Consolidated Review of

Identifiability of Link Metrics Based on End-to-end Path Measurements

1. Strengths:

Rigorous approach to answer fundamental questions related to network tomography. Nice theoretical constructions to say something quite strong about the identifiability of links when probing from exterior monitor nodes.

In addition to the substantial theory, the paper keeps sight of the practical problem to be solved and points out algorithms for testing whether the network topology meets the criteria for observability, and performs a study on several real ISP topologies. Their algorithm for optimal monitor placement is carefully evaluated using simulated and real network topologies.

It presents sound and fairly deep technical results, but with very clear exposition making it easy to follow.

2. Weaknesses

Paper makes strong assumptions that are not necessarily practical. The model left out some practical considerations, e.g., communication overhead between monitors, delay, robustness against noise/failures.

The constructions really rely on the ability to direct packets along specific paths. The source routing assumption is huge, and threatens to make the work irrelevant from a practical standpoint. It is not clear why explicit routing is allowed but paths with cycles are not. In the Internet, if routers honor the loose source routing option, then cycles are possible. If they do not honor that option then explicit routing is not possible.

Evaluation is on dated topologies.

3. Comments

On the whole, I enjoyed reading this paper. It takes a well-studied domain - network tomography - and provides a nice theoretical basis for some of the under-studied aspects of the domain. In doing so, it advances the domain significantly. Yet, I feel that the results are of theoretical interest and are not really applicable to some practical situations.

I like that the necessary and sufficient conditions are first established for two monitors and then extended to the case of k>= 3 monitors by converting the problem to identifying interior graph using two monitors. Very neat. I also like that the minimum monitor placement (MMP) algorithm is based on rules extracted from the necessary & sufficient conditions.

This is a good (though quite dense) paper that extends a long line of work in network tomography. The key difference with much prior work and the present paper is that prior work took the network routing as a given, and the present work assumes that packets can be directed along any link using source routing. I wish there was some discussion of this assumption (e.g., motivating why it is reasonable, or why the resulting research is relevant), since source routing is generally unavailable in the Internet, except for certain controlled environments. The end result is that the work is quite nice, but potentially quite limited.

The paper relies on the use of source routing to measuring the end-to-end path. Whether this is a reasonable assumption or not depends on the setting you are considering. If you are talking about measuring an overlay path that is stitched from multiple independent path segments between overlay waypoints then this is a fine assumption to make. In that scenario, you are taking end-to-end measurements of a collection of overlay multi-hop paths and using that to figure out the metrics for individual overlay links, perhaps in this way you end up using much less than the O(N^2) worst case measurements you may need. This use case is quite practical - it applies to a variety of overlay based systems. You would need to do two things to bring this scenario in: (a) rewrite your paper with an explicit overlay measurement focus and (b) use a different data set, e.g., measurements from RON, to validate your approach.

If we think about measurement properties of network links, then I would think this assumption about source routing does not hold true. Unfortunately however your writing and evaluation are both focused on this specific use case to which, I'm afraid, your approach does not apply, generally speaking. Perhaps one could imagine an ISP allowing source routing internal only for its network administrators, which admins can use to run your approach and measure their networks. It would be good to see a clear elucidation of such a use case, if this is indeed what you have in mind. Even that, your evaluation is flawed and needs to be updated to use more recent ISP topologies (Rocketfuel is > 10 years old).

The evaluation study (Section 7) includes considerable amount of experiments using random graphs, which is less interesting to me. Instead, I would have liked to see some of the practical issues raised above being investigated using realistic ISP topologies.

One major concern I have is that the monitor placement problem completely left out practical considerations, e.g., communication overhead between monitors & central server (where the inference is made), delay, and robustness against noise/failures. Therefore, although the resulting placement is 'optimal' in terms of minimal monitors required to ensure coverage, it is unclear what the performance would be like in real deployment.

I was very surprised not to see some citation and discussion of much prior work in the monitor placement problem. This is a very well-studied area, and there was basically no discussion of the prior and relevant work along these lines. One reason that it might not be cited is that prior work assumes an a priori network routing matrix/function, but that's the same assumption made by much prior work in tomographic studies, and certainly those works were cited and discussed.

Proof of Theorem 3.1: To prove an upper bound on the rank of R, it is necessary to show that R can be constructed from R' (not vice versa as done), or to show that each row operation preserves rank.

Last paragraph before Corollary 4.1: Perhaps give the intuition that if one link could be identified then all could. Is this related to the fact that the null space of R' is not perpendicular to any coordinate axis?

Section 4.2: If cycles are allowed, then it seems reasonable that the self-loop m_1 to m_1 should be allowed. If that is possible, then the argument in 4.2 seems to break down. (Technically, this is not a path between two monitors, and so the Section 4.2 is correct.)

It would be useful if all required lemmas from [22] are stated without proof in this paper. Some space could be saved by reducing the height of figures 12-14.

It would use useful if footnote 5 defined what is meant by "a link separating its endpoints". Does it mean a link whose removal would make the graph disconnected? If so, why say "separating its endpoints", since it is not merely the endpoints that are separated.

Section 5.2.2: Since the appendix provides no proofs, I think the statements of the lemmas should appear in the body of the proof. Lemma A.1 is central to understanding the proof, and the purpose of an appendix is to contain material that a reader could skip. (Also, it is wrong to say "we first show in Lemma A.2-(b)", since no proof is given.)

Section 6.1 is not a proof.

