Cross-Domain Fault Localization: A Case for a Graph Digest Approach

William Fischer Geoffrey Xie Joel Young

Department of Computer Science
Naval Postgraduate School

Network Fault Localization

Fault localization: integral part of network management/troubleshooting

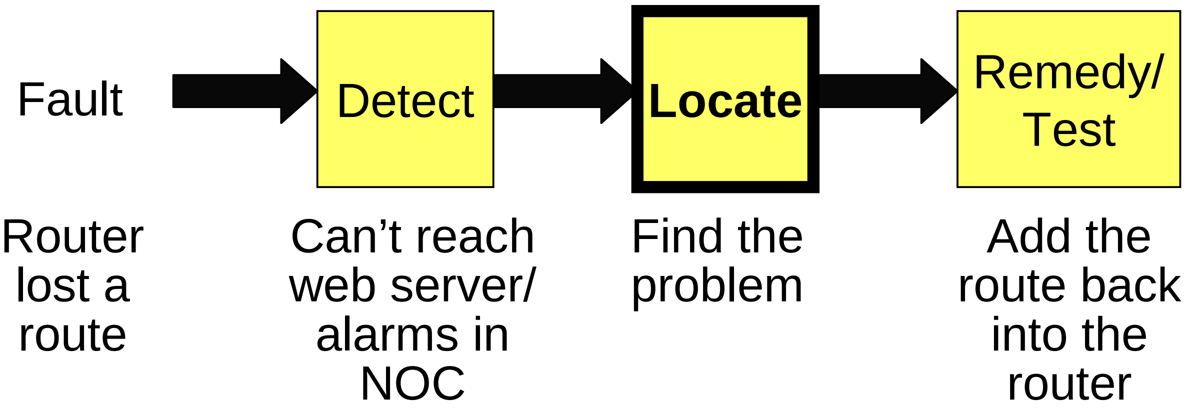

Not always easy to locate a fault

- Large number of devices
- Stale topology databases
- Human-introduced errors tough to find

Cross-Domain Fault Localization

Networks are highly connected

- Some faults can affect many domains
    - E.g. DNS failure, link congestion
- **Correlating observations across domains intuitively increases accuracy of locating these types of faults**

Domains can represent fault propagation with a graph

- Common Ground

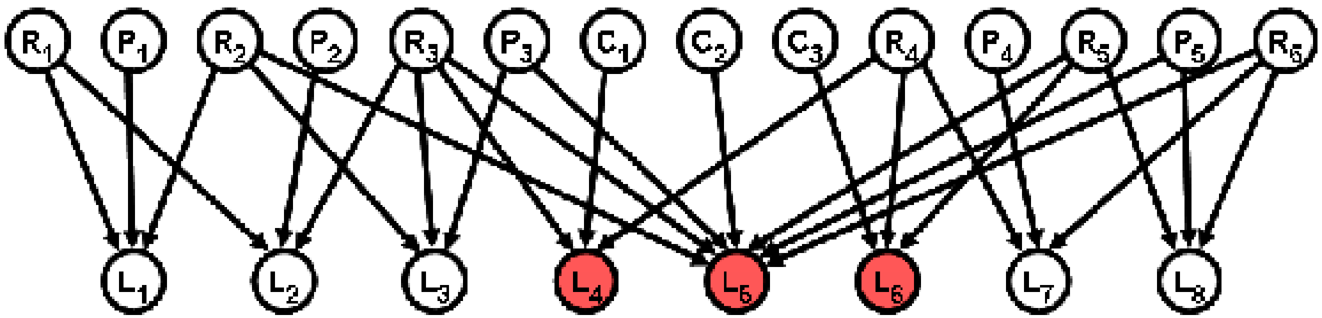

Challenges and Assumption

Challenges

- Domain managers are reluctant to share internal information
  - Topology, state, sensitive properties, etc
  - May even view other domains as *adversaries*
- Accuracy versus Privacy: competing goals
- Scalability: measured by the size of the aggregated inference graph

Assumption

- Domain managers will want to participate if
  - Privacy preserved
  - Fault localization accuracy improved

Related Work

Intra-domain: Significant recent advances in fault localization

- SHRINK [Kandula et al., 2005] and SCORE [Kompella et al., 2005]
  - bipartite causal graph model
- Sherlock [Bahl et al., 2007]
  - Multi-level causal graph model

Cross-domain: under-researched

- End-to-end approach for hierarchical organizations [Steinder et al., 2008]
  - Constrained environment

Our Graph Digest Approach

What is a graph digest?

A reduction of a fault propagation model (eg causal graph) to a digest representation of nodes and edges

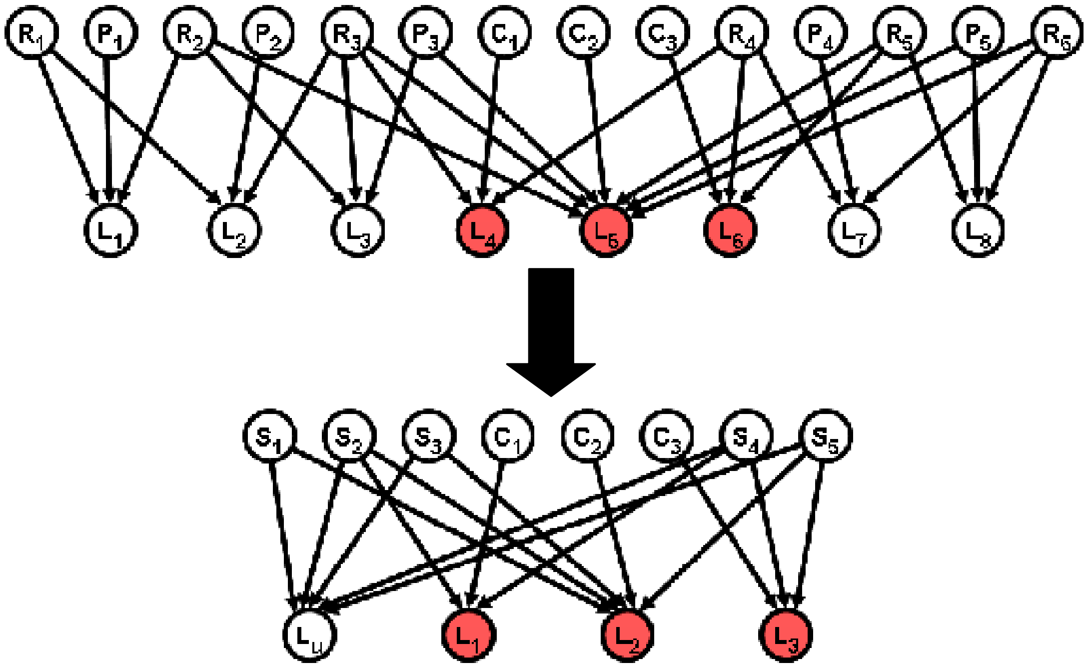

Our Graph Digest Approach

$$\mathcal{G}^{j} = \left(\biguplus_{i \neq j}^{n} f(G_i) \right) \biguplus G_j$$

