GTE Communications Transmission Systems

DEMODULATOR

JULY/AUGUST 1981

Satellite RTD - Parts 1 & 2

RTD—PART 1
VOICE TRANSMISSION

This issue of the Demodulator considers how including a satellite link in a transmission path affects communications. The round trip delay (RTD) satellite relayed signals encounter is described and how this delay affects transmission characteristics including echo is discussed.

A relatively new device for reducing echoes called an Echo Canceller is described and compared to echo suppressors. The article is presented in two parts. Part one discusses satellite transmission of voice signals. Part two discusses satellite transmission of data.

Some background knowledge of satellite communications is useful in understanding how a satellite link affects communications. The bibliography lists several past issues of the Demodulator which discuss satellites. The following paragraphs are extracted from the January 1979 issue.

"At an altitude of 22.270 miles, a satellite orbits the earth in exactly the same time period as the earth's rotation — just under 24 hours. Such a satellite is called a geosynchronous satellite because its orbit is synchronized to the earth's rotation. A synchronized satellite, placed in orbit directly above the equator on an eastward heading, appears to be stationary in the sky so its orbit is sometimes referred to as the geostationary orbit. Virtually all commercial communications satellites are geosynchronous.

The advantages of a geostationary satellite can be emphasized by viewing it as a fixed microwave repeater, on a very high tower costing less than $2.00 per foot. The disadvantage is the time delay, which equals about 0.5 seconds round trip for a two-way conversation. The delay also aggravates echo problems."

Satellite Vs Radio

Figure 1 shows alternative transmission paths that might be used to provide a long distance connection between subscribers 4,000 miles apart. One path contains a terrestrial microwave radio link and the other contains a satellite microwave radio link.

One obvious difference between the two paths is that while the microwave link contains several repeatered sections, the satellite is the only repeater in its link. Another difference is that a ray representing the terrestrial link's propagation path is horizontal relative to the earth's surface. A ray representing the satellites propagation path is vertical.

The two differences described above are advantageous to the satel-

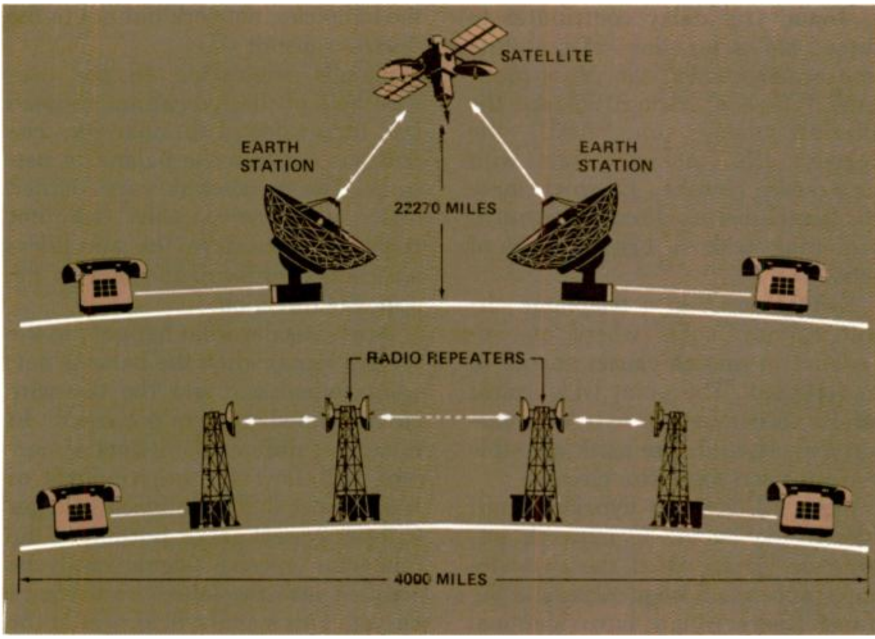

Figure 1. Satellite Vs Terrestrial Connection.

lite system. Fewer repeaters means less exposure to possible transmission impairments. The vertical path is relatively immune to fading caused by atmospheric conditions.

A third difference between the two paths is the distance the signal travels. The round trip delay (RTD) accrued by signals traversing the satellite path is more than 10 times as long as the delay over the terrestrial system.

Referring to Figure 1, the distance from an earth station to a satellite varies between 22,270 and about 26,500 miles, depending on the geographic location of the earth station. A message originating at one end of the path travels to an earth station, up to the satellite, back down to the second earth station and on to the far end. A reply to the message travels the same route in reverse. Note that, neglecting the ground portion of the path, the distance traveled is four times the distance from the earth stations to the satellite.

