MEANINGFUL METRICS FOR RETURN PATH MONITORING

Bill Morgan and Craig Chamberlain
Hewlett-Packard

Abstract

The deployment of advanced data services and the decrease in homes passed per node in order to accommodate more data traffic is putting more pressure on developing a cost effective solution to return path monitoring. This paper will discuss some of the efforts made to date, identify metrics that are currently available for quantifying the performance of the return path, and discuss several new measurements which support this objective. The goal of this paper is provide quality information to the system operator and enable the delivery of higher performance return path services.

What Have We Learned?

Over the last several years, we have participated in several attempts to monitor and ultimately improve the performance of the return path. Large quantities of data have been gathered during this period providing insight about the types of data that directly contribute to improved return path performance. In collaboration with a software partner and a major customer, a monitoring system was developed which used a standard laboratory spectrum analyzer optimized to collect data on a single node as quickly as possible. The software used multi-levels of alarm detection and provided feedback to the user both graphically and statistically. The goal was to plan maintenance for the network based on this feedback and reduce the number of alarms. This was intended to improve the overall quality of the network.

There were several drawbacks with this approach; one was the cost of the system, per node monitored. There is a general belief that if the system had indeed provided a higher quality network, this cost could have been justified. The second and more severe problem with this approach was the low probability of a scanning spectrum analyzer capturing a burst noise event. This probability has been compared to the probability of capturing meaningful data about a rainstorm by placing a shot glass in a parking lot for an instant. A short period of time provides a very poor indication of the amount of rain falling, but after a long period of time the shot glass will provide a good indication of the average rainfall.

Shot Glass in a Rainstorm

To better visualize this, it is important to understand the process a spectrum analyzer uses to scan a frequency band. Because a spectrum analyzer is a narrowband receiver, at any one instant in time it is unable to look at a frequency span wider than its resolution bandwidth. To monitor a wide span of frequencies, the analyzer tunes to a sequential series of frequencies making a measurement at each step. This series of frequencies can be anywhere from 200 to 600 points depending on the analyzer being used.

If this were the only issue, sweeping 400 points the analyzer will be at each specific frequency point for 0.25% of the time. In reality, there are other overheads that reduce this time even further. When an analyzer finishes a sweep it has to retrace to the start frequency. This time is variable depending on the analyzer, but can be from 10% to 20% of the total sweep time. There is also overhead at each data point while the analyzer deals with phase locks and display issues. These overheads can chew up as much as 50% of the time at each data point. So in the final analysis, an analyzer sweeping 400 points at a rate of 20 mSec (an excellent sweep rate for a modern analyzer) may be spending less than 25 μSec , or about 0.1% of its time at any one frequency. This gets reduced even further as the same analyzer is combined with a switch to monitor multiple nodes. It starts to look like a shot glass in a parking lot, doesn't it?

The end result of this effort, which lasted for over two years, was a hard drive full of spectrum response data and no measurable improvement in the quality of the network. We followed this effort very closely and as a result have attempted to draw conclusions and improve the monitoring approach. As a result of this, collecting large amounts of data as quickly as possible is not the correct way to predict future network performance. In addition, spectrum scan data is especially difficult to interpret / correlate and very time consuming to use as a performance monitoring tool. The meaningful information collected from these initial efforts was long-term trend data from measurements made over days, months and seasons.

What is the Proper Perspective?

Another flaw with current approaches to return path monitoring is the assumption that the network is faulty and the monitoring is responsible for documenting these faults. It has been our experience that a properly aligned return path is typically well behaved and the monitoring effort needs to focus on meaningful metrics and performance trends, not fault documentation. It is absolutely critical to any network monitoring effort that the operator begin with a properly aligned network. It is surprising to us how often this is not the case. We need to be in the mode of proactively planning maintenance on the return path rather than reactively responding to faults.

An Alternate Approach to Monitoring

An alternative approach to monitoring is to establish a user-defined set of baseline performance metrics and monitor against these metrics, identifying when measurable performance degradations occur and enabling scheduled maintenance to resolve problems before they impact customers. Ideally, this needs to be coupled with metrics that are available from the terminal and headend equipment, enabling quick response to catastrophic problems requiring immediate attention.

For the purpose of this discussion, it is convenient to group artifacts which impact the performance of the return path into two broad categories.

Short-term events ($< \approx$ 200 mSec) are typically characterized as burst noise and are good indicators of problems in the network. They are usually effectively handled by the error correction algorithms in the data communications channel and do not directly impact the user, but will degrade data throughput.

Long-term events ($> \approx$ 200 mSec) are typically characterized as broadband noise or interfering carriers and in sufficient level can be catastrophic to the network and completely disrupt communications. Narrowband interfering carriers can be addressed by frequency hoping modems if these are available, but there is no place to hop to when the interference is broadband.

