A Longitudinal and Cross-Dataset Study of Internet Latency and Path Stability

Mosharaf Chowdhury
Rachit Agarwal
Vyas Sekar
Ion Stoica

Electrical Engineering and Computer Sciences
University of California at Berkeley

Technical Report No. UCB/EECS-2014-172
http://www.eecs.berkeley.edu/Pubs/TechRpts/2014/EECS-2014-172.html

October 11, 2014

A Longitudinal and Cross-Dataset Study of Internet Latency and Path Stability

Mosharaf Chowdhury, UC Berkeley
Rachit Agarwal, UC Berkeley
Vyas Sekar, Carnegie Mellon University
Ion Stoica, UC Berkeley

ABSTRACT

We present a retrospective and longitudinal study of Internet latency and path stability using three large-scale `traceroute` datasets collected over several years: Ark and iPlane from 2008 to 2013 and a proprietary CDN's `traceroute` dataset spanning 2012 and 2013. Using these different "lenses", we revisit classical properties of Internet paths such as end-to-end latency, stability, and of routing graph structure. Iterative data analysis at this scale is challenging given the idiosyncrasies of different collection tools, measurement noise, and the diverse analysis we desire. To this end, we leverage recent big-data techniques to develop a scalable data analysis toolkit, Hummus, that enables rapid and iterative analysis on large `traceroute` measurement datasets. Our key findings are: (1) overall latency seems to be decreasing; (2) some geographical regions still have poor latency; (3) route stability (prevalence and persistence) is increasing; and (4) we observe a mixture of effects in the routing graph structure with high-degree ASes rapidly increasing in degree and lower-degree ASes forming denser "communities".

1 Introduction

Understanding Internet path properties is a fundamental requirement for many applications including server placement (e.g., [37]) and selection (e.g., [26]), fault detection and diagnosis (e.g, [16]), and analyzing key inefficiencies in routing protocols (e.g., [13]). Measuring Internet paths is almost as old as the discipline of Internet measurement itself [35]. Starting from the seminal studies by Vern Paxson [36], there have been several efforts to analyze latency (e.g., [19,31,32]), route predictability (e.g., [12,38]), as well as numerous tools and datasets (e.g., [1,7,30]).

Our work follows in the spirit of this rich body of work in routing measurement. Our main contribution here is to systematically analyze key characteristics of the Internet paths using multiple large-scale `traceroute` datasets collected over several years. Specifically we use a six-year dataset (2008-2013) from iPlane [30] that uses over 1000 vantage points, another six-year (2008-2013) dataset from Ark that uses over 80 vantage points [1], and a proprietary dataset from a large CDN with `traceroutes` from over 1800 vantage points over 2012-2013. Together, these datasets provide a panoramic and longitudinal view of Internet routing behavior to more than 25000 destination ASes spread across 200 countries.

Even though our work does not provide new active measurement techniques or new datasets, we believe that there is value in this retrospective analysis on several fronts. First, it provides a historical and longitudinal perspective of Internet path properties that are surprisingly lacking in the measurement community today. Second, it can help us revisit and reappraise classical assumptions about path latency and stability used in designing Internet-scale systems. Third, such a cross-dataset analysis can shed light on potential "blind spots" and potential biases in our understanding even with large-scale datasets.

Specifically, we analyze the following dimensions of Internet path properties:

- Has the end-to-end *latency*[1] changed over this measurement period? Have specific geographic regions improved more than others or have some regressed to worse connectivity? (§4)
- Do classical assumptions about *route stability* in terms of route persistence and route prevalence [36] still hold? How have these evolved over this time period? (§5)
- Has the *routing graph structure* changed significantly over the last several years? (§6)
- Do different "lenses" provide complementary, consistent, or contradictory views into the above routing characteristics? (§4–6)

Performing such a study over diverse and large-scale datasets raises a number of practical scalability challenges in terms of preprocessing (e.g., cleaning missing traces and converting traces to AS-granularity) and extracting meaningful information — the combined raw `traceroute` data in our study amounts to roughly 1 terabyte of uncompressed data. Conventional data analysis techniques (e.g., custom scripts) are simply not scalable or sustainable for such iterative analysis. Our contribution here is a systematic data analysis toolkit, Hummus, implemented on top of an in-memory distributed data processing system called Apache Spark [42]. Hummus enables rapid and iterative analysis of large-scale `traceroute` measurement datasets that would otherwise be infeasible or tedious. For instance, several analysis tasks performed in this paper that would have otherwise taken several hours using conventional techniques can be com-

[1]Since `traceroute` datasets allow us to only measure round-trip times (RTT) between any source-destination pair, we use RTT as the metric for latency and use the terms RTT and latency interchangeably in this paper.

pleted within six minutes using Hummus. We cannot stress the value of such a capability enough—it simplified our analysis workflow and enabled iterative analysis, typical of large-scale measurement studies, to extract interesting information. We have made our code public[2] for future measurement studies to benefit from the scalability and rapid iterative analysis enabled by Hummus.