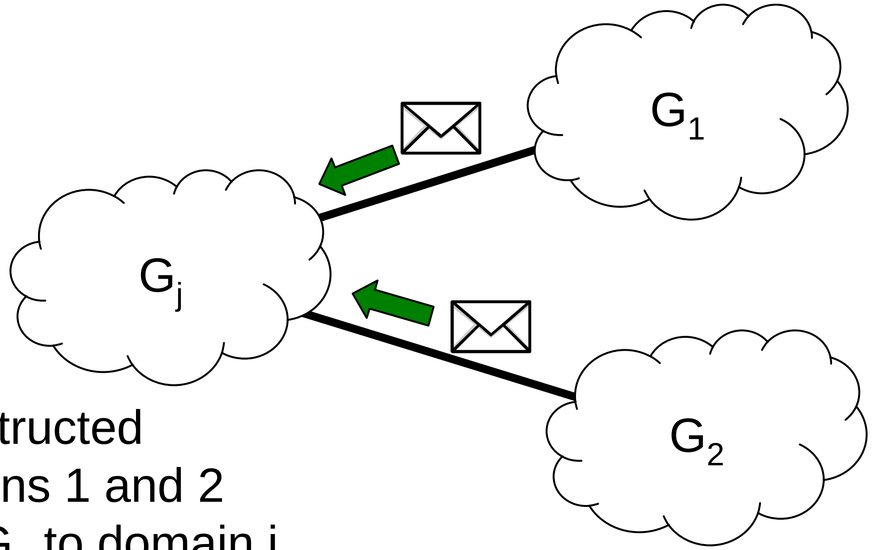

1. Fault detected in j^{th} domain, G_j constructed
2. Domain j asks for digest from domains 1 and 2
3. Domains 1 and 2 send digests G_1, G_2 to domain j
4. Domain j imports the digests and runs inference

A general framework for creating and using a digest that explicitly models the inference accuracy and privacy requirements

Performance Criteria

Heart of the approach is quantifiable metrics for accuracy and privacy

Accuracy

$$h = \frac{|B_d \cap B_u|}{|B_d|} \qquad c = \frac{|B_d \cap B_u|}{|B_u|}$$

B_u: Best explanation using undigested graphs

B_d: Best explanation using 1 undigested graph, 1+ digests

$$\alpha = \frac{2 * h * c}{h + c} \text{ for } h + c > 0, \text{ otherwise } \alpha = 0$$

Privacy: let *S* model adversary's knowledge of sensitive property

$$KL((S \mid digest),\ S) = \sum_{x \in X} \Pr(S = x \mid digest) \log_2 \frac{\Pr(S = x \mid digest)}{\Pr(S = x)}$$

Practical Privacy Metrics

KL Distance an ideal metric for a privacy criterion, but hard to quantify

In this work we look at protecting a domain against causal graph attacks

From causal graphs, can infer

- In-degree, out-degree, path lengths, reachability, etc

Sample privacy metrics:

- Reachability, number of routers, maximum node degree, diameter

Illustration
SHRINK

Assumptions

- Bipartite model
- Independent SRG failures
- No more than 3 simultaneous failures

$$\arg \max_{\langle S1,\dots,Sn \rangle} \Pr(\langle S_1,\dots,S_n \rangle \mid \langle L_1,\dots,L_m \rangle)$$

Adds “noisy” edges to form complete bipartite graph

- d = .0001 edge strength for a noisy edge $\Pr(L_j \mid S_i) = .0001$
- Subtract d from any edge strength of 1.0 $\Pr(L_j \mid S_i) = .9999$

Returns most probable explanation for the observations

Illustration
SHRINK

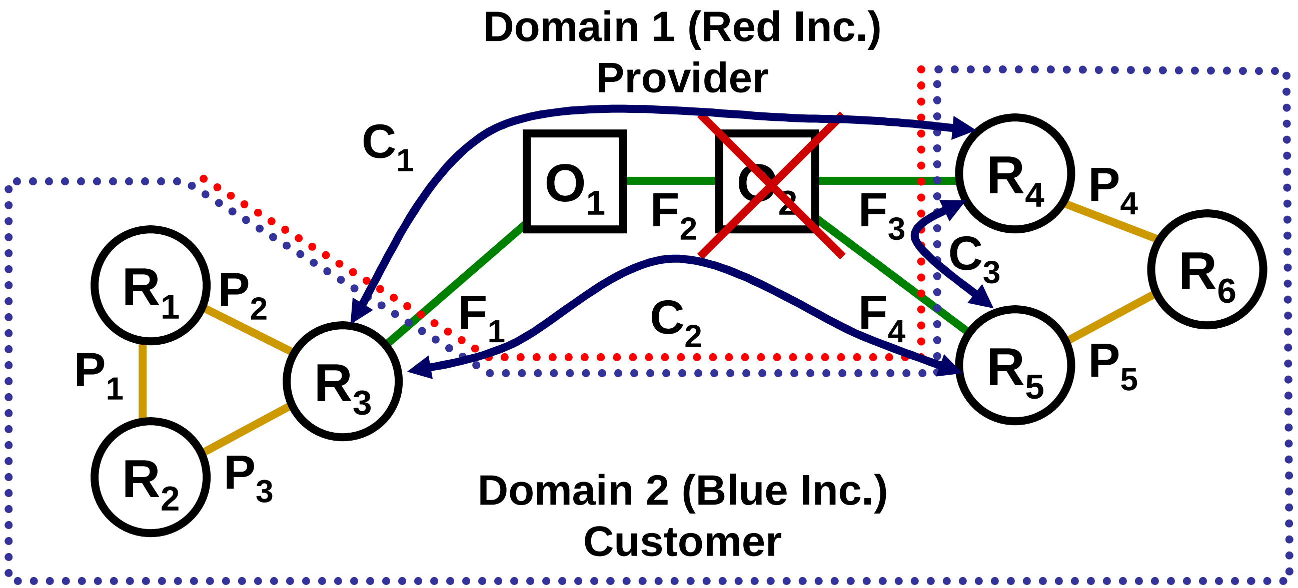

Topology

Identifiers:

R* IP Router
P* Point to Point link
C* Leased Optical Circuit
F* Optical Fiber
O* Optical Switch

