An Exploration of Exact Methods for Effective Network Failure Detection and Diagnosis

Auguste Burlats Cristel Pelsser
Pierre Schaus
UCLouvain, ICTEAM
e-mail: firstname.surname@uclouvain.be

Keywords: Boolean tomography, Network supervision.

Computer networks form the backbone of modern digital communication, and their reliability is crucial for maintaining seamless connectivity across various sectors. Failures within these networks can have significant consequences, leading to service disruption and potential financial loss. As such, it is essential to develop efficient and accurate methods for detecting and diagnosing network failures, enabling swift recovery and minimizing the impact on end-users.

In this study, we focus on Boolean network tomography, a research field that holds great promise for enhancing the resilience of networks. Boolean network tomography combines end-to-end measures with inference algorithms to estimate the state of different elements in the network. Its advantage is that it only requires a subset of nodes to be monitors and supervise an entire network.

With this approach, monitors send messages to each other through *measurement paths*. When a failure occurs on a node, all paths that cross it fail. Thus, the failure can be detected by observing if some measurement paths are not working. If the set of failed measurement paths forms a unique signature, then it is even possible to identify the failed node.

Our investigation focuses on minimizing the number of designated monitor nodes while ensuring some level of quality of network monitoring. This is crucial for minimizing monitoring costs without compromising the network's overall health and performance. We explore two critical monitoring challenges: the *cover problem*, which seeks to detect failures, and the *1-identifiability problem*, which requires pinpointing the exact failing node.

Conceptually, a node failure in a network results in the disruption of all paths traversing it. These affected paths collectively constitute the *symptom* associated with the failing node. A network is covered if there is a non-empty symptom for each node. Additionally, a network is considered 1-identifiable if every node possesses a unique, non-empty symptom, thereby serving as an identifier for the node in the event of a failure. By compiling a comprehensive list of these identifiers, one can efficiently diagnose a failure by simply observing the disrupted paths and cross-referencing a precomputed table that maps the failed paths (symptoms) to the corresponding node.

An important assumption of the considered networks in this study is that the routes between any pairs of nodes are imposed by the routing protocol and

known by the planning tool that will select the monitors. A pair of monitors is only able to verify the status of those routes. In practice network operators usually configure link (IGP) weights to influence where the traffic flows in the network assuming they follow shortest paths (see for instance [1] for optimizing IGP weights). Alternatively, other protocols such as segment routing or MPLS [2, 3, 4] make it possible to introduce deviations or explicit route set-ups between pairs of nodes, deviating from shortest paths. For all these protocols, the monitors are able to determine which data paths between them are affected by a failure.

We introduce and compare an integer linear programming, a constraint programming and a MaxSAT model to solve both the node cover and 1-identifiability NP-hard problems. To reduce the search space, we propose redundant constraints and problem reductions exploiting the network topologies. Using 625 real network topologies, we compare the solutions returned by each model to a specialized version of a greedy algorithm called MNMP [5], tailored for 1-identifiability. Our findings reveal that the introduced models reduce the amount of monitor nodes compared to the greedy approach for 35 instances of the cover problem and 157 instances of the 1-identifiability problem, paving the way for more robust and reliable telecommunication networks.

References

[1] Bernard Fortz and Mikkel Thorup. Internet traffic engineering by optimizing ospf weights. In *Proceedings IEEE INFOCOM 2000. conference on computer communications. Nineteenth annual joint conference of the IEEE computer and communications societies (Cat. No. 00CH37064)*, volume 2, pages 519–528. IEEE, 2000.

[2] Renaud Hartert, Pierre Schaus, Stefano Vissicchio, and Olivier Bonaventure. Solving segment routing problems with hybrid constraint programming techniques. In *Principles and Practice of Constraint Programming: 21st International Conference, CP 2015, Cork, Ireland, August 31–September 4, 2015, Proceedings 21*, pages 592–608. Springer, 2015.

[3] Renaud Hartert, Stefano Vissicchio, Pierre Schaus, Olivier Bonaventure, Clarence Filsfils, Thomas Telkamp, and Pierre Francois. A declarative and expressive approach to control forwarding paths in carrier-grade networks. *ACM SIGCOMM computer communication review*, 45(4):15–28, 2015.

[4] Youngseok Lee, Yongho Seok, Yanghee Choi, and Changhoon Kim. A constrained multipath traffic engineering scheme for mpls networks. In *2002 IEEE International Conference on Communications. Conference Proceedings. ICC 2002 (Cat. No. 02CH37333)*, volume 4, pages 2431–2436. IEEE, 2002.

[5] Liang Ma, Ting He, Ananthram Swami, Don Towsley, and Kin K. Leung. On optimal monitor placement for localizing node failures via network tomography. *Performance Evaluation*, 91:16–37, September 2015.