CableMon: Improving the Reliability of Cable Broadband Networks via Proactive Network Maintenance

Jiyao Hu, Zhenyu Zhou, and Xiaowei Yang
jyhu@cs.duke.edu, zzy@cs.duke.edu, xwy@cs.duke.edu
Department of Computer Science, Duke University

1 Introduction

Broadband access networks play a crucial role in modern life. They help narrow digital divide, enable e-commerce, and provide opportunities for remote work, study, and entertainment.

- However, much measurement study has shown that broadband networks have poor reliability [2].
- A recent study shows that the average availability of broadband Internet access is at most two nines (99%), much less than the minimum FCC's requirement (four nines 99.99%) for the public switched telephone network (PSTN).
- The cable industry has long recognized this problem and developed a platform called Proactive Network Maintenance (PNM) to improve the reliability of cable broadband networks [1]. PNM enables a cable ISP to collect a set of performance metrics data from each customer's cable modem.
- However, existing tools in the public domain use a set of PNM metrics and static thresholds to detect network faults, which lead to a prohibitive high false positive rate.

03/13/2019	04/09/19	06/25/19	07/15/19	08/15/19	Eight-month
24.95 %	25.45 %	27.16 %	27.07 %	27.38 %	26.15 %

Table 1: **The percentage of cable modems that need to be repaired if an ISP were to follow one of the CableLabs' recommendations (MTR < 18dB).**

2 Motivation

- The values of PNM data correlate with how frequently customer trouble tickets are created.
- The average number of customer tickets created in a unit time is defined as the ticketing rate.

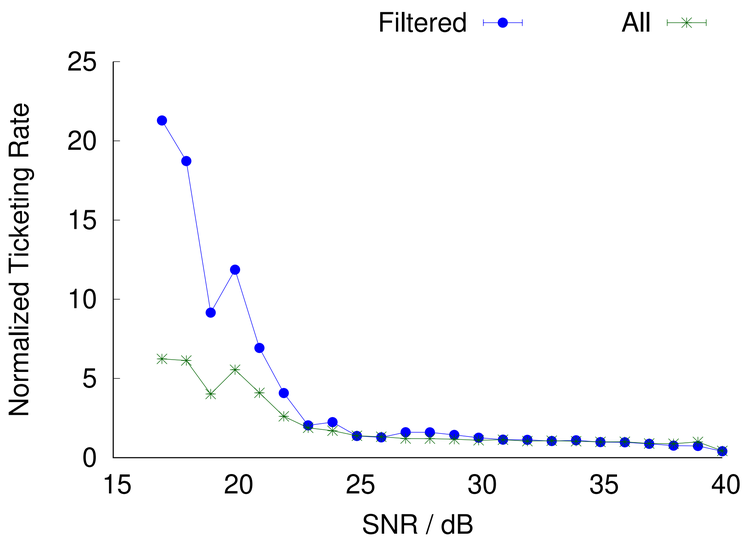

Figure 1: **This figure shows how the customer ticketing rate varies with the values of SNR. Ticketing rate tends to increase when SNR values are low.**

3 Design

Goals

1. *High precision and moderate recall.* Avoiding false positives is practically more important than repairing all faults proactively. Many tickets are not related to cable network problems.
2. *No manual labeling.* We aim to develop CableMon without manually labelled data.
3. *No extensive parameter tuning.* We aim to release CableMon as an off-the-shelf-tool at cable ISPs.
4. *Efficient.* We aim to detect network faults in real-time.

Main idea

Divide the feature value space into an abnormal sub-space and a normal sub-space to maximize the ratio of ticketing rates between the sub-spaces (Ticketing Rate Ratio, TRR).

