Robust Network Tomography in the Presence of Failures

Stikar Tati, Simone Silvestri, Ting He, Thomas La Porta

Network tomography has been proved to be a valuable tool for network monitoring and management in large scale networks. However, previous work in network tomography do not consider the occurrence of failures, while designing the network measurement framework. In this work, we study the problem of selecting paths to improve the performance of network tomography applications in the presence of network element failures.

Under the assumption that the failure distribution is known, we propose an algorithm with guaranteed approximation ratio and feasible run-time; it is called RObust MEasurements (RoMe). We also consider the case in which the failure distribution is not known, and propose a reinforcement learning algorithm to solve our optimization problem, using RoMe as a subroutine.

Network Tomography

- Goals:
  - Determine the state of internal elements using e2e measurements
  - Estimate the complete set of e2e measurements from an incomplete set

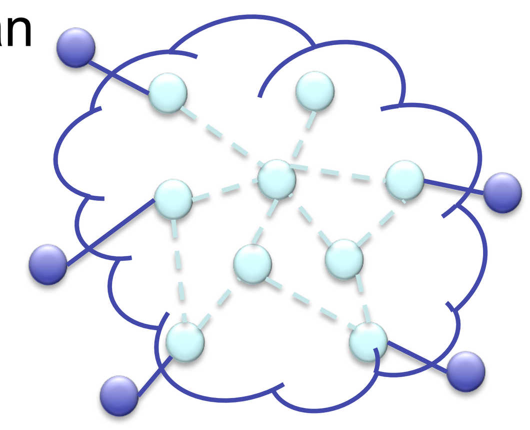

- Additive metrics (e.g. delay, log of loss rate)
- Solve a linear system: Ax = y (x: Link metrics; y: e2e meas.; A: Path matrix)
- Reduce overhead by probing a subset (basis) of the available paths
- Previous approaohces **do not consider failures** of network elements

Impact of failures

Motivational example:

- Monitors: m1, m2,..., m6

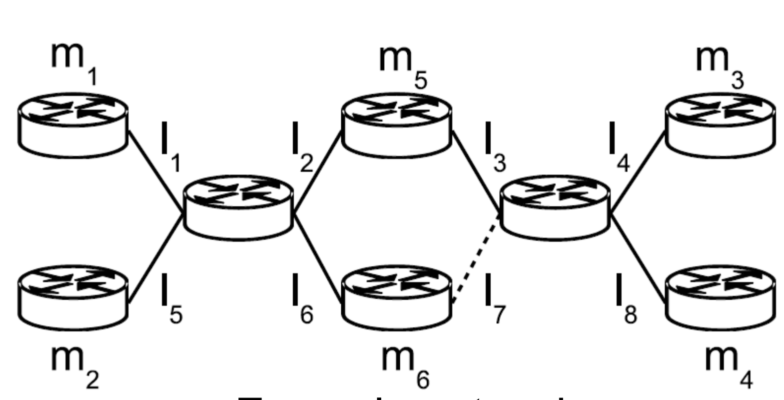

Example network

- Considered bases
  - $R_1 = (q_1,q_2,q_4,q_{11},q_{15},q_5,q_6,q_7)$
  - $R_2 = (q_5,q_6,q_7,q_8,q_9,q_{10},q_{11},q_{12})$

- Probing R_1 when l_7 fails:
  - Provides rank 3
  - Does not cover links l_3 and l_6
  - Results in a linear system with infinite solutions

- Probing R2 when l7 fails:
  - Provides rank 6
  - Enables to determine that l7 has failed
  - Results in a linear system that has a solution for each link

Motivational experiments:

- Realistic ISP topologies from the Rocketfuel Project [1]

