

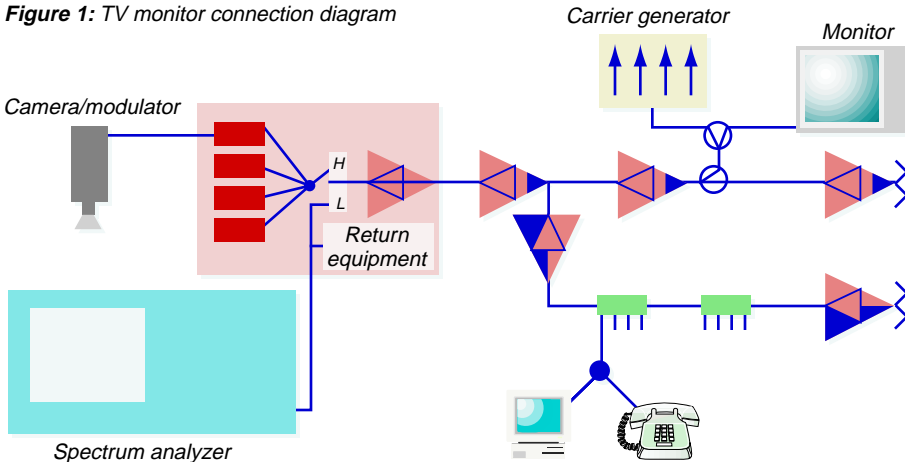
Insights into proper return path alignment

Digital troubleshooting is a whole new ballgame

By Bill Morgan, R&D Project Manager, Hewlett-Packard

Editor's note: As the cable industry rushes to get into high-speed data, telephony and perhaps other services requiring real-time transactions, a critical network component is a reliable working return path. But that's easier said than done, many are finding. This article is part one of a three-part series on cleaning up the return path in preparation for advanced, two-way services. This initial installment focuses on return path alignment techniques; future articles will focus on ongoing maintenance as well as return path noise and ingress performance.

Figure 1: TV monitor connection diagram



The signing of the Telecom bill, coupled with consumer demand for interactive services in the home, has created the potential for cable television companies to explore new sources of revenue. They now have the opportunity to provide their subscribers with two-way services like Internet communications, interactive TV and telephony. But there's a catch—cable operators need to successfully activate their return paths before they can deliver these potentially lucrative new services. As the industry begins to activate the return path, a host of problems are surfacing. Many of these problems can be avoided if an accurate alignment procedure is followed.

Table 1: TV monitor with 4-carrier generator

Pros	Cons
Utilizes readily available hardware	Technician carries multiple boxes
Ingress and gain balance on same screen	Requires 6 MHz forward bandwidth for each return monitor
Real-time ingress response simplifies troubleshooting	Limited frequency resolution

Table 2: Return sweep alignment

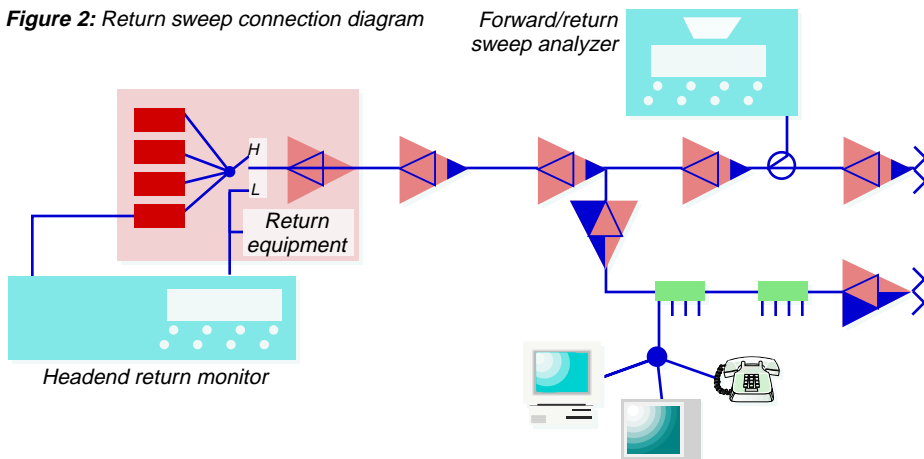
Pros	Cons
Uses minimal forward bandwidth	Slower ingress response
Provides best (variable) frequency resolution	Requires functional return communications (for sweep)
Technician carries one box for all measurements	

This article provides insight into some of the problems that have been encountered while working with operators on alignment techniques.

Customer impact

Over the years, cable operators and technicians have become familiar with the picture artifacts related to poor alignment in the forward path, but the artifacts related to poor return alignment are new. Techs are no longer dealing

Figure 2: Return sweep connection diagram



with picture artifacts, but communications degradation. The ability to identify system problems by the type of degradation is in its infancy. These problems may include many of the following:

- 1) Source levels from modems end up higher than practical.
- 2) Excessive return levels cause clipping of the return laser, affecting all the signals on one return.
- 3) Communications throughput is reduced (poor response times).
- 4) Telephone calls are dropped or service is delayed.
- 5) IPPV requests respond slowly or intermittently.

◇ RETURN PATH ALIGNMENT

6) Customers experience service outages.

7) All of the above may be intermittent.

Problems in the forward path have typically been diagnosed by visually observing the degradation, but the digital communications inherent in return services makes this method of troubleshooting impractical. It also places even more importance on the quality of the initial return path alignment. One unique difference between maintaining the forward and return paths is that each test location in the forward path is affected only by the amplifiers closer to the headend, but amplifiers in either direction can affect the current location in the return path. For example, when sweeping the forward path, if the noise in the system increases, the technician knows the problem is between the current test-point and the headend. However, in the return path, the technician cannot be sure. The source of the problem could be on a different feeder or trunk altogether. In this case, not only is the technician faced with the challenge of finding the noise/ingress problem, he or she may also be faced with the problem of not being able to complete the testing of the current amplifier.

