New Analytic Methods For Determining Network Performance Issues And Predicting Service Disruptions in Cable Networks

A Technical Paper prepared for the Society of Cable Telecommunications Engineers
By

John Holobinko
Independent Communications Consultant
&
(former) VP, Strategy and Business Development, ARRIS
jholobinko@comcast.net

Joe Jensen
Executive Vice President - Cable & Telecommunications
Block Communications, Inc.
Toledo, Oh
jjensen@cablsystem.com

David Bowler
Director of RF and Video Development
Arris Group, Inc.
Lowell, MA
david.bowler@arrisi.com

Jim Wolsiffer
Director of Technical Operations
Buckeye CableSystem
Toledo Ohio
jwolsiffer@cablesystem.com

Brian Basile
Director of Software Development
Arris Group, Inc.
Lowell, MA
Brian.basile@arrisi.com

Overview

Analytics - definition: *the method of logical analysis* - Merriam Webster Dictionary

An ongoing challenge for cable operators is network diagnostics. The costs of having incomplete or inaccurate data are: multiple service calls for the same problem; costly and unnecessary replacement of cable modems, set tops and drop cables; and in many cases poor data and voice service, or system outages.

Virtually all of today's cable network "Analytic" tools are misnamed in that they only collect and report large volumes of data primarily relating to network and device operating parameters such as SNR, power levels, and similar data. Typically, the analytics, ie. *the logical analysis* these tools provide is minimal. The process of analyzing and correlating vast amounts of data to identify specific network locations and specific causes of network issues is cumbersome and left to the burden of the operator.

This paper presents a significantly different approach for determining network performance issues and predicting service disruptions with dramatically improved network problem detection and resolution. The solution employs a comprehensive data repository which includes customer addresses, network topology and measured data from cable modems, and cable modem termination systems (CMTS's). True analytics based on extensive set of algorithms are used collaboratively to analyze this data repository and thereby derive the locations and causes of cable performance issues. The system approach is designed not only to pinpoint existing network problems, but importantly also to provide a means to preemptively address previously undetectable issues before they can become service affecting. The algorithms also rank network problems in order of their severity to enable operators to focus on the most significant issues in their networks in the order of their potential impact on services delivery and performance.

The paper includes a description of an actual cable system implementation where the software tool has been used to diagnose network problems and predict locations of future potential outages.

Contents

Introduction

The costs of hybrid fiber coax (HFC) network performance issues are if anything, underestimated. While the most attention is placed on truck rolls, the financial impact of performance degradation goes far beyond that measure. As we have seen from implementing the network analytics solution that this paper addresses, network performance issues cause increased costs with unnecessary replacement of customer premise set tops and cable modems, unnecessary replacement of drop cables, and dispatch of service personnel to wrong locations. Periodic network sweeps are another operational expense. But perhaps the greatest impact to the cable operator's business is on (reduced) subscriber penetration . In serving areas where problems cannot be resolved or more importantly, where the operator is not aware that a long term performance problem exists, customers begin to cancel their data and video services and seek alternative providers. This becomes a social behavior in which the experience of one or more neighbors impacts the consumer choices across an entire optical service area or service group.

Today it is common to define "proactive maintenance" as the act of repairing a problem before customers begin to report outages, as opposed to taking a corrective action to prevent a network problem before it can impact performance. It is questionable whether any commercially available tool can provide the operator with a clear view of system deterioration to enable true proactive network maintenance via interdiction before performance degradation or a system outage occurs. Existing tools provide overwhelming volumes of data around carrier to noise ratios (CNRs), signal levels, and bit error rates. It is up to highly experienced radio frequency (RF) technical experts to interpret this data and attempt to correlate it to a specific network performance problem. This is a prime example of a "big data" issue. The task is so daunting, that it is humanly impossible to assimilate and correlate all of this data. As a result, cable systems employing these tools for the most part still rely on trouble tickets to determine where to prioritize and to dispatch technicians to address system problems.

In 2010-2011, Cable Television Laboratories released the "Proactive Network Maintenance Using Pre-equalization Best Practices and Guidelines" (see bibliography at the end of this paper) in an effort to educate vendors and operators on how the pre-equalization data could be used to identify network faults. This guideline provides a wealth of information on pre-equalization and defines metrics on how one could use these coefficients in order to determine if the amount of correction being applied to an individual modem is potentially indicative of an issue. While this document is a valuable resource, locating the root cause of a network fault in a full scale, live HFC network can be a different matter entirely.