A good monitoring system should provide data on both categories of interference. The goal is to develop measurements that augment, but don't duplicate, data that is already available from the terminal devices and headend equipment. The balance of this paper will discuss several new measurements developed to meet these requirements, including carrier-to-noise on TDMA digital channels (Time Domain Multiple Access - signals in the return path emanating from multiple sources and sharing the same frequency spectrum), burst event counting and a dual dwell frequency table sweep.

Carrier-to-noise (C/N) is not a new measurement, but developing a good carrier-to-noise measurement in the return path requires the ability to accurately measure the channel power of a TDMA digitally modulated carrier. This capability has evolved from the ability to measure digital channel power on continuous carriers. The evolution of these measurements began with the development of the digital channel power measurement that was first available in spectrum analyzers in the mid 1990's.

Measuring Digital Channel Power

The NCTA Recommended Practices define digital channel power as "the power level as measured by a power meter which uses a thermal couple as a transducer. That is … the average power in the signal, integrated over the actual occupied bandwidth of the signal." The NCTA has provided several methods for making a channel power measurement from a single sample in the center of the channel with a correction for channel bandwidth, to full sampling across the band with integration of results. The integrated measurement is the preferred approach since the shape of the digital channel will not affect the accuracy of the result. It is important to understand how this measurement is made in order to understand its evolution.

Because the spectrum analyzer measurement bandwidth (resolution bandwidth or RBW) can be set narrow relative to the channel bandwidth, it enables the measurement of a single digital channel in the presence of adjacent channels. But because of this narrow bandwidth (see Figure 1), multiple measurements must be made across the bandwidth of the carrier and the results integrated and adjusted to arrive at the accurate channel power.

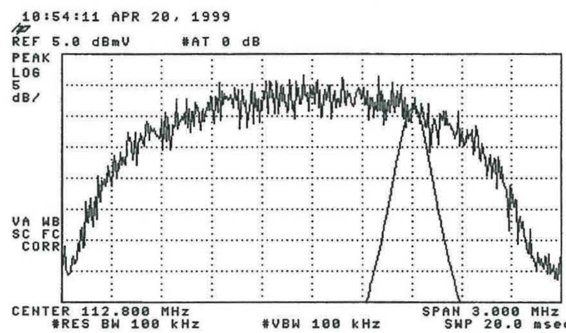

Figure 1 - Digital Channel Spectrum and RBW Filter Response

Using the noise equivalent bandwidth of the spectrum analyzer to define the measurement bandwidth of a single sample, the contribution of each sample to the total power of the channel may be calculated. This allows a simple integration of all the samples to calculate the total power of the channel.

Measuring TDMA Channel Power

The next requirement is to apply the same measurement capability to TDMA digital carriers. Figure 2 is an example of a TDMA digital carrier with a good demonstration of the alternate carrier and noise floor time slots. In order to measure the power of the channel while it is present, a threshold sampling approach is used to capture the TDMA carrier during its ON time. A properly set threshold for sample control will make sure samples are captured only while the carrier is present.

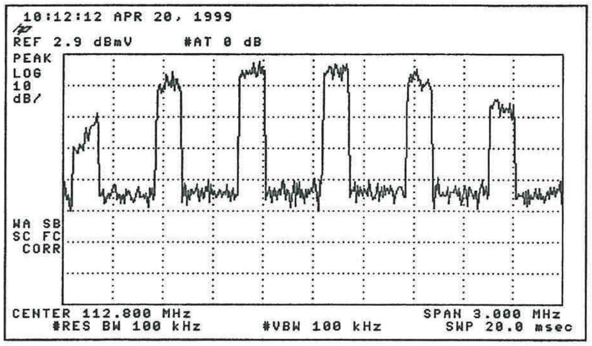

Figure 2 - TDMA Digital Channel Spectrum Response

Once the capability of measuring TDMA carriers using a threshold was developed, it became a natural migration to reverse the sampling algorithm and keep only samples below the threshold. This enabled the measurement of the average noise power in the TDMA channel bandwidth. The combination of these measurements provides the tools required for monitoring C/N in an active return path, providing an extremely flexible measurement for monitoring the performance of a return channel.

<u>Limitations / Variations</u>

There are several limitations to this measurement. The first is that in a TDMA channel slot, multiple transmitter sources are talking and a sampled approach for this measurement captures samples from many of the sources. Therefore, the result is an average of the power from multiple sources. If we assume a properly aligned network these sources which are controlled by a long loop AGC should be operating close to each other in level. The individual transmitter source level is typically available from the headend terminal equipment and this data may be used to monitor individual sources that are operating out of range.