If a well-documented alignment procedure is followed, and the technicians understand the inter-relationships of how the return operates, these problems can be minimized.

Alignment methods

There are currently many methods being used to align the return, but only the two most prevalent are discussed here because they do a good job of representing the range of capability:

1) **TV monitor** and portable 2- or 4-carrier generator in the field with a headend spectrum

analyzer and video modulator to send the response downstream;

2) **Return sweep** generator with headend sweep receiver and ingress monitor. Tables 1 and 2 provide a short summary of the pros and cons of these two methods (see page 42).

One of the differences between the two alignment methods is frequency resolution. Quite

often, the narrow bandwidth of the return path is used to justify lower frequency resolution. In reality, a 5 MHz to 42 MHz return path still consists of more than three octaves of bandwidth, only one octave less than a state-of-the-art forward path. Many of the problems found and resolved while aligning the system will be repeated over octaves and may be as significant in the lower frequency octaves as in the upper.

The sweep response traces shown in Figures 3 through 5 provide an example of the benefits of improved frequency resolution. Figure 3 is the response of a return cable system (11 amplifiers deep) indicative of a significant reflection which could push the error correction in a digital communications link to its limit. This sweep response was taken using 135 kHz of frequency resolution.

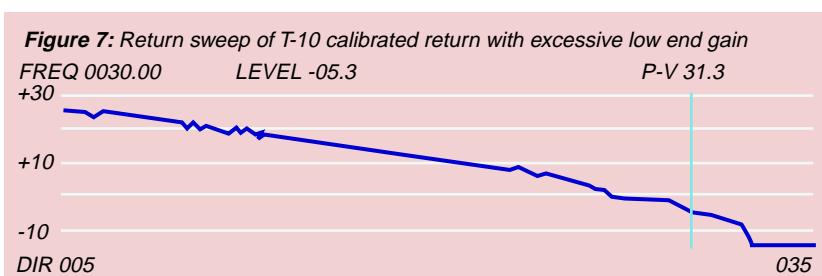


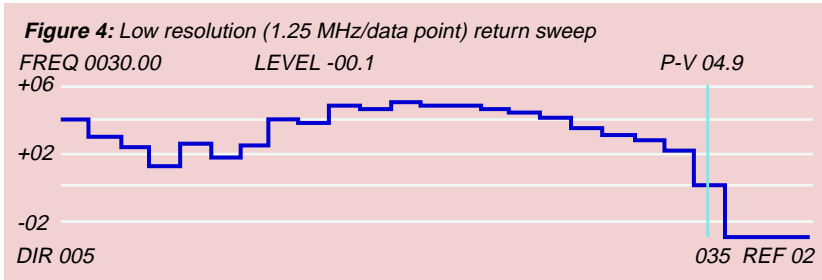

Figure 4 shows the same system response using only 1.25 MHz of frequency resolution. Note that the reflection which is obvious in Figure 3 could be missed in Figure 4. Figure 5 shows the same system response using a simulated four-carrier approach.

It is important to be aware of the tradeoffs being made when test equipment that provides less frequency resolution is selected. It is also important to be familiar with the passives in the system when using a carrier generator for alignment.

Roll-off in the passives may be compensated for by misadjusting the amplifier slope. Some of the newer 1 GHz passives roll-off below 10 MHz, so when using the carrier generator approach, the carriers should be placed at frequencies that are flat through the passives. In summary, a return sweep system with high frequency resolution has several advantages:

1) Flatness discontinuities and suckouts can be seen.

2) Roll-offs at the band edges are visible,



and diplexer problems may be eliminated.

3) Reflections and return loss problems show up as ripples in the sweep response and can be repaired.

4) Modern sweep systems with short duration sweep pulses can be used in the presence of carriers with minimal interference and don't take up the bandwidth required by CW carriers.

Ingress problems

If the return ingress is extremely high, repairs may be required prior to starting the alignment process.

Experience shows that 70 percent of ingress problems occur in the home, 25 percent in the drop, and only five percent in the coaxial trunk and feeder itself. It is also becoming apparent that a major contributor to the ingress in the coaxial trunk is actually common path distortion. Excessive ingress can interfere with the sweep systems, and may drive the laser into compression, causing the output levels to be in error.

In order to follow the

Figure 8: Spectrum scan of properly calibrated return





Figure 9: Return sweep of properly calibrated return



Once the proper pads and equalizers were installed, small amounts of noise caused no problems

process described in this article for return system alignment, one needs to be able to start from the fiber node and proceed through the network (one visit per location being the goal). Return path "blockers," or some alternate methods of disconnecting the return input to the amplifier currently being tested, need to be available to establish proper set-up from the current location back to the headend. This will be an important step in reducing return path test time and meeting the one visit per location goal.