As we began dialog with operators over 18 months ago, a recurring issue voiced was that their existing diagnostic tools were telling them that their system was fine and that all operating parameters were within the normal operating limits. Yet, certain geographic areas were generating extraordinary numbers of tickets and there was no visibility into the underlying causes. A common example of this phenomenon is an area in which there are customer complaints and suspicions on the part of the operator that there must be a data problem in the return path based on the customer's experience with slow response and poor data transfer rates, but the existing tools report back that all return path modems in the optical serving area have robust RF output levels and CNR's all in the higher part of the range. Digging deeper, the modems of the customers reporting problems forward error correction (FEC) metrics have very significant pre-FEC and post-FEC error rates, which are resulting in enormous throughput issues. The explanation for this is that FEC error rates do not correlate directly to an equivalent impairment in the data rate. The effective data rate may be much worse. In some cases, a small performance degradation can result in the loss of a significant number of packets, which translates to slow response times for the customer and low data throughput. Therefore, resolving symbol errors in the physical link can provide a very substantial improvement in throughput by eliminating both the first order effect of dropped packets and the second order effect of latency due to retransmission.

Without adequate insight into the cause of a performance issue, technicians are often repeatedly dispatched to the wrong location in an attempt to solve a customer issue. While one or more customers at the end of a line may be experiencing problems, a dented cable, leaking directional coupler or improperly aligned equipment further up the line may be causing the problem, but the performance of customers located closer to the fault is not degraded. Technicians may become frustrated by repeated calls, changing out customer premises equipment (CPE), replacing drop cables and taps and doing other replacements out of desperation, still without being able to solve the problem. Ultimately, if the fault cannot be corrected, frustrated customers often disconnect resulting in a loss of revenue, in addition to the costs of truck rolls and equipment replacements. This loss of revenue is far more negatively impacting in the long term than the cost of truck rolls.

Another problem caused by lack of visibility into a problem is the dispatch of more than one team to address the same problem. Based on the trouble ticket, each team believes that it knows the cause of the problem and independently goes to solve it. This increases maintenance costs, and worse in some cases, the first technician that arrives and solves the problem can have his/her work reversed by a second technician that arrives and believes that s/he is solving the problem without knowledge of the corrective action that the first technician has taken.

One of the strongest arguments that passive optical network (PON) advocates use against HFC network advocates is the far greater yearly cost of maintaining an HFC system compared to a PON system. The argument put forward is that the greater cost of PON equipment compared to HFC equipment is more than offset by operational expenditures (OpEx) savings down the road. We believe that the cost differential of operating an HFC network compared to a PON network can be reduced significantly by implementing a network analytics solution.

Analytics Solution Objectives

We began by establishing the following functional goals for the analytics solution:

The solution must provide the following:

1. An easy means to identify problems and impending problems, including the ability to visualize specific problems in the network through the use of a mapping system (e.g. Google maps or equivalent) upon which the cable network and cable customers are displayed as an overlay.
2. A simple visual means to determine the status of the cable network and the ability to visually determine the locations (based on both geo-coordinates and street addresses) of any faults or parts of the network where performance is degrading, which can be securely accessed from anywhere by an authorized user. Ability to selectively view the network from a high level with adjustable granularity all the way down to the individual street and individual customer level.
3. Ability to identify the location of the fault (in addition and separate of identifying the locations of customers impacted by the fault). Rank potential causes of the fault.
4. Ability to graphically display via color code, the status of every service area, every device and every customer, and to display the actual associated data and performance numbers.
5. Ability to set variable thresholds on various defects and performance parameters, and ability to correlate alarms, as a means of prioritizing problems and managing the level of alarms both initially when the solution is implemented and after network performance is stabilized.

Implementation Goals

1. The functional goals should be achieved using no new equipment whatsoever in the network; with a bare minimum of additional equipment in the headend/hub that is necessary for data gathering.
2. Data gathering should put an absolute minimal load on the CMTS and not impact any data or video services performance
3. Initial input of all network data and customer location and services data should be accomplished automatically via simple interfaces to the existing mapping system and customer data bases, including a means to insure that any changes are rapidly and automatically reflected in the system. In the case where a network design is not available from a mapping system, the solution should have the ability to impute the network design automatically and superimpose this upon the graphical maps.

4. Processing shall be accomplished via a cloud-based model in which the customer does not have to have any additional computational capacity in their network.

Example System Design and Methodology

Figure 1 shows the architecture of the network analytics solution explored in this paper. The system host is cloud based and sits in a central location. One or more system collectors sit in the customer's network. The network topology and customer information is input to the collector via a simple file transfer function. Notably, customer adds/deletes and network equipment and topology changes can be fed automatically into the system collector, insuring that the analytics database is always up to date. Network performance data from cable modems and CMTS's are fed to the data collector so that there is little or no increased data load on the network. The system software creates a representative mapped solution of all the equipment, network topology and performance data in a hierarchical format.