Illustration
SHRINK

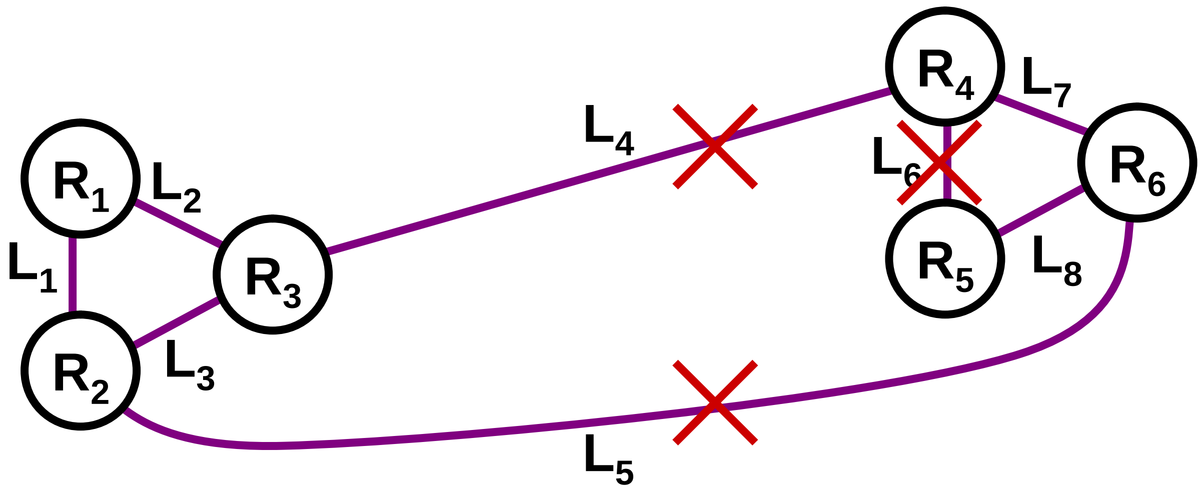

IP View

Illustration
SHRINK

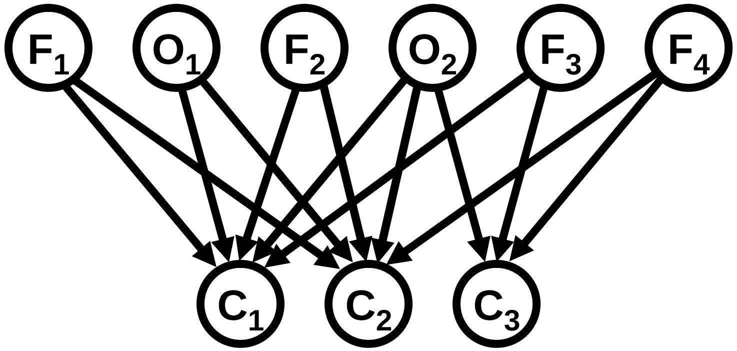

Provider Causal Graph
(with respect to customer)