Appendix B, case (1)(c): "if v_1,v_2 \in \{u_1,u_2\}" The comma could mean either "and" or "or". Do you mean "if \{v_1,v_2\} = \{u_1,\u_2\}" or "if \{v_1, v_2\} \cap \{u_1,\u_2\} = \emptyset"?

Typo: Section 2.3: The paths from m_2 -> m_3 are actually paths m_3 -> m_2.

4. Summary from PC Discussion

In the PC meeting discussion, the reviewers were somewhat divided on this work. On the one hand, they felt that the theoretical contributions were solid. On the other hand, there was significant discussion on the fact that the theoretical constructions ignore many important practical concerns, not least of which is the fact that source routing is not generally enabled in the internet. The reviewers suggested that the authors consider either casting the work in terms of an overlay network in which routes can be explicitly specified, or in terms of tunneling within a single ISP, where paths can also be explicitly chosen (although this latter scenario is potentially of no interest, since an ISP would almost certainly already be able to measure any link-level characteristic of interest). In the end, the reviewers decided on acceptance of the paper, with some members still voicing concerns about the practical importance and appeal of the work.

5. Authors' Response

Responses to main concerns:

As the reviewers correctly recognize, our results depend on the assumption of (cycle-free) source routing. We would like to emphasize that the nature of this work is fundamental research, and our results apply to any type of network that allows the monitoring system to control routing of measurement packets under the cycle-free constraint. While source routing may not be widely supported in the current Internet, it is generally feasible in overlay networks and single administration networks as the reviewers have pointed out. Moreover, our assumption also holds true in Software Defined Networks (SDN), where we can select an arbitrary path for each flow through the SDN controller, with the only constraint being that the path does not contain cycles. For these networks, our contributions are: (i) developing explicit, verifiable conditions for the network administrator (or owner of monitoring system) to test identifiability of internal link states from external end-to-end measurements, and (ii) developing an efficient algorithm to select locations for placing monitors (e.g., by deploying monitoring client or turning on monitoring functionality) so as to achieve identifiability with the minimum number of monitors. We have added a discussion of how our results are applicable to some practical network environments in our revised Section 1.

Meanwhile, we note that it is difficult to cover all applicable network scenarios mentioned above in evaluations, due to limitation in space and time, and most of all, access to suitable data sets. We have carefully studied the RON data set pointed out by Reviewer 4 in the hope of evaluating our solutions on the topology of an overlay network, but found that this data set does not contain topology information. Hence, we leverage random graph models to evaluate our solution on a variety of network topologies, as presented in Section 7.3.1. Moreover, to evaluate the performance of our proposed algorithm (MMP) in more updated Autonomous System topologies (which belong to single administration networks), we added an evaluation section in Section 7.3.2 based on the publicly available topology data set CAIDA, which is released in Apr. 2013.

Detailed responses:

In addition to the discussion of applicable network scenarios and further algorithm evaluations mentioned above, we also addressed a number of reviewer comments as follows:

1. *"I was very surprised not to see some citation and discussion of much prior work in the monitor placement problem. This is a very well-studied area, and there was basically no discussion of the prior and relevant work along these lines."*
We have included a discussion of existing work on monitor placement in Section 1.1.

2. *"Proof of Theorem 3.1: To prove an upper bound on the rank of R, it is necessary to show that R can be constructed from R' (not vice versa as done), or to show that each row operation preserves rank."*
We have added arguments showing that R can be reconstructed from R' in the proof.

3. *"Last paragraph before Corollary 4.1: Perhaps give the intuition that if one link could be identified then all could"*
We have added this intuition to Corollary 4.1 and updated its proof.

4. *"Section 4.2: If cycles are allowed, then it seems reasonable that the self-loop m_1 to m_1 should be allowed. If that is possible, then the argument in 4.2 seems to break down."*
We have clarified the condition for our conclusion in Section 4.2, which excludes the counter-example raised by the reviewer.

5. *"It would be useful if all required lemmas from [22] are stated without proof in this paper."*
We have moved all lemmas and propositions used in this paper from the tech report to the appendix.

6. *"It would use useful if footnote 5 defined what is meant by "a link separating its endpoints". Does it mean a link whose removal would make the graph disconnected?"*
We have rephrased the definition of bridge in footnote 6.

7. *"Section 5.2.2: Since the appendix provides no proofs, I think the statements of the lemmas should appear in the body of the proof."*
We have moved the lemmas used in Section 5.2.2 to the main text.

8. *"Section 6.1 is not a proof."*
We have restructured Section 6 and renamed Section 6.2 (originally Section 6.1) to rigorously reflect its content.

9. *"Appendix B, case (1)(c): "if v_1,v_2 \in \{u_1,u_2\}" The comma could mean either "and" or "or". Do you mean "if \{v_1,v_2\} = \{u_1,\u_2\}" or "if \{v_1, v_2\} \cap \{u_1,\u_2\} = \emptyset"?"*
We have modified this condition to eliminate the confusion. Note that in the revised version, Appendix B (proof of optimality of MMP) has been moved to the technical report due to space limit, and its conclusion is stated in Theorem 7.1 in this paper.

10. *"Typo: Section 2.3: The paths from m_2 -> m_3 are actually paths m_3 -> m_2."*
We have corrected the typo.

Moreover, we feel that the issues (e.g., communication overhead, robustness against noise/failures) raised by Reviewer 2 are very interesting considerations for practical monitor deployment. However, we believe the investigation of these issues are out of scope of the current paper, which focuses on network identifiability and monitor placement to achieve identifiability.