Using the speed of the light as the propagation speed and assuming the distance from the earth stations to the satellite is 25,000 miles, we can calculate the RTD.

$$\text{RTD} = \frac{(\text{Distance in miles})}{(186{,}000)} \text{ sec.}$$

$$= \frac{(4 \times 25{,}000)}{(186{,}000)} \approx .540 \text{ sec.}$$

Since RTD is generally expressed in milliseconds, we multiply the result by 1,000 to give 540 ms.

The RTD for the 4,000 mile terrestrial path is

$$\text{RTD} = 1{,}000 \frac{(2 \times 4{,}000)}{(186{,}000)} \approx 44 \text{ ms.}$$

Round trip delay contributes to three characteristics of satellite transmission which can cause problems. These characteristics are the absolute delay time itself, the delayed echo and problems with suppressor circuits. Each of these characteristics are discussed in turn, beginning with a brief review of echo.

Echoes occur in a telephone circuit at any point where an impedance mismatch causes energy to be reflected. The point where most echoes occur is at hybrid transformers. Hybrids are used to couple two-wire and four-wire circuits.

Figure 2 shows a hybrid circuit. Ideally the balancing network impedance would equal the two-wire impedance at all frequencies and the transformers would have identical electrical characteristics. Under these ideal circumstances, signals on the "receive" side of the hybrid are coupled to the two-wire circuit and the balancing network but not to the 4-wire transmit side.

Signals originating on the two-wire side of the hybrid are coupled into both sides of the four-wire circuit but not into the balancing network. The signals which are coupled into the receive side are not transmitted because the amplifiers are connected for signal flow in the opposite direction.

Now consider what happens to the receive signal when the balance network impedance and the two-wire circuit impedance are not equal. In this case, different amounts of current will flow in the windings of transformer 2. The induced voltages will not cancel and a portion of the four-wire receive signal will be coupled into the four-wire transmit circuit. This signal will appear at the other end of the four-wire line as an echo.

Two factors, loudness and delay, determine how objectionable an

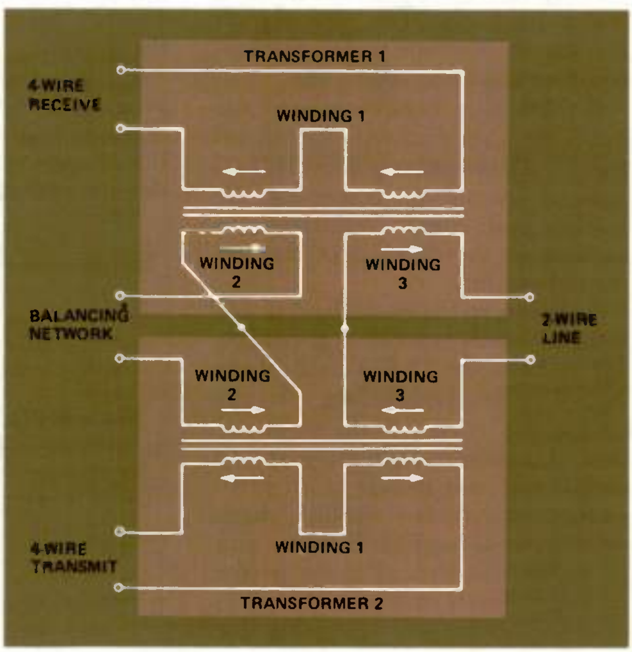

Figure 2. Two transformers may be connected to form a high-quality hybrid. Arrows show how current flow in the four-wire receive branch induces secondary current flow in the two-wire line and the balancing network. No current flows in four-wire transmit branch because the opposite potentials, induced across the two halves of the winding, cancel each other.

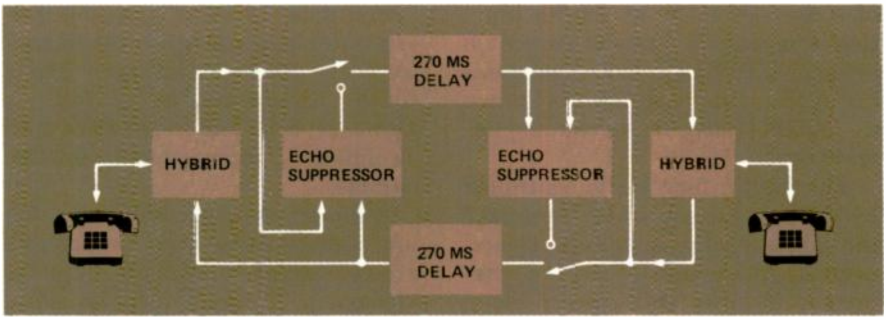

Figure 3. Split Echo Suppressor.

echo is. It is apparent that an echo that is inaudible will not be annoying, no matter how long it is delayed. However, studies show that if an echo is loud enough to be noticeable it becomes more objectionable with increasing delay between it and the signal. Given the long round trip delay, it is apparent that echo can be a serious problem when a satellite link is in the transmission path.

