



King : Estimating latency between arbitrary Internet end hosts


Krishna Gummadi, Stefan Saroiu

Steven D. Gribble

University of Washington

A Question:

What can we do with a tool that could estimate latency between any two arbitrary hosts accurately?



C can measure latency between A and B



Potential uses of such a tool

- ◆ Building topologically sensitive overlays
- ◆ Selecting a close replica server
- ◆ Scaling wide-area measurement studies involving latency estimation
 - Detour, IP2Geo study etc.,
 - current state of the art techniques allow at most a few hundred end hosts to be measured



Current state of the art

- ◆ Use techniques like IDMaps and GNP
 - inaccuracy in estimates: We need a tool that can *measure* latency rather than *compute* it
 - issues with deployment: IDMaps requires additional infrastructure;
- ◆ Share pre-collected data sets
 - e.g., Skitter data from CAIDA
- ◆ Use shared measurement infrastructure
 - e.g., trace-route servers, PlanetLab, NIMI



King: a new measurement tool

- ◆ Estimate latency between arbitrary end hosts
- ◆ Requires no additional infrastructure
 - leverages existing DNS infrastructure enabling a large fraction of Internet hosts to be measured
- ◆ Provides highly accurate latency estimates
- ◆ Fast and light-weight
 - requires only a few DNS queries per estimate

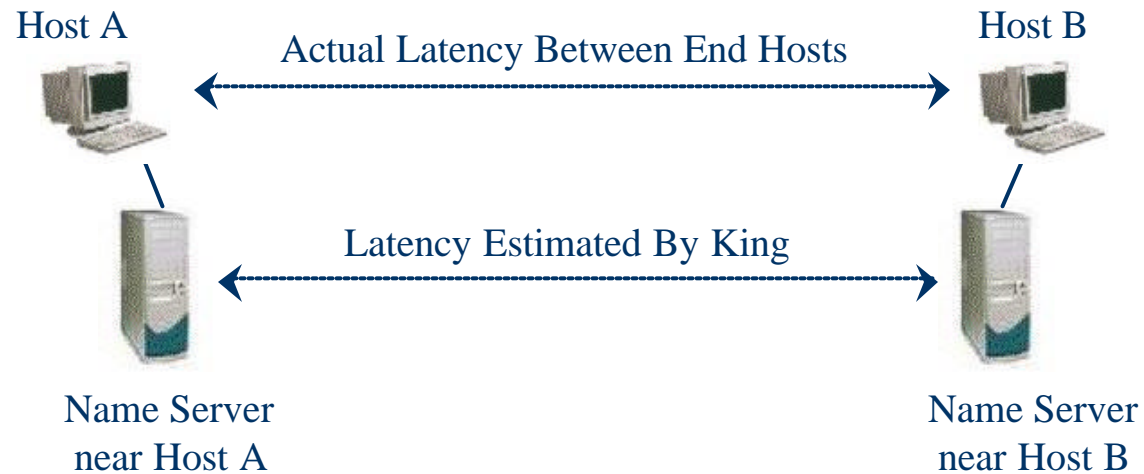
We hope that King will be used in many unanticipated ways like in the case of Ping and Traceroute