Illustration
SHRINK

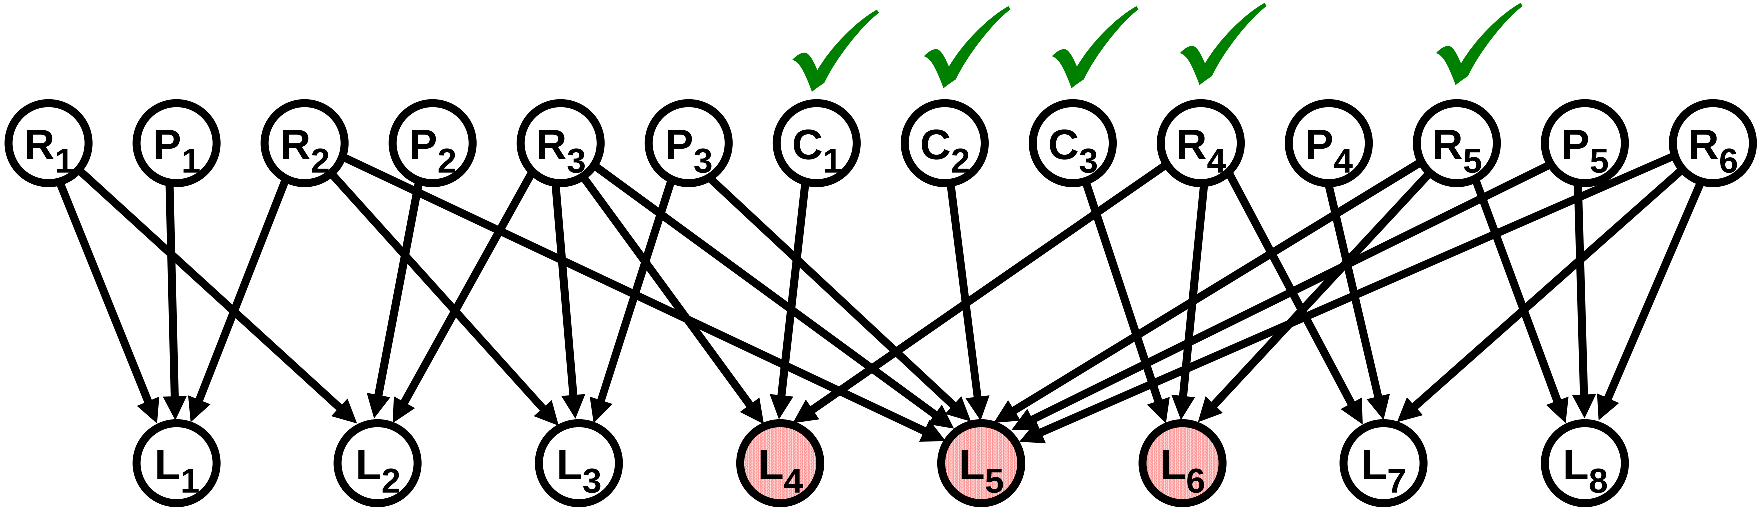

Customer Causal Graph
with observation state


Best explanation

Illustration
SHRINK

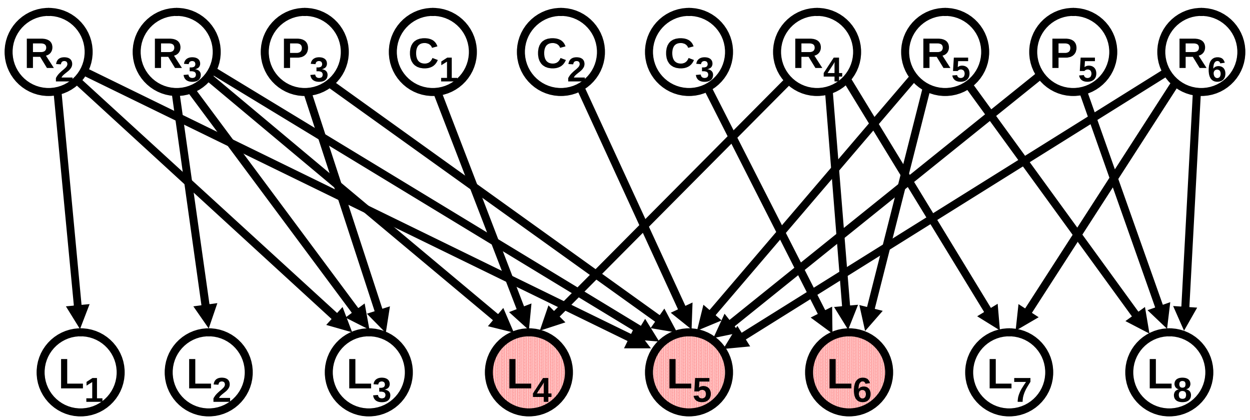

Nodes R_1, P_1, P_2, P_4 pruned

Illustration
SHRINK

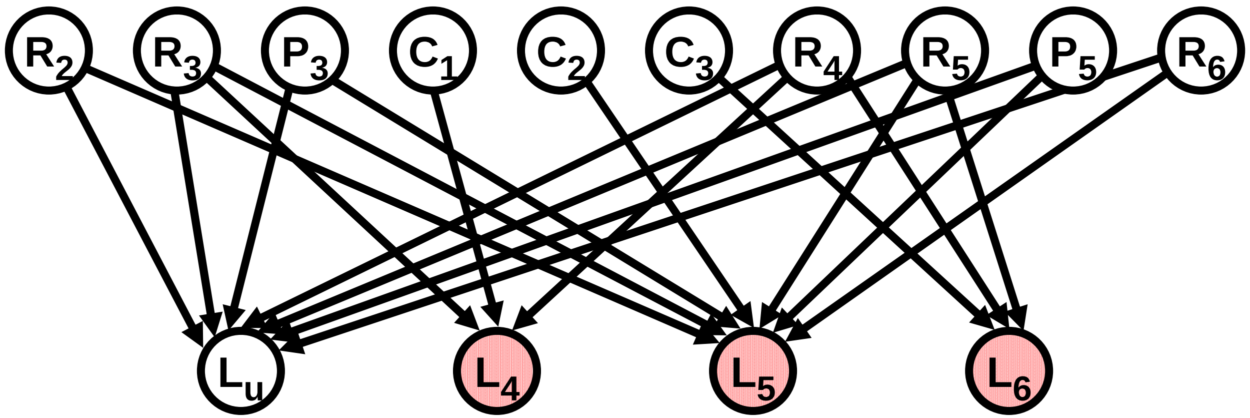

Combing all “up” observations

Illustration
SHRINK

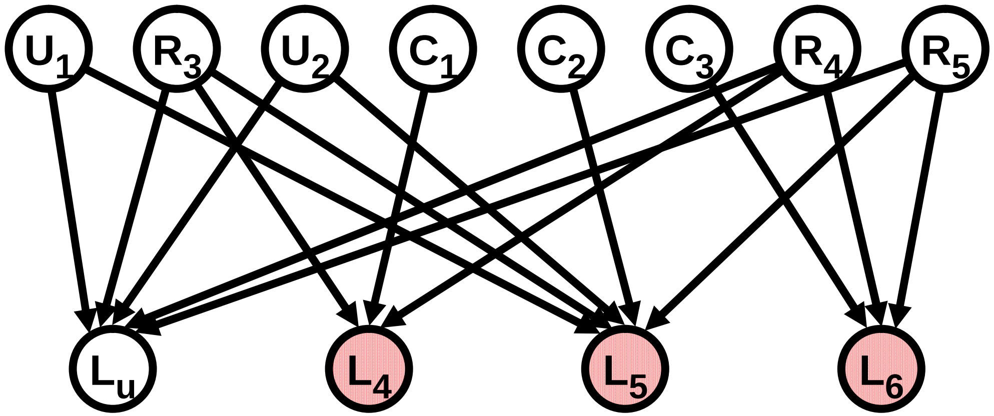

U1 = R2,R6
U2 = P3,P5

Aggregating nodes that are indistinguishable in the causal graph

Illustration