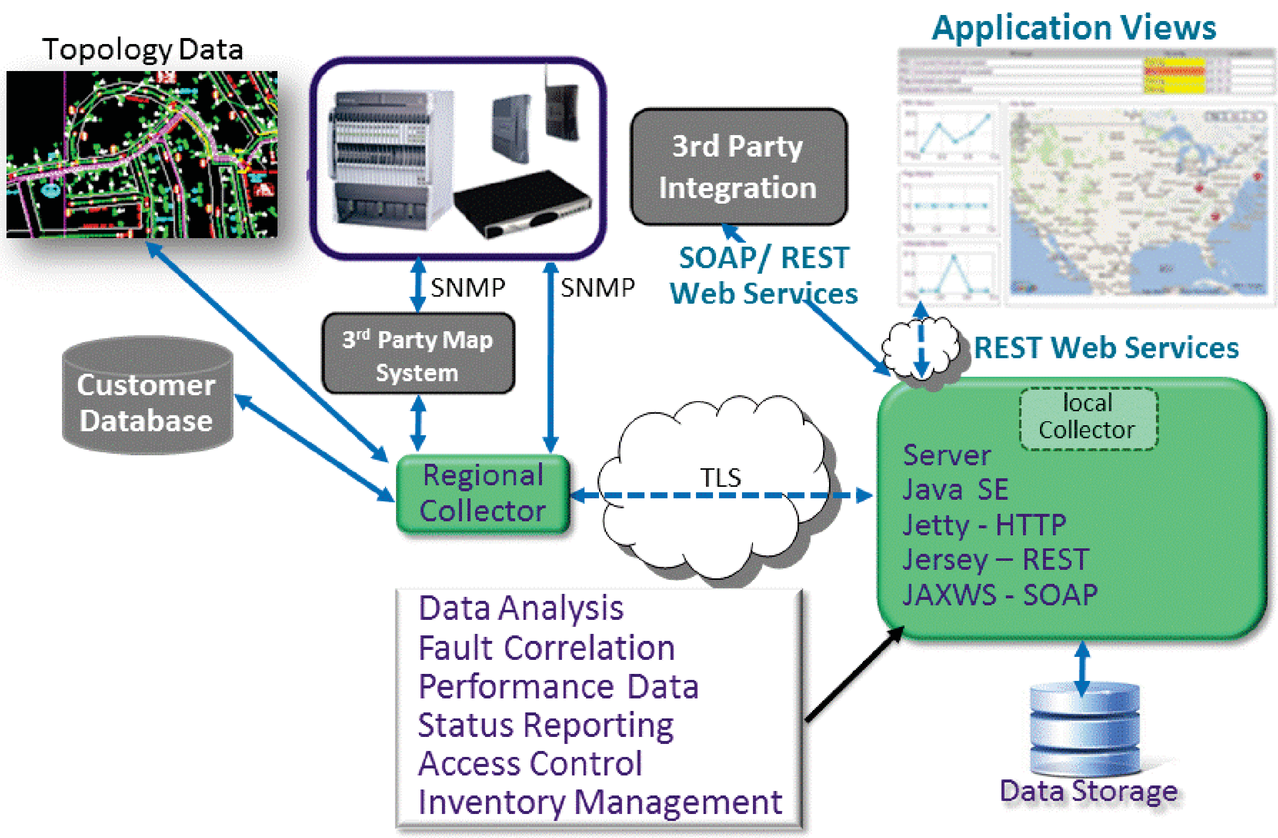

Figure 1: Example System Architecture

Fault determination and location is based upon the implementation of a series of algorithms which examine data returned by the CMTS and data modems in each service group, and correlate this data to the geographic address and network topology information in the database. While the use of algorithms to attempt to determine the location and cause of network faults has been attempted previously, the accuracy of such algorithms has not been sufficient to provide a consistently accurate analysis. To determine location of faults requires implementation of a number of expert algorithms which consider factors such as power levels, modulation error ratio (MER), signal to noise ratio (SNR), equalization coefficients and plant topology information. We began with a blank slate, and developed a new set of algorithms based on achieving the objectives defined earlier in this paper.