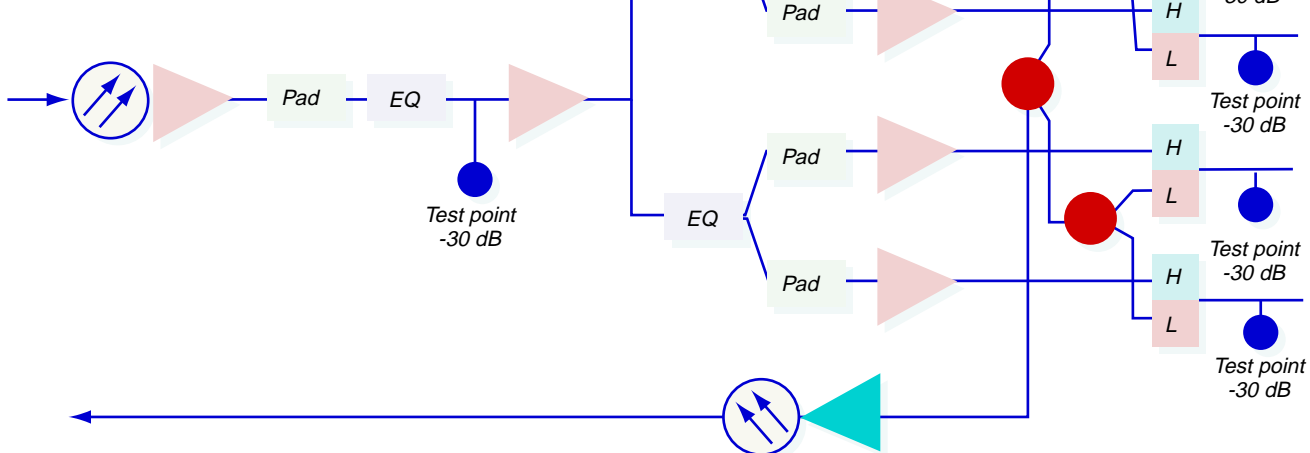
Blockers are particularly use-

ful on new systems where customers are not installed. If a new section is connected into an existing system, an ingress problem in the new section may degrade or disable the existing system. There have been many articles written discussing the source of ingress and solutions, and we have listed some of them in the reference section (see page 53).

Figures 6 and 7 provide an extreme example of how poor alignment techniques can aggravate ingress problems. This particular return path was "aligned" using the level of the T-10 carrier as a reference, but ignoring the slope of the return. Because of the excessive gain at the low end of the spectrum, the return amplifiers were pushed into compression by small amounts of burst noise below 10 MHz. These bursts of compressed noise affected the entire return spectrum. Once the proper pads and equalizers were installed, the same small amounts of noise caused no problems.

The traces shown in Figures 8 and 9 are of a properly aligned return and the associated well-behaved noise response of the return. Once again, it is important to have good frequency resolution in the headend spectrum analyzer to effectively identify and troubleshoot sources of ingress.

Figure 10: Typical fiber hub



◇ RETURN PATH ALIGNMENT



Return path alignment using the HP CaLan Sweep/Ingress Analyzer.

Alignment process

The approach to alignment of the return path is similar to the forward path in the sense that it should be aligned for unity gain. In the return path, unity gain is referenced to the input of the amplifiers. In the forward path, unity gain is referenced to the output of the amplifiers. In each case, the gain of the amplifier is compensating for the loss of the section of cable between the previous amplifier (closer to the headend) and the current amplifier. Attention to detail in the return path is critical to successful

alignment. A poorly aligned amplifier farther out in the trunk may make alignment impossible because of excessive noise in the communications path.

Reference output

The first step in the alignment process is to measure the output level at the headend for each return path using a given reference level input to the fiber node return laser. When choosing this reference input level, consideration must be made for low-end optical and RF noise floors, as well as high-end clipping

resulting from overdrive. A typical manufacturer's specification for optimum input level to the return laser is +20 dBmV. This level assumes standard video carriers, and it is becoming common practice to use this as the reference level. Modern sweep systems are designed to operate 10 dB or greater below optimum carrier levels, so +10 dBmV is used as the reference input level in this discussion.

An accurate and flat input to each laser is necessary to establish the proper headend reference, and this input level must be maintained for the return system to operate properly. This input may be provided by either of the methods discussed earlier. The sweep system with higher frequency resolution has the advantage of allowing a tech to see problems in the return frequency response during the alignment.

It is important to be familiar with the amplifier and fiber optic block diagrams (see Figure 10). Internal coupling and test point variations determine the loss between the sweep insertion point (IP) and the input to the amplifier or laser. In our experience, not having this information has been a major contributor to alignment problems.

We recommend developing a level matrix for your equipment which technicians can refer to when setting the source level. A sample matrix is provided in Table 3 (page 53). It may be necessary to contact the manufacturers of the specific hardware to verify the configurations. This table should provide a concise summary of the internal losses in the hardware and the level required from the return sweep source to provide a known level to the active device. It is also necessary to have block diagrams of the hardware with the available return test points.

The need for a readily available concise summary of this information should not be underestimated. It is often the case that sweep technicians seldom know the configuration of the return test points, or what the losses are to the active devices.

Normalize outputs

With the correct input level to the fiber node or first return amplifier, the output at the headend is measured with a spectrum analyzer or sweep receiver. Because the output from each return will vary by different lengths of return fiber or coax, these outputs should be normalized to the lowest level return by attenuating the higher level returns. This step creates a common output at the headend for all returns, assuming +10 dBmV input to the laser or first return amplifier. This common output level is referred to here as the "X" level.

Table 3: Source level matrix

$$\text{Sweep input level} = \text{carrier level} - 10 \text{ dB} = +20 - 10 = +10 \text{ dBmV}$$

	Type of hardware			
	Laser hub	Line extender	Trunk amp	Bridger
Sweep input level	+10 dBmV	+10 dBmV	+10 dBmV	+10 dBmV
Internal coupling loss	4 dB	1 dB	5 dB	13 dB
Test point loss	30 dB	20 dB	20 dB	20 dB
Total insertion point loss	34 dB	21 dB	25 dB	33 dB
Source level	+44 dBmV	+31 dBmV	+35 dBmV	+43 dBmV

Source level = sweep input level + IP = +10 + 34 = +44 dBmV

Check and align sweep response

All subsequent amplifiers should be adjusted to re-establish the X level output at the headend with the same +10 dBmV at each amplifier input. The amplifier is adjusted using the plug-in pad and equalizer for coarse adjustment and the gain and slope controls for the fine adjustment. The alignment should proceed from the fiber node or first return amplifier out, making sure each amplifier is calibrated properly before moving on.

