Profile Management Informed Proactive Network Maintenance

Comparison of RxMER Per Subcarrier, Bit Loading, and Impairment Driven versus Measurement Variability

A Technical Paper prepared for SCTE•ISBE by

Jason Rupe
Principal Architect
CableLabs®
858 Coal Creek Circle, Louisville, CO 80027
303.661.3332
j.rupe@cablelabs.com

Jingjie Zhu
Senior Engineer
CableLabs
858 Coal Creek Circle, Louisville, CO 80027
303.661.3312
j.zhu@cablelabs.com

Table of Contents

Title **Page Number**

List of Figures

Title **Page Number**

List of Tables

1. Introduction

With the introduction of DOCSIS® 3.1 technology, profiles and profile management became important to operators who wish to get the most out of their network capacity. With the resiliency advantages of DOCSIS 3.1 technology, network impairments impact service through a loss of capacity but not failure of service until these impairments become severe. Therefore, there is an opportunity to use lost capacity as a way to measure the severity of impairments, and to prioritize proactive network maintenance (PNM) work as well. But with profiles being limited in number on many cable modem termination systems (CMTSs), the theoretical maximum capacity can't always be obtained. So, profiles need to be considered to truly measure the impact of proactive maintenance operations.

In this paper, we present some competing methods for prioritizing PNM work in terms of optimal possible profiles, as well as the measures involved in setting these profiles (RxMER per subcarrier, and bit load). Depending on the conditions of the operator's network and PNM tool capabilities, some solutions will be better than others. So, we will help operators decide an approach that works best for them. The solutions we compare will be offered as workers in our Proactive Operations (ProOps) platform as well, so that anyone interested can conveniently try them for themselves.

Identifying PNM opportunities is one challenge, but prioritizing them is another, and selecting the important problems to work on is yet another. All these steps must be done right for PNM to be effective because repair resources are limited.

Profile management is the practice of optimizing bitrates overall for the cable modems (CMs) on a CMTS; Profile Management Application (PMA) is the name of the application that optimizes profiles, say to maximize overall bit rate across all CMs subject to a limited number of profiles to share.

Operators express that they are reluctant to do proactive maintenance in part because they do not know which problems to address, or which problems will impact service in the future. But the way DOCSIS resiliency mechanisms work, bitrate is sacrificed for service reliability in the face of impairments. So, while one impairment may not itself impact service, a future one will because the impacts pile up until the customer notices. All impairments impact bitrate in some way, be that through codewords that need correction, data that has to be re-sent, or subcarriers that have to be avoided, for examples. This suggests that the impact to bitrate may be a good candidate for prioritizing and selecting important proactive maintenance projects to tackle.

But how do you determine the bitrate impact of impairments? That is where a look at PMA can be informative. For downstream PMA, downstream RxMER per subcarrier is collected and used to determine bit loading, which in turn informs bit load, which informs how the profiles are set. Therefore, there are three potential candidates to consider: RxMER per subcarrier, bit loading, or the profiles assigned.

Therefore, we investigate each of these approaches as methods for assigning PNM work. For some of these methods to work, however, anomaly detection is all but necessary. We also look at possible methods that include anomaly detection using a machine learning approach, and methods that avoid use of an anomaly detector (AD) when possible, for cases where AD is not available. Statistical methods can be used to detect impairments as well; for example, operators and vendors who are members of CableLabs can obtain a copy of Spectra which detects anomalies in spectrum data, including methods that were contributed back by Comcast.

After outlining methods for consideration, we test these methods with an available data set, and compare the merits of the approaches. While we don't have access to operator networks to test the methods in deployment, the comparisons we have made are informative enough to narrow the options down and inform which may be tried first in various situations. We encoded some of the better options into the ProOps platform so we can work with operators to test and tune the methods for their use. These solutions are available to CableLabs member operators for free, and provided in some form to vendors as well.

2. Approaches Studied

Next, we introduce the approaches we defined and compared for prioritizing PNM work. Each of these can be updated after information changes, such as the completion of maintenance.

2.1. Prioritize by worst profile group without anomaly detection

The target in this method is the worst performing CMs in the worst profile. If your PMA calculation uses anomaly detection within it, then anomaly detection is used indirectly but its information is not directly used in this case.

1. For each CM, calculate its profile from PMA based on the configured number of profiles in its CMTS (Assigned Profile, as is typically done with PMA), and the best profile it could have been assigned with unlimited profiles possible. (Possible Profile is a profile it would be assigned if it were in a cluster of one). Note this is before consideration of impairments that can be removed.
2. Calculate the bit rates (bit loading) these CMs should achieve with these two profiles (BR_AP and BR_PP). [1]
3. Cluster CMs by profile. Given the current PMA approach is to cluster before setting profiles, the clustering comes essentially for free in step 1. Order the clusters by bit rate from lowest to highest so that the worst profile group is selected for consideration first.
4. <u>Start with the profile with the lowest BR_AP</u>; find the CMs in this profile where BR_AP = BR_PP. If none exists, find the CMs with the lowest absolute difference. Call this CM subset CM_Cluster_i for i=1. Assign to this CM subset CM_Cluster_i a severity measure calculated as the difference between BR_AP and the average BR_PP of all CMs sharing the profile but have a higher BR_PP than the group. So, we are assigning to each CM a cluster number and severity number.
5. Remove this CM_Cluster_1 set of CMs from the set of CMs and recalculate the profiles using PMA (returning to step 1 above).
6. Repeat the steps above to find CM_Cluster_2, and repeat again, until all CMs are clustered into CM_Clusters_i for i = 1 to n for an arbitrary n. Note that n will be larger than the number of profiles. The severity will be used to assign importance for maintenance of the CMs in question, and the CM clusters too.

The work is thus clustered. Sum the severities for the CMs in the cluster. The resulting severities are the priorities to sort by for the work.

Note that the number of profiles has to be discovered from a trusted source such as the engineering team. While the CMTS has a maximum number of profiles possible, operations engineering teams can decide to

---

[1] In a different version, one could calculate the bit rate for each CM if it could be assigned an optimal profile (max bit rate achievable given how the CM was provisioned) (BP_OP). This could also be defined as the target profile based on provisioning targets.

use a smaller number of profiles when implementing. This setting, therefore, will need to be configured in ProOps or your own PNM application environment for each CMTS.

Measuring the severity of CM clusters in this way presents a tradeoff. Clustering by CMs sharing a profile may correlate with geography but maybe not enough to be efficient with truck rolls. And there is not a guarantee that one cause leads to this CM cluster to share a profile. The theory behind this approach is that one or a few related problems are leading to the lowest profile to bring down a group of CMs. If one or two problems are impacting the group, then this approach should highlight the group of CMs experiencing the problem. But if that is not the case, then the clustering may not work as intended.

Due to the nature of profiles, and the small number of profiles to share among all CMs, one change of a CM's performance can cause re-optimization of the mix of profiles; while some CMs get a better performing profile, others can get a worse one. This makes comparison and prioritization difficult.

If a group of CMs is working fine at a low profile due to being at the end of a line, with low RxMER average values, a profile may be selected in this approach which has no addressable issues. Therefore, human intervention will be needed, and perhaps additional rules or checking is needed to remove false positives.

The need for information from experts to inform the profile management, as well as the fact that this method may indicate false positives, suggests it will be difficult to manage.

This method can easily be augmented with anomaly detection. By adding anomaly detection to the calculation, selecting the worst performing CMs, removing the anomalies, and recalculating the profiles to see the impact, the profile information can be used to prioritize proactive work. Using anomaly detection should reduce the expected false positive rate significantly, and make the method worth considering.

2.2. Prioritize by worst profile modem without anomaly detection

This method is one which we encoded and will provide comparisons later in this document. This procedure below aligns with our measurement labeled Severity(PMA).

The target in this method is the worst performing CMs over all profiles in terms of their impact on the profile. If your PMA calculation uses anomaly detection within it, then anomaly detection is used indirectly but its information is not directly used in this case.

1. For each CM, calculate its profile from PMA based on the number of profiles possible in its CMTS (Assigned Profile, as is typically done with PMA), and the best profile it could have been assigned with unlimited profiles possible (Possible Profile, a profile it would be assigned if it were in a cluster of one). Note this is before consideration of impairments that can be removed.
2. Calculate the bit rates (bit loading) these CMs should achieve with these two profiles (BR_AP, BR_PP).
3. Cluster CMs by profile. Given the current PMA approach is to cluster before setting profiles, the clustering comes essentially for free in step 1. Order the clusters by bit rate from lowest to highest so that the worst profile group is selected for consideration first.
4. <u>Searching over all profiles</u>, find the CMs where BR_AP = BR_PP (plus or minus a small delta). Or, if none exist, find the CM(s) with the smallest difference in BR_PP-BR_AP. Call this CM subset CM_Cluster_i for i=1. Assign to CM_Cluster_i a severity measure calculated as the difference between BR_AP and the average BR_PP of all CMs sharing the profile but have a higher BR_PP than the group. So, we are assigning to each CM a cluster number and severity number.

5. Remove this CM_Cluster_1 set of CMs from the set of CMs and recalculate the profiles using PMA (returning to step 1 above).
6. Repeat the steps above to find CM_Cluster_2, and repeat again, until all CMs are clustered into CM_Clusters_i for i = 1 to n for an arbitrary n. Note that n will be larger than the number of profiles. The severity will be used to assign importance for maintenance of the CMs in question, and the CM clusters too.

The work is thus clustered. Sum the severities for the CMs in the cluster. The resulting severities are the priorities to sort by for the proactive work.

Just as in the previous approach, this approach has the same tradeoff due to the clustering utilizing profiles, and using profile impact as the measure of severity.