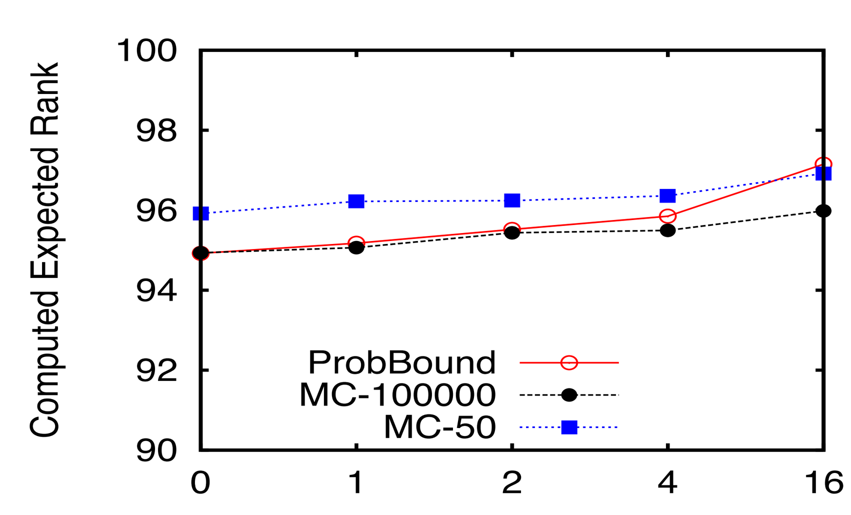

Problem formulation

- Problem setting:
  - Heterogeneous probing cost
  - Limited budget
  - Set of *candidate paths* R_M
- Objective function: Expected Rank (ER)

$$ER(R) = \sum_{v \in \{0,1\}^{|E|}} r(R_v)P(v)$$

- Optimization problem:

Find $R^* \subseteq R_M$

$$R^* = \underset{R^* \subseteq R_M}{\textbf{argmax}}\ ER(R)$$

$$\text{s. t.}\quad PC(R) \leq B$$

(PC: Probing cost; B: Budget)

Known failure distribution

- RoMe
  - Approximation factor $(1-1/\sqrt{e})$
  - Exponential complexity in computing ER value
- Approximating ER value to achieve polynomial complexity

$$E(Z_{\hat{R}_q}) \leq \sum_{w \in R_{ind}} E(X_w) + \sum_{q \in R_{dep}} E(D_q)$$

- Optimal with additional constraint of linear independence in paths

Algorithm 1: the RoMe algorithm

```
Input: Set R_M, budget B, cost function PC : 2^{R_M} → ℝ+,
       objective function ER : 2^{R_M} → ℝ+
Output: A subset of R_M
R̄ = arg max_{q∈R_M} {ER({q}) : PC({q}) ≤ B};
R_out = ∅;
R = R_M;
w_q = (ER({q}) − ER(∅)) / PC({q}), ∀q ∈ R_M;
while R ≠ ∅ do
    q_max = arg max_{q∈R} w_q;
    if PC(R_out) + PC({q_max}) ≤ B then
        R_out ← R_out ∪ {q_max};
    R = R \ {q_max};
    // Weight update
    forall the q ∈ R do
        w_q = (ER(R_out ∪ {q}) − ER(R_out)) / PC({q});
if ER(R_out) > ER(R̄) then
    return R_out
else
    return R̄
```

Unknown failure distribution

- Reinforcement learning approach
  - Learning with Submodular Rewards (LSR) algorithm
  - We learn path availabilities through probing
- LSR has performance guarantees under certain conditions