Again, this may be done with a carrier generator approach, or a sweep system. Care must be taken if using the carrier generator approach because the frequency resolution is limited, and flatness problems may be missed. Once again, it is critical that the proper source levels be used. The level matrix created earlier minimizes the errors in this step.

The successful delivery of interactive services to subscribers is dependent upon the proper alignment of the return path. The key points are:

- 1) Attention to the alignment process is

absolutely critical;

- 2) Unity gain in the return path is as important as in the forward path; and

- 3) Good frequency resolution in both the return sweep and spectrum monitoring test gear can help identify many problems before they become customer complaints.

In the November issue, ongoing maintenance of the return path will be addressed.

References

1. CableLabs, *Digital Transmission Characterization of CATV Systems*, Nov. 1994.
2. Farmer, James O., Managing the Return Spectrum to Optimize Interactive Revenue Opportunities, *1995 NCTA Technical Papers*.
3. Farmer, James O., Paul Gemme, Charles Cerino, and Mark Millet, Two-Way Cable Plant Telephony and Data Experience, *1996 SCTE Emerging Technologies Proceedings Manual*.
4. Johanson, Brian, Bob Chamberlain, and Aravanan Gurusami, HFC Return System: Management of Subscriber-Induced Noise,

1996 SCTE Emerging Technologies Proceedings Manual.

5. Kim, Albert J., Two-way Plant Characterization, *1995 NCTA Technical Papers*.

6. Natarajan, Raja, and Paul Vilmur, Providing High Reliability Telephony and Data Services over HFC Networks, *1996 SCTE Emerging Technologies Proceedings Manual*.

7. Paff, Andy, Implications for the Introduction of Telephony Services on the HFC Architecture, *1996 SCTE Emerging Technologies Proceedings Manual*.

8. Prodan, Richard S., Majid Chelehmal, Tom Williams, and Craig Chamberlain, The Cable System Return Channel Transmission Environment, *1995 NCTA Technical Papers*.

9. Staniec, Thomas J., Making it Work: Return Systems 101, *CED*, August 1995.

10. Stonebeck, Dean A., and William F. Beck, Designing the Return System for Full Digital Services, *1996 SCTE Emerging Technologies Proceedings Manual*. **CED**



Proactive return path maintenance

It's time to sweat the small stuff

By Bill Morgan, R&D Project Manager, Hewlett-Packard

Editor's note: This article is part two of a three-part series on the return path. The final installment will appear in the February issue.

In last month's issue, we discussed return path alignment procedures and solutions to some of the problems that may be encountered while turning on the return path. Activating the return path is only the first step toward offering subscribers interactive services. If the quality of the service you provide is not superior, your customers will look to your competitors. To keep the competitive edge, you will need to be able to find and fix problems before they are visible to your customers. In this article, we will discuss the importance of a return path preventive maintenance program and look at issues surrounding proactive maintenance of the return path.

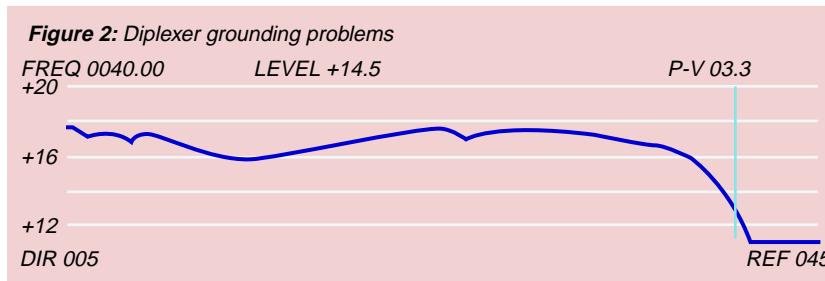

As fiber is installed deeper into today's cable systems, and the number of cascaded amplifiers is reduced, the tendency is to pay less attention to keeping the remaining amplifiers operating at their optimum performance. But with the activation of return path services, optimum performance is even more critical. Attention to detail is mandatory, regardless of the number of cascaded amplifiers and the age of the plant. Technicians need to be obsessed with details and investigate all discrepancies. In other words, don't ignore the small stuff!

Any maintenance program, forward or

return, requires a schedule to make it happen. You need a list of daily, weekly, monthly, semi-annual and annual inspections and procedures to be performed. Establish a checklist and keep it updated as the system architecture changes. Information from the routine tests should be recorded and kept in a database as a history of the active network.

Return vs. forward

Preventive maintenance for the return path



is similar to many of the current maintenance programs used in the forward path. These methods include periodic system sweep, end-of-line monitoring and periodic performance tests. Because the return path is configured differently and carries different services, some of the return path tests have changed. As an example, a carrier-to-noise measurement in the forward path can be made relative to a measured visual carrier because the forward path carriers are always active. Quite often, return path carriers are only on when transmitting

data, so a carrier-to-ingress measurement in the return path must be made relative to an intermittent data carrier or its expected level.

Cable operators are discovering that maintaining the return path is much more difficult than the forward path. The cost of maintaining a bi-directional system is currently two to four times greater than the cost of maintaining the forward path alone. Much of the increase in labor is related to multiple trips to the same location repairing self-inflicted problems. The goal of return path maintenance is to minimize trips to a given site by maintaining precise gain alignment. This can be accomplished by adding return path sweep testing and ingress monitoring to the normal forward path test program.