Like the previous method, this method can easily be augmented with anomaly detection to allow selection by anomaly that potentially would be removed.

2.3. Prioritize by profile bitrate impact using anomaly detection

This method is one which we did not encode for comparison as we expect it to be somewhat like the previous one which we did. However, we think it may be one worth exploring with operators.

The target in this method is the anomalies in the signal, so anomaly detection is explicitly required. But we still use profiles to determine which anomalies to dispatch for.

Assume we have a list of impairments generated for a set of CMs on a given node. Like impairments across CMs, they need to be clustered so that each impairment is associated with a list of CMs that show the impairment, and each CM could be associated with any non-negative number of impairments.

1. For a set of CMs on the same node, and impairments found to appear on the set of CMs, order the impairments found by severity on the CMs and create an association list for each impairment containing the CMs that show the impairment.
2. Find the profiles for each CM on the node using PMA and call these profiles the Assigned Profiles.
3. Take the most severe impairment on the ordered list of impairments, as measured by the amount of impact to RxMER on the subcarriers that it impairs, using a method such as those that follow. Adjust the RxMER per subcarrier for each CM impacted by the impairment as though the impairment was removed (methods like LMS (Cole 1990)), sliding window average or median, EWMA (bi-directional averaged), or FFT-smooth-IFFT methods may be useful here). One reasonable approximation is to assign new RxMER values over affected subcarriers such as the statistical average of the non-affected subcarriers, or a target value for the subcarriers, or for impairments like waves to assign the average of affected subcarriers. Find new profiles for the new set of CMs on the node using the adjusted RxMER per subcarrier.
4. Take the difference in the capacity gains of the newfound profiles and the previously calculated profiles, then assign the delta value as a score to the impairment. Depending on desired approach, cluster the CMs by shared anomaly, profile, or not at all.
5. Update Assigned Profile with the newfound profiles, for all CMs on the node, and repeat the process above (return to step 2) for the next impairment until no more impairments are on the list. All impairments have a severity assigned now.

The default is to cluster by impairment found, as that is given in the process. Alternately, we could sum up severities by CM or profile and assign work based on the severities summed.

By using RxMER impact or bit load as the measure of performance, prioritization is straightforward. There are impacts in these measures due to clustering, as profile changes can have the same impact as discussed above, and removal of an impairment can impact more than one CM. But, the measures of RxMER impact (dB or linear) or bit loading are closely linked to network capacity and service bandwidth, so are excellent candidates for prioritizing work. With AD, a focus on the impairment, and therefore the issue to address, reduces the need to cluster. The clustering only helps to measure the impact of the impairment.

Note: A possible adjustment to this method is to take the square of the difference of RxMER values by subcarrier for each CM and sum those square values so that deviations are measured as a squared loss function instead of linear.

2.4. Prioritize by bit-loading or RxMER per subcarrier impact using anomaly detection

This method is one which we encoded and will provide comparisons later in this document. We label the measurements for this process as Severity(MER) and Severity(Bitload).

The target in this method is the anomalies in the signal again, so therefore anomaly detection is explicitly required here too. But this time we just use RxMER per subcarrier or the bit loading possible, avoiding explicit use of PMA.

Assume we have a list of impairments generated for a set of CMs on a given node. Like impairments across CMs need to be clustered so that each impairment is associated with a list of CMs that show the impairment, and each CM could be associated with any non-negative number of impairments.

1. For a set of CMs on the same node, and impairments found to appear on the set of CMs, order the impairments found by severity on the CMs and create an association list for each impairment containing the CMs that show the impairment.
2. Find the RxMER per subcarrier for each CM on the node and call these the real values.
3. Take the most severe impairment on the ordered list of impairments, as measured by the amount of impact to RxMER on the subcarriers that it impairs, using a method such as those that follow. Adjust the RxMER per subcarrier for each CM impacted by the impairment as though the impairment was removed (methods like LMS, sliding window average or median, exponentially weighted moving average (EWMA) (bi-directional averaged), or FFT-smooth-IFFT methods may be useful here). Note this is a clustering step as well, resulting in CMs that are clustered around an anomaly. One reasonable approximation is to assign new RxMER values over affected subcarriers that are the average of the non-affected subcarriers, or a target value for the subcarriers, or for impairments like waves to assign the average of affected subcarriers. Call this adjusted RxMER per subcarrier the adjusted values. Apply this step to all impacted CMs in the group.
4. Sum the improvement in bit rate (best possible bit loading) that is gained by removal of the impairment, over all impacted CMs, and assign this result as a score to the impairment. Alternately, one can use the improvement in RxMER per subcarrier values summed over subcarriers.
5. Update the real RxMER values with the adjusted RxMER values for all subcarriers, for all CMs on the node, and repeat the process above (return to step 3 with the updated RxMER per

subcarrier values) for the next impairment until no more impairments are on the list. All impairments have a severity assigned now.

The CMs are clustered by anomaly, so that can be used easily by default. Individual CMs can drive the work too, depending on how the user wants to apply the severity and sort for work.

As this method uses AD and measures the impact of the removal of the anomaly, the advantages of the previous method apply here as well.

Note: A possible adjustment to this method is to take the square of the difference of RxMER values by subcarrier for each CM and sum those square values so that deviations are measured as a squared loss function instead of linear.

2.5. Prioritize by bit-loading or RxMER per subcarrier without anomaly detection

This method targets poor performing CMs without the use of AD or even PMA or profile information of any kind. It instead goes to the source of profile information, the downstream RxMER per subcarrier values for the CMs involved. We encoded this measurement for comparison later, indicating the measurements as NormalScore(MER) and NormalScore(Bitload).

This simple approach is offered as a method to begin using and developing from. The PMA application is shared in limited ways with vendors, and the CableLabs AD is shared with vendors as an executable only. Not every operator will implement PMA. Such operators are able to make use of this method here too, but some of the methods using PMA and AD are expected to perform better for many operator applications.

1. For a set of CMs on the same node, find the RxMER per subcarrier for each CM on the node, and call these the real values.
2. For each CM, adjust the RxMER per subcarrier for each CM through smoothing (methods like LMS, sliding window average or median, EWMA (bi-directional averaged), or FFT-smooth-IFFT methods may be useful here). One reasonable approximation is to assign new RxMER values over affected subcarriers that are average of the non-affected subcarriers, or a target value for the subcarriers, or for impairments like waves to assign the average of affected subcarriers. Call this adjusted RxMER per subcarrier the adjusted values.
3. Take the sum of the squares or absolute value of the differences in each subcarrier between the real and adjusted RxMER values and call this the severity. A bit loading calculation can be done here as well, and the overall bit loading may be used as the measure instead (inverse of bit load, or difference in bit loading from the real and smoothed RxMER per subcarrier results).
4. Prioritize and sort the CMs by their severities.

GIS information would be needed to cluster work without alternatives such as by anomaly. CM performance may inform that clustering too, as experts have suggested.

3. Details, Discussion, and Evaluation

3.1. A note on using these results in a PNM framework

To use these approaches as workers in ProOps, and to have a solution that turns data into actionable information, we need to have elements that collect data and calculate statistics (observe), assign priorities and cluster work (orient), sort the clusters and choose which are opportunities (decide) on which to act (act). The elements in these flows sometimes do not follow an order for OODA, but the elements are represented. Some of the elements which may be separate workers when coded for ProOps, and the roles they may provide, are described next.

- Observe
  - RxMER per subcarrier
  - AD
  - Profile
- Orient
  - RxMER improvement or severity
  - AD severity
  - Profile improvement
  - Anomaly matching
- Decide
  - Severity assignment
  - Sorting results
- Act
  - Dashboard
  - Work list

For an initial configuration, here are the general configuration elements and steps.

- Collect RxMER per subcarrier every 4 hours, and profiles if relevant, then immediately run through any AD, or statistics calculators.
- Calculate any severities or scores needed.
- If the severity calculation did not require clustering yet, then execute on any clustering needed.
- Execute next on any added data to collect for future AD or impairment detection training, etc.
- Calculate any statistics needed based on the clustering, including calculating the cluster severities.
- Sort the clusters or single entities by severity, and calculate any needed data for the dashboard and prioritized work list, which may have a threshold value applied as criteria for inclusion in the work list.
- Present results in a desired file format or data base, along with updating the dashboard.

3.2. A note on clustering work by anomalies

As we examine anomalies found in CMs, we will calculate their impact on RxMER per subcarrier or on profiles and assign a value to the CMs where the impairment appears. But in doing this, we want to also combine CMs by impairment in many cases. Combining by profile is done in some methods by default. But working by anomalies may not result in clustered CMs unless the anomalies are compared. An untested approach we suggest is as follows.

1. Cluster anomalies by their features:
   a. For waves, match by cavity wall size first (distance between peaks or troughs), then severity.
   b. For slope, by slope value.
   c. For other impairments, by subcarrier frequency or data ranges.
2. For each cluster, add the net severity scores of the MAC addresses in the clusters; call this the cluster severity.

Several methods may be useful to consider here. For waves, normalize the slope to 0, then take an IFFT to find the frequency values most prominent, and finally correlate the transformed data. Likewise, correlate the untransformed data. Methods to correlate these data include k-means, difference tests like paired-t and Chi-Squared, and some statistical parameter tests such as moment tests.

There are several challenges with comparing anomalies to cluster work, and we leave this work for future consideration. We may start with investigating the IFFT approach with checking for correlation of the transformed and untransformed tests for waves, and use difference tests for frequency bound anomalies (matched by anomaly type or not).

3.3. What to monitor

Several statistics may be worth monitoring including the following.