This sampled measurement approach, although extremely accurate, requires a variable amount of time to capture samples, since it is dependent upon how often and for what duty cycle the channel is present. During high traffic time, it may take quite a bit of time to capture enough noise samples, and during late night hours when there is no traffic, it may require a similar length of time to capture channel power samples. But in a system that is designed to deliver long-term trend analysis, this is not a limitation.

In addition to this method, there are several others used for measuring carrier-to-noise in a TDMA environment. A reasonable approximation can be made of TDMA channel power by capturing a single sample and adjusting the result for the channel's bandwidth with a known correction factor. As discussed earlier, the accuracy of this result is affected by the difference in the shape of the channel measured and the shape of the channel used to calculate the correction factor. But it does provide a much faster measurement result since the instrument does not have to wait for the carrier to turn on multiple times. It is also possible to measure the average noise power using a single sample with bandwidth compensation, or measure the noise at a clear area of the spectrum offset from the channel. Both of these are also limited since they do not provide the true integrated noise power under the carrier.

Monitoring carrier-to-noise is an effective way to track longer-term events that affect the return path performance. Small changes in system gain or broadband noise performance show up quickly and can be repaired before network performance is affected. In addition, by measuring each individual channel, the user has the capability of setting different performance parameters on each service dependent on the quality of service required. This tool becomes very powerful as operators are required to provide a higher level of service to different customers (small business, etc.). This methodology has the flexibility to grow as new advanced measurements are defined and implemented.

Addressing Short-term Events

The averaging measurement used to monitor channel power and noise does not do a good job of capturing burst events. To address short-term events, two new measurements have been developed. An additional sampled measurement takes samples at a single frequency and measures the # of occurrences and duration of events that occur above a threshold. Knowing the number and duration of the events over a sample time indicates the "error free" time we could expect from the network. This result may be converted to an estimate of percent availability, a metric familiar to the telecom environment with value for any data communications environment. It is an estimated result since the sampling approach is not capable of capturing every burst transition. Percent availability is an excellent metric to use for trend analysis and also works well for qualifying a new frequency slot prior to the deployment of new services.

Other uses have been conceived for burst event counting. Intentional monitoring of specific operational frequencies can provide information about the use of these frequencies and enable the user to build a picture of usage. For instance, the user will know the approximate timeframe modems "talk" and also the amount of time converters require to request PPV events or to respond to a billing request. This can also be used to give an indication of mode usage and impending bandwidth limitations. In addition, having the duration of the events available provides a valuable troubleshooting tool to the cause of the events.

Another new measurement developed is a multiple dwell table scan measurement which scans a specific frequency table and measures both a short (100 μSec) peak detector dwell result and a long (9 mSec or variable length) peak detector dwell result. The short measurement has a tendency to provide the value of the ambient noise floor and is a reference point. The longer measurement captures the peak level of burst events, since it dwells at the frequency with an armed peak detector for a longer period.

When the two measurements are observed relative to each other, an indication of the increased "noise" from burst events is provided and this can represent an approximation of noise floor to short-term interference ratio. This measurement also gives the user more control over the traditional scanning approach of the spectrum analyzer. By sweeping a frequency table instead of an entire span, the speed of the measurement may be optimized for the needs of the user.

Another use of this measurement is the intentional monitoring of occupied frequencies in the spectrum. This provides rough data indicating if the path has amplitude stability, if the transmitting devices are operating within their designated ranges and if the system has failed during the test period. A lack of signal or even a substantial drop in noise level could indicate an actual failure. Using this information in an alarm window can alert the user to changes in noise and intermittent failure events.

Summary

The combination of these new measurement capabilities may be used in several different modes. Using feedback from the CMTS (Cable Modem Termination System) in the headend, these measurements can be focused on the nodes or channels that have been exhibiting packet or byte errors. This is more of a reactive mode, and is a mode that we prefer to stay out of but improves the efficiency of the monitoring. By making periodic measurements on a schedule, the baseline performance of the nodes may be used as a reference for watching trends. This is a preferred proactive mode that will allow the operator to schedule maintenance as nodes degrade, but before quality of service is impacted. In addition, the frequency of measurements on a node can be adjusted by the relative performance of one node to other nodes.

The data services being deployed in today's CATV network provide several methods of actively monitoring the performance of the return path. It is our goal not to verify or duplicate the results from the terminal equipment, but instead find a way to augment and begin to correlate these results. This can be accomplished by using measurements, which allow the tracking of trends in the performance of the network and predict network problems. This feedback allows maintenance to be scheduled so problems may be fixed before system performance is affected.

Bibliography

Gorin, Joe, Spectrum Analyzer Measurements and Noise - Measuring Noise and Noise-like Digital Communications Signals with a Spectrum Analyzer, Hewlett-Packard Application Note #1303

NCTA Recommended Practices for Measurements on Cable Television Systems, Second Edition, National Cable Television Association, Science and Technology Department