SHRINK

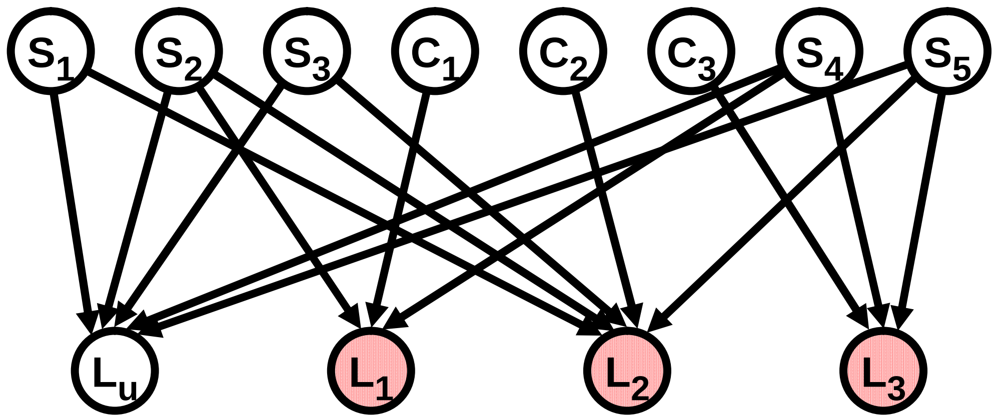

Renaming all but the “shared attribute” nodes and L_u

Illustration
SHRINK

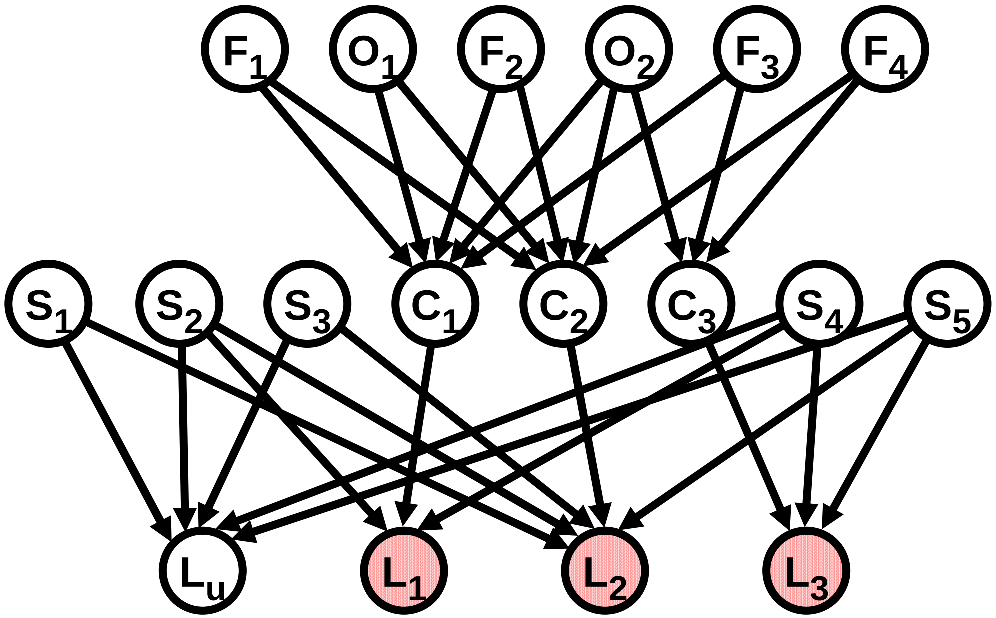

Multilevel union

Illustration
SHRINK

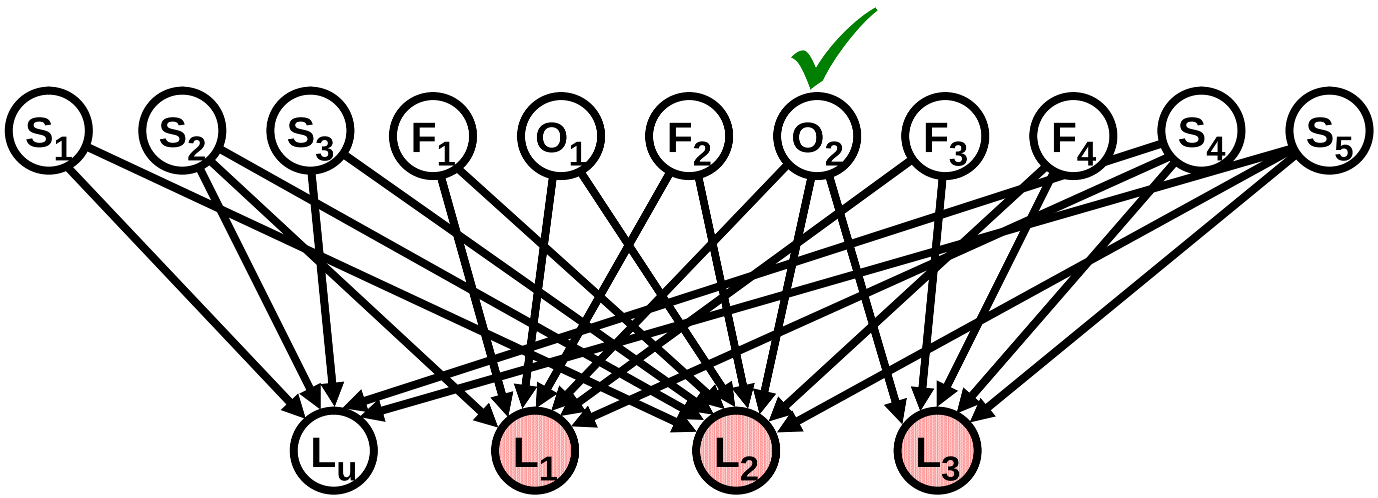

Model Specific Bipartite Union

✓ Best explanation

Continuing Work
Customer Network Physical Topology

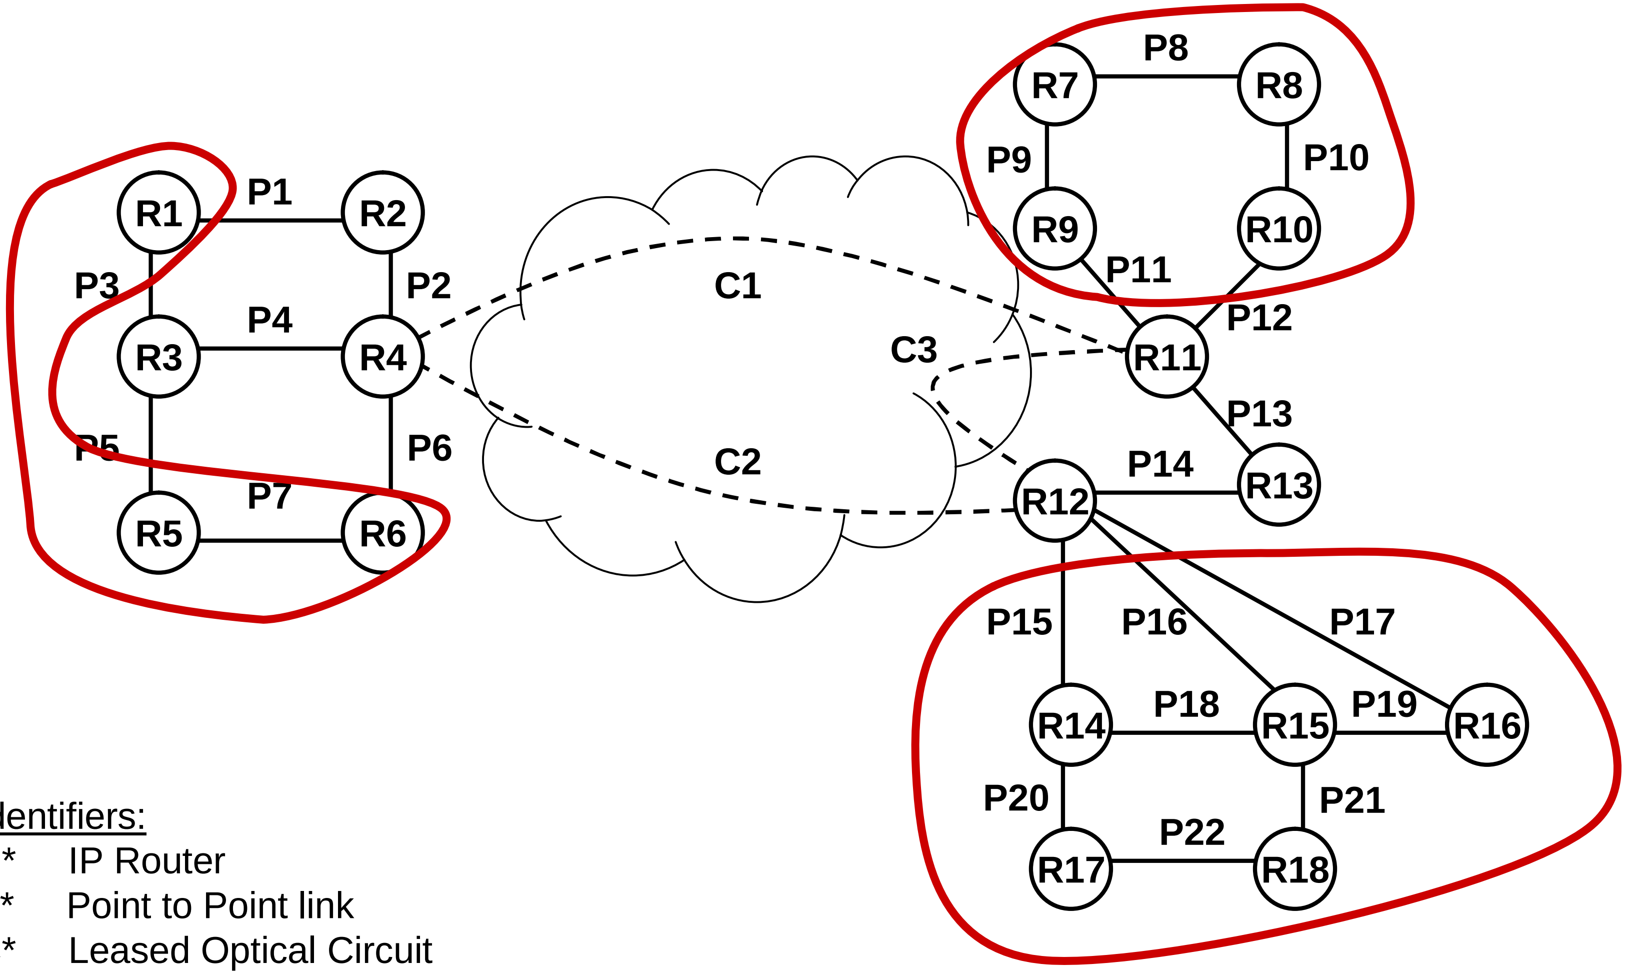

<u>Identifiers:</u>
R* IP Router
P* Point to Point link
C* Leased Optical Circuit