Sweep testing

Routine sweep testing of the forward and return paths gives the technician many advantages over other alignment methods. Forward path sweep response data is valuable information for the technician as he is troubleshooting the return path. The improved frequency resolution inherent in today's return path sweep systems enables the technician to find and fix problems before they become customer complaints. Some examples of return path problems identified by sweep testing are pictured in Figures 1 and 2.

Maintenance costs will be reduced if the technician sweeps the forward and return paths at the same time. Combining the two tests eliminates a second trip to the same location and minimizes housing openings which can aggravate ingress and leakage problems. In addition, many

amplifier designs have return signal paths on the forward amplifier motherboard. Therefore, if the forward amplifier is replaced, the return amplifier may also be affected. Faulty grounding of the modules in the amplifier can cause ingress problems in the return, but only have a minor effect on the forward path. Checking both forward and return paths at the same time can help eliminate future problems in both directions.

The sweep tech needs to understand the importance of return path unity gain, the prop-

er levels to insert into the return, and what levels to expect at the headend.

In last month's article (see October issue, page 42), we discussed using a level matrix to set the sweep source level. This matrix is just as important for maintenance as it is for alignment. If these levels aren't well controlled, all the work done during initial alignment can be wasted. Routine sweeping of the return should include sweeping the fiber link and reconfirming the normalized levels at the headend (the "X" level) to verify the performance of the optics.

Some manufacturers' distribution equipment will recommend different optimum input levels. This is sometimes found when comparing the recommended input to the return amplifier gain block and the fiber node return laser. This is a deviation from the unity gain theory. The gain block (return amplifier and cable loss) can have either gain or loss in order to set the signals to the correct level for the fiber node. But the philosophy that each amplifier in the return path is adjusted to compensate for the

loss of the section of cable following (in this case, the cable and passives between the fiber node and the first amplifier) still holds true. If your sweep gear uses an insertion point loss variable to adjust the actual sweep source level,

The safest approach is to sweep and align the new return before connecting it

this loss may be adjusted to compensate for the variation.

It is important to maintain a history of amplifier serial numbers, pad and equalizer values, and performance data for identifying trends. Keeping a record of the previous sweep results in the amplifier housing, or any place easily accessible to the technician, will also help to prevent or track down problems. Technicians should compare the current and previous results and question any discrepancies. Don't just change pad or equalizer values if the sweep insertion level is correct, and the levels at the headend are incorrect. Find the problem! The goal is to continually improve the performance of the system. Small incremental improvements each time the location is visited will keep the system at peak performance.

Table 1: Source level matrix

	Type of hardware			
	Laser hub	Line extender	Trunk amp	Bridger
Sweep input level	+10 dBmV	+10 dBmV	+10 dBmV	+10 dBmV
Internal coupling loss	9 dB	1 dB	5 dB	13 dB
Test point loss	30 dB	0 dB	20 dB	20 dB
Tap loss		24 dB		
Cable loss		8 dB		
Total insertion point loss	39 dB	33 dB	25 dB	33 dB
Source level	+49 dBmV	+43 dBmV	+35 dBmV	+43 dBmV

If it is necessary to sweep from a tap in the return, you can use this same level matrix. Determine the new source level by ignoring the amplifier test point loss and adding the loss from the tap to the amplifier output (return input). The level matrix in Table 1 illustrates an example of this new calculation. This same approach can be used for testing the return path from any return insertion point in the system.

Turning on new sections

Connecting new forward plant to an operational forward plant usually presents little risk to the existing system, with the exception of the power demands and tap insertion loss. In contrast, connecting new return plant into an existing system can create major problems. In Figure 3, if the last three amplifiers were connected into the system with zero attenuation, and the gain and slope controls set to the maximum output level, this could result in a significant increase in noise at the laser and disable the return.

It is extremely important to check the pad, equalizer, gain and slope values before connecting the new equipment.

The operator has a couple of choices when connecting new return plant to operational sections. The safest approach is to sweep and align the new return before connecting it. This method requires a return sweep system with a portable sweep receiver placed at the node and connected to the return feeder under test. When using this approach, the sweep source level is determined using the same source level matrix dis-

cussed earlier. The return amplifier pad and equalizer are then adjusted to achieve the required level at the return input to the node. This approach allows the technician to locate and fix any ingress problems before connecting the new section to the existing plant. After the new section is aligned, it can safely be connected to the existing return and a final check made for overall performance.

A more expedient approach is to maximize the pads in the new return section or disable the return amplifiers at the time of installation and then connect the new section to the existing trunk. Now the new section can be aligned in the same manner as the initial alignment using a return sweep receiver at the headend. The operator needs to be aware that using this approach, an ingress problem in the new section may impact the existing return before the ingress can be located and repaired.

Ingress monitoring

Some of the test equipment currently used for sweeping the return path makes it easy to check the headend return ingress while sweeping. Monitoring ingress as part of the sweep routine enables the technician to continually lower the composite ingress distortion as part of the routine maintenance. In addition, having access to headend ingress while sweeping provides an excellent troubleshooting tool when problems occur. By comparing the ingress at the




Figure 3: Typical system diagram

◇ RETURN PATH MAINTENANCE

Figure 4: Return path ingress performance





Figure 5: Percentage of return path availability



current field test location to the ingress at the headend, the technician can isolate the source.

Continuously monitoring ingress will also help the operator predict problems in the system by viewing slight changes in the system's performance over time. This process helps the operator direct maintenance technicians to potential trouble locations before system performance degrades.