Outline

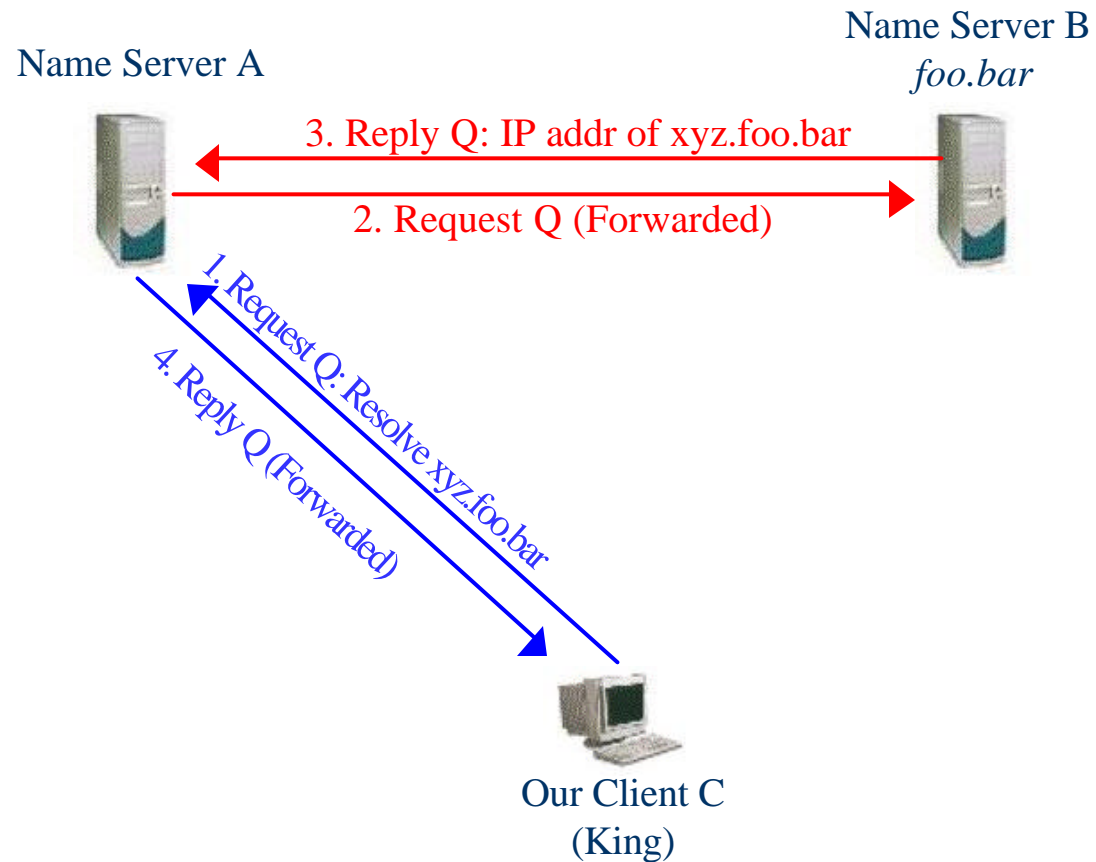
- ◆ Motivation
- ◆ How King works
- ◆ Evaluation of King
- ◆ Conclusions

How King Works: The Basic Idea



- ◆ Challenge 1: How to find name servers that are close to end hosts
- ◆ Challenge 2: How to estimate latency between two name servers

Challenge 2: How do we estimate the latency between name servers?





Success of Recursive DNS

- ◆ For King to work, name servers must support recursive queries
 - in a large random sample, > 75% of name servers supported recursion
 - translates to > 90% success rate given a pair
 - as we can measure from A->B, or B->A



Challenge 1: How to find DNS servers nearby the end hosts

- ◆ Assumption: Authoritative name servers for the host are closeby (topologically and geographically)
- ◆ This assumption may not always hold, but our evaluation shows that it is true in general
 - e.g., AOL is an exception
- ◆ To find an authoritative name server
 - given host name, use forward name resolution
 - given host IP, use reverse lookup in in-addr.arpa domain



Selecting a close name server

- ◆ When multiple authoritative name servers are deployed, how do we choose a close one?
 - select the server with longest matching DNS suffix and IP prefix with end host



Outline

- ◆ Motivation
- ◆ How King works
- ◆ Evaluation of King
- ◆ Conclusions



Evaluation of King

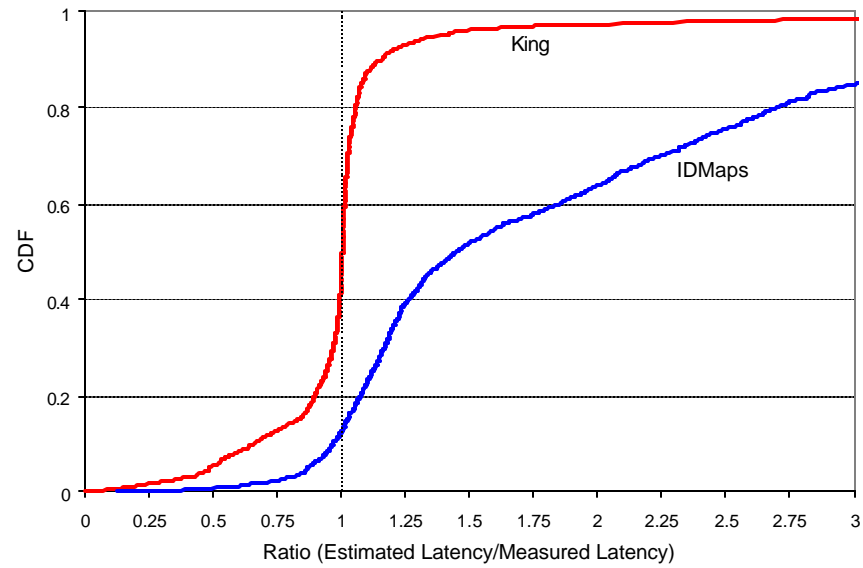
- ◆ How accurate is King?
- ◆ What are the causes of inaccuracy?
- ◆ Can King identify its own inaccurate estimates?
- ◆ Does King preserve order among its estimates?



