King: Estimating Latency between Arbitrary Internet End Hosts *

Krishna P. Gummadi, Stefan Saroiu, Steven D. Gribble
Department of Computer Science & Engineering
University of Washington Seattle, WA, USA, 98195-2350
{gummadi,tzoompy,gribble}@cs.washington.edu

The ability to estimate network latencies accurately between arbitrary Internet end hosts would enable wide-area measurement studies and applications, such as investigating routing path inefficiencies on a wide-scale [1] or constructing topologically sensitive overlay networks. In this poster we present King [2], a tool that accurately, quickly and cheaply estimates the latency between arbitrary end hosts by using existing DNS infrastructure in a novel way. Our approach is inspired by tools like Sting [3] and T-BIT [4], which show that it is possible to use existing protocols in unanticipated ways to obtain results that were previously intractable.

Our technique relies on two observations. First, given a pair of end hosts to be measured, in most cases it is possible for King to find their authoritative DNS name servers that are topologically close to the end hosts. Second, given a pair of DNS name servers, King can accurately estimate the latency between them using recursive DNS queries. Thus, King is able to use the measured latency between the name servers as an estimate of the latency between the end hosts. Although this estimate will inevitably suffer from inaccuracies, our extensive evaluation demonstrates that this estimation error is small (less than 20% error in over three-quarters of generated estimates).

Compared to previous approaches like IDMaps [5] and GNP [6], King has two primary advantages (1) deployability and (2) accuracy.

1. *Deployability:* Unlike IDMaps, King does not require the deployment of additional infrastructure, and unlike GNP, King does not require end hosts to agree upon a single set of landmarks or share the coordinates with other end hosts. Because King uses existing DNS infrastructure, it scales naturally in terms of the number hosts that could be measured. Our evaluation indicates that King can be used to estimate latencies between arbitrary IP hosts more than 75% of time in the Internet today.

2. Accuracy: Unlike IDMaps and GNP the accuracy of King is not affected by the placement of measurement points (tracers or landmarks) in the Internet. King's estimates are based on direct online measurements rather than extrapolation from latencies of a subset of paths measured offline. Our evaluation shows that the error in King latency estimates between end hosts is significantly smaller than that of IDMaps, while the error in its estimates between name servers is negligible. Further, we show that King can identify those of its own estimates that are likely to be inaccurate. These characteristics make King an ideal tool for a number of wide-area measurement studies that involve measuring latencies accurately between a large number of geographically distant end hosts.

King is also fast and lightweight requiring the generation of only a few packets to produce an estimate. Our evaluation also demonstrates that (a) King preserves order among its latency estimates; a feature useful in identifying a close server from a group of server replicas and (b) authoritative name servers are located topologically close to their end hosts in a majority of scenarios (i.e., latencies as measured from a far away host to both the end host and its name server are very similar); this explains the high accuracy of King. Finally, we test the feasibility of using King for measurement studies by performing Detour study [1] on a larger scale. The results we obtained are consistent with the earlier Detour study. We were able to obtain our results with far greater ease as, unlike the previous study, we were not restricted to the publicly available traceroute servers and their usage policies. We also believe that the simplicity of our King tool will enable researchers to use it in myriads of unanticipated ways like the more popular Ping tool.

*This work was supported by National Science Foundation under ITR grant No. CCR-0121341. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

References

[1] "The Detour web page." www.cs.washington.edu/research/networking/detour/.

[2] "The King web page." www.cs.washington.edu/homes/gummadi/king.

[3] "The Sting web page." www.cs.washington.edu/homes/savage/sting/.

[4] "The TBIT web page." www.icir.org/tbit/.

[5] "The IDMaps web page." idmaps.eecs.umich.edu/.

[6] "The GNP web page." www-2.cs.cmu.edu/~eugeneng/research/gnp/.