Proper monitoring provides a measure of the availability of the system and a look at the performance of the return path. Problems in the return path can cause an increase in the ingress signals received at the headend. By sampling, storing and time stamping these measurements, the performance of the system can be viewed over time. These measurements can be compared to current measurements for troubleshooting purposes. We may be able to

**Return path
maintenance
requires
more diligence
than the
forward path**

One method of tabulating ingress data is to store peak and average values which give an indication of long term performance. An example of this type of data collected over a two-hour period is shown in Figure 4.

Another method is to record the number of times ingress is detected above a given threshold. This method can be used to provide a measure of percent availability. An example of this type of data also collected over a two-hour period is shown in Figure 5.

The NCTA engineering committee is in the process of developing recommended measurement practices for the return path, and among these efforts is radiated and conducted ingress. Results from this effort should be available early in 1997 and will provide guidelines for these measurements.

Summary

The successful delivery of two-way services to subscribers requires a higher reliability system. An effective proactive maintenance program can help achieve this level of performance while reducing resource demands overall. The key points to remember are:

1) Return path maintenance requires more diligence than the forward path

2) Routine sweep testing of the return will prevent problems and minimize repeat trips

3) Ingress management requires continuous monitoring of the return path. **CEC**

References

1. CableLabs, "Digital Transmission Characterization of CATV Systems," November 1994.
2. Farmer, James O., Paul Gemme, Charles Cerino, and Mark Millet, "Two Way Cable Plant Telephony and Data Experience," 1996 *SCTE Emerging Technologies Proceedings Manual*.
3. Kim, Albert J., "Two-way Plant Characterization," 1995 *NCTA Technical Papers*.
4. Natarajan, Raja and Paul Vilmur, "Providing High Reliability Telephony and Data Services over HFC Networks," 1996 *SCTE Emerging Technologies Proceedings Manual*.
5. Paff, Andy, "Implications for the Introduction of Telephony Services on the HFC Architecture," 1996 *SCTE Emerging Technologies Proceedings Manual*.
6. Prodan, Richard S., Majid Chelehmal, Tom Williams, and Craig Chamberlain, "The Cable System Return Channel Transmission Environment," 1995 *NCTA Technical Papers*.
7. Staniec, Thomas J., "Making it Work: Return Systems 101," *CEC*, August 1995.



Noise and ingress performance in the return path

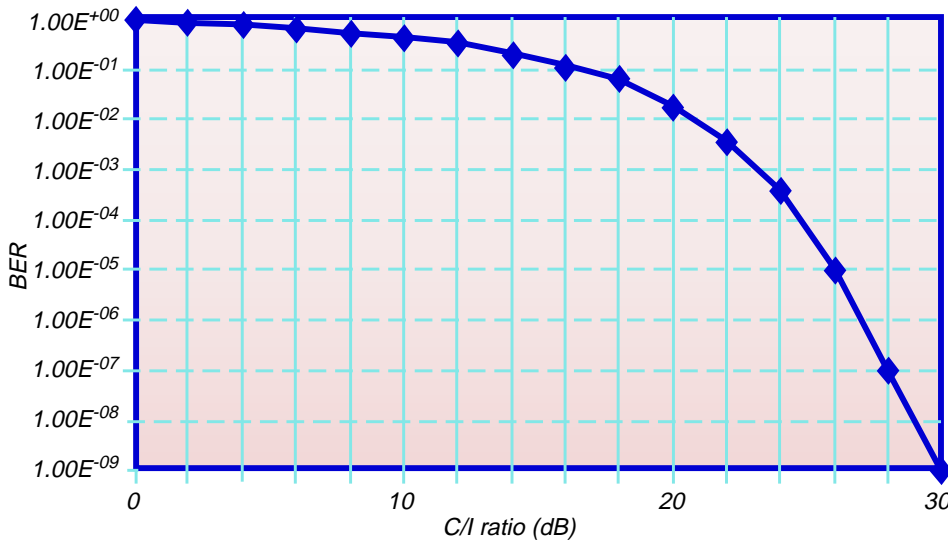
Starting a monitoring program

By Bill Morgan, R&D Project Manager, Hewlett-Packard

The previous two parts of this series (in the October and November 1996 issues of *CED*) discussed the proper procedures for alignment of the return path, and recommended proactive maintenance practices to keep the return path operating effectively. Cable operators are finding that ingress is the biggest roadblock to offer-

trum analyzer to the return path and check it daily or whenever a problem is reported. This can hardly be classified as a proactive maintenance approach, but it is better than waiting for the customer's complaint. As a first level of monitoring, this same spectrum analyzer can be connected to a computer and trace data stored periodically to create a history of your return path's performance. The analyzer can be configured to keep the sweep rate as fast as

Figure 1: BER vs. C/I ratio



ing the two-way services that customers are demanding. This article will recommend several ingress measurements to help quantify the return's performance, discuss the relationship between ingress and BER, and provide some suggestions for data handling. Prior to starting an ingress monitoring program, it is critical that your return path be properly aligned and your technicians understand the importance of maintaining proper gain balance. For more detailed information on this subject, we have provided references at the conclusion of this article.

Spectrum monitoring

The most common method of return path monitoring in use today is to connect a spec-

trum analyzer to the return path and check it daily or whenever a problem is reported. This can hardly be classified as a proactive maintenance approach, but it is better than waiting for the customer's complaint. As a first level of monitoring, this same spectrum analyzer can be connected to a computer and trace data stored periodically to create a history of your return path's performance. The analyzer can be configured to keep the sweep rate as fast as

possible and still have the frequency resolution necessary to identify problems. The data processing should include comparing the data to thresholds and identifying alarm conditions, as well as keeping a running average of the data to characterize the performance of the return path over long periods of time. By comparing the data samples which fall below a given threshold to the total test time, the user is able to derive a rough approximation of percent availability vs. frequency. There are many things that can be done with the results from this type of test, and there is a potential for generating massive amounts of data. Because of this, it is important to establish good methods for compressing the data

into usable archives. Programs designed for this type of monitoring are currently available.