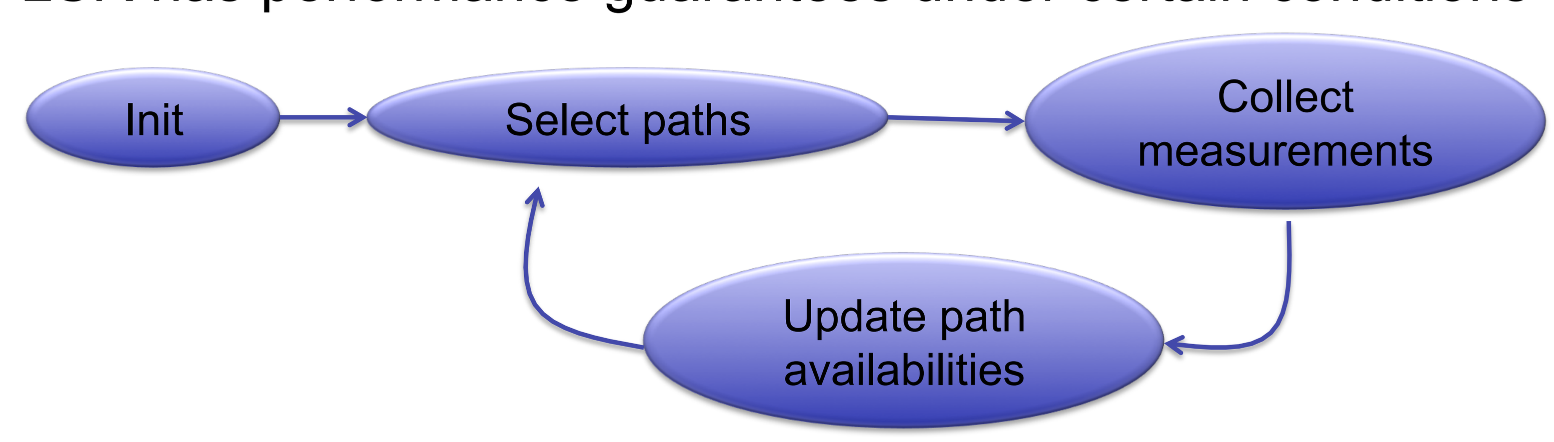

Results

- Realistic ISP topologies from the Rocketfuel Project [1]
- Realistic failure model [2]
- We compared to SelectPath [3]

