Aggregating Noisy and Unsynchronized Measurements for Network Tomography

Ítalo Cunha[†‡] Renata Teixeira[♮‡] Nick Feamster[♯] Christophe Diot[†]

† Thomson ‡ UPMC Paris Universitas ♮ CNRS ♯ Georgia Tech

1. INTRODUCTION

Internet paths can experience failures that are not reported or automatically solved by network equipment. These failures are often called “blackholes” and can happen due to wrong router configurations, software bugs, or errors in the interaction between multiple routing layers [3]. Blackholes are of special interest because they can cause persistent packet loss or complete disconnectivity, and they persist until a network operator identifies and fixes the problem. Moreover, because blackholes do not raise any alarms, end-to-end probes are usually the only form of detecting them. As a result, it is not uncommon for network operators to deploy end-to-end monitoring on their networks.

Correlating end-to-end probe failures with the network topology to infer the failure location (e.g., link or router) is a technique referred to as *binary tomography*. In a tomography system, as exemplified in Fig. 1, *monitors* send end-to-end probes along pre-selected paths to determine their status (step 1). These statuses are then reported to a *coordinator* (step 2). The coordinator aggregates measurements from all monitors into a *reachability matrix*. This reachability matrix and the network topology are used as input to a tomography algorithm that builds a hypothesis of failed links that explain the observed status of paths (step 3). For example, if monitor A detects that its path to destination C is down and monitor B detects that its path to destination C is up, the coordinator would infer that link 3 has failed.

Although promising, network tomography systems have not been widely deployed due to practical issues [2]. For example, in our experiments in PlanetLab and in an enterprise network, applying naïve techniques to aggregate measurements into reachability matrices results in a large number of alarms: a tomography system would raise one alarm per minute in PlanetLab and one alarm every three minutes in the enterprise network. This amount of alarms is too large to be useful in practice and the majority of them are probably not real failures (i.e., *false alarms*).

These false alarms are due to errors in the reachability matrices used by tomography algorithms. We identify two sources of errors when building reachability matrices: **(1)** synchronizing the time probes go through a link is impossible and they can experience different network conditions, and **(2)** transient packet losses not related to failures (e.g., due to congestion) can result in *detection errors* (i.e., wrong measurement of the status of paths). To understand how Type 1 errors affect tomography, suppose monitor A probes C while link 2 is down and that monitor B probes C successfully after link 2 has come back up. Type 2 errors affect tomography similarly: suppose in the example above that link 3 is working, the probe from A to C is lost due to congestion, and the probe from B to C succeeds. In both cases, inferring that link 3 has failed is a false alarm.

PAM2009, April 1-3, 2009, Seoul, Korea.

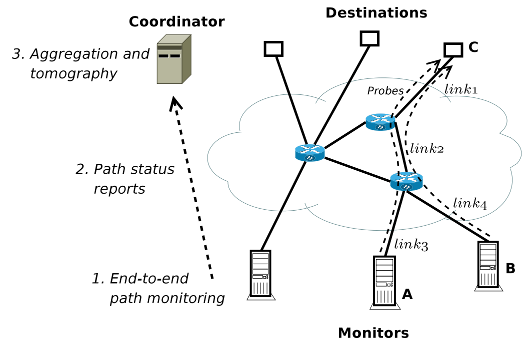

Figure 1: Overview of a tomography system.

In this abstract we analyze three different *aggregation strategies* that combine measurements from different monitors into a reachability matrix. The strategies trade speed for accuracy, waiting for multiple measurements to gather more information and then build a more accurate reachability matrix. Our experiments on PlanetLab and on an enterprise network show that better aggregation strategies can significantly reduce the number of alarms raised by tomography systems.

2. AGGREGATION STRATEGIES

This section describes three strategies by which we combine the set of path measurements into a reachability matrix (step 3 in Fig. 1). An inconsistent reachability matrix could result in false alarms or inaccurate identification of failures. Existing work on network tomography presumes the existence of such a reachability matrix.

Basic aggregation. The BASIC aggregation strategy works as follows. First, detect that a path status changed from up to down. Then, wait for a full *measurement cycle*, i.e., the period it takes to probe all monitored paths once. Finally, build a reachability matrix by combining the statuses reported in the latest cycle. This simple strategy assumes that if a link ℓ fails, all paths that traverse ℓ will be confirmed as down after the measurement cycle.

However, this assumption is not always true. If the failure is shorter than one measurement cycle, synchronization errors (Type 1) can happen: some paths that traverse ℓ may be probed while the link is down and others when ℓ is up. Also, detection errors (Type 2) add wrong information about the status of paths in the reachability matrix. To summarize, the BASIC strategy is not robust to errors of Type 1 and 2.