BER vs. carrier/ingress

Although spectrum monitoring is a good first line of defense, it is important to monitor other key parameters in the return path and repair problems before customers are affected by slow response times or loss of service. Bit Error Rate (BER) is often discussed as one of the key parameters to use for monitoring return path communication channels, because this is a true measure of the performance delivered to the customer. Each of the services (cable modems, telephony, VOD, LAN, etc.) require varying levels of performance, anywhere from one error per 10,000 bits (10^{-4}) to one error per 10 million bits (10^{-7}). BER is one of the measurements that has traditionally been used in digital communications networks to monitor performance, and in a completely digital network, BER is an excellent metric to monitor. But in a mixed RF/digital network such as the cable TV environment, BER has several disadvantages.

Figure 1 is a typical example of the variation of BER vs. the carrier/ingress ratio (C/I) for a given type of modulation and error correction. The horizontal axis will shift depending on the type of modulation used, and the sharpness of the knee will vary depending on the type of error correction. The important characteristic to notice is the rapid increase in the BER as the C/I degrades near the knee of the curve. With only a 4 dB drop in C/I from 28 dB to 24 dB, the BER increases from 1×10^{-7} to 4×10^{-4} . In an RF network, the advanced warning received as the C/I degrades is limited if BER is the chosen parameter to monitor.

In addition, most of the digital services being carried by the return path today do not carry the reference bit stream required to make a BER measurement. The time required to make a BER measurement can be significant, and is defined by the equation:

$$\text{Time} = \frac{1}{\text{Data Rate} \times \text{BER}}$$

Using this equation, we can calculate that on a typical QPSK return carrier with a data rate of 1.5 Mbps, the time to measure 1×10^{-8} BER will be almost 67 seconds. Because the carriers are typically TDMA, the time required to capture enough data is lengthened even further. Another drawback to measuring BER is that it requires test equipment not currently found in most cable systems.

Another method often used to measure data transmission performance is monitoring packet errors. This is a measure of data packets (as few as 64 or as large as several thousand




Figure 2: 64 QAM constellation with ingress.




Figure 3: Wideband noise measurement.

bytes) which were re-sent because of errors detected by the receiver. Monitoring packet errors has the advantage of measuring the performance of an operating link without requiring embedded test code sequences. But packet errors also have a similar sharp transition from low errors to high errors as the C/I decreases.

A simpler and more effective approach to monitoring data transmission performance in an RF network is monitoring the C/I instead.

Once the type of digital modulation and the error correction are defined, the relationship between the measured C/I and potential BER performance can be predicted. If necessary, a specific digital link can be characterized by measuring the BER as the C/I is varied. This data will normally be provided by the test equipment vendor, but should be verified anyway. The sensitivity and advanced warning gained by monitoring C/I is significant.

How do we measure it?

The C/I ratio can be calculated directly by first measuring the average power of the carrier, and then the average power of the noise in the same frequency span. When measuring the return path noise, you are actually measuring a combination of many artifacts. They are, in addition to noise from amplifiers, common path distortion generated from the forward path signals, second- and third-order distortion generated from excessive signal levels, and impulse noise and inducted interference coming from the drop and home. In general, all of these artifacts degrade the performance of the RF/digital communications link.

Figure 2 is an example of a 64 QAM carrier constellation in the presence of noise. As the noise approaches the knee of the BER vs. C/I curve, the points on the constellation spread far enough that they become indistinguishable from each other, and data errors occur.

The biggest difficulty encountered when monitoring noise in the return path is the presence of carriers. One option is to measure the noise in a vacant portion of the return band. Unfortunately, the magnitude of the noise in a typical return path can vary quite a bit with frequency. This approach also misses narrow-band signals which fall on the carrier. But measuring noise offset from the carrier will help identify noise problems which cover wider bands, typical of impulse noise.

Gated spectrum monitoring

What is needed is a method for measuring noise in the presence of TDMA carriers. This is an ideal application for a measurement approach developed specifically for bursted carriers. The video gate was developed for the spectrum analyzer specifically for measuring intermittent events. In this case, the video gate can be triggered by the envelope of the detected data pulse. An in-depth discussion of gated measurements is beyond the scope of this article, but gated measurements are discussed in several articles referenced at the end of this text. In addition, gated measurement application notes are available from the spectrum analyzer manufacturers.

By using a gated measurement approach, and changing the video gating from the rising to the falling edge of the burst detector, the spectrum analyzer video can be gated both when the carrier is on for the carrier measurement, and when the carrier is off for the noise measurement.

Using currently available test equipment, there are several other tests which can be performed to help quantify the magnitude of the ingress in the return path.

Wide bandwidth average power

An average power measurement set up to measure the entire return band will provide a single quantity result (the ultimate in data reduction) for the average noise performance of the return path. This measurement algorithm is currently available in some spectrum analyzers for measuring the average power of digital carriers. To measure digital carrier power, the analyzer samples the signal in pre-determined increments across the bandwidth of the carrier and integrates the samples to arrive at the average power of the carrier.

This same measurement can be used to measure the average power of wide bands of noise simply by increasing the measurement span to the full bandwidth of the return path. Figure 3 is

becomes a frequency selective voltmeter with the selectivity determined by the analyzer's IF resolution bandwidth, when it is tuned to a frequency of interest and set to 0 MHz span. The analyzer's video output can be sampled at a high speed by a digital oscilloscope. The samples can then be compared to a user-defined threshold to locate impulse or burst events in the return path which exceed the threshold. If you keep a log of the number, amplitude and duration of the events, you will have a good picture of the burst noise performance of your return path. Figure 4 is an example histogram from the results of a 24-hour test.