On long paths, echo suppressors are commonly used to control echoes. Figure 3 is a simplified diagram of a split echo suppressor which might be used on our satellite example path. Basically the echo suppressor is a pair of voice actuated switches. The suppressor on the right opens the bottom path when it detects speech energy in the top path. The suppressor on the left opens the top path when it detects speech energy in the bottom path.

An actual suppressor is somewhat more sophisticated, as shown in Figure 4 which expands the right hand side of Figure 3. The bottom path carries the transmit signal from the telephone. It also carries any echo signal generated at the hybrid.

Most of the time the people at either end of the line speak in turn, so the echo can be stopped by allowing the switch to open the path while the far end party is talking. The switch is controlled by a voice detector. However, interruptions and exclamations are part of the dynamics

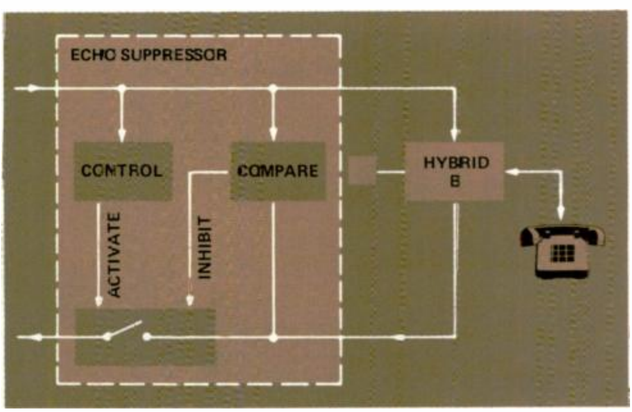

Figure 4. Simplified Echo Suppressor Logic.

of conversation. So an inhibit function is included to prevent the switch being opened if both parties are talking at the same time. This is known as "double talk" in telephone terminology.

The inhibit function is controlled by the comparator which compares the speech energy in both paths and overrules the switch if there is a component of near end speech along with the echo. Although this action makes the telephone conversations more "natural" it prevents echo suppression during double talk.

Also, the switching decisions require a finite amount of time and cannot be made with absolute accuracy. So, some amount of speech clipping is inevitable, particularly on the first part of an interruption. Experience shows that echoes during double talk and speech clipping are more annoying when they are coupled with satellite delay.

Another problem with suppressors arises from the fact that voice signals are not continuous. There are gaps between phrases, words and even syllables. If the switches acted instantaneously, the path would be constantly cutting in and out. So, there is a .1 to .2 second built-in delay before the near end transmit path is cut in.

For these reasons, echo suppressors are not considered satisfactory for U.S., domestic, full-hop circuits (both directions by satellite). However, echo suppressors are used on full-hop, international, satellite circuits.

Although domestic half-hop satellite circuits using echo suppressors are considered satisfactory from the viewpoint of transmission quality, the satellite circuit's potential for use is drastically reduced. A relatively new device called an echo canceller provides a solution to this problem. This device synthesizes a replica of the echo and subtracts the replica from the real echo. The echo is effectively cancelled. Since the path is not interrupted, clipping does not occur. Full-hop satellite transmission is satisfactory.

Figure 5 is a block diagram of an echo canceller. As shown in the figure, samples of the incoming signal are connected to the input of a tapped delay line. Each tap position adaptively attenuates the sample reaching it by a weighted amount. The weighted samples are accumulated to form an echo estimate. This estimate is subtracted from the actual echo in the transmit path.

If the impulse response of the actual echo path has decayed to zero

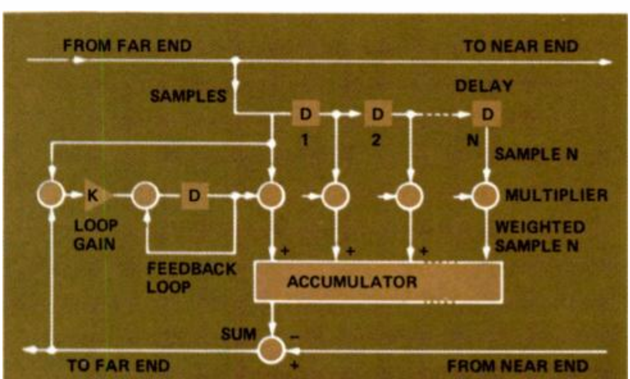

Figure 5. Echo Canceller.

by the time of the last tap and the weights are good approximations of the echo response, the echo estimate will be equal to the actual echo. The echo will be cancelled. If the canceller signal is not a good approximation of the echo, the difference between the two signals is fed back to the input of the delay line as an error signal.