Coping with lack of synchronization. To address synchronization errors (Type 1), we propose multi-cycle aggregation strategies that improve consistency by using extra time to gather more measurements and build a more consistent reachability matrix. Similar to BASIC, the *multi-cycle* (MC-ALL) strategy starts aggregation upon a path status change. Instead of building the reachability matrix after one measurement cycle, MC-ALL waits for n cycles with identical measurements of all paths to build the reachability matrix. MC-ALL avoids building reachability matrices when measurements are not stable.

MC-ALL solves synchronization issues (Type 1) by building reachability matrices only for failures that last more than n measurement cycles. Failures shorter than that can never result in n identical measurement cycles and are filtered. MC-ALL is also robust to detection errors (Type 2). To be aggregated in a reachability matrix, detection errors need to happen in the same set of paths for n consecutive cycles. However, MC-ALL may be too conservative. If there are frequent detection errors, aggregation is delayed until n consecutive matrices are identical. This additional delay impacts the number of failures that MC-ALL can build a matrix for. Ultimately, this compromises its usefulness if detection errors are frequent.

Coping with detection errors. We propose an extension to MC-ALL called *multi-cycle noise-tolerant* (MC-PATH), which is less restrictive than MC-ALL and more tolerant to detection errors. This strategy also waits n cycles to build the reachability matrix. However, instead of requiring the statuses of all monitored paths to be identical, each path is considered individually: paths that are down for n consecutive cycles are added as down in the reachability matrix, while others are added as up. This approach focuses on stable paths and eliminates unstable ones.

Detection errors are filtered on a per-path basis and do not delay the aggregation of failed paths. MC-PATH has a higher probability of adding detection errors in a reachability matrix than MC-ALL, however, it may be an interesting trade-off in scenarios where detection errors are frequent.

3. WIDE-AREA EXPERIMENTS

We apply the aggregation strategies described above on measurements taken from PlanetLab and an enterprise network during 12 days. PlanetLab measurements are collected from 200 nodes probing almost 40,000 paths every minute, and measurements on the enterprise network are collected from 8 monitors probing 56 paths every 5 seconds.

In these experiments, we do not have ground truth about detection errors or real failures. So, we compare the absolute number of alarms that would be raised by a tomography algorithm [1] when using each aggregation strategy without taking into consideration if they are false or not. Tab. 1 shows the number of alarms for each aggregation strategy.

These results emphasize the contrast between the highly dynamic PlanetLab environment and the more stable enterprise network. The application of tomography to raw measurements from both networks (represented by BASIC) would lead to thousands of alarms as every change in the status of paths could result in a reachability matrix being built. PlanetLab would have one alarm per minute and the enterprise network one alarm every three minutes.

	PlanetLab	Enterprise
BASIC	16,260	6,256
MC-ALL	—	13
MC-PATH	251	27

Table 1: Number of alarms in 12 days.

We configured the multi-cycle strategies to filter all failures shorter than 10 minutes by choosing large values of n. As discussed in Sec. 2, MC-ALL is too conservative for dynamic environments; it never builds a reachability matrix for PlanetLab. In the enterprise network, MC-ALL is able to detect some failures, but not as many as MC-PATH. MC-PATH's flexible aggregation strategy works for both deployments. MC-PATH triggers 135 alarms in PlanetLab (approximately 12 alarms per day); this number is more than a hundred times smaller than using BASIC. In the enterprise network, MC-PATH reduces the number of alarms to 2 per day. It is more realistic to expect that an operational team will deal with two alarms per day than one alarm every three minutes. We could also configure our mechanisms to filter longer failures, which would reduce even more the number of alarms.

4. DISCUSSION AND FUTURE WORK

We have discussed three strategies to aggregate measurements from multiple monitors in a reachability matrix for use with tomography algorithms. These strategies address errors resulting from (1) the impossibility of synchronizing probes and (2) detection errors caused by transient packet loss. Our evaluation of these strategies using measurements from PlanetLab and from an enterprise network shows that adequate aggregation strategies can significantly reduce the number of alarms raised by tomography algorithms.

As future work, we believe that analytical modeling and controlled experiments (for ground truth) can give further insight on the impact of aggregation strategies into a tomography system's accuracy and limitations. We also plan to study techniques that would reduce the probing overhead of tomography systems when maintaining an updated snapshot of the network topology.

Acknowledgments

This work was supported by the European Community's Seventh Framework Programme (FP7/2007-2013) no. 223850.

5. REFERENCES

[1] A. Dhamdhere, R. Teixeira, C. Drovolis, and C. Diot. NetDiagnoser: Troubleshooting Network Unreachabilities Usind End-to-end Probes and Routing Data. In *Proc. CoNEXT*, New York, NY, 2007.

[2] Y. Huang, N. Feamster, and R. Teixeira. Practical Issues with Using Network Tomography for Fault Diagnosis. *SIGCOMM CCR*, 38(5):53–58, 2008.

[3] R. Kompella, J. Yates, A. Greenberg, and A. Snoeren. Detection and Localization of Network Blackholes. In *Proc. IEEE INFOCOM*, Anchorage, AK, 2007.