Our key findings are:

- Overall, most countries are *improving* in terms of latency from different vantage points. Countries in Africa, which currently have the highest latencies, are also the ones showing the most improvements over the last six years. Some countries are surprisingly regressing in the Ark dataset; we identify them as anomalies and point out the likely causes.
- We observe that both routing prevalence (i.e., how frequently the dominant route is used) and persistence (i.e., how stable are routes across consecutive measurements) are *increasing* both at AS- and city-granularity. However, we find the absolute numbers to be significantly lower than those observed in classical studies. In general, we observe that prevalence/persistence are inversely correlated with the in-degrees of destination ASes.
- The AS-granularity Internet *routing* graph is getting denser and more clustered, which is correlated with the improvements in latencies, route lengths, and stability. Following a "rich-getting-richer" phenomenon, the high-degree ASes are connecting to increasingly larger fraction of ASes, and lower degree ASes are forming more tightly-knit communities.
- We do find different datasets to be largely in *qualitative* agreement and having good coverage. But we observe several subtle differences (e.g., latency anomalies in Ark) and some non-trivial coverage gaps (e.g., many countries in Africa and Asia often do not have a sufficient number of measurements).

These results have important implications both for measurement research as well as the design of Internet-scale systems. For instance, we find that our visibility into significant geographical regions (e.g., Africa, Asia) is quite limited, motivating the need for more careful selection of targets. Similarly, we find that most stable vantage points are in the US and Europe providing a very US- and Euro-centric view of Internet connectivity. While this is very valuable, it does indicate where our future gaps will be. Third, we find that there is significant potential for latency improvement for large portions of the Internet. Given the critical role that latency plays in user quality-of-experience (e.g., [2]), content providers will be well advised to expand their vantage point presence to these emerging regions [10].

[2] https://github.com/mosharaf/hummus

	Ark	iPlane	LargeCDN
Period	2008-2013	2008-2013	2012/2013
Sampling	1 cycle/month	1 day/month	No
`traceroute`	Paris [7]	Normal	Normal
Traces	2.52 billion	1.85 billion	0.83 billion
Valid Traces	187 million	449 million	104 million
Vantage IPs	80	1051	1895
Vantage Countries	35	40	89
Vantage ASes	72	226	815
Dst Countries	219	226	222
Dst ASes	36238	27243	25876

Table 1: Dataset details.

2 Datasets and Methodology

In this section, we describe the `traceroute` datasets used in our study (§2.1) as well as the methodology we follow to preprocess and clean the datasets, followed by a discussion of potential biases and limitations of this study (§2.2).

2.1 Datasets

We use three `traceroute` datasets from Ark [1], iPlane [30], and a large commercial CDN (LargeCDN). The Ark and iPlane data spans six years from 2008 to 2013, while the LargeCDN dataset spans 2012 and 2013. Table 1 provides a quick summary of each dataset. `traceroutes` from all three datasets included only IPv4 addresses. While all three datasets cover almost all of the geographical regions and ASes, both Ark and iPlane have their vantage points located in a significantly smaller number of locations than LargeCDN. Each dataset is collected by (periodically) issuing `traceroutes` from a fixed set of vantage points to some destinations based on dataset-specific criteria. For completeness, we briefly summarize how these individual `traceroutes` were collected and refer readers to the above references for more details.

- In iPlane, each vantage point daily probes a selective list of IP prefixes (120K prefixes out of 450K operational prefixes [25]) concentrated on the core Internet.
- Ark ensures that in each "probing cycle" `traceroutes` are sent to randomly selected destinations in every routable /24 IP prefix.
- For the proprietary LargeCDN dataset, the destination selection process information is unavailable; however, we observed significantly higher "coverage" over the IPv4 address space relative to Ark and iPlane.

Due to the massive scale of these datasets and practical limitations in downloading the entire raw data, we subsample the traces for Ark and iPlane. Specifically, we downloaded one cycle per month for Ark and iPlane — for Ark, the cycle length is two to three days, whereas it is one day for iPlane. For LargeCDN, we use the entire dataset without sampling. Since only Ark and iPlane datasets span six years from 2008 to 2013, we rely on them for longitudinal analysis. We primarily use the LargeCDN dataset to compare the observations

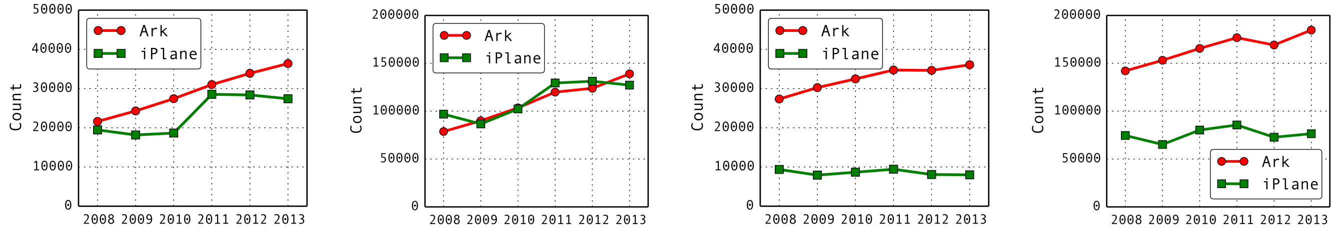

(a) Number of ASes (b) Number of AS-AS links (c) Number of cities (d) Number of city-city links

Figure 1: Number of observed ASes (cities) and AS-AS (city-city) links in Ark and iPlane over time.

from traceroutes collected by Ark and iPlane in 2012/2013, especially because of LargeCDN's larger spatial coverage.