- Table of CMs by severity, ordered by severity
- Table of CMs by hours in the top worst (configurable) by net severity
- Plot of clusters by number of CMs in the cluster and severity
- Plot of the top worst (configurable) CMs or clusters by severity, over time
- Overall node health score by severity total

3.4. Training data for anomaly detection

Training on new data of the AD needs a human supervisor to be involved to label samples or make corrections. Labeling samples means manually adding labels to the new samples. We can also use the AD to perform detection on the new samples, and the human supervisor makes corrections to the localization and classification results. Even if the prediction seems precise enough, small corrections are still necessary to make the labels even more precise. This work is needed because the feed forward neural networks are deterministic in terms of making predictions on the same input data. If we use the same data with the label that's generated by the AD itself for training, it's only reinforcing the model to what it has learned and not helping the model to learn from the new samples.

Each type of spectrum data will use a separate training process to train a separate model using labeled datasets, because different spectrum data may reflect impairments differently. Also using separate models can help stabilize the model and improve efficiency and accuracy.

3.5. Bit loading

Bit loading is handled differently depending on whether the solution is based on PMA or AD (including non-AD RxMER calculations).

With PMA, the bit loading gain is represented by the J value. PMA calculates profiles, then calculates a relative capacity gain by using dynamic profiles over 256-QAM flat profiles as the J value. The J value can be used to find an overall channel capacity gain in bit load. By calculating a set of modulation profiles, making an adjustment to the data, then calculating a new set of profiles, the differences in total bit loading can be compared. But realize that by looking at profiles and making changes based on those, the gain in bit loading depends on the other CMs on the impacted channels, so making a change to a CM or one or more channels impacts multiple CMs, and potentially more than one profile, so an overall impact has to be considered.

With AD, anomalies are first detected by the AD. As we know their type and frequency range from the AD, we first calculate the average bit loading from the average MER value from all subcarriers, then calculate the difference between subcarrier (within the anomaly range) bit loading values and the average bit loading. So, for each anomaly we have a total bit loading loss calculated. These scores can be weighted by anomaly types and added together. This is how we calculate the bit loading in these approaches using AD; RxMER-based bit loading calculations are similar.

3.6. Measuring performance

Because these methods have a few very different methods for measuring severity, it is difficult to compare the performance of some of these approaches. Likewise, combining these methods is possible but a weighting scheme for that might need to be developed. Further, there really are two ways to think about the operational performance of these approaches: the alignment of severity to actual problems in service and plant, and the alignment of severity to financially responsible repair work. The methods informing these severity measures are separately comparable by their ability to collect valuable information to support manual troubleshooting and repair. In addition, the calculation time required (performance) of the methods is a consideration too.

For the methods we encoded, we can provide computational performance comparisons. But without field studies to compare effectiveness, we can't provide any useful guidance on the performance of one approach over another in terms of the effectiveness of the work performed as a result. We can, however, provide a comparison of the information that each approach provides to the technicians, and our risk assessment of which methods should most closely align with effective work. We leave operational performance comparisons to the operators, and we look forward to working with any operator willing to work with us to optimize these methods.

3.7. Comparison of measurements

We applied several of our methods to a large data set of RxMER per subcarrier values and calculated the output measurements of each. Where we could apply different measurements to the same CM, we did so; the results are plotted in the figures that follow.

In these figures, when referring to bit load (labeled Bitload) and RxMER (labeled as MER), severity score and normal score are different ways to calculate similar statistics, and thus potentially used in the same way.

- **Severity score** is the sum of separate anomaly scores (calculated against the sample average in this case), so these two severity scores are from the AD involved methods.
- Alternately, the **normal score** calculation doesn't include AD. It only calculates the "area" below the sliding median across the sample. These normal scores are therefore a measure of variability. We can use them as a severity measure like severity score.

However, both severity and normal scores are from the perspective of the single CM.

Alternately, the channel score is the sum of all CM scores on the same channel; the optimal total capacity of the channel with a certain number of profiles and other constraints can be estimated using profiles generated by PMA, so is the important measure for PMA consideration. In the figures below, we refer to channel score as Severity(PMA).

Note that Severity(Bitload) would be equivalent to Severity(PMA) if there were unlimited profiles (or as many or more profiles as CMs, so that each CM can have its perfectly matched profile).

See Figure 1 for the relationship between Severity(MER) and NormalScore(MER). Note all scores are positive, and we see a cone shape relationship. While the two approaches correlate, this graph does suggest that one measure will prioritize work differently than the other. The top few for one measure are different than the top few for the other.

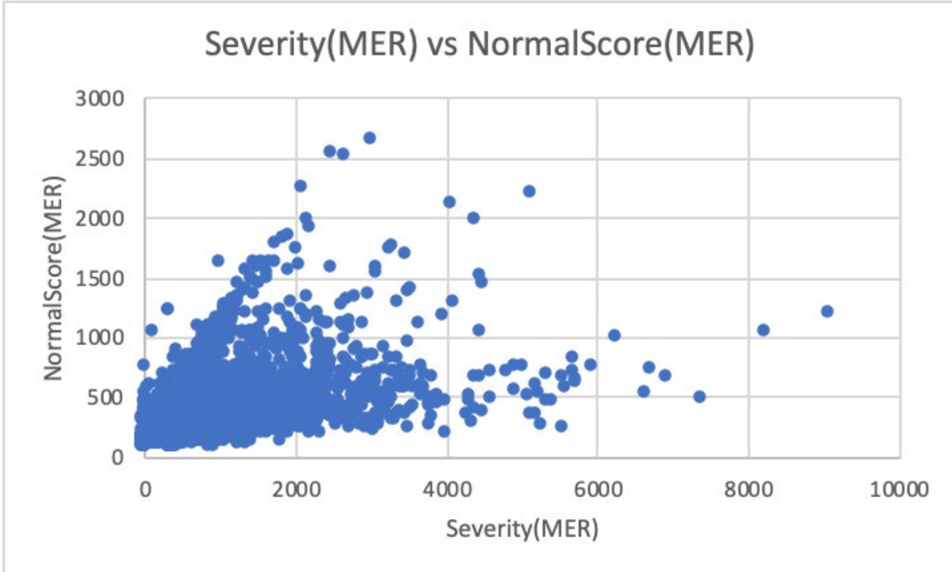

Figure 1 - Severity versus normal score from RxMER

Figure 2 shows the same comparison between severity and normal score, but for bit loading. Again, all measurements are positive. But there are several observations that score differently on one measure with little change in the other. Again, taking the top opportunities with one measure and the top for the other measure may not have any overlap, or very few. The two approaches will suggest different work to consider.

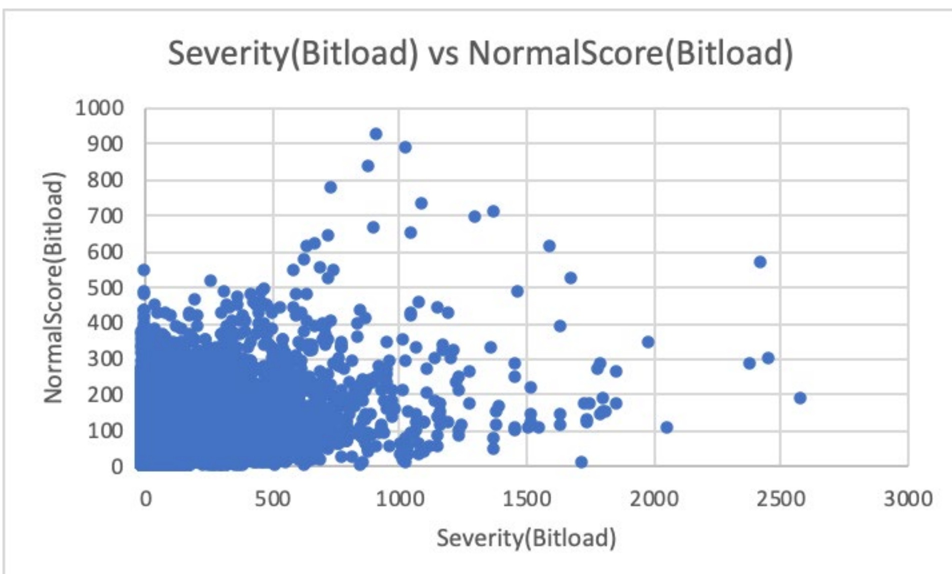

Figure 2 - Severity versus normal score from bit load

Figure 3 compares the severity measures of MER and bit load. Notice a high degree of correlation. This graph suggests that using either measurement is sufficient. And due to the way that bit loading is calculated, this relationship is expected.