The weight of each tap is periodically adjusted by a small correction factor which is proportional to the product of the error signal and the corresponding tap signal. The adjustment is controlled by a separate feedback loop for each tap position. An echo canceller can also be described as an adaptive transversal filter.

The canceller has a limited bandwidth determined by the per section time delay of the tapped line. This time cannot exceed the period of 1/2 cycle of the highest frequency to be cancelled. For a 4 kHz voice channel, the maximum delay time is 125 microseconds:

$$t = \left(\frac{1}{f}\right)\left(\frac{1}{2}\right) = \frac{1}{2 \times 4000} = 125 \text{ ms}$$

This is the period of an 8 kHz wave. Since this is the Nyquist frequency for PCM sampling of a 4 kHz channel, 125 microseconds is also referred to as the Nyquist interval.

A Bell Laboratory's echo canceller, described in the October 1980 issue of the IEEE Spectrum, is an integrated circuit with 128 taps and therefore 126 feedback loops. The number of taps and the loop gains determine how closely the canceller approximates the echo signal.

A greater number of taps and smaller incremental changes in loop gain will provide a closer approximation of the echo but it will take a longer period of time for the canceller to adapt to a new echo path. Cancellers with 256 taps have been designed. Each time a call is originated, a talker may hear an echo of his first syllable or two. However, this is considered less objectionable than the degradation in echo estimates that would result if the adaptation time were decreased to prevent the initial echo.

The total number of filter segments also determines how far from the hybrid a canceller can be installed. Figure 6 shows a canceller installation with 32 milliseconds delay between it and the hybrid. This canceller requires:

$$\frac{32 \times 10^{-3}}{125 \times 10^{-6}} = 256 \text{ sections}$$

A 128 section canceller could be placed up to 16 milliseconds from the hybrid.

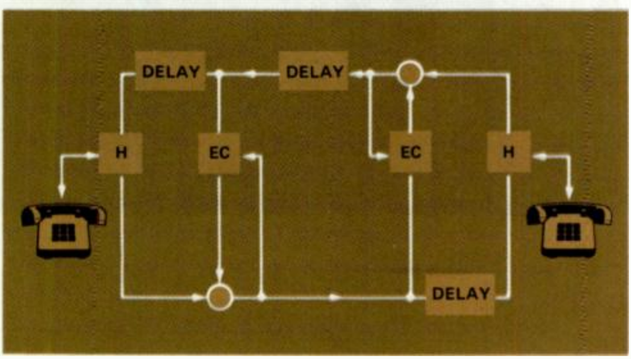

Figure 6. Echo Canceller Installation.

The latest canceller designs use integrated circuits for filter implementation because discrete component filters of the required complexity are bulky and costly.

Even the best canceller cannot achieve perfect cancellation. There is always some residual, uncancelled output due to circuit noise, nonlinear distortion and echoes which are dispersed beyond the length of the filter.

These small, undesired signals can be blocked by a center clipper which only passes signals with amplitudes above a preset threshold value. Since the undesired signals are very small, the threshold is quite low and easily overcome by double talk signals.

Double Talk Operations

During one way conversation, the principal input to the echo canceller will be the reference signal, the echo and some noise. The noise will cause some slight pertubations in the tap weights. However, this noise is not correlated with the reference waveform. Therefore, the tap settings are subject to small fluctuations but are correct on the average.

In the case of double talk, the second talker's signal appears to the canceller as a huge component of additive noise. This would upset the echo estimate and cause a large variation in tap weights, if the adaption was allowed to continue during double talk. To prevent this, echo estimating is discontinued and the tap weights remain at the values present at the time double talk is detected.

The double talk detector (Figure 7) must be very fast acting to prevent the canceller being "swamped" by the strong second talker's signal. Also their is a possibility that the echo characteristics might change during the double talk period. However, this seldom occurs because double talk is normally of short duration.

Figure 7. Echo Canceller Showing a Double Talk Detector and a Single Tap Adjustment Circuit.