Coverage: Figure 1 shows that the coverage of Ark has consistently been increasing over the years both in terms of the number of ASes observed and geographical span. For iPlane, there was a sudden increase in 2011 in number of ASes observed, but the geographical span is essentially consistent across the years. That said, we do observe that there are non-trivial blind spots in the Ark and iPlane datasets in terms of country-level coverage relative to the LargeCDN dataset, suggesting the need to revisit some of the destination selection strategy. More importantly, the blind spots occur in precisely the regions of the world where connectivity tends to be poor and where there is a huge scope for improvements (§4.1).

2.2 Preprocessing and potential biases

Many traceroutes have some "failure" where one or more of the intermediate IP hops do not respond (Table 1). Our focus in this paper is on analyzing end-to-end latency and stability; hence, we conservatively ignore such traceroutes.[3] After this preprocessing step, we had roughly 740 million traceroutes where all hops responded.

Our goal in this paper is to analyze the routing properties (latency, stability, structure) for different *source-destination* (SrcDst) pairs at different granularities: AS, country, and city. To perform geographical and AS-level analysis, we map the observed IPv4 addresses to their corresponding geolocations (city, country) and ASes. For this mapping, we used a proprietary conversion table obtained from LargeCDN because it had significantly higher coverage relative to other public sources such as Maxmind [4]. We do so after confirming that there is indeed a very high match rate (88.51%) with the MaxMind dataset [4] for the IP addresses common across the two conversion tables. For our observations to be statistically meaningful, however, we need to ensure that each SrcDst pair (at AS-, country-, or city-granularity) has a sufficiently large number of observations. Unless mentioned otherwise, we ensure that each SrcDst pair has at least 100 observations.

Potential biases of traceroute-based analysis: There are three known biases of using traceroute measurements that have been pointed out in prior work. For completeness, we highlight these and also describe why our analysis is not impacted by these biases. First, traceroutes may be inaccurate if there is some inherent load balancing or multi-path routing in place (e.g., [12]). Since our focus is mostly on the city-, AS-, and country-level paths and on long-term stability patterns this does not induce a significant bias. Second, the IP-level routes inferred from traceroute can be misleading due to aliasing where the routers use different interface IPs in sending the ICMP responses (e.g., [24]). Most antialiasing techniques require some form of active probing. Since our analysis is based on historical traceroute data, it is not meaningful to resolve the aliases retrospectively and thus we do not perform any IP-/host-level analysis. As such, our stability and structure analysis is at much coarser granularity and is less impacted by router IP aliasing. Third, topologies inferred from traceroute-like shortest path routes can show spurious relationships in topology inference (e.g., [28]). Again, our goal in this paper is not topology inference per se; we use the routing graph structure mainly to explain the latency and stability results and this is not a significant bias.

3 Hummus: Large-Scale Traceroute-Data Analysis Toolkit

The raw datasets from the previous section amount to roughly 1 terabyte in uncompressed form. Analyzing such large-scale datasets raises a number of practical scalability challenges in terms of preprocessing (cleaning missing traces, mapping IPs to desired granularity, etc.) and multiple round of iterations to extract meaningful information. Conventional analysis tools such as custom perl or python scripts require impractically large analysis time due to the associated compute and I/O bottlenecks.

To overcome these challenges, we developed Hummus — a tool that exploits state-of-the-art data processing tools (specifically, Apache Spark framework [42]) to enable rapid iterative analysis of traceroute datasets. In this section, we

[3] For end-to-end analysis, we can potentially use a few more traceroutes where the destination responds but intermediate hops fail. For consistency across the paper, we do not report results for the end-to-end analysis from such traceroutes and use only the "clean" set.

give a brief description of Hummus, functionalities enabled by Hummus, and a quick performance benchmark suggesting that future studies can exploit Hummus to significantly reduce the data preprocessing and analysis time.

Hummus comes with two key advantages. First, similar to tools such as MapReduce [15], it can run on distributed clusters and take advantage of data and CPU parallelism. Moreover, Hummus makes the "distributed" aspect (using multiple machines and available cores, handling machine failures and task scheduling, etc.) transparent to the user who essentially writes a single-machine code, as in conventional tools. Our own analysis, for instance, uses a cluster of 20 m2.4xlarge machines on Amazon EC2 with a total of 160 CPU cores and 1.3 terabytes of main memory. Second, and perhaps more importantly, Hummus leverages the Spark framework to provide support for avoiding disk I/O overheads by caching datasets and intermediate results into main memory. This leads to significant speed-ups since workloads in our measurement study are extremely I/O bound due to the inherent exploratory and iterative nature of our analysis tasks.

We emphasize that the actual code we had to write for our analysis is almost the same as what we would have written for a local, single-threaded script or program to analyze the data. We have made Hummus public at https://github.com/mosharaf/hummus for future measurement studies to quickly analyze and extract information from large datasets. We describe the functioning of Hummus using an example below.

