King: Estimating Latency Between Arbitrary Internet Hosts

Krishna P. Gummadi, Stefan Saroiu and Steven D. Gribble

Motivation and Related Work

King: A technique (similar to Ping) to estimate latency between any pair of Internet hosts, from any other Internet host.

Applications:

- Scaling wide-area measurement studies such as Detour.
- Building topologically sensitive overlay networks for Distributed P2P systems (e.g., Gnutella, Chord, Kazaa).

Problems with existing techniques like IDMaps and GNP:

- Deployability: Additional infrastructure needed for IDMaps. End hosts estimate and share their coordinates in GNP.
- Accuracy of estimates: Errors due to offline extrapolation of latencies assuming that Internet topology resembles a metric space.

How King works

Observations:

- A large number of IP hosts are topologically close to their authoritative name servers.
- Latency between any two name servers can be *accurately* measured by using Recursive DNS queries.
- Latency between end hosts can be *approximated* as the latency between their name servers.

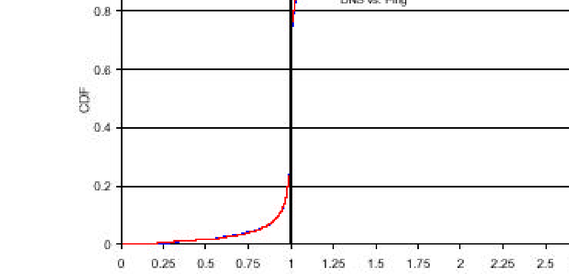

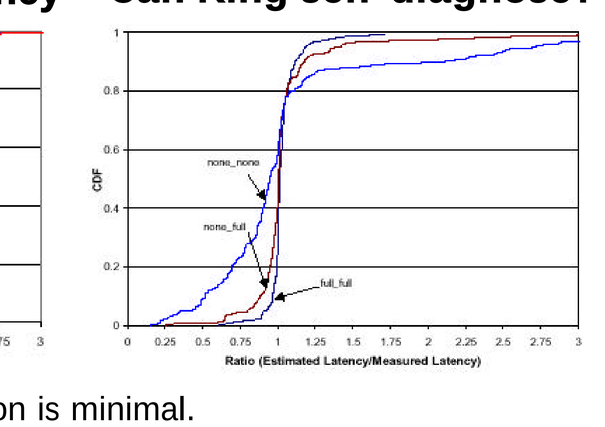

Evaluation

Accuracy of King

Estimates to P2P clients

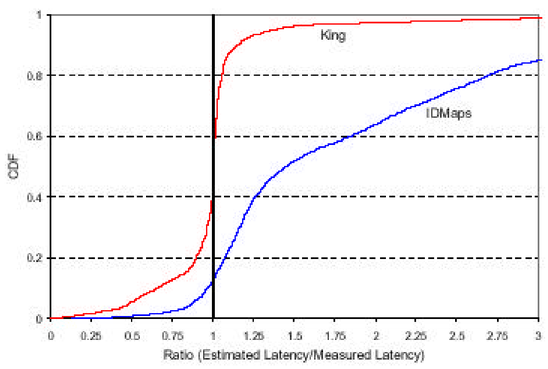

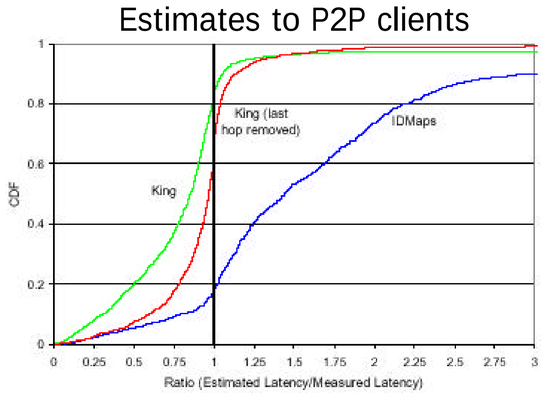

- King estimates are significantly better than those by IDMaps.
- King tends to underestimate latencies for hosts with large last hop latency.

Does King preserve order?