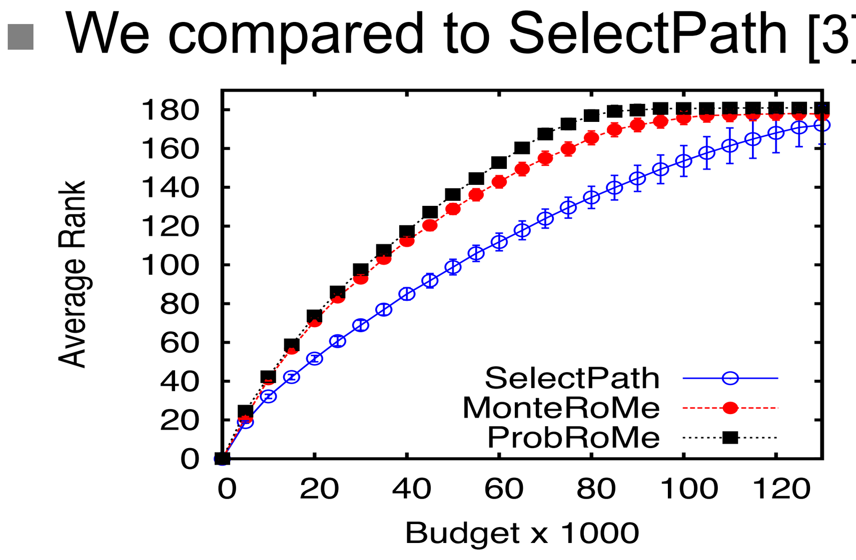

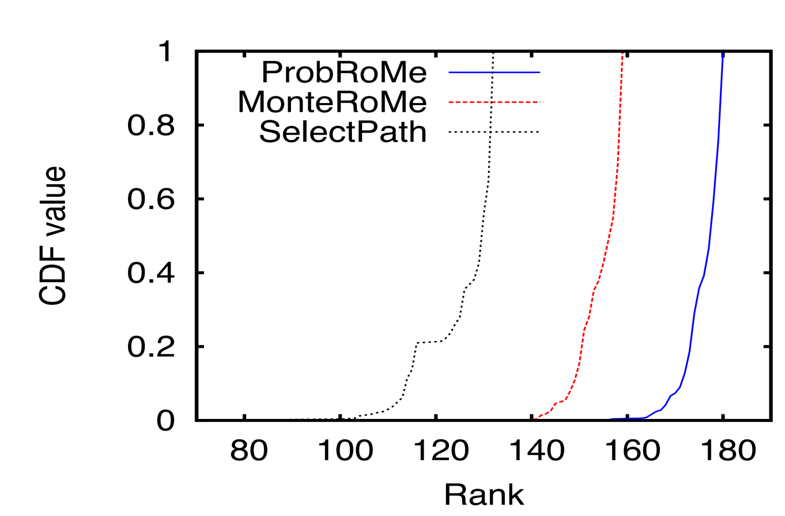

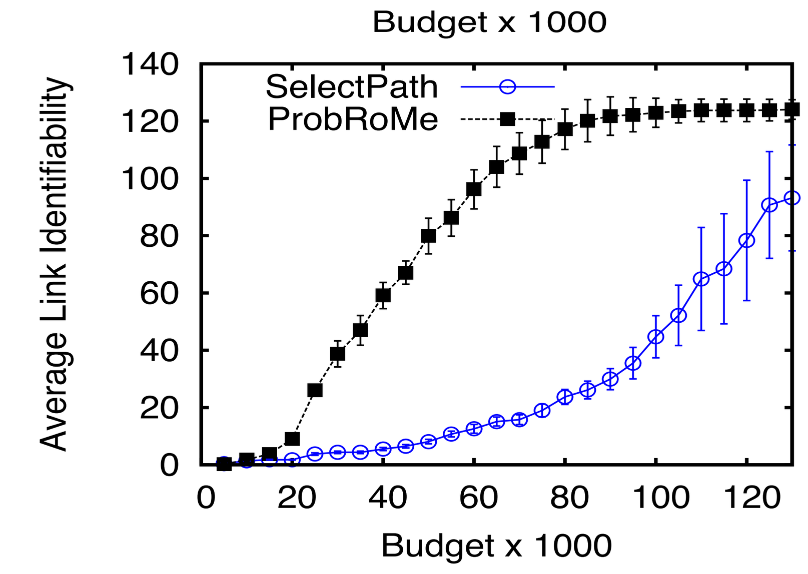

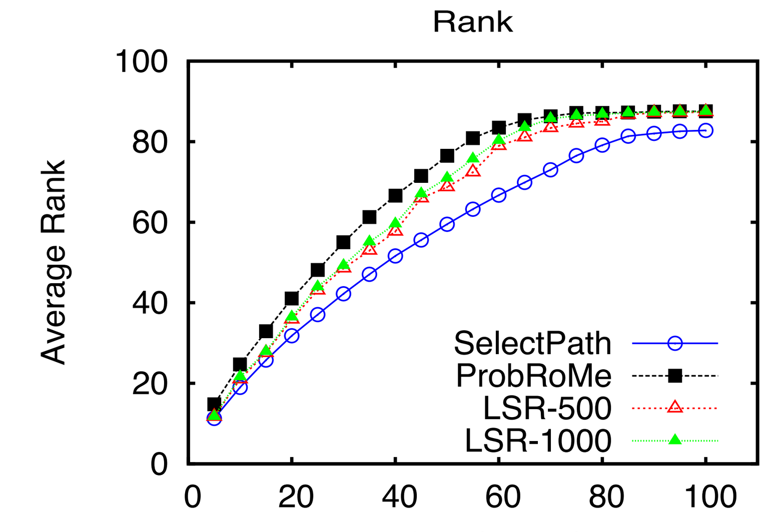

References

[1] “Rocketfuel Project: Internet topologies,” http://www.cs.washington.edu/research/networking/rocketfuel/.

[2] A. Markopoulou, G. Iannaccone, S. Bhattacharyya, C.-N. Chuah, and C. Diot, “Characterization of failures in an IP backbone,” IEEE INFOCOM, 2004.

[3] Y. Chen, D. Bindel, H. Song, and R. H. Katz, “An algebraic approach to practical and scalable overlay network monitoring,” ACM SIGCOMM Comp. Com. Rev., 2004.

[4] S. Tati, S. Silvestri, T. He, T. La Porta, “Robust Network Tomography in the Presence of Failures”, IEEE ICDCS 2014

Sponsored by Army Research Lab and UK Ministry of Defence - ITA Program