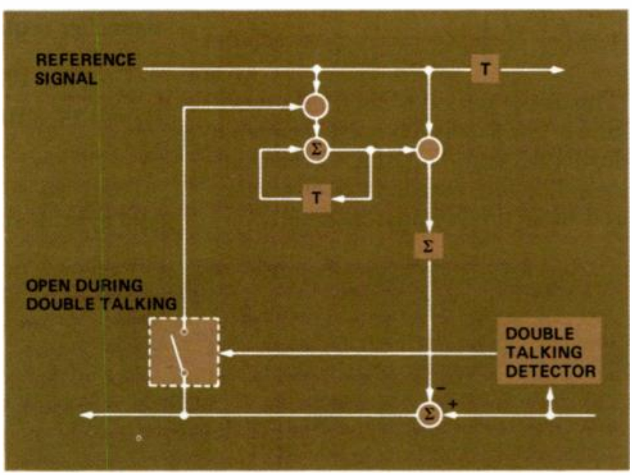

RTD—PART 2
DATA TRANSMISSION

Part one discussed the affects of a satellite link on voice transmission. This second part discusses the affects of a satellite link on data transmission.

A satellite link does not present serious problems to a four-wire data system. In the first place such a system is a full duplex channel — there are no two-to-four wire junctions so echo should not occur.

If full duplex protocol is used, the system should not experience a serious decrease in throughput. It is apparent that the long round trip delay will decrease throughput, if stop and wait protocols are used.

There are three areas where a satellite link may affect two-wire data transmission. These are, data set operation, line protocol and throughput.

Data Set Operations

There are two ways to operate over a path containing echo suppressors. The first is to disable the echo suppressors. The second is to operate with the echo suppressors actively in the circuit.

Any data set that sends an answer tone will initially disable echo suppressors because the answer tone is within the echo suppressor disabling frequency range. Once the suppressors are disabled, there must not be a gap greater than 100 milliseconds in the energy on the line. If such a gap exists, the suppressors will be reactivated.

Full duplex data sets always have energy in both directions, without gaps. Reverse channel data sets maintain continuous energy on the line by turning on low speed transmission whenever high speed transmission is not occurring.

A third type of data set uses holding tones to keep the echo suppressors disabled. The holding tone output is present whenever the data set is not transmitting or receiving. The tone's purpose is to make sure the echo suppressors remain disabled.

Figure 8 is a timing diagram of data transmission between two stations, A and B. As shown in the figure, the data sets operate with holding tones to keep the echo suppressors disabled.

Referring to the diagram, the called station A puts out an answer tone followed by a holding tone. At a

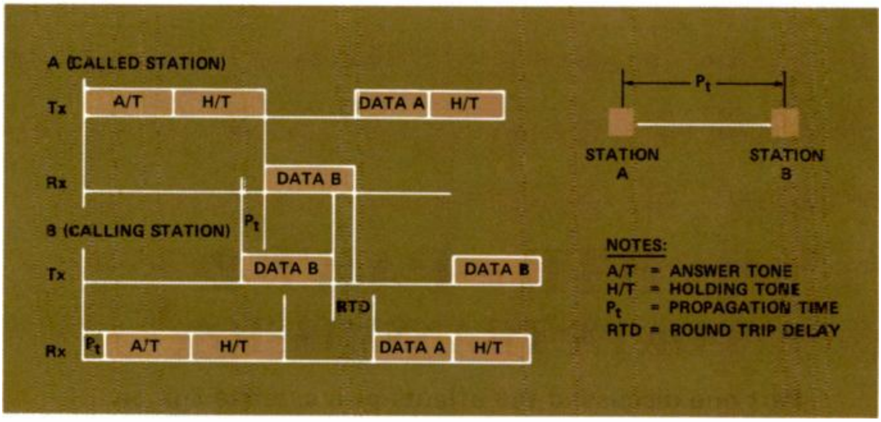

Figure 8. Operation with Echo Suppressors Disabled.

later time. equal to the propagation time, station B receives these signals and at some subsequent time, begins to transmit data.

One propagation time later, B's data is received at station A which immediately turns off its holding tone. At the end of B's transmission, its carrier is turned off. It takes one propagation time before A is aware that B's carrier is off. If A begins transmitting immediately, another propagation time elapses before that signal is received at B. In other words, there is a total elapsed time of one round trip delay between the time B stops transmitting and the time A's signal arrives at B. During this time, there is no energy on the line in either direction at station B. This can be seen on station B's Tx and Rx lines in figure 8.

If a satellite is in the path, the RTD will greatly exceed 100 milliseconds. The echo suppressors at station B will be reenabled. Now lets assume that B begins transmitting a data block as soon it detects the end of A's data. Note that a holding tone and therefore energy is still present in the receive direction. B's transmission must now contend with the delay built into the echo suppressor. As much as 50 milliseconds of transmit data may be clipped and therefore possibly not received at station A.

