

Failure Diagnosis with Incomplete Information in Cable Networks

Yun Mao
University of Pennsylvania

Hani Jamjoom
IBM T.J. Watson Research

Shu Tao
IBM T.J. Watson Research

1. INTRODUCTION

Cable network has become one of the most popular ways of high-speed Internet access for homes and small businesses. For broadband cable providers, managing their large-scale network infrastructure is highly challenging because these networks are geographically dispersed and contain a large number of devices. A single administrative area typically serves hundreds of thousands of end-customers (or cable modems) with thousands of intermediate distribution devices that operate on different protocol layers. Such devices include routers, Cable Modem Termination Systems (CMTS's), fiber nodes, repeaters, etc. It is, thus, critical for the providers to monitor the health of their infrastructure and perform quick failure diagnosis. However, there are many challenges in failure diagnosis. In this paper, we focus on two challenges that are motivated by input from a large U.S. cable provider: *missing device status information* and *incomplete topology*.

To highlight the significance of both challenges, consider a simplified cable network in Fig. 1. The first challenge is that it is infeasible for the Network Operation Center (NOC) to constantly monitor the status of all devices, because (1) some devices are passive and do not respond to diagnostic packets or signals,¹ and (2) the cost of probing all devices is too expensive.² Therefore, the NOC has to rely on limited monitoring data obtained from only the accessible devices (e.g., CMTS's and cable modems) to diagnose failures spanning the entire network. The second challenge is that in many cases, a complete network topology is not available. This typically happens because topology information is spread across disparate plant mapping applications, which may not integrate back into monitoring and management applications. As a result, the NOC often needs to deal with an incomplete topology like the one exemplified by Fig. 1, where the exact number of repeaters and their connections to fiber nodes and cable modems are unknown.

With these two challenges in mind, the goal of this paper is to develop an approach that helps identify the root cause of failures, even though failures may be caused by devices that are neither remotely measurable by the NOC, nor visible from the topology standpoint. Previous works on network failure diagnosis have depended on the availability of complete topology (with possibly incorrect) information [2, 1]. In these works, network topology is used to understand the dependency between network devices (e.g., routers) and observed failures (e.g., IP link failures). Here, we study the problem of failure diagnosis with *incomplete information* about device status and network topology. To this end, we adopt the concept of *failure group* (FG)—a group of end components (i.e., cable modems) that are likely to share the same risk of failure. The failure of a FG dictates the failures of all its components. The idea is that if we can infer FGs that directly reflect the missing part of the topology, we can then use FGs to localize failures, diagnose root causes, and even detect misconfigurations in the existing topology. For instance, both a repeater and a fiber node may form FGs that

consist of all their downstream cable modems. If the corresponding FGs can be identified, when a similar group of cable modems are observed in failure, we will be able to distinguish whether the root cause is the repeater or the fiber node. Following this idea, we develop an approach that can infer failure groups by collecting and analyzing historical measurements on cable modems. Our study shows that this approach can accurately infer the failure group association of the monitored nodes, hence provide meaningful assistance to failure diagnosis in cable networks.




Figure 1: A typical cable network topology consisting of routers/switches, CMTS's, fiber nodes, repeaters, and cable modems, which are organized in a tree structure. Here, connectivity between end modems and fiber nodes is missing.

2. APPROACH

We use binary status of cable modems (i.e., 1 if faulty, 0 if not) as the primary information source to infer failure groups. Due to the high overhead of probing, it is uncommon for management applications at NOC to periodically check the status of all cable modems. Instead, active probing to all cable modems is only invoked if a higher-level device (e.g., a CMTS interface) issues an alarm. For instance, an alarm could be triggered if the number of *live* modems registered at a CMTS interface is lower than a threshold.

The above measurement yields a *failure instance matrix* $X = [x_{ij}]_{n \times d}$ for each CMTS interface, where n is the number of measurements initiated for the interface, and d is the number of modems probed in each measurement. x_{ij} represents the status of the j -th modem during the i -th measurement: $x_{ij} = 1$ indicates that the modem is faulty; $x_{ij} = 0$ means the modem is functioning.

Our goal is to identify the FGs and determine their failure status based on the failure instance matrix X . Note that the number of FGs, r , is typically much smaller than the number of modems, d , due to the high-density tree topology of the network. Ideally, the status of FGs and the association of cable modems to different FGs jointly determine X . Specifically, suppose we know *a priori* the compositions of all FG's and their status during the measurement, we can construct two binary matrices, a *failure explanation matrix*, $U = [u_{ij}]_{n \times r}$, and a *failure group matrix*, $V = [v_{ij}]_{r \times d}$. Each row of V represents an FG: $v_{ij} = 1$ iff the j -th modem is associated with the i -th FG. Each column of U represents the status of an FG: $u_{ij} = 1$ iff the j -th FG fails during the i -th round of measurement. Thus, the product of the two binary matrices should equal

¹For example, the repeaters that connect cable modems and fiber nodes operate on the physical layer. Their status is invisible to the NOC.

²For example, it is very costly to probe all cable modems in an administrative area with hundreds of thousands of customers.

the original failure instance matrix X :

$$X = U \times V \quad (1)$$

In practice, however, the failure instance matrix collected by probes are not always accurate. This is caused by the fact that (1) probes are unreliable; (2) monitored cable modems may be powered off when the probes are sent; and (3) not all failure events should be characterized by failure groups (e.g., isolated failure events should be considered as measurement noise). As a result, X might not exactly equal $U \times V$. We need to find the right decomposition of X such that U and V best represent the actual causes of the failures, despite the existence of noise.

