases, the receive terminal circuit (RTC) itself provides a meaningful measure of recovery time. If the RTC is unable to frame on the data it receives for about 3 ms, it inserts a "blue signal" to satisfy downstream equipment. When the RTC is able to reframe, the blue signal is removed, and data transmission is resumed. We define maintenance span recovery time as the time from the end of the interruption to the removal of the blue signal by the RTC. One limitation of this definition is that only recovery times in response to interruptions of 3 ms or longer can be measured. Otherwise, the blue signal is not inserted. This is not a serious limitation, since the terminal circuits should have no effect on the recovery time for interruptions less than 3 ms.

The expected recovery time for a maintenance span is simply the line recovery time plus the reframe time of the RTC. This expectation was confirmed by recovery time measurements. The maximum average reframe time for the RTC is 1.5 ms. Therefore we expect a recovery time less than 5.5 ms for a maintenance span of 20 regenerators for interruptions of 3 to 10 ms. For interruptions less than 3 ms, we expect maintenance span recovery times less than 1.0 ms (the line recovery time).

6.4 System recovery for maintenance spans in tandem

Sufficient system elements were available to arrange three maintenance spans in tandem for a total transmission distance of 61 km (38 miles). With this arrangement it was possible to observe the effect of signal interrupts in one maintenance span on subsequent maintenance spans. These tests served to confirm the failure sectionalization capabilities of the receive terminal circuit (RTC) and uncovered no anomalies in system operation. If, during an interruption, all downstream main-

tenance spans have had time to reframe on the blue signal inserted by the RTC of the span experiencing the interruption, then the total system recovery time is simply the recovery time of the interrupted maintenance span. This also holds true, of course, if the interruption is less than 3 ms since all the RTCs start to reframe simultaneously when the data resume.

The system behavior is more complicated in the situation where downstream RTCs do not have time to frame on the failed span's blue signal during the interruption. In the worst case, the total system recovery time is the recovery time of the interrupted maintenance span plus the reframe times of the downstream RTCs.

The system recovery performance described here should be adequate to virtually insure that calls will not be dropped during normal maintenance activity.

VII. DC POWERING EXPERIMENT

The +5.0-Vdc and −5.2-Vdc supply voltages for the regenerators and terminal circuits were provided by dc-to-dc converters operating from 48 Vdc. In this experiment, we monitored system performance while inducing transients on the common power distribution bus. The results led to some modifications of the regenerator circuit.

Normal activity on a bay of equipment includes the connection and removal of regenerator units. (A regenerator draws about 150 mA from the +5.0 Vdc supply and 390 mA from the −5.2 Vdc supply.) The resulting transients appearing on the common power distribution bus must not impair the performance of working equipment in the bay.

Tests were made on a maintenance span in which each regenerator had a detected optical power between −52 and −45 dBm. It was found that the load change of connecting and removing one regenerator caused the system to make a significant number of errors. The sensitivity was traced to the regenerators and not the terminal circuits.

Further tests were made with even larger load changes to insure that worst-case situations would be covered with adequate margin. A load change of 790 mA on the −5.2-Vdc supply caused about 100 errors. A load change of 720 mA on the +5.0-Vdc supply would induce RTC out-of-frames.

Some waveforms accompanying this out-of-frame response are shown in Fig. 18. The voltage on the +5.0 Vdc supply bus shows a 200-mV spike and step change of 70 mV. This transient is a function of not only the current load change but also the bus impedance and the dc-to-dc converter characteristics. Figure 18 also shows the signal at the decision point in the regenerator. The 1-v transient during the recovery of dc level was the immediate cause of the errors.

One regenerator was modified to improve the filtering of the supply

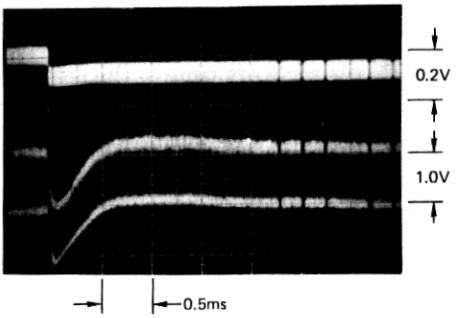

Fig. 18—Transients accompanying a step load change of 720 mA on the +5.0-Vdc supply. Upper trace: voltage on +5.0-Vdc supply buss. Lower trace: voltage at decision point in regenerator.

voltages in its linear channel. The above tests with load changes of about 720 mA were repeated on the modified regenerator. No errors resulted. Subsequently, all regenerators were similarly modified.

Another source of interference on the supply busses is from the dc-to-dc converters themselves. However, the specification for the power units guarantees 30 mV peak-to-peak or less ripple at their outputs. The tests on the modified regenerator assured us that transients up to 100 mV peak-to-peak cause negligible degradation. The conclusion is that the modified regenerator will tolerate all power supply transients and noise routinely incurred in an operating bay.