Example: The following scala snippet finds the number of distinct vantage points in a dataset using the Spark[4] shell:

```
var lines = sc.textFile("hdfs://ark-traces/")
var traces = lines.filter(l => !l.contains("-")).cache
var vps = traces.map(t => t.split(" ")(0)).distinct()
```

In the snippet above, the first statement reads the entire dataset (the ark-traces directory in this case) from a distributed file system (HDFS [9]) into the variable called lines. The second statement performs the cleaning of the dataset as described in previous section — it "filters out" the traceroutes that exhibit negative latency values. The ".cache" at the end of the second statement ensures that Hummus caches traces in main memory. Finally, the third statement takes the first field of each line (the source IP) and counts the number of distinct vantage points. To get the actual set of unique vantage points, we could use a simple vps.collect() statement following the above snippet. To count the number of traces instead, we write:

```
var numTraces = traces.count()
```

Because traces was cached earlier, this query will avoid reading from the disk and run many times faster. For example, the former query took 98.5 seconds and the latter 237 milliseconds in our cluster. The corresponding queries took 94 minutes and 37 minutes on a local machine. We have found the ability to perform quick, ad-hoc analysis to be the biggest advantage of Hummus, saving hours of analysis time and dramatically speeding up our analysis workflow.

[4]Spark natively runs on scala, but it provides language bindings for java and python.

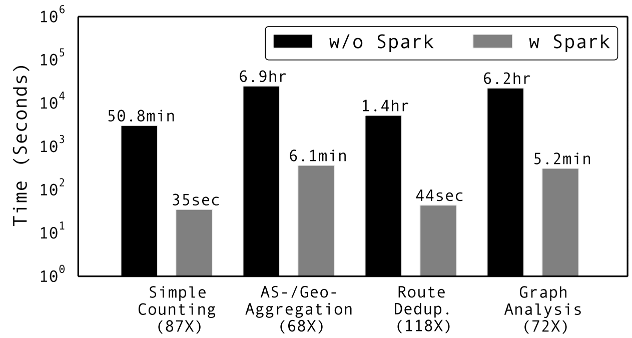

Figure 2: Example analytical tasks in the paper and corresponding durations on a Spark cluster vs. local machine for the Ark dataset. Speedups are in parentheses. Note that the y-axis is in log scale.

Speedup: The speedups we can achieve using Hummus depend on the type of analytical task we want to perform as well as on the fraction of data that can be cached and reused. We benchmarked the runtime of four typical analytical tasks on our 20-machine cluster and compared it with corresponding runtime on a single machine (Figure 2). The benchmark was performed for the Ark dataset in the following sequence with aggressive caching between steps.

1. *Simple counting:* First, we load and cache the entire Ark dataset into memory to perform simple counting tasks, similar to the ones used to produce Table 1.
2. *AS-/Geo-Aggregation:* Using this task, we convert each traceroute to a sequence of ASes/geolocations, and then group them by location, for example, to calculate aggregate statistics like the latencies presented in Section 4. We cache the AS-/Geo-converted form for the next task.
3. *Route deduplication:* In this case, we used the cached AS-granularity traces to determine the unique AS-level routes for each SrcDst pair to perform route stability analysis (Section 5). Because of caching in the last step, it runs much faster than other tasks.
4. *Graph analysis:* Finally, we convert each AS-granularity routes to AS-AS links and calculate characteristics of the AS graph, e.g., density and degree distributions discussed in Section 6.

Since the above tasks occur frequently in traceroute studies, we believe that Hummus will be a useful tool for future large-scale measurement studies.

4 Country-Level View of Latency

In this section, we analyze end-to-end Internet latency at the granularity of *countries*. We focus on the country granular-

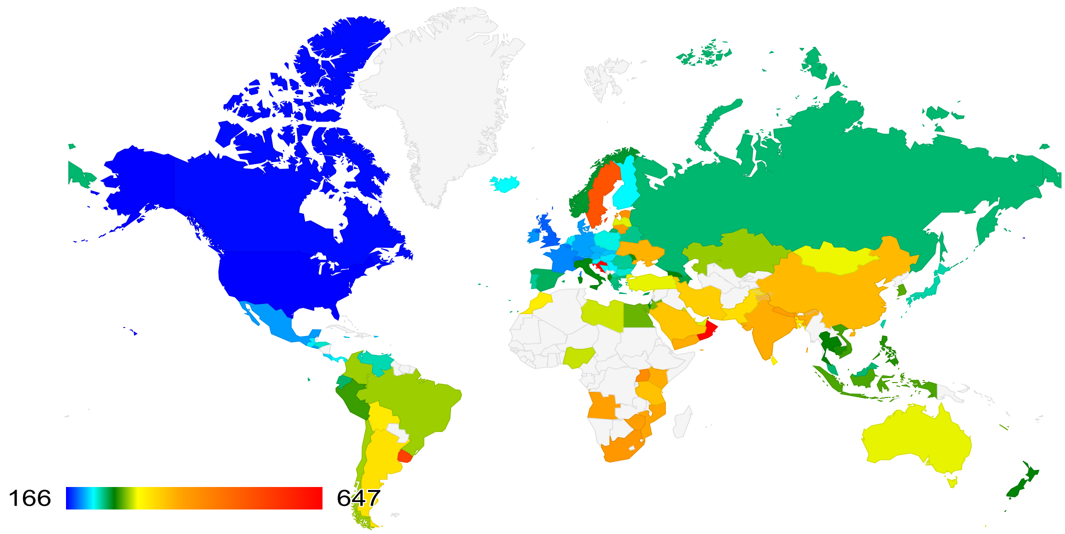

Figure 3: Average latency across the world in 2013 based on combined Ark, iPlane, and LargeCDN datasets.