To find the right decomposition, we propose an algorithm based on the *Non-negative Matrix Factorization* (NMF) method [3]. Suppose the number of FGs r is given.³ NMF can decompose X into two non-negative real matrices $U'_{n \times r}$ and $V'_{r \times d}$, such that the *derived failure instance matrix*, $X' = [x'_{ij}] = U' \times V'$, is a good approximation of the original failure instance matrix X . i.e., NMF aims at minimizing the reconstruction error function

$$\delta = \|X - X'\|^2 = \sum_i \sum_j (x_{ij} - x'_{ij})^2 \quad (2)$$

Minimizing this error function can be viewed as minimizing the gap between the hypothesis and the reality. After U' and V' are obtained, we further normalize them and apply a threshold-based algorithm to derive the binary matrices U from U' and V from V' .

We have also extended our algorithm to work without given the number of FGs r . The idea is to use an estimated lower bound of r as the starting point, and gradually increase it until the diagnosis result is satisfied according to several heuristics. The details are not included in this paper due to space constraints.

3. EVALUATION

As discussed earlier, the status information collected from cable modems is not always accurate. Therefore, it is important for our approach to be robust against noisy input. We have evaluated our algorithm in both simulation and experimental settings based on real data. Here, we present a set of simulation results to demonstrate its effectiveness.




Figure 2: Simulated topology

In this study, we simulate 200 cable modems in a randomly generated balanced tree topology similar to that of Fig. 2. Both level-1 and level-2 nodes are hidden from our algorithm. Here we show how we can identify level-1 nodes as major FGs. We have also tackled complex multi-level hidden topologies, which is not part of this paper. In each time epoch, the nodes in the topology can fail independently with deterministic failure probabilities. We set the failure probability of level-2 and leaf nodes to 10% as the noise to the experiment. The failure probability of the level-1 nodes is set

³For instance, we may know the number of repeaters in the network, but not the connections between cable modems and these repeaters.

to 3.3%. All failure events are simulated independently. When the ratio of failed leaf nodes exceeds 15%, active probes are launched and the status of all leaf nodes are collected as a measurement sample. We keep the simulation running until 150 measurement samples are collected. The algorithm has no prior knowledge of either the topology or the numbers of nodes on level 1 and 2.

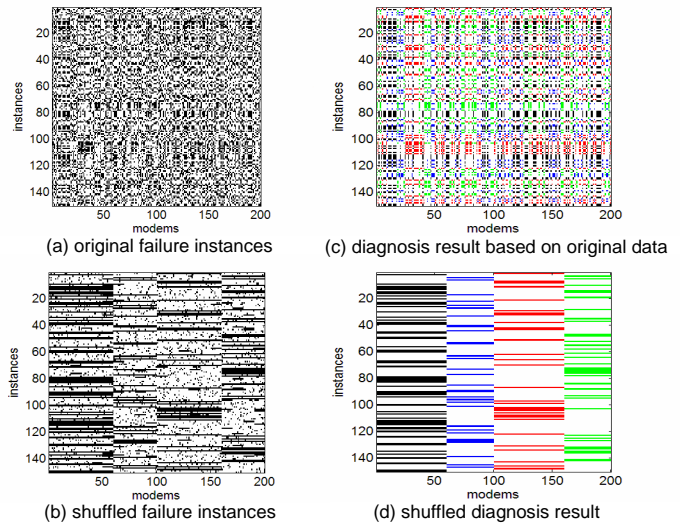


Figure 3: Results of the simulation

Fig. 3(a) shows the simulated failure matrix. A black dot in the figure indicates a failure observed on a modem in a probing instance. Note that the failure groups associated with the 4 level-1 nodes are hardly visible, unless we re-order the columns according to the modems' association with these 4 nodes (see Fig. 3(b)). Using the original failure matrix as input, our algorithm infers 4 failure groups and computes the derived failure matrix (X'). The latter is shown in Fig. 3(c) with each FG marked with a different color. If we reorder the columns in Fig. 3(c) based on the modems' association to the derived FGs as in Fig. 3(d), we can see that the algorithm successfully filters out the noise and accurately identifies the dominant FGs that are represented by the level-1 nodes. Note that although the high level nodes have a lower failure probability than the low-level ones, the high level nodes have a bigger impact on the leaf nodes once they fail. This explains why only the 4 FGs associated with the 4 level-1 nodes are identified here by our algorithm. We have extended our algorithm to iteratively find the more fine-grain FGs. Other simulation and experimental results consistently show that the proposed approach is effective in real network settings.

4. REFERENCES

- [1] S. Kandula, D. Katabi, and J. P. Vasseur. Shrink: A tool for failure diagnosis in ip networks. In *Proceedings of MineNet workshop*, 2005.
- [2] R. R. Kompella, J. Yates, A. Greenberg, and A. C. Snoeren. IP fault localization via risk modeling. In *Proceedings of NSDI*, 2005.
- [3] D. D. Lee and H. S. Seung. Algorithms for non-negative matrix factorization. In *Proceedings of Neural Information Processing Systems (NIPS)*, pages 556–562, 2000.