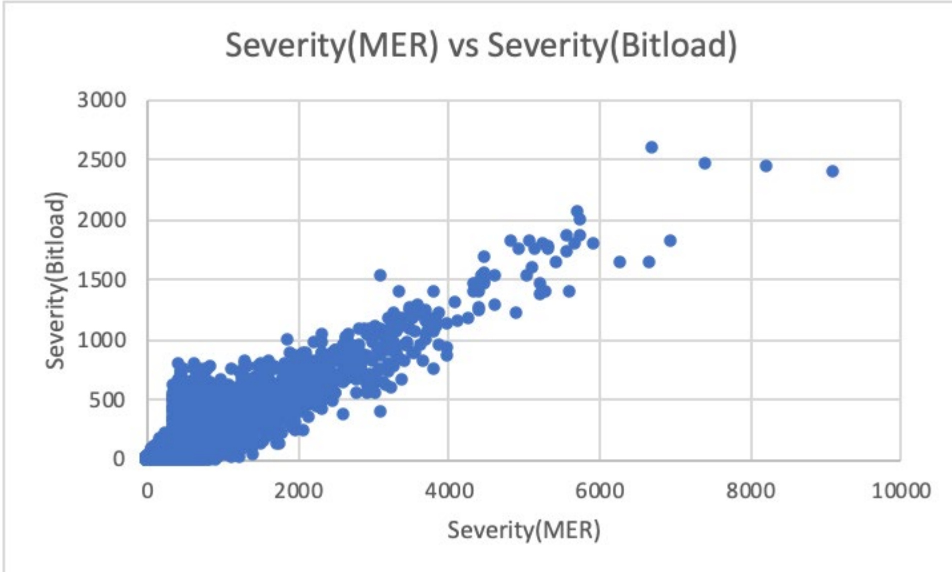

Figure 3 - Severity from RxMER value versus severity from bit loading value

With profile management added to consideration, we can calculate the severity of the PMA measurement and compare it to the other severity measurements. Figure 4 shows the MER severity compared to the PMA severity. As expected, there are clusters where addressing the problem may have negative impact on the overall profile, and those do have smaller MER severity. Certainly, if removing an impairment will result in a net overall lower performance from PMA, that might not be work you want to prioritize (there are other conditions that may change that decision, and there may be different approaches to PMA that avoid the conflict, as we discuss later).

However, note that at least for this set of data, those with negative impact after considering profiles are less severe. Therefore, if we use RxMER per subcarrier to prioritize PNM work, the risk may be low of having a negative or small impact after considering profiles. It may be sufficient just to check the profile impact when in doubt, if using RxMER per subcarrier as the basis for selecting proactive work.

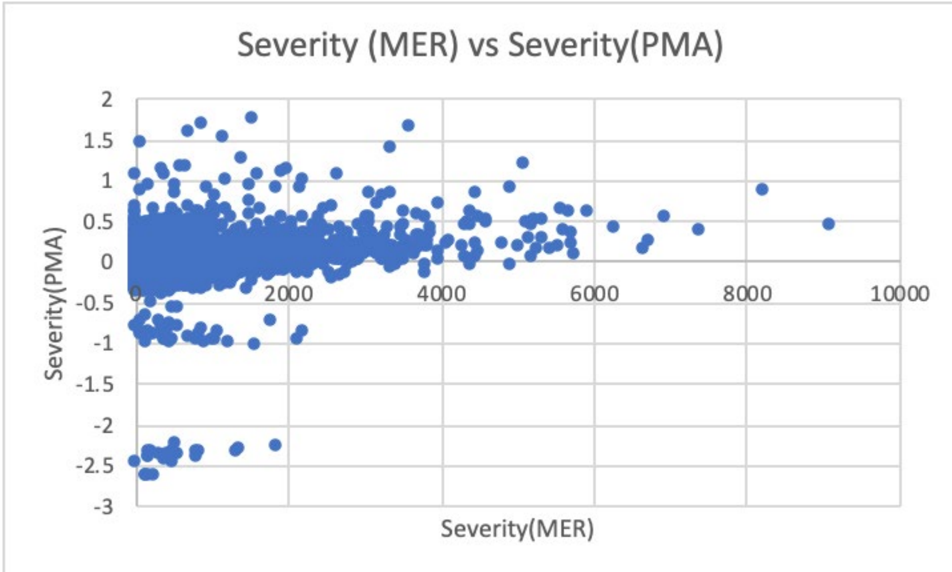

Figure 4 - Severity from RxMER value versus severity from PMA value

Figure 5 compares severity of bit loading and PMA. Like Figure 4, we see the same relationship and the conclusions are the same, as expected.

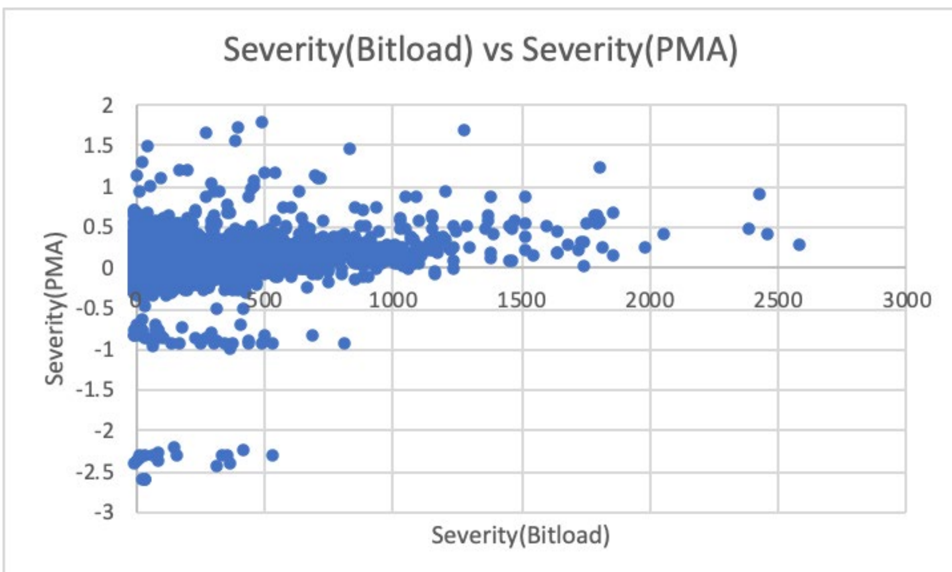

Figure 5 - Severity from bit loading versus severity from PMA value

Figure 6 relates the normal scores of MER and bit loading. Like Figure 3, we see a strong correlation. But this time there is a large number of high normal scores for bit loading that have small MER normal scores. This suggests that, without the ability to find impairments in an automated fashion, there is good reason to look at both measurements and perhaps look deeper to determine which is best. But this is also another reason for using an AD.

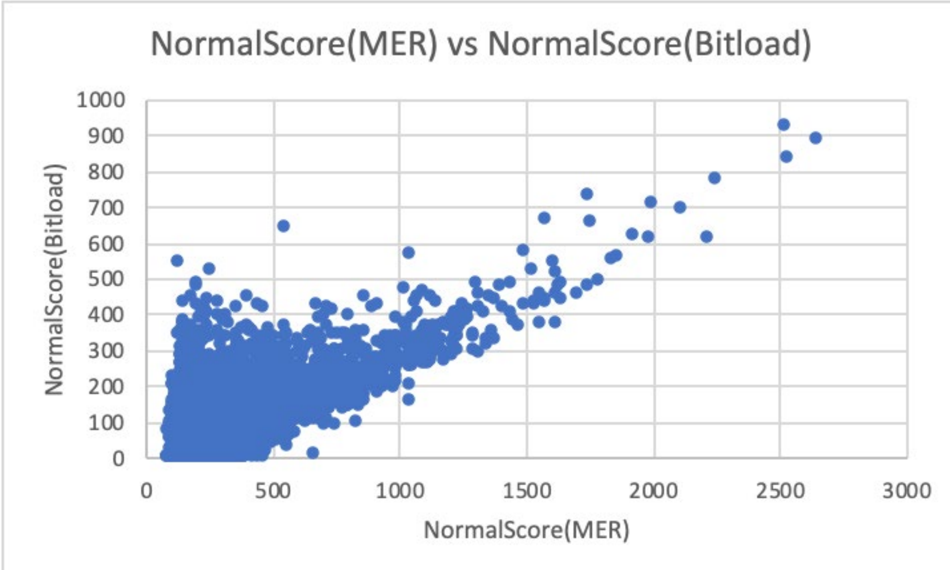

Figure 6 - Normal score from RxMER value versus normal score from bit loading value

Figure 4 and Figure 5 suggest looking at severity for MER and bit loading for only positive PMA severity. Figure 7 shows the conditional relationship, and it strongly resembles Figure 3. This result suggests that the dynamics of PMA cannot be well predicted by just looking at improvements in MER or bit loading, so some consideration of profile impact is warranted for selecting PNM work.

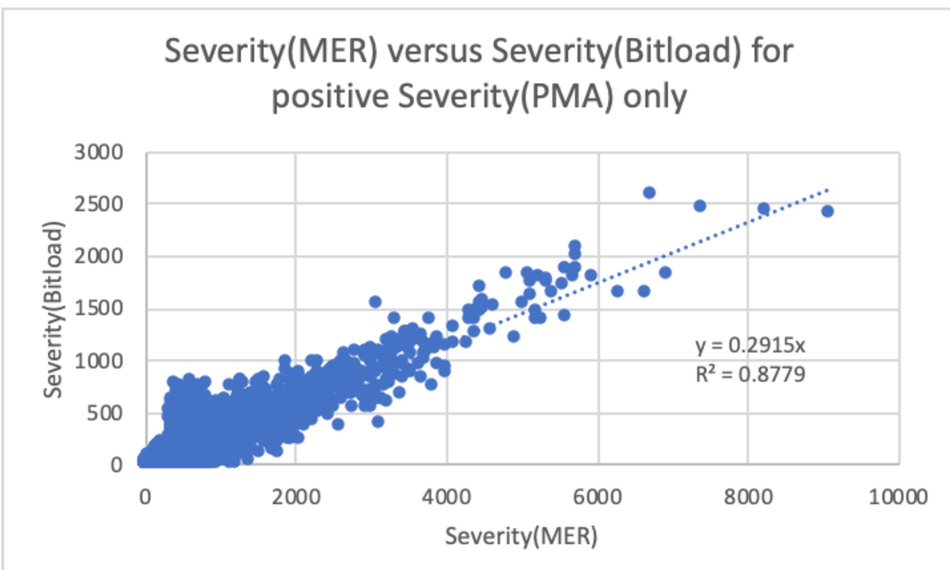

Figure 7 - Severity for RxMER value versus bit loading for CM that score a positive PMA severity

Next, we want to look at individual anomaly types to see if some types of anomalies exhibit different relationships. If we find that some anomaly types appear to score generally higher, or have different relationships with profiles, then we may want to treat them differently. But as we see in the figures that follow, there seems to be no difference in these relationships by anomaly type.