VIII. CONCLUSIONS

The Atlanta Fiber System Experiment was begun with high expectations for the system performance. There was some uncertainty due to lack of experience with lightwave systems; the Atlanta system was the Bell System's first complete lightwave system available in conditions approximating field environment. These experiments provided us with the needed experience, and in all cases the system met or exceeded our expectations.

Measurable crosstalk between fibers proved to be rare and to have only a negligible effect on system performance when it did occur. The effect of fiber dispersion is small compared to that of fiber loss; at 44.7 Mb/s the regenerator spacing is essentially loss-limited, not dispersion-limited. The jitter accumulation in a line of 50 regenerators should be well within the capability of the terminal circuits. System recovery time will be a fraction of interruption times caused by protection switches. This should virtually assure that no calls are dropped. The modified regenerator tolerates power supply transients and noise such as would be expected in a field application.

Through experience with the Atlanta lightwave system, we gained confidence in our technology. We were assured that the system and

equipment design will meet the conditions of a practical Bell System installation.

IX. ACKNOWLEDGMENTS

Much of the test circuitry used in the experiments was constructed by H. W. Proudfoot. We gratefully acknowledge his effort as well as that of M. J. Buckler, S. S. Cheng, V. J. Mazurczyk, and C. M. Miller in assisting with some of the tests. D. D. Sell provided valuable aid in designing the experiments and interpreting their results.

REFERENCES

1. R. S. Kerdock and D. H. Wolaver, "Performance of an Experimental Fiber-Optic Transmission System," Conference Record of National Telecommunications Conference, November 29 to December 1, 1976.
2. T. L. Maione and D. D. Sell, "Experimental Fiber-Optic Transmission System for Interoffice Trunks," IEEE Trans. on Communications, *COM-25*, No. 5 (May 1977), pp. 517–523.
3. I. Jacobs, "Lightwave Communications Passes Its First Test," Bell Laboratories Record, December 1976, pp. 291–297.
4. M. I. Schwartz, "Optical Fiber Cabling and Splicing," Technical Digest of Topical Meeting on Optical Transmission, Williamsburg, Virginia, January, 1975, pp. WAZ-1 to WAZ-4.
5. C. M. Miller, "A Fiber-Optic-Cable Connector," B.S.T.J., *54*, No. 9 (November 1975), pp. 1547–1555.
6. P. K. Runge and S. S. Cheng, "Demountable Single Fiber Optical Fiber Connectors and Their Measurement on Location," B.S.T.J., this issue, pp. 1771–1790.
7. T. L. Maione, D. D. Sell, and D. H. Wolaver, "Practical 45 Mb/s Regenerator for Lightwave Transmission," B.S.T.J., this issue, pp. 1837–1856.
8. R. G. Smith, C. A. Brackett, and H. W. Reinbold, "Optical Detector Package," B.S.T.J., this issue, pp. 1809–1822.
9. P. W. Shumate, Jr., F. S. Chen, and P. W. Dorman, "GaAlAs Laser Transmitter for Lightwave Transmission Systems," B.S.T.J., this issue, pp. 1823–1836.
10. M. J. Buckler and C. M. Miller, "Optical Crosstalk Evaluation for Two End-to-End Lightguide System Installations," B.S.T.J., this issue, pp. 1759–1770.
11. S. D. Personick, "Receiver Design for Digital Fiber Optical Communication Systems, I," B.S.T.J., *52*, No. 6 (July–August 1973), pp. 843–874.
12. S. D. Personick, P. Balaban, and J. H. Bobsin, "A Detailed Comparison of Four Approaches to the Calculation of the Sensitivity of Optical Fiber Systems Receivers," IEEE Trans. Communications, *COM-25*, No. 5 (May 1977), pp. 541–548.
13. A. G. Chynoweth, "The Fiber Lightguide," Physics Today, May 1976, pp. 28–37.
14. M. R. Santana, M. J. Buckler, and M. J. Saunders, "Lightguide Cable Manufacture and Performance," B.S.T.J., this issue, pp. 1745–1757.
15. C. M. Miller, "Loose Tube Splices for Optical Fiber," B.S.T.J., *54*, No. 7 (September 1975), pp. 1215–1225.
16. O. E. DeLange, "The Timing of High-Speed Regenerative Repeaters," B.S.T.J., *37*, No. 6 (November 1958), pp. 1455–1486.
17. C. J. Byrne, et al., "Systematic Jitter in a Chain of Digital Regenerators," B.S.T.J., *42*, No. 6 (November 1963), pp. 2692–2714.
18. D. L. Duttweiler, "The Jitter Performance of Phase-Locked Loops Extracting Timing from Baseband Data Waveforms," B.S.T.J., *55*, No. 1 (January 1976), pp. 37–58.