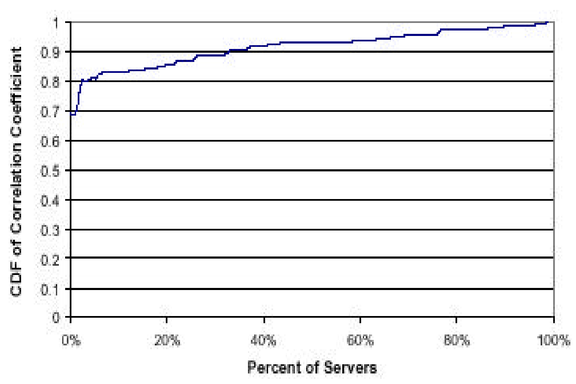

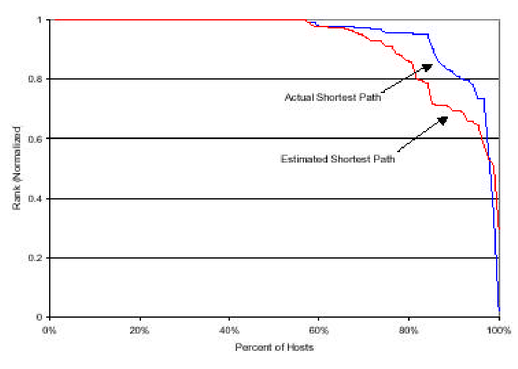

- Very high correlation between the orderings of estimated and true latencies.
- King selects a close server in a vast majority of cases.

Causes of error in King estimates

Are end hosts close to their authoritative name servers?

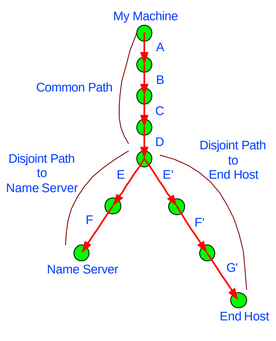

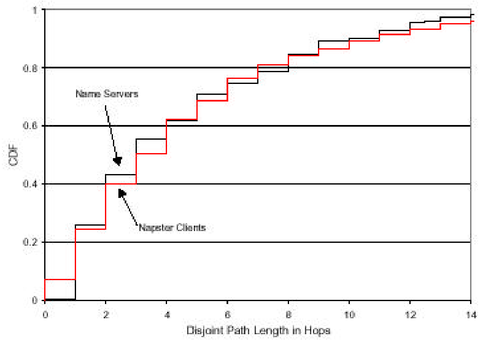

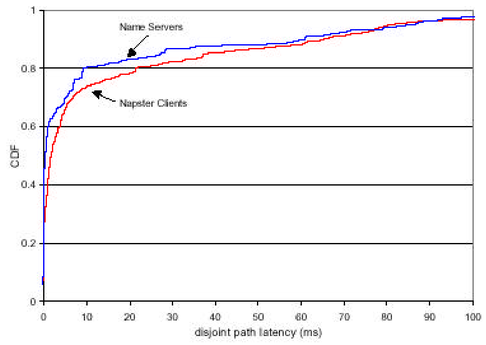

- The divergent path latencies show that end hosts are close to their name servers in a majority of scenarios.
- Earlier studies considered divergent path hop counts only.

DNS query resolution latency

- Latency due to DNS query resolution is minimal.
- King can predict accuracy of its estimates matching the domain names and IP addresses of end hosts and their name servers.

Can King self-diagnose?

Potential Limitations of King

At least one end host must have a name server that resolves non-local recursive queries

- ~75% of name servers in the Internet support such queries.

How do we choose a right server when there are multiple authoritative name servers?

- An Improved version of King can pick the right ones.

End hosts using modems have a large last hop latency

- But ordering of latencies is unaffected by ignoring last hop.

Conclusions and Future Work

Our evaluation indicates:

- Latency estimates by King are fairly accurate.
- King estimates are very good at preserving order.
- King can rank its own estimates on their likelihood to be inaccurate.

Use King in wide-area measurement projects

- Evaluating efficiency of Akamai's server selection.
- Generating proximity maps of P2P overlays such as Gnutella.

A detailed evaluation of the King technique is to appear in SIGCOMM IMW 2002.
http://www.cs.washington.edu/homes/gummadi/king