The CW Tester™, developed by CableLabs, is a variation on this test that measures the interfering signal's impact on the amplitude

burst noise in the return path is just beginning. The NCTA Engineering Committee is in the process of writing an addendum to the Recommended Measurement Practices which will cover the return path. The goal is to provide measurement procedures which rely on currently available test equipment, when possible.

What should I do with all of this data?

The goal of this effort is to provide different methods of ingress measurement which can be used as required, depending upon the nature of the anticipated problems. Ideally, a combination of all of the above should be used to give the broadest picture of the return path's performance. If these measurements are made on the return prior to service activation, you can select where

Figure 4: Burst event histogram results at 16.3 MHz.



an example of this measurement, with the start and stop frequencies set to 5 and 30 MHz, respectively. If carriers are active in the measured band, they will be included in the average power result, but this measurement can be an effective method for comparing the performance of multiple returns before there are active carriers. This measurement provides a single number for the result which is a good indication of the average noise performance of the measured path. This is not the best measurement for finding impulse problems or very narrowband interference. But it will help spot 2 dB or 3 dB variations in the average noise performance over time, which are difficult to identify by looking at a number of spectrum analyzer traces.

Time domain burst event counting

There are several variations on tests which tabulate burst events. A spectrum analyzer

and phase of a reference CW carrier. The storage of data by the CW Tester™ is triggered by sensing when the reference carrier is shifted far enough (in phase or amplitude) to cause a data bit error. This additional phase information provides a unique signature of the burst, which may make identification easier. There are additional tests going on in which a high-speed oscilloscope is being used to monitor the return without frequency selectivity. This approach also has merit because the high-speed impulse noise in the return quite often has a wide frequency spectrum. Tom Staniec of Time Warner's Excalibur Group has even gone so far as to capture these impulses with a high-speed digital scope and re-create them with an arbitrary function generator in the lab to simulate return path disturbances.

Stay tuned to these trials, because the evolution of the methods for capturing impulse and

carriers are assigned in the return band to provide the best performance for that carrier. Some carriers, such as IPPV, can tolerate higher levels of noise and can be placed in portions of the band with higher noise/ingress. Telephony carriers, which have the lowest tolerance to impulse noise, can be placed in the portion of the band where the impulse events are at a minimum.

The amount of data captured in a good return path monitoring program has the potential of filling even the largest computer disk drive in a short period of time. It is important to implement data compression which will prevent storing massive amounts of data which will most likely never be accessed. The spectrum analyzer scans can be stored as the average, peak and minimum of the data over specified periods of time.

Another approach to use when analyzing spectrum scan data is to use a waterfall display, which adds a third dimension to the graphs for

◇ INGRESS MANAGEMENT

easier visualization of performance vs. time. A good example of this was provided by John Mattson and Joe Pendergrass of Arris Interactive in a paper presented at the 1997 SCTE Emerging Technologies Conference. But once again, the amount of data stored is significant.

The wide-band average power data could be stored as a graph of the average, maximum and minimum over time, with time as the horizontal axis in this case. This approach makes it easy to see trends in the data which may indicate gradual performance degradation. The best method for displaying the burst event data is in a histogram, with event length in the horizontal axis and number of events in the vertical axis.

Archival data from any of these measurements can be used to compare the performance of the plant as it grows, and subscribers are added. With a historical perspective of the return plant performance, gradual degradation will not go unnoticed.

The successful delivery of interactive services to your subscribers is contingent upon an efficient ingress management strategy. Know your options for measuring ingress and be prepared to

be creative. Effective ingress management, along with proper alignment techniques and a proactive preventive maintenance program, will assure that you are ready to provide the new services that your customers are demanding. **CED**

Acknowledgements

The author would like to thank Beth Armantrout, Syd Fluck and Jerry Green of Hewlett-Packard for their help in creating these three articles.

References

- 1) Staniec, Thomas J., "Return Systems 102: What goes around...", *CED*, December 1996.
- 2) Kim, Albert, Rogers Engineering and CableLabs, "Two-way Cable Television System Characterization," April 1995.
- 3) Prodan, Richard, Majid Chelehmal, Tom Williams and Craig Chamberlain, "The Cable System Return Channel Transmission Environment," *1995 NCTA Technical Papers*.
- 4) Prodan, Richard, Majid Chelehmal and Tom Williams, "Analysis of Two-way Cable System Transient Impairments," *1996 NCTA*

Technical Papers.

5) Raskin, D., D. Stoneback, J. Chrostowski and R. Menna, "Don't Get Clipped on the Information Highway," *1996 NCTA Technical Papers*.

6) Chen, Helen and Don Gardina, "Testing Digital Video on Cable TV Systems," *1996 SCTE Cable-Tec Expo Proceedings*.

7) Hranac, Ron and Thomas J. Staniec, "Making Two-way Work," *1996 SCTE Cable-Tec Expo Proceedings*.

8) Mayes, Rembrandt and Jeffrey Sauter, "HFC Return Systems—A Tutorial," *1996 SCTE Cable-Tec Expo Proceedings*.

9) Huff, John L., "Time-Selective Spectrum Analysis: Extending Spectrum Analyzer's Usefulness," *1983 NCTA Technical Papers*.

10) Mattson, John and Joe Pendergrass, "A Cost-Effective Approach to Return Path Monitoring," *1997 SCTE Emerging Technologies Conference Proceedings*.

11) Product Note #8590-2, available from Hewlett-Packard, discusses gated measurements for the HP 8591C cable TV analyzer.



800-452-4844 x HPTV
www.hp.com/go/catv