Lab Validation

To validate the example network analytics system employing the algorithms, we initially mapped a service area adjacent to our facility and reproduced this service area in the lab, using the amplifiers, equivalent cable lengths, DCs and tap values as in the system we were duplicating. Next, to insure that any results measured were not skewed to a single vendor's equipment, we acquired a variety of modems from a retail store. Once the test network was operational, we introduced a series of faults representing real world issues in networks and tested the ability of the solution to locate and determine the cause of each fault within the test network. Defects tested included water in taps, dented cables, poor connectors and poor shielding, improper levels, defective amplifiers, improperly installed or configured equipment and a variety of other defects. We also invited RF experts into our lab and provided them with any test equipment that they preferred to see if they could find the location and cause of the fault using traditional testing techniques. The results showed that that the analytics software was superior both in speed and accuracy compared to the RF experts tested.

Actual Network Validation

Next, we implemented the analytics solution within a portion of Buckeye CableSystem's live cable network. After first successfully validating the solution in Buckeye's lab, the initial network implementation covered an area consisting of eleven nodes with a total of approximately 3000 households passed. These nodes were served by two CMTS's, each from a different manufacturer so as to remove any possibility of the CMTS impacting results. The criteria used to choose the specific nodes for the implementation was based upon those nodes having a high number of customer tickets written against the serving areas, coupled with the fact that other installed diagnostic tools reported no problems within these serving areas. Customer premises equipment included cable modems, multimedia terminal adapters (MTAs) or both as well as digital set tops for video services, from more than one manufacturer. To gather data from the network a 1RU server with collector software was installed at the cable operator's headend where the CMTS's were located. This was the only piece of equipment required in the operator's network. The actual analytics application software was cloud-based, residing in our facility. To insure that there was no way for the customer's network to be compromised by hacking or other surreptitious means, the server had to initiate sessions with the cloud application, and was not able to respond to any external requests. Communications between the server and the cloud based application software was accomplished via an SSL link which was fully encrypted.

Upon start up, the analytics system rapidly detected a number of problems which heretofore had gone undetected at the network level and therefore were unsolved, causing poor or spotty performance in the downstream and/or the upstream, and intermittent service drops to some customers.

- Previously unknown upstream performance problems affecting customers connected to multiple different nodes were identified for the first time and made highly visible via the analytics program. The reported CNRs and RF levels of these modems were all excellent, and did not indicate any performance issue. In each node, the analytics program displayed a large number of data subscriber homes with either a red or yellow status. The analytics solution determined that the problem was not flatness or level, but rather attributed to group delay associated with the highest frequency return path channel on the system. The program reported poor post-FEC performance in the return path and determined that error levels were extremely high indicating that these modems lost sessions or operated very slowly the majority of the time. The problem was attributed to a return path DOCSIS® channel operating too close to the highest cutoff frequency where the diplex filter was introducing sufficient group delay as to cause the problems reported.

- A downstream performance issue was indentified in one section of a node which causing significant post FEC errors. Closer analysis revealed modems with RF levels that were suspect but not at alarm levels. However the system identified that an amplifier feeding this neighborhood was the common point of all customers affected. This amplifier was checked and although RF levels were close to spec it was found to be creating high levels of forward errors, and was replaced thereby correcting the performance issue.

- The system was also instrumental in identifying downstream performance issues in a node with a serving area that includes a large percentage of multiple dwelling units (MDU’s) and where the plant design was made based on typical residential requirements, This meant that in specific locations where conditions are "just right", system levels at these devices can be at the upper edge of specifications. The analytics system determined that the likely cause of poor data throughput in these MDU's was related to the one specific leg of the node serving this area. A sweep of this section of the node was performed to bring levels into spec and the performance issue was corrected.

- For a single customer, the system identified a single modem that exhibited downstream FEC errors on just two of the eight bonded channels in use and therefore relatively poor overall throughput. A search of that customer's history revealed they had not had a trouble call for quite some time. The analytics system pointed to the drop cable as the problem source. A technician was then sent to investigate. While testing the technician found significant downstream errors at the ground block and noted what appeared to be some “not-to-spec” splice work done on the customers drop. The drop itself was found to have been ripped apart and repaired by twisting the conductors together and wrapped with electrical tape. After confirming that the FEC errors were not seen at the tap, the drop was replaced and checked for proper operation at the ground block. Checking with the customer the tech was told they had seen issues of slow loading web pages and speed checks that were not consistent with the service level they had subscribed to but had not taken the time to call in. They were very thankful that Buckeye Cablesystem had identified a problem and corrected it for them.

In every case the analytics solution was able to find problems that were causing trouble tickets to be generated and customer dissatisfaction, and that were not identified by other systems. Visually, the solution displayed customer penetration across the geography, making it visually easy to correlate high problem areas with low customer penetration.