Figure 8 through Figure 13 are repeats of Figure 1 through Figure 6 respectively, but for only LTE anomalies. Note the patterns are respectively very similar, leading to the same conclusions for just LTE anomalies.

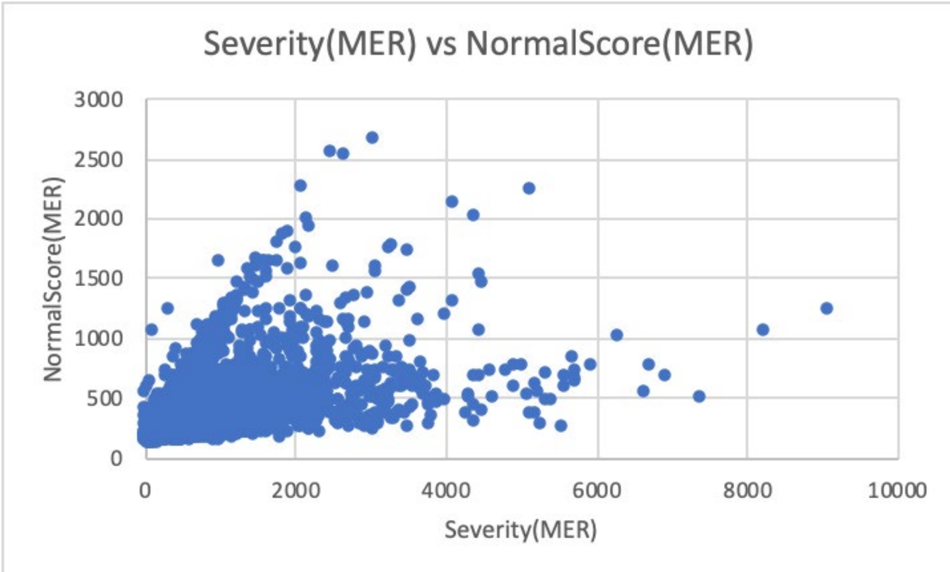

Figure 8 - Severity versus normal score from RxMER for LTE ingress impairments

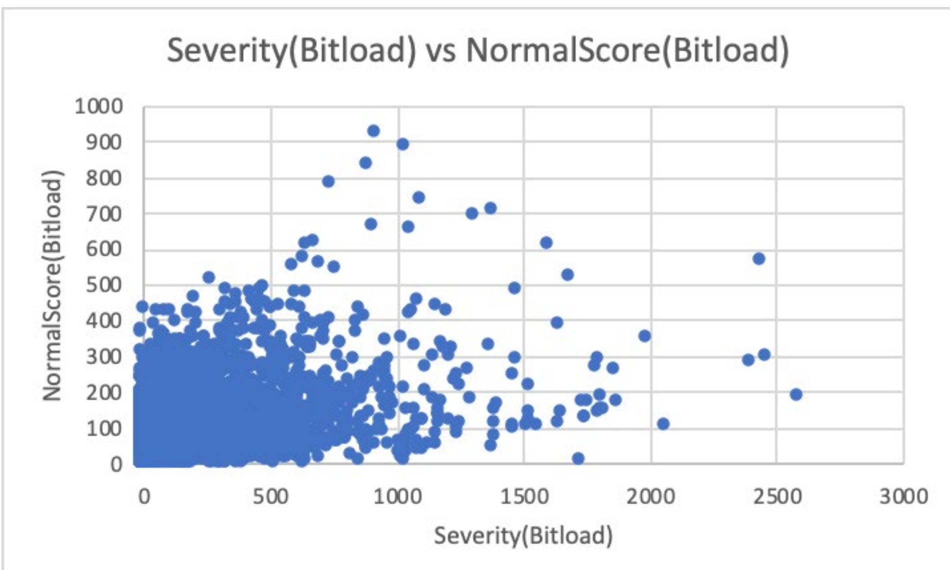

Figure 9 - Severity versus normal score from bit load for LTE ingress impairments

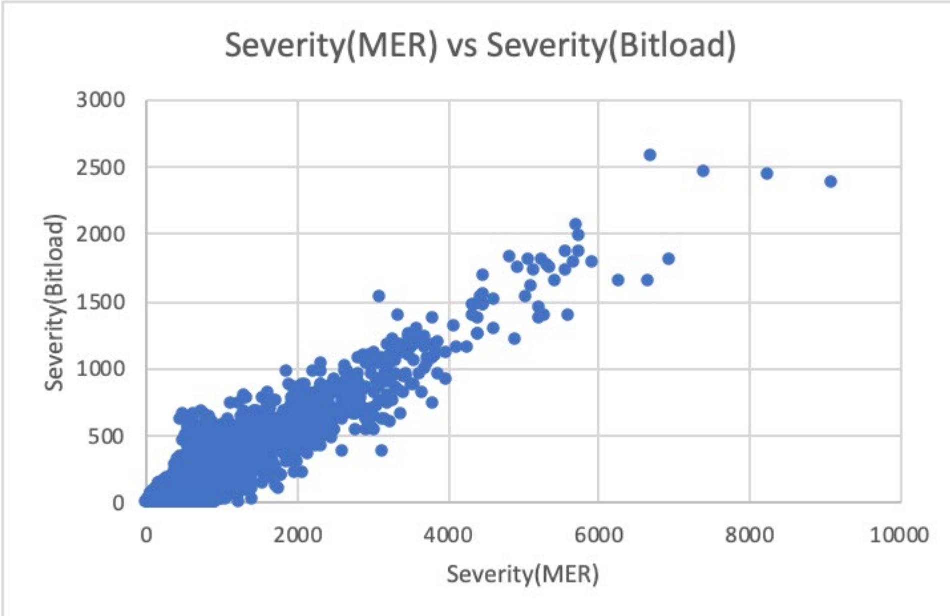

Figure 10 - Severity from RxMER value versus severity from bit loading value for LTE ingress impairments

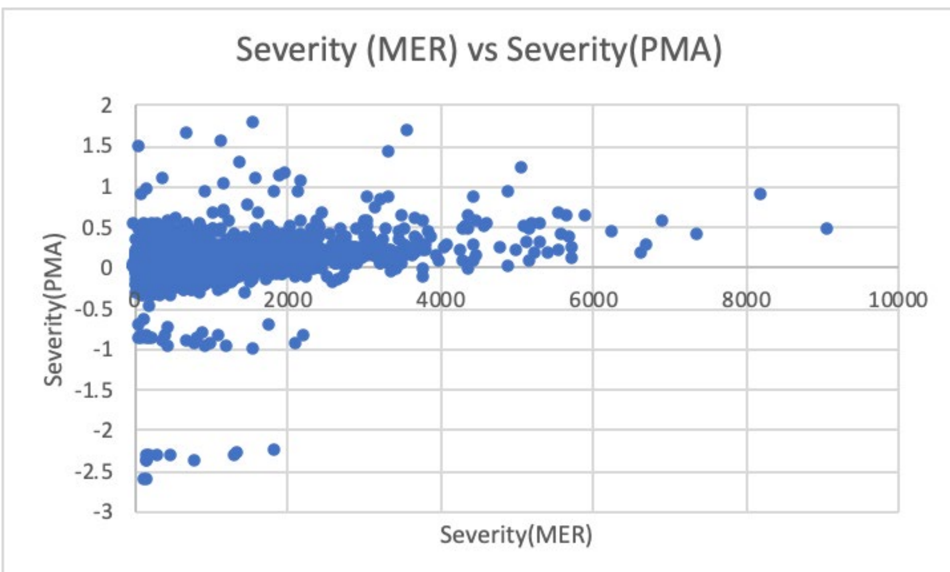

Figure 11 - Severity from RxMER value versus severity from PMA value for LTE ingress impairments

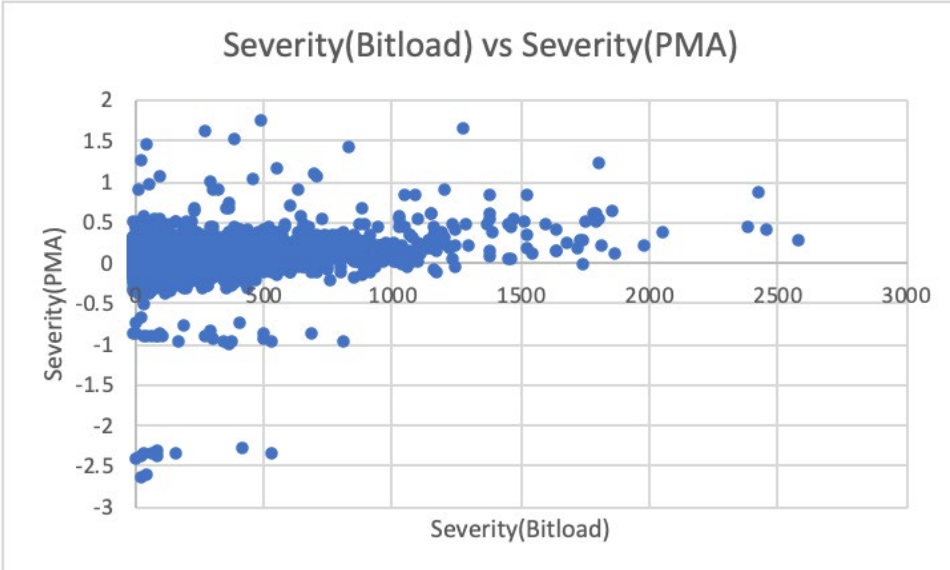

Figure 12 - Severity from bit loading versus severity from PMA value for LTE ingress impairments

