iTop: Inferred Topologies with Partial Path Information

B. Holbert S. Tati S. Silvestri T. La Porta

Full knowledge of the routing topology of the Internet is useful for a multitude of network management tasks. However, the full topology is often not known and is instead estimated using topology inference algorithms. Many of these algorithms use Traceroute to probe paths in the network and then use the collected information to infer the topology. In practice some router configurations may severely disrupt the operation of Traceroute and cause it to only provide partial information about paths in the network. iTop is an algorithm for inferring the network topology when only this partial information is available.

iTop constructs a virtual topology, which overestimates the number of network components, and then repeatedly merges links in this topology to resolve it towards the structure of the true network. iTop has been compared to other inference algorithms through extensive simulations. Results show that iTop significantly outperforms previous approaches and its inferred topologies are within 5% of the original networks for all the considered metrics. Additionally, the topologies inferred by iTop significantly improve the performance of fault localization algorithms when compared to other approaches.

Network Model

- **Monitors**
  - Nodes that participate in information collection process
  - Partial paths determined using Traceroute between monitors
  - Estimate path lengths between monitors (e.g., UDP TTL field)
- **Observed Network Components**
  - Nodes and links reported by Traceroute
  - Nodes classified as Responding (R)
- **Non-Cooperating Nodes (NC)**
  - Anonymous Nodes (A)
    - Do not reply to Traceroute probes
    - Do forward Traceroute probes
  - Blocking Nodes (B)
    - Do not forward or reply to Traceroute probes
    - Result in hidden nodes (HID) between two blocking nodes in a path
      - True type of hidden nodes cannot be discerned from traces

	R-R	R-A	R-B	R-NC	A-A	A-HID	NC-NC	NC-HID	HID-HID	A-NC	B-NC	A-B
R-R	-	-	-	-	-	-	-	-	R-R	-	-	-
R-A	-	**R-A**	-	**R-A**	-	R-A	-	R-A	R-A	-	-	-
R-B	-	-	**R-B**	**R-B**	-	-	-	R-B	R-B	-	-	-
R-NC	-	**R-A**	**R-B**	**R-NC**	-	-	-	R-NC	R-NC	-	-	-
A-A	-	-	-	-	A-A	A-A	A-A	A-A	A-A	A-A	-	-
A-HID	-	R-A	-	-	A-A	A-HID	A-NC	A-HID	A-HID	A-NC	A-B	A-B
NC-NC	-	-	-	-	A-A	A-NC	NC-NC	NC-NC	NC-NC	A-NC	B-NC	A-B
NC-HID	-	-	-	-	A-A	A-HID	NC-NC	NC-HID	NC-HID	A-NC	B-NC	A-B
HID-HID	R-R	R-A	R-B	R-NC	A-A	A-HID	NC-NC	NC-HID	HID-HID	A-NC	B-NC	A-B
A-NC	-	-	-	-	A-A	A-NC	A-NC	A-NC	A-NC	A-NC	A-B	A-B
B-NC	-	-	-	-	-	A-B	B-NC	B-NC	B-NC	A-B	B-NC	A-B
A-B	-	-	-	-	-	A-B	A-B	A-B	A-B	A-B	A-B	A-B

Compatibility Matrix

iTop Algorithm

- **Phase 1: Construct Virtual Topology (VT)**
  - Start with observed network components
  - Add virtual components to represent unobserved ones
    - Inferred from path lengths between monitors
    - Real components may be duplicated in multiple virtual ones
- **Phase 2: Identify Merge Options**
  - Merge options indicate VT links that can be combined
  - Check set of rules to determine valid merges
    - Links cannot be observed in the same path in VT
    - Merging two links must preserve observed shortest distances between monitors
    - Link types must fit compatibility matrix
      - Matrix indicates what links can be combined based on the types of their endpoints
- **Phase 3: Merge Links**
  - Iteratively combine valid pairs of links
    - Always choose links with fewest possible valid merges
    - Valid merges reduced after each merge is performed
  - Recheck compatibility matrix before each merge
    - Some link endpoints change due to previous merges
  - Repeat until no valid merges remain
  - Resulting topology is the iTop topology

