ARCHIPELAGO
Internet Measurement Infrastructure
http://www.caida.org/projects/ark/

caida
www.caida.org

Why should my network host an Ark node?

When your network hosts a measurement node that participates in CAIDA's Archipelago (Ark) infrastructure[1], it broadens the view of the global Internet for the network research community. Network researchers use CAIDA topology data to conceive, develop, and test their models and methods. Participating networks agree to our Memorandum of Cooperation (MoC)[2] and donate rack space, power, and cooling to host a 1u (virtual) server dedicated to Ark measurement.

Once deployed, the Ark node conducts continuous measurements of the routed IPv4 (and IPv6 where available) address space. Ark aggregates all of the resulting data on a server located at the San Diego Supercomputer Center on the campus of University of California, San Diego.

Each additional node contributes to and increases the completeness and accuracy of data representing the topological structure of the Internet core.

Based on the local view of the network topology that each Ark node collects, CAIDA generates regularly updates reports to illustrate the dynamics of connectivity and performance from the measurement source hosts to the rest of the routed IPv4 Internet.[3] These reports support a variety of research, engineering, and operational interests in promoting a robust, scalable global Internet infrastructure.

Other experiments running on the Ark measurement platform have investigated relative effectiveness of different path probing methods (IMC'08), evaluating the efficacy of deployed Internet source address validation filtering (IMC'09), measuring the impact of certain causes of missing hops in traceroute paths (CCR'11), and a comparison of public and commercial geolocation databases (NMMC'11). This range of scientific experiments has successfully demonstrated our vision of a metaphorical distributed measurement ``operating system'' to support empirical Internet science.

Figure 1: AS Dispersion by AS hop

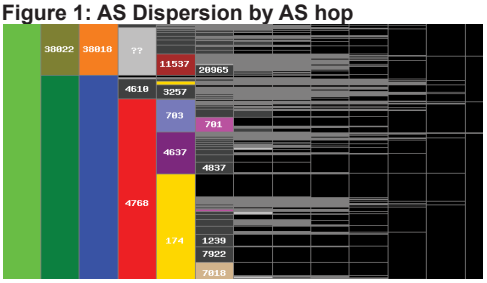

This graph shows a view of AS peering relationships observed near the monitor, useful for a hosting organization to compare against a known list of their upstream transit providers. Typically, the leftmost column will be all one AS (i.e., the local network that the monitor is in), with the second or third hop starting to split into different ASes as probes go to different destinations.

Figure 2: RTT density vs. geographical distance

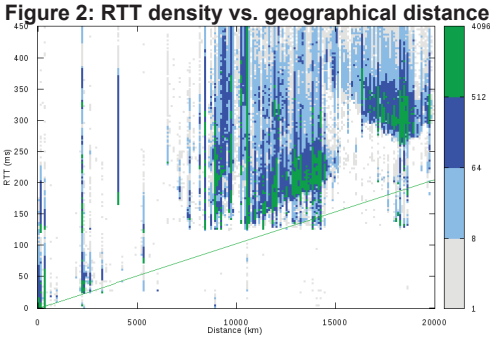

This graph shows the density of round-trip times (RTTs) as a function of geographical distance from the probe source. The RTTs and distances are binned (to every 3 ms and 100 km); dot color reflects the number of hops seen a given (RTT,distance) combination. The straight green line represents the theoretical best RTT (given the speed of light in fiber) for a distance.

Figure 3: RTT quartiles vs. hop distance

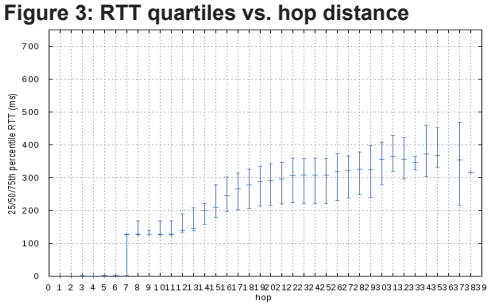

This graph shows the quartiles (25th, 50th, and 75th percentiles) of round-trip (RTTs) found at each hop in a trace. These distributions can provide insight into topological issues that contribute to high network latency from a given host to a broad cross-section of the Internet.

Figure 4: Median RTT per country

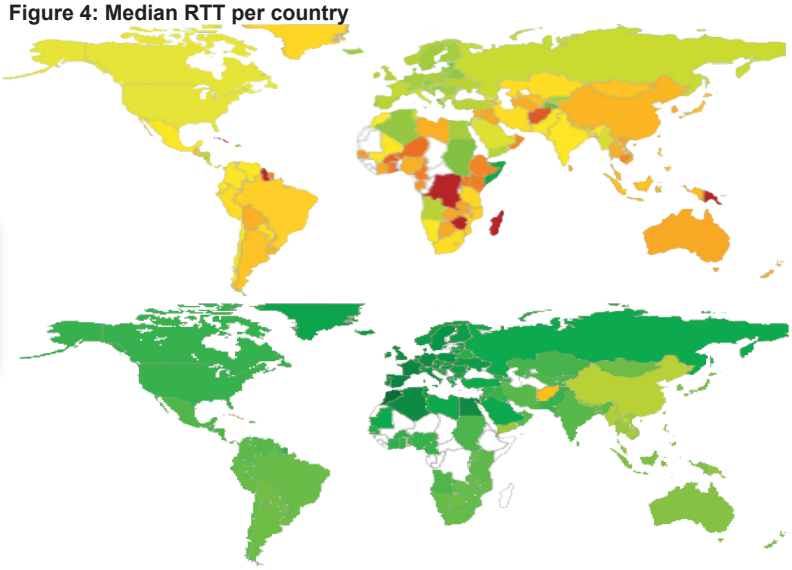

These graphs show the median round-trip times (RTTs) from the probing monitor to destinations in various countries (and US states). These two snapshots were taken from the Morrocan monitor prior to (Sept 2010, top) and then following (Oct. 2011, bottom) a new fiber installed down the West coast of Africa.

If interested, please contact
ark-info@caida.org

References:

1. Archipelago Measurement Infrastructure
http://www.caida.org/projects/ark/

2. Archipelago Memorandum of Cooperation (MOC)
http://www.caida.org/projects/ark/moc/

3. Archipelago Monitor Statistics
http://www.caida.org/projects/ark/statistics/

Ark statistics per monitor:

* Total number of responding ASes
* Total number of responding Prefixes
* Median RTT per country and US state (Fig. 4)
* AS Path Dispersion Graphs (Fig. 1)
* IP Path Dispersion Graph
* Path Length Distributions
* RTT Distribution Graphs (Fig. 3)
* RTT vs. Distance Graph (Fig. 2)

UC San Diego SDSC
San Diego Supercomputer Center, University of California San Diego
9500 Gilman Drive, mc0505, La Jolla, CA 92093-0505 | 858-534-5000 | www.caida.org