ity to expose interesting geo-political and geo-economic effects on latencies across the world. We begin by looking at current "state of the world" – highlighting countries with the best and the worst latencies – using the Ark, iPlane and LargeCDN datasets (§4.1). We then analyze how latency has evolved over the past six years, and identify the top "gainers" and "losers" (§4.2). For the latter, we use only the Ark and iPlane datasets as the LargeCDN dataset is available only for 2012-2013. Next, we analyze the correlation between different datasets highlighting their differences and commonalities (§4.3). We close this section by summarizing our key observations (§4.4).

4.1 State of the world in 2013

The goal of this subsection is to identify "well-connected" and "not-so-well connected" countries in the current Internet and to validate whether the three datasets conform to similar set of observations.

Methodology: For this analysis, we restrict ourselves to the `traceroutes` collected by Ark, iPlane, and LargeCDN in 2013. We map each `traceroute` latency result to a particular source-destination (SrcDst) pair at the country granularity; i.e., with countries being the source and the destination. To obtain a unified cross-dataset view, we pick source countries that appear in *all three* datasets (there are 10 such countries, denoted as S^{2013}_{common} in Table 2). To avoid selection bias, we consider all destination countries in each dataset that received ≥ 30 `traceroutes` from *each* of these ten source countries; this reduced the number of destination countries to 125, 145 and 163 for Ark, iPlane, and LargeCDN datasets, respectively. Finally, a unified cross-dataset view is obtained by considering only those destination countries (with $\geq$30 `traceroutes` from each source) that appear in *all three* datasets; there were 111 such countries.

We acknowledge that there are potential sources of bias here; for instance, source countries like the United States may adversely affect our observations due to the sheer geographical expanse. However, our results suggest that the former effect is minimal on the aggregated result of the ten sources. Moreover, studying latency at the granularity of the country does not provide with information about latency of specific regions within a country. We provide such finer-grain information while discussing our observations as potential reasons that affect the country-wide latency profile.

In terms of summarization strategies, we performed analysis using median values as well (as opposed to average values used in final presentation). We chose to show the averaged results, because (a) our goal is to get a high-order understanding connectivities across countries, and averages suffice for this purpose; and (b) the results using median as the summarization strategy were qualitatively similar with an obvious disadvantage of less visibility into interesting outlier cases.

Best- and worst-connected countries: Figure 3 shows the geographical distribution of the average latency of the 111 destination countries as described above. At a high level, we see that the countries roughly fall into three tiers: (1) the best-connected countries with an average latency of around 200 ms that are largely in North America and Western Europe; (2) countries in South America, North Africa, Eastern Europe, Southeast Asia, and Australia form the middle tier with average around 300 ms; and (3) the rest of Africa, Middle East, South Asia, and China in the bottom tier with >400 ms average latency. That said, we do observe some anomalous hotspots on closer inspection (see §4.3).

Consequently, we compute the average ranks of countries based on their ranks in individual datasets; the resulting top-10 highest and lowest ranked countries are shown in Table 3. The top-10 list is also consistent with the three-tier observation from earlier. We find most of the developed countries to

		Source Countries
§4.1	S^{2013}_{common}	United States, Canada, Brazil, France, Germany, Great Britain, Hong Kong, Singapore, Taiwan, Australia
§4.2	$S^{\text{all}}_{\text{Ark}}$	United States, Canada, Brazil, Great Britain, Ireland, Netherlands, Spain, Japan, South Korea, Philippines
§4.2	$S^{\text{all}}_{\text{iPlane}}$	United States, Canada, Brazil, France, Germany, Poland, Russia, Switzerland, Japan, Taiwan

Table 2: Common source countries for latency analyses.

Low Latency (Lowest-to-Higher)	High Latency (Highest-to-Lower)
United States	Uganda
Canada	Angola
Great Britain	Zimbabwe
Belgium	South Africa
Bermuda	Mozambique
Switzerland	Kenya
France	Tanzania
Liechtenstein	Nepal
Mexico	Uruguay
El Salvador	Yemen

Table 3: Combined rank of countries based on their ranks in Ark, iPlane, and LargeCDN datasets.

have very good connectivity across all vantage points while some of the countries in Africa still have very poor latency characteristics. However, we will see in the next subsection, many of these African countries have in fact been improving in connectivity quite significantly over the last six years.

Number of IP and AS hops: In all datasets, we observed positive correlations between end-to-end latency and the number of IP/AS hops in `traceroutes`. Specifically, Pearson's r values were between 0.2 and 0.3 for latency-vs-IP hops and between 0.3 and 0.5 for latency-vs-AS hops.