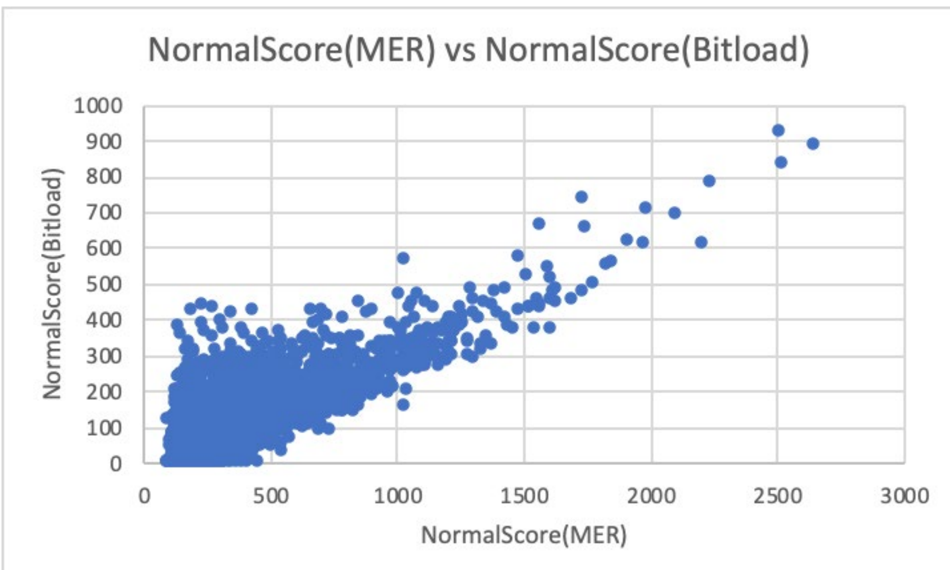

Figure 13 - Normal score from RxMER value versus normal score from bit loading value for LTE ingress impairments

With five separate impairment types to identify, and several modems showing multiple impairment types, showing all graph combinations would fill several pages with little value. Therefore, we are showing a few sample graphs for other single impairments, and for those modems with multiple impairments in Figure 14 through Figure 18, which follow.

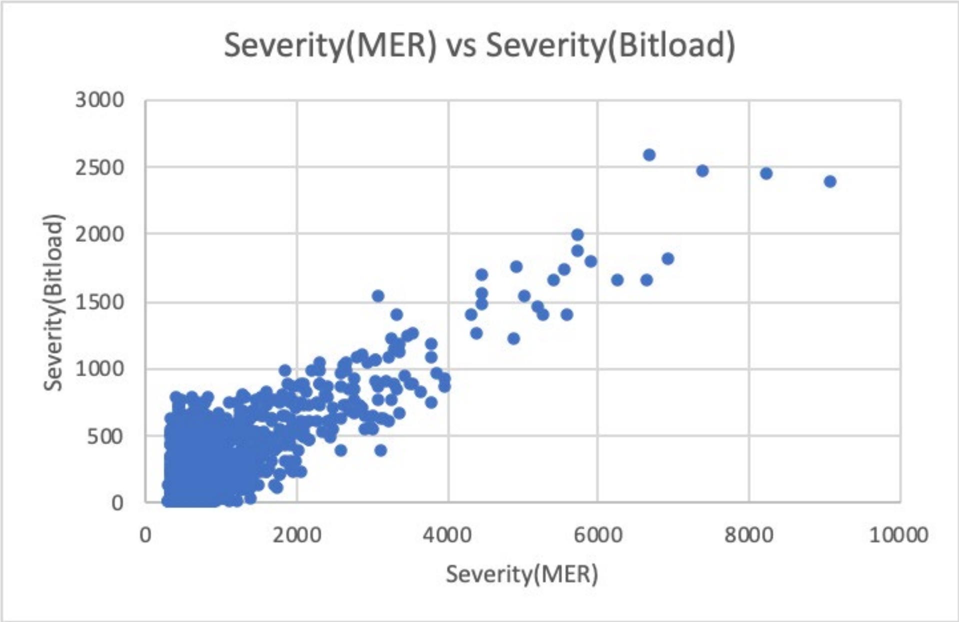

Figure 14 - Severity from RxMER value versus severity from bit loading value for wave impairments

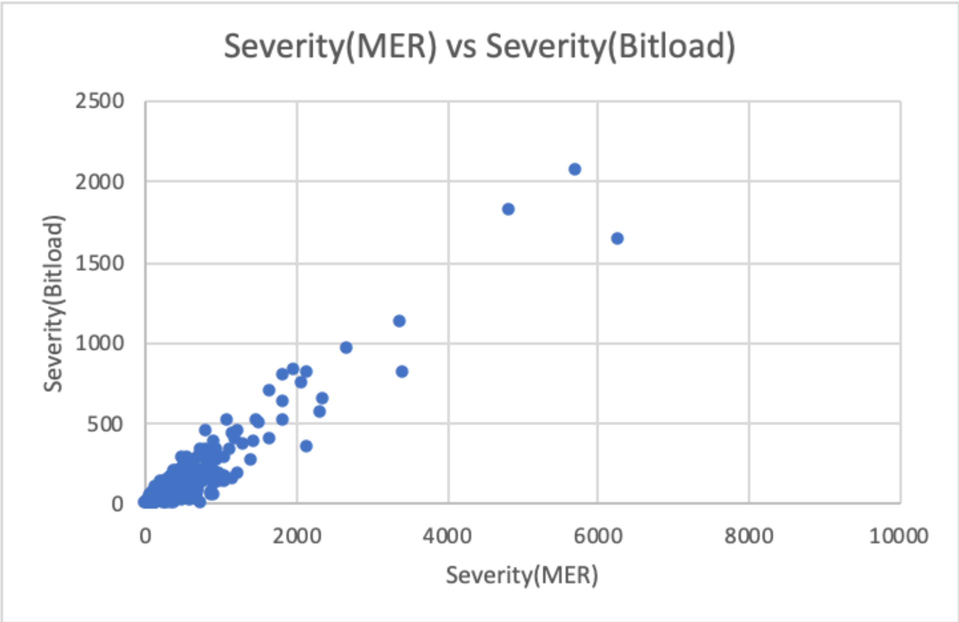

Figure 15 - Severity from RxMER value versus severity from bit loading value for rolloff impairments

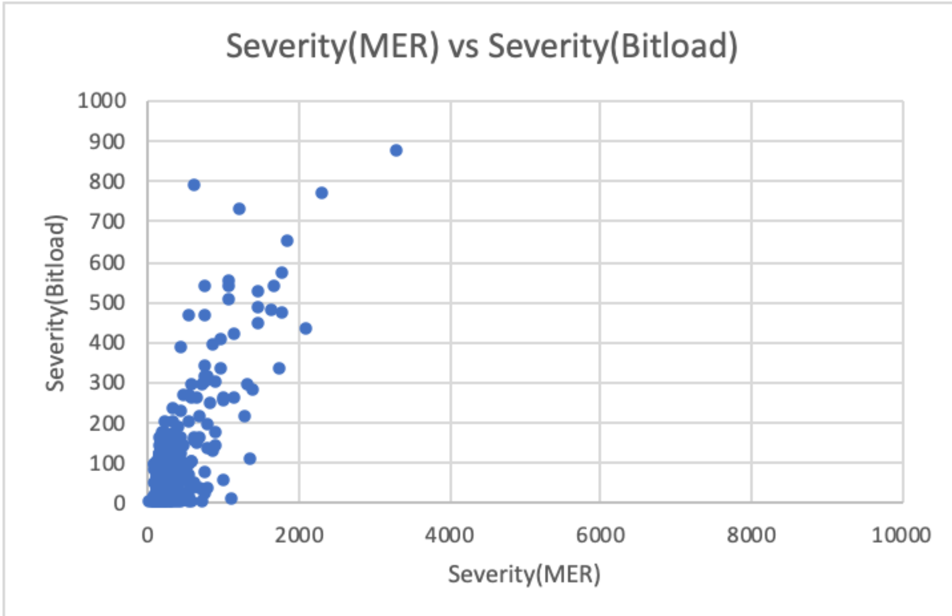

Figure 16 - Severity from RxMER value versus severity from bit loading value for suckout impairments

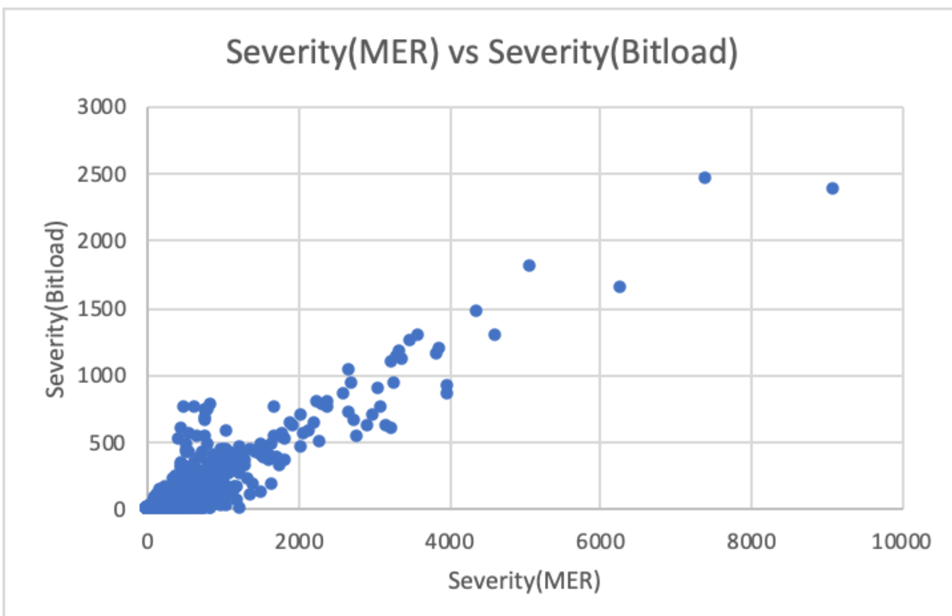

Figure 17 - Severity from RxMER value versus severity from bit loading value for spike impairments