Operations with Active Suppressors

Energy gaps greater than 100 milliseconds occur during the operation of many half duplex data sets. Some of these sets do not produce an answer tone, so they operate with active echo suppressors. Data sets operating with active echo suppressors are protected against echo. However, they must transmit sacrificial carrier to compensate for the delay time built into echo suppressors which prevents the path cutting in and our during pauses in voice transmissions. This provision holds suppression in the transmit path for a short while after the receive energy ceases. Consequently, loss of data could occur if return transmission began immediately. Therefore, half duplex data sets, which operate with active echo suppressors, should transmit a sacrificial, carrier-only signal for about 150 ms. before transmitting data.

Figure 9 is a timing diagram of data transmission between two stations with active echo suppressors. The answer tone originating at station A is received at station B after a lapse of time equal to the propagation time (P_t). This tone disables the echo suppressor at B.

The figure shows a gap between the answer tone and the first sacrificial carrier on (Cs station A). If the gap is long enough, the suppressor will be reenabled. Assuming this is the case, the message, Data A, takes one propagation time to arrive at B. The reply from B requires one propagation time to get back to A. Note that a part of B's sacrificial carrier is used in overcoming the echo suppressor delay. Since there is a quiet time equal to one RTD at station A the echo suppressor is active when B's transmission is received.

Assuming there is no gap between the answering tone and A's transmission, there will be continuous energy on the line until the end of A's transmission. There will not be any echo protection for this first transmission.

It takes one propagation time for the end of A's transmission to reach B and an additional propagation time for B's reply to return to A. Therefore, there is no energy on the line at A for a time equal to at least one RTD. The echo suppressor will be active at A.

Subsequent transmission from A to B will also be subject to RTD so the echo suppressor at B will be active. In other words, all transmissions after the first would be protected by echo suppressors. If an echo occurs on the first transmission, line protocols may be affected.

The Echo Canceller with Data

Data transmissions may encounter problems if echo cancellers are included in the path. In the case of full duplex and reverse channel operation, energy is continuously maintained in both directions. Since currently available echo cancellers do not train in the double talk mode, it is likely that the canceller will not

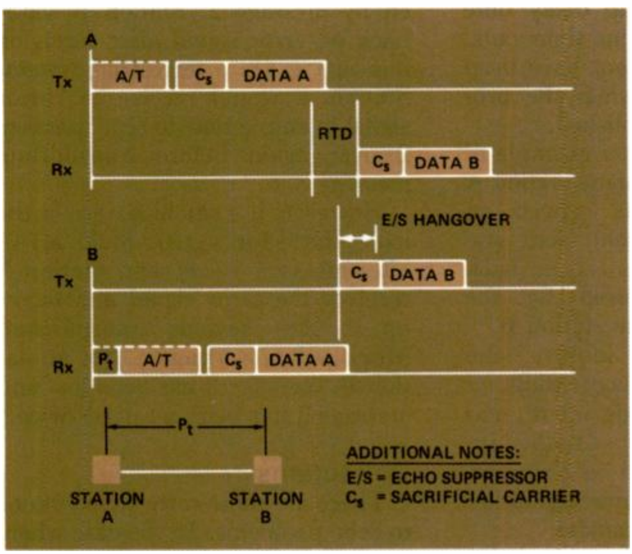

Figure 9. Operation with Echo Suppressors Active.

properly adapt to the echo path. This is not a particular problem with full duplex data via satellite because the echo is an out-of-band signal. However, reverse channel operation could have problems because the echo is inband.

Echo cancellers will provide protection against echo for half duplex operations. Also, echo cancellers will protect all data sets which operate with echo suppressors enabled. Data sets that operate with holding tones will be protected, even though there is always energy on the line, because it is in only one direction at a time. The echo cancellers can train on the holding tone signals.

Protocol Problems

There are four ways that a satellite link may cause protocol problems. The first problem could result if the data set mistakes its own echo as a control character and assumes the other terminal is trying to gain control. In other words, the demodulated echo simulates a contention. Echo can also be seen as an incorrect response.

The long round trip delay time may exceed the system time outs. Also, the RTD may not have been taken into account when the protocol timing was established.

An inquiry loop is an example of the first problem. Assume station A sends an inquiry and expects an acknowledgement signal from station B. The inquiry is echoed back to station A, followed by the acknowledgement from station B.

Station A sees the inquiry echo first, assumes there is contention, ignores the acknowledgement and sends another inquiry. Another inquiry echo is received in response. This process can continue indefinitely and prevent data transfer.

The second problem occurs if the RTD exceeds the protocol time out interval. For example, assume a control station (A) has a one-half second time out interval. The station sends an inquiry via satellite to station B. Assuming an RTD of 540 milliseconds, B's acknowledgement will not be received at A until the time out has expired and A has transmitted a second inquiry.

When A does receive the acknowledgement to his first inquiry, he interprets it as a response to his second inquiry and transmits his message. The next signal received at A is the real acknowledgement of this second inquiry. This is out of sequence. At this point A may give up and send an end of transmission or may send another inquiry to begin another cycle of confusion.