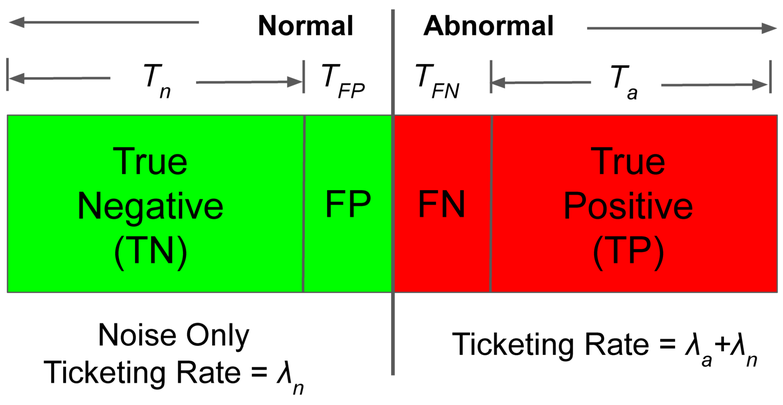

Figure 2: **Analysis of Ticketing Rate Ratio.**

Key Observation

- TRR is monotonous w.r.t. false positives and false negatives.
- TRR is maximized when both false positives and false negatives are zero.

Goal ③ ✓

Challenges

1. *Noisy tickets.* A customer may or may not call when there is a fault. Meanwhile, a customer may make a call when there are no network outages as well.
2. *Noisy PNM data.* An added noise may make a PNM metric take an abnormally low or high instantaneous value.

1. Filter the tickets with their fields to only consider networking tickets. Challenge ① and Goal ② ✓
2. Apply time series models onto the PNM metrics to reduce noise. Challenge ② ✓
3. Train a model to maximize TRR that can detect anomalies in real-time. Goal ④ ✓

4 Evaluation

Establishing Evaluation Metrics

- Use a sliding window algorithm to combine the point-wise detection results into anomaly events.
- Estimate the best window parameters according to TRR, precision, and recall.

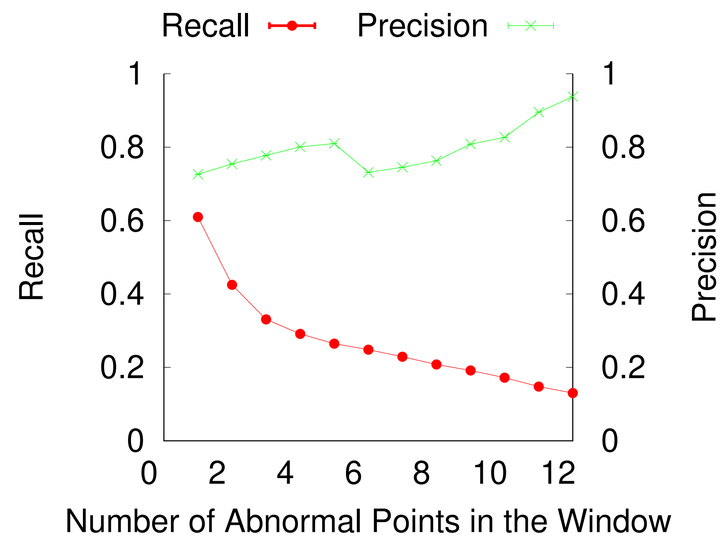

Figure 3: **Precision and Recall of Different Window Parameters.**

- High ticketing rate ratio, together with a relatively high precision ∼ 80%, and a recall around 20%. Goal ① ✓
- Compared to existing tools, CableMon detects more severe and dispatched tickets, which take longer to resolve and shorter for the customer to report.
- Compared to existing tools, the anomaly events detected by CableMon have a moderate length, indicating fewer false positives.

References

[1] DOCSIS CableLabs. Best Practices and Guidelines, PNM Best Practices: HFC Networks (DOCSIS 3.0). Technical report, CM-GL-PNMP-V03-160725, 2016.

[2] Aaron Schulman and Neil Spring. Pingin'in the Rain. In *Proceedings of the ACM SIGCOMM Conference on Internet Measurement Conference*, pages 19–28, 2011.