Continuing Work
Customer Network IP View

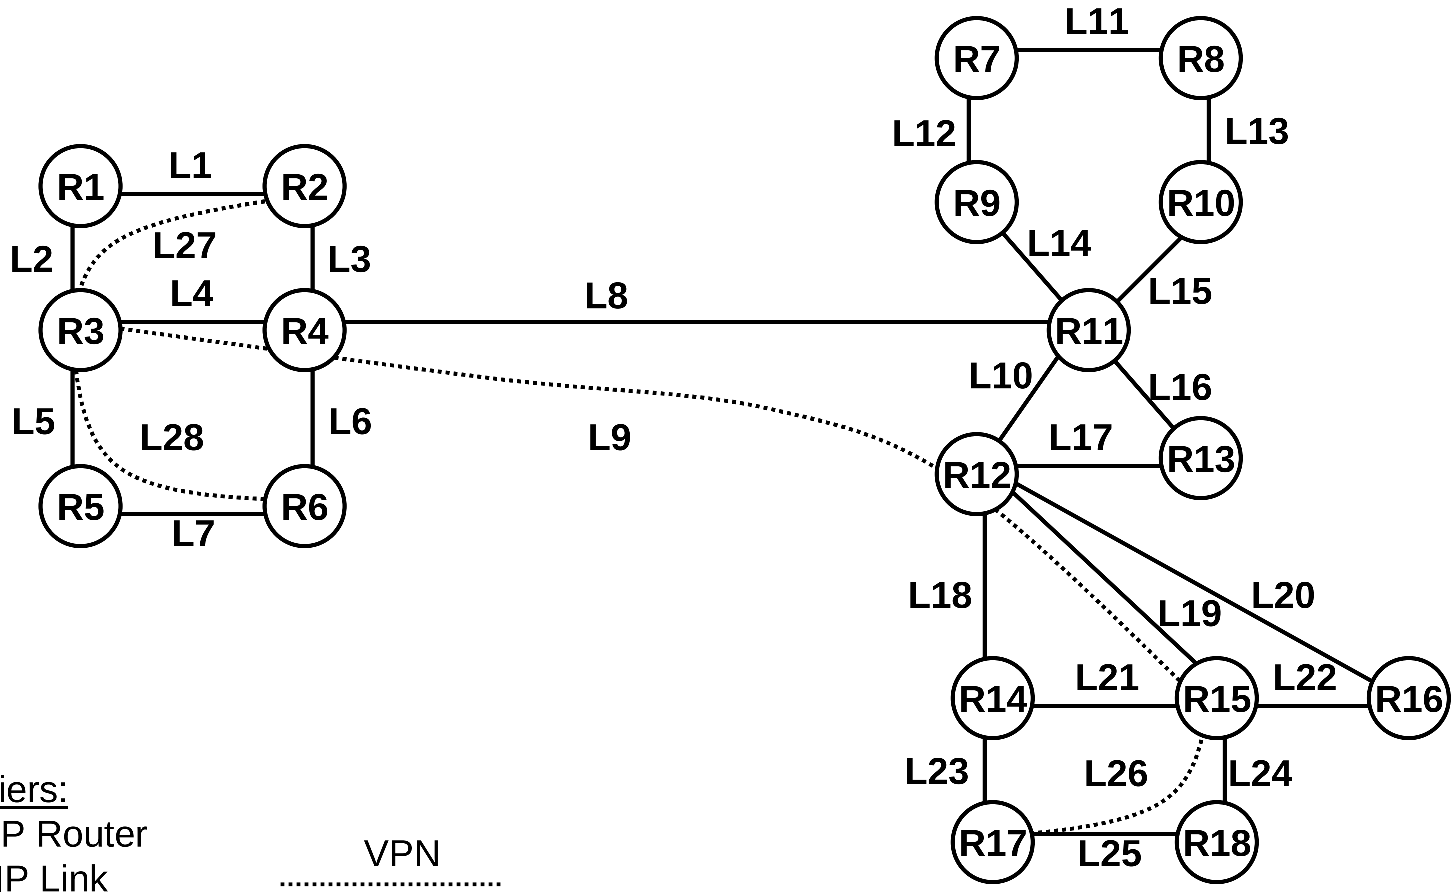

Identifiers:
R* IP Router
L* IP Link

VPN ....