4.2 Latency evolution

While the previous subsection focused on the current state of the Internet latency, here we dive deeper to understand how Internet latencies have evolved over the past six years to the current state.

Methodology: We analyze Ark and iPlane datasets for the last six years; we ignore the LargeCDN dataset since it provides only a nine month snapshot (§2). Moreover, for this subsection, we do not perform cross-dataset normalization since there were very few source-destination pairs left after normalizing across datasets *and* across years. Instead, for each individual dataset, we choose ten source countries that appear every year ($S^{\text{all}}_{\text{Ark}}$ and $S^{\text{all}}_{\text{iPlane}}$ in Table 2), and pick destinations that appear across all six years and have at least 30 traces to each source country. In the end, we get 93 destination countries each year for Ark and 139 destination countries each year for iPlane.

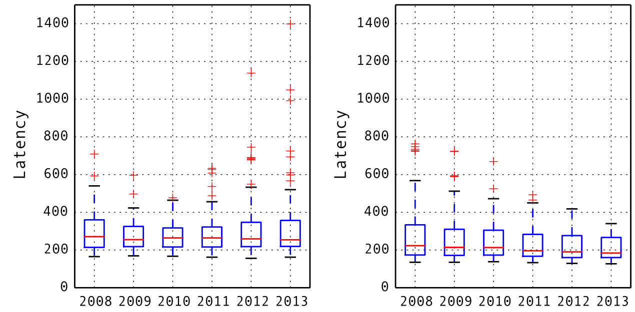

(a) Ark (b) iPlane

Figure 4: Box plots showing the median, 25-th, and 75-th percentiles across countries of the average `traceroute` latencies to destination countries over time. Whiskers are at $1.5\times$ the inter-quartile range, and the red crosses represent outliers.

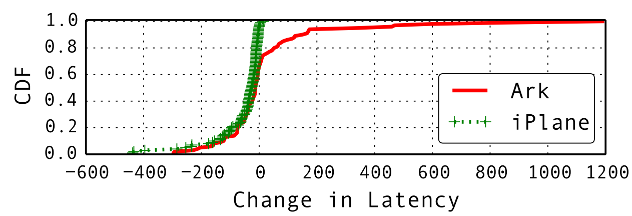

Figure 5: Change in average latencies to destination countries from 2008 to 2013 in Ark and iPlane.

Latency trends over years: Figure 4 shows the latency trends over years for the Ark and iPlane datasets. For the former, while the median latency was consistent over the years, the maximum latency suddenly started increasing after 2010. We discuss this in more depth in §4.3. For the iPlane dataset, the average, median, and maximum latencies decreased over the years.

Next, we focus on the change in average (over the sources) latency of *individual countries*. Figure 5 shows the CDFs of the change in average latencies in Ark and iPlane between 2008 to 2013. Interestingly, we see that the average latency of almost all countries decreased for the iPlane dataset. However, for the Ark dataset, almost 30% of the countries experience some increase in latency! Again, we discuss this in more depth in §4.3.

Biggest losers and gainers: Figure 5 shows that with the exception of 30% of the countries in the Ark dataset, the latency has reduced between 2008 and 2013. However, it does not provide insights into whether this change in latency occurred rather gradually or all of a sudden. To understand this (both latency improvements and regressions), we plot the latency of five most-improved and most-regressed countries from 2008 to 2013 for each dataset in Figure 6.

For both Ark and iPlane datasets, the countries whose latency improved the most belong to the third tier, as iden-

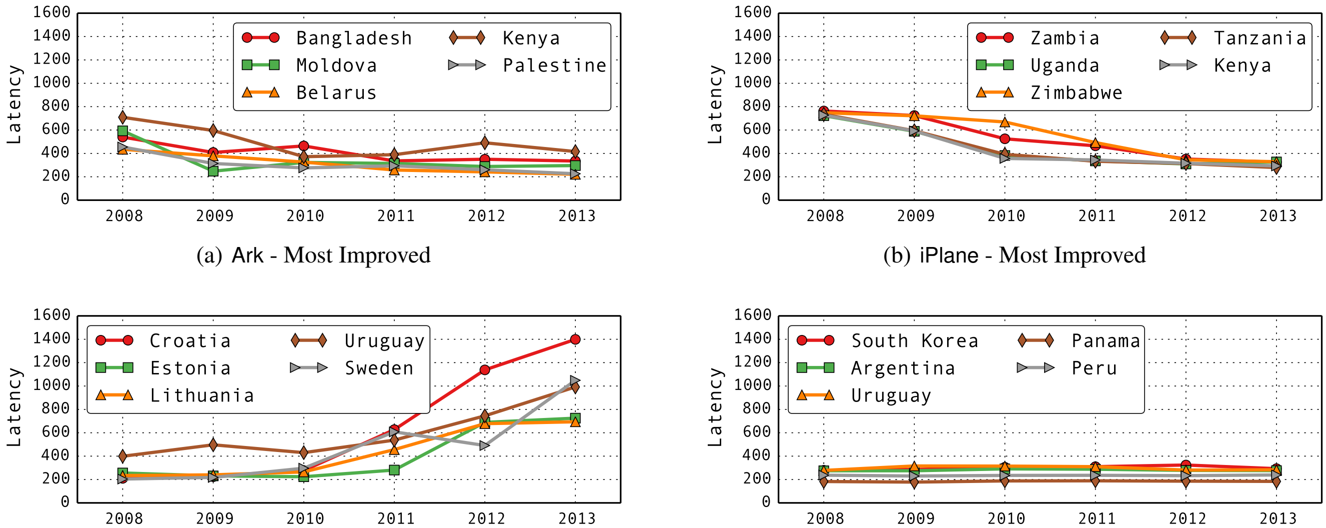

(a) Ark - Most Improved

(b) iPlane - Most Improved

(c) Ark - Most Regressed

(d) iPlane - Most Regressed

Figure 6: Changes in average end-to-end latencies for the most improved and most regressed five countries.