Simulation Results

- **Network Measures**
  - iTop compared to two other pre-existing approaches
    - Merge Nodes (MN) [2]
    - ISOMAP [3]: Merges nodes using distance thresholds
  - Use random networks and real networks from CAIDA [4]
  - iTop topologies best reflect ground truth (GT) topologies
    - Within 5% of measures for GT topology
    - Inferred numbers of edges and nodes (shown for real networks)
    - Node degree and joint degree distribution

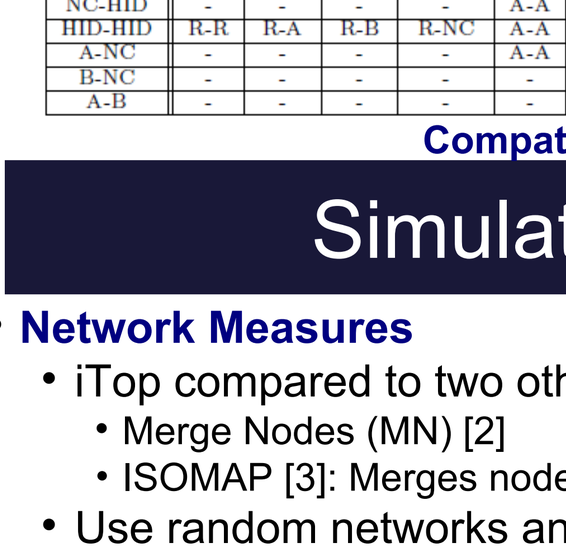

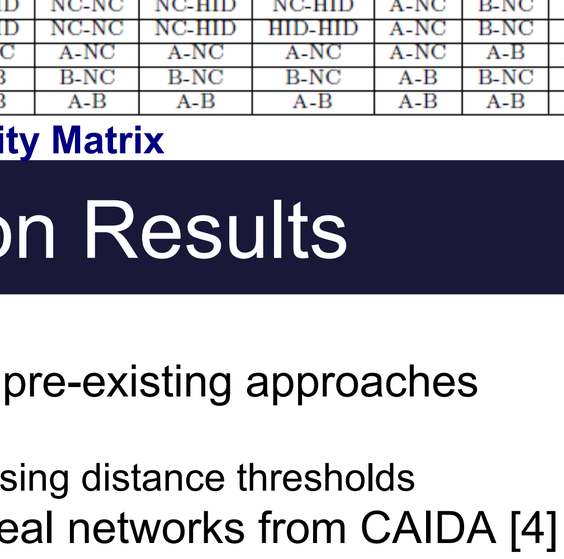

- **Fault Diagnosis**
  - Faults identified using Max-Coverage algorithm
  - Hypothesis lists modified to account for merging
  - Better inference in iTop results in better diagnosis results
  - Higher accuracy, fewer false positives

Conclusions

- **Merging**
  - iTop infers networks which are closer to the real network than networks inferred by MN and ISOMAP
  - iTop scales well with the percentage of non-cooperating routers
- **Fault Diagnosis**
  - Networks inferred by iTop result in better performance of fault diagnosis algorithms
  - Easier to find link(s) to repair in real-world scenarios

References

[1] B. Yao, R. Viswanathan, F. Chang, and D. Waddington, "Topology inference in the presence of anonymous routers," *IEEE Infocom*, 2003.

[2] X. Jin, W.-P. Yiu, S.-H. Chan, and Y. Wang, "Network topology inference based on end-to-end measurements," *IEEE Journal on Selected Areas in Communications*, vol. 24, no. 12, pp. 2182–2195, 2006.

[3] T.C.A. For Internet Data Analysis (CAIDA), "Macroscopic internet topology data kit (itdk)," 2013.

Related Publications

- B. Holbert, S. Tati, S. Silvestri, T.F. La Porta, A. Swami, " Effects of Partial Topology on Fault Diagnosis," IEEE MILCOM 2013.

Project Sponsored by the Defense Threat Reduction Agency