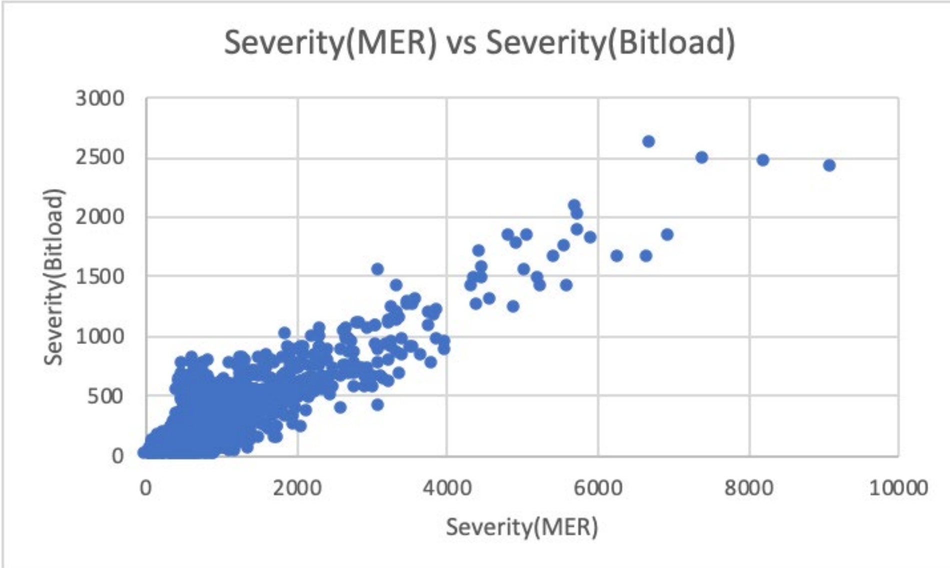

Figure 18 - Severity from RxMER value versus severity from bit loading value for modems reporting multiple impairment types

These plots suggest a few important results which inform maintenance decisions.

A reliable method for identifying anomalies in an automated way is important. Without analyzing the tradeoffs between statistical methods and machine learning techniques for this problem, we only suggest operators use anomalies to identify the opportunities to address because they are indications of cable plant problems.

Further, we suggest using either RxMER per subcarrier or bit loading as a measurement basis, with a measurement of the difference in either as a way to prioritize work. Equally, a plant quality measurement could be used if it is a comparison between the actual performance and intended performance.

Profile management must be considered when selecting proactive maintenance work. Looking at the provided figures, it is conceivable that only the most impactful impairments are addressed, and those appear to have positive impact on overall profiles. But not all PNM activities will yield positive net profile impact, so some consideration should be given here. We explore this issue a bit more next.

3.8. Manual simulation of proactive maintenance impacts

As an exercise to further demonstrate these relationships, we next conduct a manual simulation of maintenance performed on a set of modems. We select a sample of modem data from the larger data set which generated the figures previously discussed. We then select the most significant impairment on a measure of performance, remove the identified impairment from the RxMER per subcarrier data of all modems where it appears by a mathematical adjustment, recalculate profiles and measurements, then repeat a few times. We then report the results of this experiment in a few forms.

While many of the methods we outline in this paper are CM or CM cluster centric, this approach demonstrates using two of the methods as a measurement but in an impairment centric maintenance approach, which we believe is an operationally effective way to perform PNM. The impact on profiles however is examined here too, as we can already see it should be considered for an overall network service impact perspective.

In the comparison done here, the method of section 2.2 is augmented to be an anomaly driven version of using Severity(Profile) as the measure, and this is compared to the method of section 2.4 where we prioritize by bit load impact from correcting the impairment (Severity(Bitload)). Therefore, looking at the difference should show something about the impact of profiles on the simulations.

We make one large assumption in this analysis worth mentioning: we are identifying anomalies in multiple CMs and correlating anomalies so that we presumably remove one anomaly from multiple CMs where it appears. Therefore, the assumption is that the anomaly is indeed one anomaly appearing in multiple CMs, whereas it may not be the same anomaly but rather more than one that coincidentally appear the same in the RxMER per subcarrier data. While the strength of this assumption is related to our ability to accurately cluster by anomaly, the use is comparative in this simulation so should not bias the results greatly. The intent is to see the impact of profiles on PNM selection, and it should still show that impact.

For this simulation comparison, we selected a set of 526 CMs on the same section of plant. Figure 19 shows the RxMER per subcarrier data of all these CMs overlaid on the same plot. Notice related impairments are indicated.

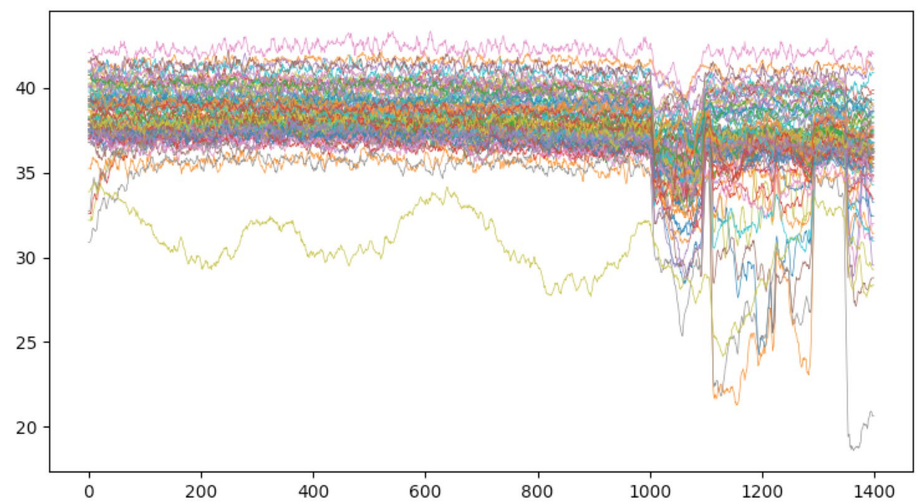

Figure 19 - 526 CMs on a common channel, many experiencing impairments that appear related

The simulation results are shown in Figure 20 and Figure 21, as well as Table 1 and Table 2. Figure 20 is a graph of the total bit load for the set of CMs at each iteration, where an iteration is the mathematical removal of an impairment from the RxMER values, with the impairment selected by the two compared methods: Severity(Bitload) and Severity(Profile). Note that for the first few impairments, there is little difference in the methods. In fact, seeing the detailed results in the two tables, the first three impairments selected by the two methods are the same. After about seven iterations, there is little difference in the two approaches. These results suggest that either approach may be sufficient for selecting PNM work. But more data and simulations would be needed to prove this hypothesis.

Looking at Figure 21, we see that the profile management results, as indicated by J value, vary more. Thus, at least for the selected PMA used here, there is impact on PNM selection and the impact of PNM on service. Of course, this impact is not likely detectable by customers or even applications in many cases, but it may be in extreme cases.

This result suggests that the profile management approach implemented should be made robust to PNM activities. Fortunately, this is easy: if each CM's RxMER is the same or improved after PNM, and a recalculation of profiles yields a solution that is further from the objective, then keep the existing profile solution set, and consider recalculating until a better solution is found. As many PMA are heuristics, you can always change the order of inputs and get a different answer. As profile calculation is fast, and profiles are often kept for long periods of time, there is little reason not to calculate several solutions and select the best according to operations and business objectives.

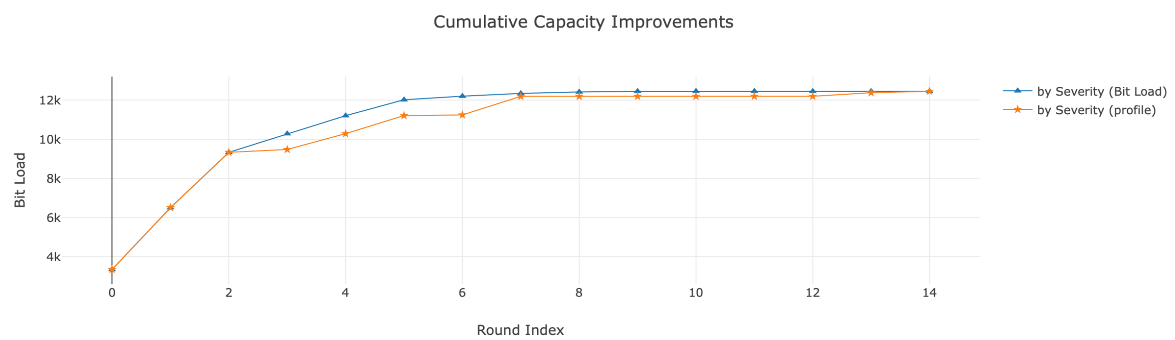

Figure 20 - Cumulative bit load improvements over multiple rounds of impairment removal for both methods