tified in §4.1. For Ark, some of the African countries from the third tier could have replaced countries that show up in Figure 6; however, these countries did not appear in 2008 or 2009, and were filtered out during cross-year normalization. Nevertheless, Figure 6 shows that the latency improvements have been gradual. An AS-granularity analysis shows two key changes in 2013 `traceroutes` from 2008: first, new ASes (with ASNs higher than 30000) in each of these countries became more frequently used. Second, AS3356 and AS3549 (i.e., Level 3) became a part of many routes.

In terms of regression in observed latencies, as Figure 5 would suggest, there is no noticeable regression for the iPlane dataset. For the Ark dataset, however, interesting trends stand out. In particular, while there was essentially no regression until 2010, the most-regressed countries are consistently becoming worse since 2011. Rather surprisingly, with the exception of Uruguay, the rest of the countries that make an appearance belong to the top-tier countries identified in §4.1. Digging deeper, we found that most `traceroutes` toward Croatia, Estonia, and Lithuania in the Ark dataset were routed through AS1257 after 2010. While it is hard to say definitively, these three countries having this commonality starting with the same year lends itself to the speculation that AS1257 may in fact be responsible for higher latencies for these countries starting 2010. Moreover, many of the `traceroute` for these countries geographically traversed through Sweden – another member of the most-regressed countries in the Ark dataset!

4.3 Differences across datasets

Figure 4, Figure 5 and Figure 6 highlight several interesting differences between the Ark and the iPlane datasets. In particular, Figure 4 shows that the Ark dataset has a large number of outliers essentially skewing the latency trend starting 2010; such outliers are not so dominant in the iPlane dataset. Figure 5 shows that while almost all countries in the iPlane dataset observed reduced latencies between 2008 and 2013, almost 30% of the countries in the Ark dataset observed significantly increased latencies. Finally, as discussed above, Figure 6 gives interesting latency regression trends starting from 2011. We observed several other interesting anomalies that we have not discussed, For instance, Croatia appears in the third tier in §4.1 with a significantly higher latency than all of its neighboring countries; moreover, the average latency for Croatia increased by almost 7× between 2008 and 2013. We believe that this is precisely due to the AS1257 anomaly discussed in reference to Figure 6.

To formally understand the differences between the datasets, we calculated Pearson's correlation coefficients across datasets to find the correlation between latencies observed across the three datasets. Table 4 shows that latencies observed by iPlane and LargeCDN show strong positive correlation; however, both are weakly correlated with the latencies observed in the Ark dataset. While anomalies like AS1257 may have adverse affect on absolute numbers, it is unclear whether such anomalies lead to different *trends*. To that end, we computed the correlation between different datasets in terms of *relative ranks* of countries based on the average latency. We find, rather interestingly, that the relative order of countries across datasets tend to be more similar (Table 4).

To summarize, our analysis suggests that the absolute numbers from any individual dataset should be used with caution; however, the temporal trends (such as those in Figure 6) and relative orderings of specific data points (such as those in Table 3) may indicate a higher level of certainty. Our

	Latency	Rank
iPlane-LargeCDN	0.70	0.72
Ark-iPlane	0.18	0.57
Ark-LargeCDN	0.27	0.66

Table 4: Pearson's correlation coefficients across datasets in terms of latencies and ranks of destination countries.

analysis also highlights the importance of a cross-dataset study as a means of providing a second-hand confirmation to the observations made using a single dataset.

4.4 Summary of key observations

In this section, we analyzed the end-to-end latency at the country granularity with a particular focus on understanding the current state of the world, the evolution of latency across years and the differences and commonalities across multiple datasets. Our key observations are:

- **State of the World:** The countries with the lowest average latency tend to be in North America and Europe while the countries with the worst average latency tend to be in underdeveloped regions of Asia and Africa.
- **Countries with improved latency:** While underdeveloped regions of Asia and Africa have the worst average latencies, these are also the ones with the most improvements over the past six years.
- **Countries with deteriorating latency:** Some datasets suggest that the countries with the most deteriorating latency belong to Eastern Europe. However, this appears to be an artifact of an anomaly caused by a single AS, specifically over the last three years. Most countries observe stagnant or improved latencies.
- **Overall latency trends:** Overall latency has improved, albeit slowly, over the last few years, with absolute numbers depending on the dataset.
- **Cross-dataset studies:** Absolute numbers from any dataset in isolation may not provide a complete picture, suggesting a need for cross-dataset analysis. However, temporal trends and relative ordering of specific data points observed in one dataset may hold across datasets.