Measuring Solution Effectiveness

A common but inaccurate measure of network problems is the measure of truck rolls. The irony is that upon implementing this solution, truck rolls will not immediately decrease. Rather, they will initially go up as the analytics solution identifies network problems that were previously undetected. The greater initial benefit with be that the number of *repeat truck rolls* to solve any specific problem will decrease markedly, and over time, the number of network problems reported by customer tickets will decrease. Therefore, we suggest the following:

1. A far more accurate measure of improvement is to track the average number of service calls necessary to resolve a trouble ticket, plus the average length of time necessary to resolve the trouble ticket.
2. Measure the number of set top boxes and cable modems that are replaced. This should go down as technicians are able to identify the root cause of the problem, as opposed to "guessing" and showing the customer that they are at least "doing something" in an attempt to solve the problem.
3. Similarly, measure the number of drop cables that are replaced over time, as historically technicians have struggled with and guessed at the root cause of problems.
4. Measure overall customer satisfaction and data service penetration. We anticipate that similar to the RF ingress problem that caused poor data performance for an individual customer in the last example above, solving problems will improve customer satisfaction and as this happens, word of mouth will result in higher penetrations. In fact, as a problem area’s network

performance issues are rectified, this could provide the basis for very targeted incentive marketing, since all customer addresses in the affected areas are known in the database.

5. While the trial at Buckeye Cable has run for a number of months, we have reason to believe (but have not yet proven categorically) that entire system sweeps can be either reduced or eliminated resulting in significant operations savings.

With the ability to determine performance issues before they reach service impacting levels, customer outages should decrease leading to an overall improvement in network reliability. We believe that improvements in network performance will lead to a gradual increase in customer satisfaction.

Conclusions

We have demonstrated via implementation in a live cable system that it is possible to create a true analytics solution that is capable of processing “big data” efficiently via algorithms to identify performance issues heretofore difficult to detect, and enable true proactive system intervention of problems before they create systems outages and to rapidly detect any outages that do occur. By creating a much higher certainty around problem determination, a specific technician with the right skill level and appropriate equipment for the specific identified issue can be dispatched to correct each problem on the first call. Guesswork can be eliminated resulting in lower replacement rates for drop cables, set tops and cable modems. System performance as perceived by customers can be enhanced through the ability to significantly improve BER rates and therefore data throughput and reductions in macroblocking. As network performance improves, traditionally lower penetration rates have the potential to rise as overall customer satisfaction increases.

As more wireless spectrum becomes utilized, the problems of interference in cable systems due to LTE signal egress will likely increase.. The challenges of keeping a network operating at peak performance and providing the best customer experience will become more difficult versus less difficult. As cable systems implement DOCSIS® 3.1, to achieve greater efficiency will require cleaner operation, or the gains of orthogonal frequency division multiplexing (OFDM) will be offset by higher error correction rates leading to little or no improvement in network throughput. Therefore, it will be important for operators to employ network analytics as a means to improve network performance. The benefits are a reduction in cable OpEx, higher customer satisfaction and the potential of higher average revenue per mile of plant.

END

Bibliography

"Proactive Network Maintenance Using Pre-equalization" DOCSIS® Best Practices and Guidelines*; CM-GL-PNMP-V02-110623; ©2010-2011 Cable Television Laboratories, Inc.*

"Broadband Cable Access Networks: The HFC Plant"; *David Large, James Farmer*

"Probability, Random Variables and Stochastic Processes"; *Athanasios Papoulis*

"SCTE Measurement Recommended Practices for Cable Systems, Fourth Edition"; Society of Cable Television Engineers; http://www.scte.org/devams/cgi-bin/msascartlist.dll/ProductInfo?productcd=TS46

Abbreviations and Acronyms

Abbreviation	Description
OpEx	Operating Expense
Post-FEC error rate	The Error Rate After Forward Error Correction
Pre-FEC error rate	The Error Rate Prior to Forward Error Correction

Acronym	Description
BER	Bit Error Ratio
CMTS	Cable Modem Termination System
CNR	Carrier to Noise Ratio
CPE	Customer Premise Equipment
DOCSIS®	A CableLabs interface specification that enables high-speed Internet services over HFC. The DOCSIS® brand for these specifications and devices built to them developed from the specifications' original name, "Data Over Cable Service Interface Specifications."
FEC	Forward Error Correction
HFC	Hybrid Fiber-Coaxial
LTE	Long Term Evolution
MDU	Multi Dwelling Unit
MER	Modulation Error Ratio
MSO	Multi-System Operator
MTA	Multimedia Terminal Adapter
OFDM	Orthogonal Frequency-Division Multiplexing
PON	Passive Optical Network
RF	Radio Frequency
RU	Rack Unit
SNR	Signal to Noise Ratio
SSL	Secure Sockets Layer