Some data systems transmit blocks of data using no response as a positive acknowledgement. When an error occurs, an error signal is returned. Assume a system where messages are sent with a 300 millisecond blank between them.

Station A sends message 1 followed by message 2. Station B sends back an error signal after receiving message 2. On a terrestrial system, Station A would receive the error signal in ample time to send message 2 over again before transmitting message 3.

However, if a satellite were in the loop, the error signal would arrive after message 3 was sent. Station A receives the error signal and backs up to the previous transmission. Message 3 is sent again. So, at station B, message 2 has been lost and message 3 has been received twice.

Corrections

There are some software solutions to echo problems. In the case where clipping occurs on the answering stations transmission, pad characters could be added to sacrifice to the clipping, thereby protecting the data characters.

When an echo is received on every transmission, different code characters may be placed at the front of the messages from calling and called stations so that signals can be separated from echoes. Alternatively, a dummy message may be sent and a timer set when a response is received. The data set is programmed to ignore all messages until the timer has expired.

In the case where the RTD exceeds the time out interval, the simplest solution is to increase the time out interval. The timing problem, where successive blocks of data are sent without acknowledgement unless there is an error, might be solved by increasing the gap between messages. However, this decreases the throughput, even on terrestrial connections. It might be better to change the protocol to wait for a response instead of using no response as an acknowledgement.

Throughput

A satellite path's long RTD has an effect on throughput efficiency. In a half-duplex protocol, such as bisync, the receipt of each message block at the receiving terminal must be acknowledged before the next message block can be transmitted. When the time spent in acknowledgement is long compared to the message time, throughput efficiency is decreased.

The most significant contributors to the acknowledgement time are the request-to-send (RTS) clear to send (CTS) interval and the round trip delay. In a terrestrial circuit the RTS/CTS is most significant although RTD is not negligible for long circuits. The reverse is true for a satellite link; RTD is most significant but RTS/CTS is not insignificant.

One way to improve the throughput efficiency is to increase the number of characters in a message block. However, as the block length is increased the error rate becomes a significant factor, so there is an optimum point. The optimum point for best efficiency in terms of characters per block is dependent upon data set speed (i.e. the amount of time taken to transmit a block). The optimum will be a smaller number of characters for low speed data sets and a larger number of characters for high speed sets.

Future Developments

The July/August 1980 issue of the Demodulator, "Fault Locating, Part 3 Datacom" discussed a possible application of echo cancellers to full-duplex, data transmission over the dial-up, switched network. The article suggests placing the canceller in the modem. The advantages and problems of this approach are discussed in the following paragraphs.

The echo canceller discussed earlier in this article is installed in the network at or near two-wire/four-wire interfaces. This canceller will not eliminate near end echoes (sidetone) because they are generated at impedance mismatches at the end terminal itself and in the circuit up to the first interface. While such echoes are actually desirable for voice transmission, they can be just as damaging as distant echoes to full duplex data transmission.

For this and other reasons, it is desirable that echo cancellers be installed at the terminals. Figure 10

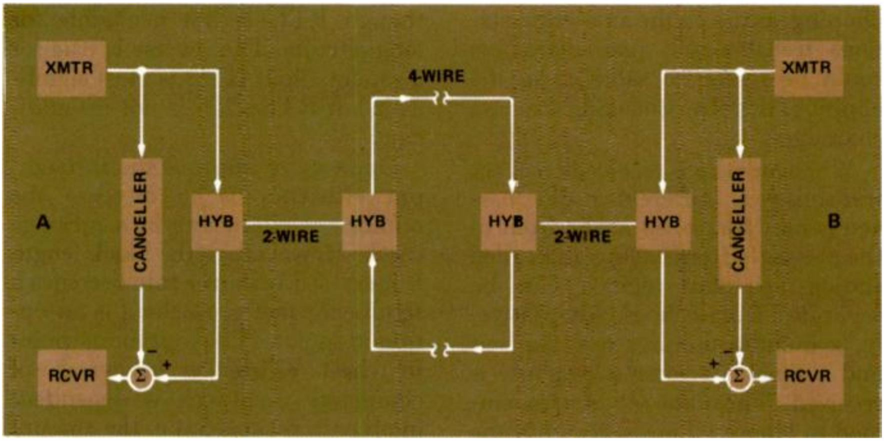

Figure 10. Echo Cancellers for Full-Duplex Data Transmission Over a Two-Wire Line.

shows a way to use hybrids and echo cancellers to provide an echo free interface between a four-wire modem and a two-wire line.