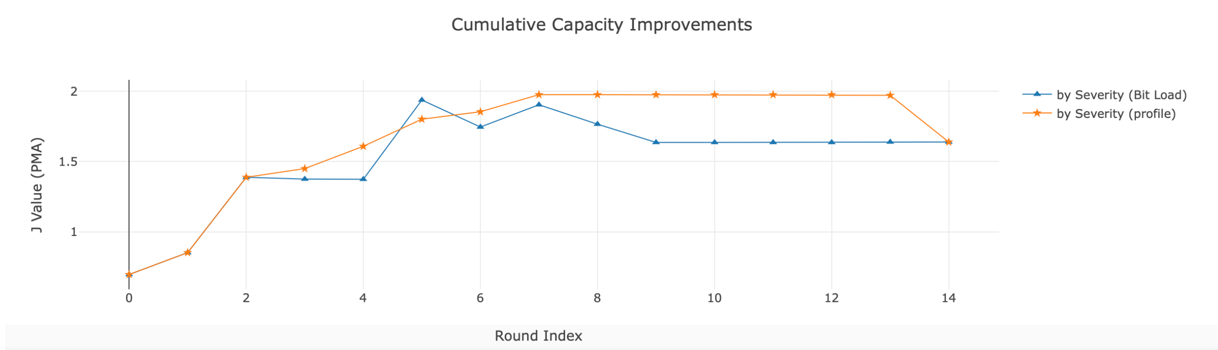

Figure 21 - Cumulative J value (PMA) improvements over multiple rounds of impairment removal for both methods

Table 1- Results prioritized by bit load (Severity(Bitload))

Round	Anomaly	Location (Subcarrier ID)	Num CMs	Severity (Profile)	Severity (Bit load)	Num CMs (Better Profile)	Num CMs (Worse Profile)
0	LTE	start: 1128, end: 1281	39/526	0.696	3356	352	174
1	LTE	start: 1010, end: 1097	90/526	0.156	3157	193	333
2	WAVE	start: 0, end: 1399	40/526	0.536	2828	178	323
3	LTE	start: 1346, end: 1376	23/526	-0.013	951	229	272
4	ROLLOFF	start: 0, end: 138	9/526	-0.001	923	357	163
5	SUCKOUT	start: 992, end: 1114	27/526	0.564	814	433	77
6	ROLLOFF	start: 1330, end: 1399	3/526	-0.192	182	220	109
7	SUCKOUT	start: 1082, end: 1318	8/526	0.158	142	90	219
8	LTE	start: 1001, end: 1028	31/526	-0.138	75	241	127
9	LTE	start: 1259, end: 1288	2/526	-0.130	33	98	78
10	SPIKE	start: 16, end: 23	2/526	0.000	2	0	0
11	SPIKE	start: 1348, end: 1354	2/526	0.000	0	0	0
12	SPIKE	start: 1108, end: 1114	3/526	0.000	0	0	0
13	SPIKE	start: 1227, end: 1234	2/526	0.003	0	2	51
14	SPIKE	start: 712, end: 718	2/526	0.000	0	0	0

Table 2 - Results prioritized by J value (equivalent to Severity(PMA).

Round	Anomaly	Location (Subcarrier ID)	Num CMs	Severity (Profile)	Severity (Bit load)	Num CMs (Better Profile)	Num CMs (Worse Profile)
0	LTE	start: 1128, end: 1281	39/526	0.696	3356	352	174
1	LTE	start: 1010, end: 1097	90/526	0.156	3157	193	333
2	WAVE	start: 0, end: 1399	40/526	0.536	2828	178	323

Round	Anomaly	Location (Subcarrier ID)	Num CMs	Severity (Profile)	Severity (Bit load)	Num CMs (Better Profile)	Num CMs (Worse Profile)
3	SUCKOUT	start: 1082, end: 1318	8/526	0.062	146	79	250
4	SUCKOUT	start: 992, end: 1114	27/526	0.159	810	335	178
5	ROLLOFF	start: 0, end: 138	9/526	0.193	923	128	125
6	LTE	start: 1259, end: 1288	2/526	0.053	33	106	23
7	LTE	start: 1346, end: 1376	23/526	0.122	951	47	305
8	SPIKE	start: 712, end: 718	2/526	0.000	0	0	0
9	SPIKE	start: 16, end: 23	2/526	0.000	2	0	0
10	SPIKE	start: 1227, end: 1234	2/526	0.000	0	0	0
11	SPIKE	start: 1108, end: 1114	3/526	0.000	0	0	0
12	SPIKE	start: 1348, end: 1354	2/526	0.000	0	0	0
13	ROLLOFF	start: 1330, end: 1399	3/526	-0.005	182	147	40
14	LTE	start: 1001, end: 1028	31/526	-0.332	75	160	154

4. Conclusion and Future Work

Profile management has a clear influence on the impact to service that results from PNM. Therefore, some consideration of profiles is important when selecting proactive work to do. It is advisable that operators at least look at the impact on profiles when determining whether to conduct PNM, and it is likely a factor in prioritizing proactive work.

Anomaly-driven PNM, finding and eliminating anomalies, appears to be the most popular approach that operators take for their PNM; but modem-driven PNM is an option to consider for some, as the data that most informs PNM comes from the CMs. Likewise, profile-driven PNM is an option to consider as well. But ultimately, each of these factors needs to be considered in some way in deciding which PNM work to take on.

Though this paper does not compare all options, an examination of the options presented here does suggest that impairment-driven identification of maintenance opportunities is a very good approach, that problem severity can be measured by impact on CMs, and that it is necessary and complimentary to PNM decision making to consider the impact on profiles when determining the capacity impact on the network for selecting maintenance. Said differently—find anomalies, investigate their impact on customers' service, estimate the impact due to profiles and on network impact expected from the work, and use these pieces of information to prioritize and select PNM opportunities. It may at times be best to let profiles drive PNM maintenance decisions, in fact.

Seeing that some network repairs can have negative impact on profiles, it suggests more work on profiles is warranted too. It may be simple enough to set a rule that says, if all CMs improve, then the recalculation of profiles can only result in a net benefit, or a benefit to each CM, or another appropriate rule, before consideration of replacing profiles. In other words, do not harm first. Alternately, it may be useful to develop algorithms that set profiles to be robust against these changes. Likewise, it may be worth investigating different objectives when setting profiles altogether.

Further computer simulations of various PNM selection criteria are warranted by this work. The two limited simulations we ran suggest that very significant PNM opportunities can have significant impact on network capacity, whereas smaller ones may not. Further, there are several measures that can be used to select PNM opportunities, and they may all be sufficient for many operations, though that should be checked with further simulation and field experiments. The limited simulations we ran also suggest that profiles recalculated after plant improvements may not always improve capacity for everyone or overall, but that can be easily addressed in software or in operations decisions.

At least, consider this work as a roadmap for analyzing PNM strategies for operations. The graphical comparisons are simple to do with your own data, and running simulations of various approaches can be very informative. Extending this work into cost analysis may be the necessary step to procure funding for operational improvements.

Our best hope is to work with operators to test some of these approaches and determine which are the most cost effective to operations and impactful to service overall. Selecting proactive work requires identifying the opportunity, measuring it, estimating the benefit of addressing it, and then executing on those opportunities with the most promise. Therefore, the best measure of success is to determine the cost and benefit of the work being done, then adjust the opportunity selection process to optimize the program on those measures. The field work that has to be done is to determine which of the approaches outlined in this paper, or alternative approaches, are best suited to optimize the maintenance program. Looking at the conclusions that follow, we already have a good idea where to start. Fortunately, for well-equipped or unequipped operators alike, CableLabs has ProOps encoded and ready to assist.

Certainly, there remains a lot of confirmation work with operators before all this can be applied optimally to operations. While identifying the anomalies and measuring their significance to service are important steps, the PNM work to remove the anomalies has to be done to measure cost and impact, so that both can be modeled better for more accurate and operator-specific decision making.

Abbreviations

AD	anomaly detection
CM	cable modem
CMTS	cable modem termination system
dB	decibels
DOCSIS	data over cable service interface specification
IFFT	inverse fast Fourier transform
ISBE	International Society of Broadband Experts

LMS	least mean square
LTE	long term evolution
MER	modulation error ratio
OODA	observe, orient, decide, act
PMA	profile management application
PNM	proactive network maintenance
ProOps	proactive operations (platform)
QAM	quadrature amplitude modulation
RxMER	receive modulation error ratio
SCTE	Society of Cable Telecommunications Engineers

Bibliography & References

[1] Cole, T.J., "The LMS method for constructing normalized growth standards," European Journal of Clinical Nutrition 44(1): 45-60, 1990

[2] Jingjie Zhu, Karthik Sundaresan, Jason Rupe, "Proactive network maintenance using fast, accurate anomaly localization and classification in 1-D data series," proceedings of the International Conference on Prognostics and Health Management (ICPHM), June 2020, 978-1-7281-6286-7.

[3] Ron Hranac, James Medlock, Bruce Currivan, Roger Fish, Tom Kolze, Jason Rupe, Tom Williams, Larry Wolcott, "Characterizing network problems using DOCSIS(R) 3.1 OFDM RxMER per subcarrier data," 2019 SCTE-ISBE and NCTA.

[4] Greg White, Karthik Sundaresan, "DOCSIS 3.1 Profile Management Application and Algorithms," 2016 SCTE-ISBE and NCTA.

[5] Jason Rupe, Jingjie Zhu, "Kickstarting proactive network maintenance with the proactive operations platform and example application," 2019 SCTE-ISBE and NCTA.

[6] Jason Rupe, "A general-purpose operations cost model to support proactive network maintenance and more," 2019 SCTE-ISBE and NCTA.