

Network Measurement Research in SPYCE

Kostas G. Anagnostakis, Michael Greenwald
 CIS Department - University of Pennsylvania
 anagnost@dsl.cis.upenn.edu

Raphael S. Ryger
 CS Department - Yale University
 ryger@cs.yale.edu


Partially supported by the CIP/SW URI "Software Quality and Infrastructure Protection for Diffuse Computing" through ONR grant N00014-01-1-0795.

Highlights

- new network tomography method and tool (**cing**) for directly measuring network-internal delays that does not require deployment of a wide-area measurement infrastructure
 - Comparison with indirect methods
 - Performance 2002 paper (best paper award)
 - Feasibility study
 - INFOCOM 2003 paper
 - Use of tool to evaluate network models
 - MASCOTS 2003 paper
 - Hybrid direct-indirect method to improve coverage
 - Work in progress
 - Use of tool for overlay adaptation and security
 - Work in progress - see demo

Direct method (cing tool)

Problem: Suppose we want to measure queuing delay between x and y on path $P = s \Rightarrow d$



Method: send m pairs of ICMP Timestamp probes to routers on head and tail of link:

$$\text{One observation: } \delta = t_y - t_x = t_{\text{queuing}} + t_{\text{link}} + O_{x,y}$$

$$\text{Two observations: } \delta_i, \delta_j : \delta_j - \delta_i = \Delta t_{\text{queuing}}$$


$$\text{After } m \text{ observations: } \delta_{\min} = \min(\delta_i)_{i=1..m}$$

$$\delta_j^i = \delta_j - \delta_{\min}^i : \text{delay estimate for } x \Rightarrow y$$

Cing: feasibility

- 96% of nodes respond to Timestamp queries
- Irregular routing limits choice of nodes


Example: Path structure, Penn to Sprintlabs



Corresponding feasible measurement partitions, Penn to Sprintlabs


Cing: feasibility (2)

- Data:** ~10k paths from 5 different sources
- Metric:** fraction of nodes usable for tomography
- Results:** ~50% nodes are usable, more difficult as distance from source increases, better when probing from multiple sources



Cing: accuracy

- Simulation study** shows direct method is more accurate on average than indirect, but much better in worst case



- Real-world measurements** show direct is more accurate than TTL-based (e.g. traceroute-like) methods

Cing applications: network modeling

Network protocol designers assume one bottleneck in path


Does model match reality?

Not always: measurement study of 50k paths shows 10% of connections go through multiple bottlenecks

Does it matter?

Yes: simulation experiment shows "congestion collapse" scenario, demonstrating tension between fairness and efficiency not foreseen by existing network models

Implications: Need to rethink simulation models, develop comprehensive benchmarking suite for protocol evaluation




Cing applications: overlay adaptation

- One can override IP-level routing by sending packets through overlay nodes
 - Useful for routing around broken or poorly-performing parts of the net (congestion, DoS, ...)
- cing** can be used for fault diagnosis to improve overlay adaptation
 - More accurate and timely fault diagnosis
 - Scalable: less probes, time to locate fault compared to end-to-end/all-to-all probing
 - Greater local autonomy: no need for cooperation, deployment of global diagnosis infrastructure

Work in progress

- Extended, hybrid direct-indirect technique**
 - Combine ICMP Timestamps, TTL-expired RTT probes and indirect inferring methods when pure direct approach cannot be used



- Preliminary results show 97% coverage

- Integration with network services**
 - Use **cing** for diagnosis and adaptation

Impact: Critical Infrastructure

- Networks are CI, and underlie other CIs
- Diagnosis, reconfiguration, recovery all important to survivable network CI
- cing** provides autonomy and accuracy in network diagnosis
- Useful, for example, in detecting existence and location of DoS and DDoS: diagnosis

Impact: Software Quality

- cing** is a novel diffuse computing approach to network measurement
 - Many cooperating components
 - Software part of SPYCElab toolkit
 - Decentralized, loosely coupled, and effective
- Approach is highly fault- and error-tolerant
- Example of "logical sandbox" characteristic of diffuse computing approach in SPYCE

Impact: DoD

- DoD has far-flung, complex networks
 - Rapid deployment, diagnosis, reconfiguration
- cing**-like approaches provide highly decentralized solutions even with partial information
- Crucial for Network Centric Operations and Warfare, as per JV 2010 /JV 2020