Accuracy Results for All Single Failures

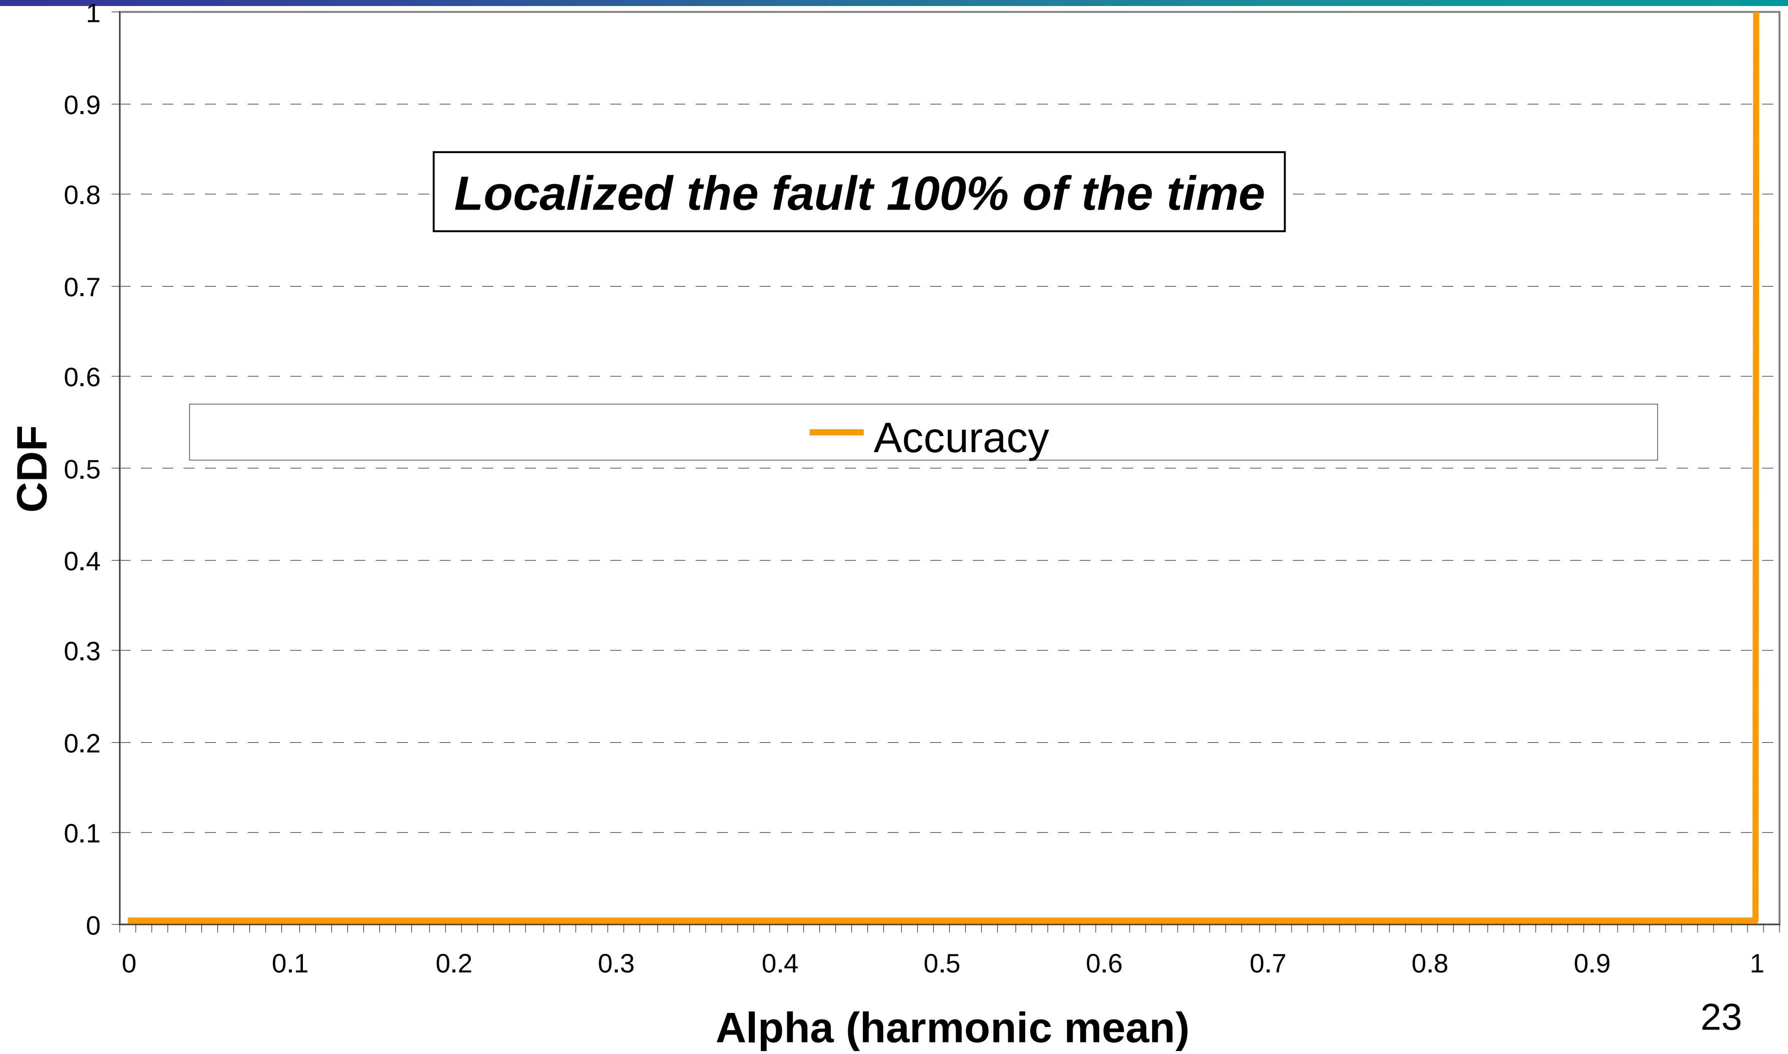

Privacy Results for All Single Failures

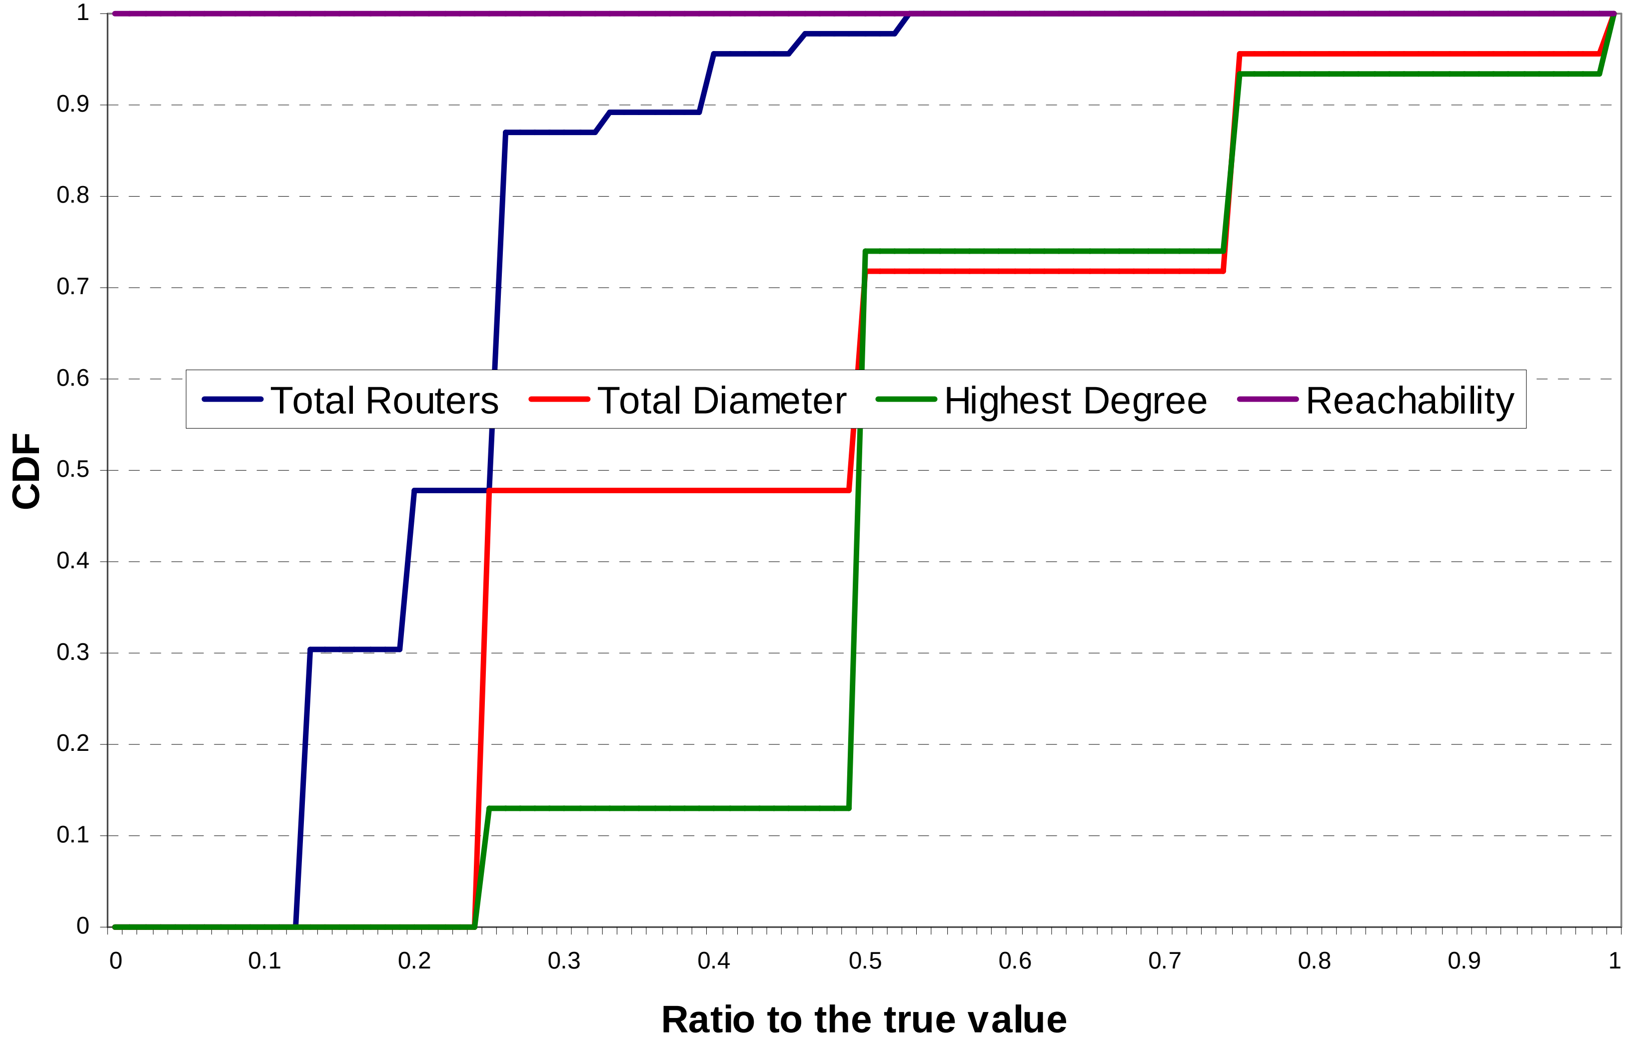

Adding Scalability Results

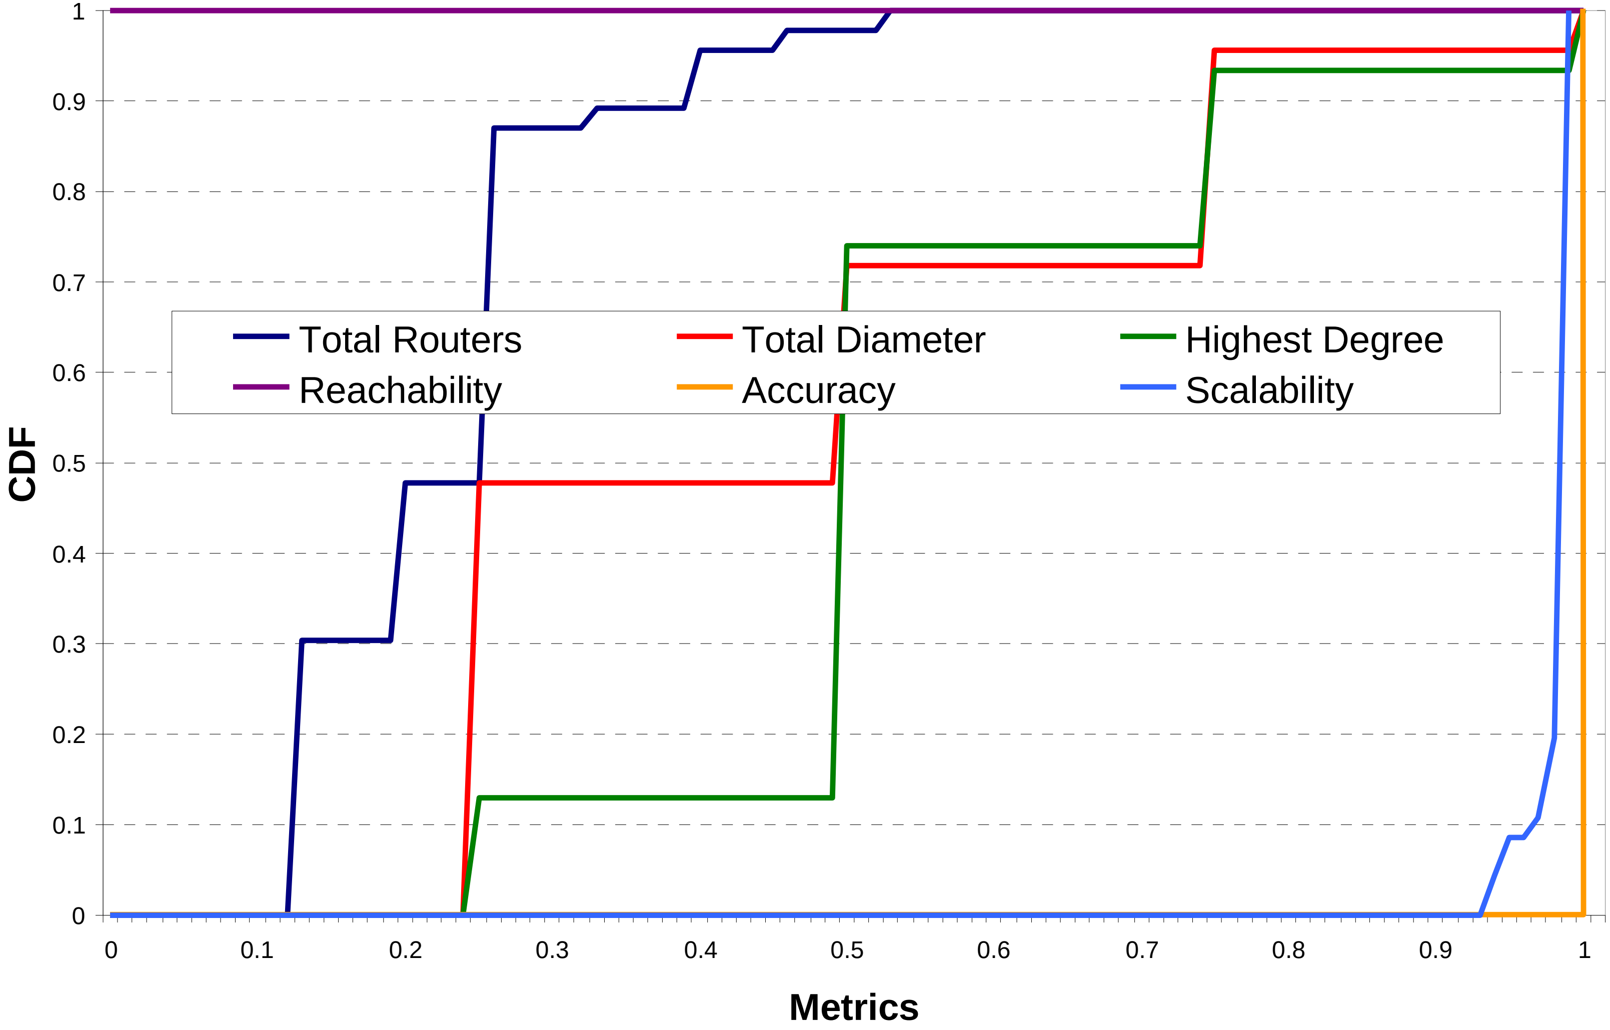

Summary and Future Work

A general framework for reasoning and discussing cross-domain fault localization

Initial results demonstrating utility of the framework

Future Work:

- Validation of the generality for the approach
- Impacts of observation and model errors must be determined
- Privacy protection given a series of digests from the same domain
- Digest-sharing format and strategies must be explored

Thank you!

Questions?