Accuracy of King

- ◆ Compare the accuracy of King with IDMaps
- ◆ Methodology
 - Measure true latency between 50 public Traceroute servers and 50 end hosts using **Traceroute**
 - Estimate latency between the same endpoints using **King** and **IDMaps**
 - Compare estimated latency with measured latency
 - Metric used : $\frac{\textit{Estimated Latency}}{\textit{Measured Latency}}$

Accuracy of King



- ◆ King is far more accurate than IDMaps
- ◆ King tends to under-estimate latencies
 - typically, name servers have higher BW and lower latency last hop links than end hosts



Evaluation of King

- ◆ How accurate is King?
- ◆ What are the causes of inaccuracy?
- ◆ Can King identify its own inaccurate estimates?
- ◆ Does King preserve order among its estimates?



Causes of Inaccuracy in King

- ◆ Authoritative name servers may not be close to end hosts
- ◆ Latency estimation between the name servers might be inaccurate
 - application level latency at DNS servers to resolve the query



Are authoritative name servers close to their end hosts?

- ◆ In a random sample, 70-80% of end hosts and their name servers are separated by less than 10-20 msec
- ◆ Our conclusion contradicts earlier studies !!
- ◆ Possible explanations:
 - We looked at more metrics
 - divergent path hop count – a misleading metric used primarily in other studies; divergent path latency – tells a different story
 - Unknown bias in our random samples



Application level latency for DNS servers

- ◆ Methodology:
 - selected a large number sample of name servers
 - measured latency to servers using **Ping** and **DNSPing** (iterative DNS query) over time
 - Query resolution latency = **DNSPing** – **Ping**
- ◆ Application level latency negligible
 - Implication: King estimates between name servers are very highly accurate



Evaluation of King

- ◆ How accurate is King?
- ◆ What are the causes of inaccuracy?
- ◆ Can King identify its own inaccurate estimates?
- ◆ Does King preserve order among its estimates?



Can King identify its own inaccurate estimates?

- ◆ Primary cause of error in King
 - authoritative name servers far from their end host
- ◆ Simple heuristics based on the lengths of DNS suffix and IP prefix match work well



Evaluation of King

- ◆ How accurate is King?
- ◆ What are the causes of inaccuracy?
- ◆ Can King identify its own inaccurate estimates?
- ◆ Does King preserve order among its estimates?



Does King preserve order among its estimates?

- ◆ Sometimes preserving order among estimates is more important than their accuracy
 - Applications like server selection
- ◆ King does very well at preserving order among its estimates
 - very high correlation coefficient (>0.8) between the orderings of estimated and true latencies
 - large latency last hops do not effect order



Summary of evaluation

- ◆ King is far more accurate than IDMaps
 - King errs more when it under-estimates due to large last hop latencies of end hosts
 - estimates the accuracy of estimates between name servers is even higher
- ◆ The primary cause of error is the authoritative name servers that are far from their end hosts
 - King uses heuristics to identify such errors
- ◆ King preserves excellent order among its estimates



Validating King's utility for wide-area measurement studies

- ◆ The Detour study
 - showed that default routes are inefficient, and alternate routes can have better latency.
 - they were limited to 35x35 data points
- ◆ We repeated study using King
 - we gathered 193x193 data points
 - The data points were name servers chosen using King's self-evaluation heuristics
 - it took less than 4-5 hours using a single machine
 - our results were consistent with those from earlier study



Conclusions

- ◆ We presented King; a new measurement tool that
 - can estimate latency between arbitrary Internet end hosts
 - does not require any additional infrastructure as it leverages existing DNS infrastructure
 - fast and light-weight
- ◆ Our evaluation of King confirms that
 - it is accurate
 - it preserves order among its estimates
- ◆ We showed that King can be used in scaling wide-area measurement studies

Questions ?

For more information visit:

www.cs.washington.edu/homes/gummadi/king