There are other problems connected with echo cancellation for full duplex data. One problem results from the fact that long periods of double talking transpire. So, it is not practical to "freeze" the canceller during double talk because changes in the echo channel may occur during a lengthty session. Adaptation can be accomplished during double talk by using an average gradient algorithm. However, adaptation is much slower than with the minimum mean-square error algorithm used in the canceller previously described.

Another problem with placing the canceller at the station is the very large delay exhibited by the distant echo returning from the far end of the circuit, particularly if the path contains a satellite link. It is impractical to use a filter with the approximately 4,000 Nyquist interval taps required to handle delays of 500+ millisecond magnitudes. A reduction by a factor of 4 can be achieved by the use of a symbol-interval transversal filter rather than the Nyquist interval type but the number of taps required for satellite routing is still too great.

A solution is to provide sectors of bulk delay between active delay sections. This approach would substantially reduce the required number of Nyquist interval taps. It would then be necessary to send a test signal at the beginning of a call, to measure the delay of the distant echo group. However, the allocation of a little time, to set-up a protocol at the beginning of a transmission, is not unusual in data communications.

In summary, if echo cancellation could be achieved at the modem at a moderate cost, the dial up network could be more useful in computer communications. Development efforts are underway to achieve this result and to produce echo cancellers on one or two microcircuit chips. The Demodulator will continue to monitor these developments and report the results in a future issue.

BIBLIOGRAPHY

1. Sondhi, M.M. and Berkley, D.A., *Silencing Echoes on the Telephone Network* — Proceedings of the IEEE, Vol. 68, No. 8, August 1980.

2. Weinstein, Stephen B., *Echo Cancellation in the Telephone Network*, IEEE Communications Society Magazine, January 1977.

3. Duttweiler, D.L., *A Twelve-Channel Digital Echo Canceller*, IEEE Transactions on Communications, Vol. Com-26, No. 5, May 1978.

4. Curtis, D.M., (American Telephone and Telegraph) *Slide Presentation on the Characteristics of Data Transmission over Satellite Facilities*. Presented to Boston Section IEEE, March 6, 1979.

5. *Satellite Communications Update Part 1*, GTE Lenkurt Demodulator, January/February 1979.

6. *Satellite Communications Update Part 2*, GTE Lenkurt Demodulator, May/June 1979.

GTE LENKURT
1105 COUNTY ROAD
SAN CARLOS, CALIFORNIA 94070

ADDRESS CORRECTION REQUESTED

Bulk Rate
U. S. Postage
PAID
San Carlos, CA
Permit No. 37

95673OBR IPO-BFJ
S
OM MGR
ACIFIC TEL CO
R. J J O BRIEN
O BOX 700
IO LINDA CA 95673

914A/PCM Subscriber Multiplex System (MXU)

GTE Lenkurt's 914A PCM Subscriber Multiplex System can extend a central office switch up to 50 miles. The 96-channel 914A is an attractive alternative to installing new cable to meet the expanding needs of small communities and is particularly appealing to companies with an immediate need to provide additional services coupled with future plans for a GTD-5 EAX installation.

The flexible 914A System can function as a 96-line digital pair-gain system operating into an analog switch and be quickly converted to function directly with Automatic Electric's GTD-5 EAX. A modular component of the GTD-5 EAX family, the 914A MXU has software commonality with that digital system.

Conversion from a pair-gain system to a switch-remote multi-

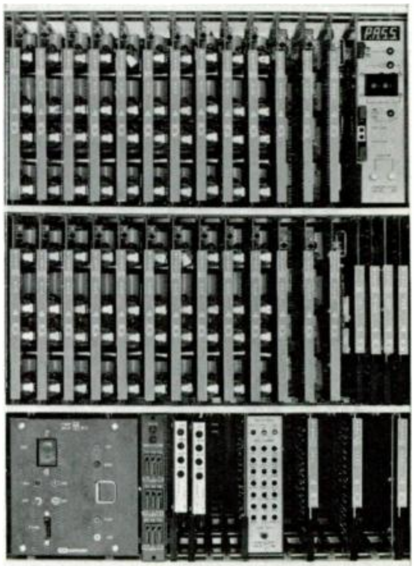

plexer is accomplished by rerouting the primary T1-type span, at the DSX, to the assigned digital trunk interface of the electronics office. The MXU software is automatically changed from 914A central office terminal operation to GTD-5 operation.

GTE LENKURT

VIDEO, VOICE & DATA
TRANSMISSION SYSTEMS

The GTE Lenkurt Demodulator is circulated bimonthly to selected technicians, engineers and managers employed by companies or government agencies who use and operate communications systems, and to educational institutions. Permission to reprint granted on request.

LITHO IN U.S.A