5 Routing Stability

In this section, we analyze routing stability [36, 38] with a particular focus on how stability varies across years and across datasets. Similar to Paxson's study [36], we focus on route stability at the granularity of ASes and cities w.r.t. two key notions of prevalence and persistence [36]. We then analyze path prevalence in §5.1 and path persistence in §5.2. We then discuss the differences observed between the Ark and the iPlane datasets in §5.3 and close the section by summarizing the key observations in §5.4.

Metrics of stability: Paxson's seminal study of Internet routing behavior [36] proposed two metrics to capture route stability: *prevalence* and *persistence*. Intuitively, prevalence refers to the probability of observing a given route, whereas persistence is the duration we expect to keep observing the same route. More formally:

- Prevalence is defined as the probability of observing the dominant route (i.e., the most frequently observed route) for any given source-destination pair. It is computed as the ratio of the number of times the dominant route is observed and the total number of observations.

$$\text{Prevalence} = \frac{\text{Dominant Route's Popularity}}{\text{Number of Observations}}$$

- Persistence is used to measure the frequency at which a particular SrcDst pair switches between different routes.

$$\text{Persistence} = \left(1 - \frac{\text{Number of Route Changes}}{\text{Number of Observations} - 1}\right)$$

Note that these metrics are distinct and cannot be inferred from one another. For instance, consider a SrcDst pair with three candidate routes: R_1, R_2, and R_3. Consider the following two sequences of observed routes across two set of measurements:

$$R_1, R_2, R_3, R_2, R_3, R_1, R_3, R_1, R_2$$

and

$$R_1, R_1, R_1, R_2, R_2, R_2, R_3, R_3, R_3$$

Note that each of R_1, R_2 and R_3 is a dominant path in the two sequences (i.e., they appear the same number of times) and indeed, the prevalence for each of the routes is 0.33 in both measurements. However, the routes in the first measurement have 0 persistence while routes in the second observation have 0.75 persistence precisely because routes are stable across consecutive observations in the second measurement.

Methodology: To perform longitudinal analysis, we restrict ourselves only those source-destination pairs that appear across each of the six years (in each individual dataset). Furthermore, we only consider source-destination pairs with at least 100 `traceroutes` each year. We do not perform cross-dataset normalization to avoid vantage point bias.

We analyze the route prevalence and persistence at two different granularities: AS-granularity (a sequence of AS hops) and city-granularity (a sequence of cities); this is similar to [36] with the difference that we do not study these metrics at the host granularity. As discussed in §2, studying routes at the granularity of the host requires retroactively "unaliasing" routers that has its own limitations. In the end, we had 5653 and 17174 AS-granularity SrcDst pairs and 9395 and 26192 city-granularity pairs for Ark and iPlane, respectively.

There are several plausible limitations in our analysis on route stability. First, imprecise AS and city resolution may lead to skew in results. However, as discussed in Section 2, we have cross-validated the mapping tables used for AS and city resolution against the Maxmind tables with a high match

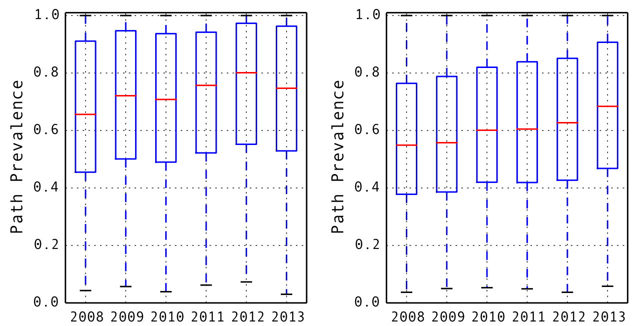

(a) Ark (b) iPlane

Figure 7: Route prevalence between AS-AS pairs is gradually increasing.

rate. Second, subsampled datasets (§2) only allow us to capture *long-term* persistence properties at a month-level granularity.

5.1 Prevalence analysis

We start with analyzing route prevalence at the AS-granularity and then move on to the city-granularity.

5.1.1 AS-granularity

Figure 7 shows the distribution of prevalence of the dominant routes over the years for Ark and iPlane datasets at the AS granularity. We observe that prevalence is increasing over the years in both datasets – median prevalence increased from 0.66 to 0.75 in Ark and from 0.55 to 0.68 in iPlane between 2008 to 2013.

Interestingly, we find the prevalence values in our analysis to be significantly higher than Paxson's study [36]— he observed "[..] In general, Internet paths are strongly dominated by a single route [..]". For example, Paxson found the median prevalence at AS granularity to be 1, which is significantly higher than our observations. We believe that the difference may be due to the source-destination pairs chosen in Paxson's study and in our study — the focus in Paxson's work was largely on university sites whereas we have a greater view of the "commodity" Internet (using Ark and iPlane datasets).

We found the route prevalence as the AS granularity to be negatively correlated with average number of AS hops and with average latency (Pearson's r values between -0.1 and -0.3) for both datasets. We speculate that this is because longer AS paths are more likely to be impacted by routing instabilities or routing convergence issues; i.e., even if just one intermediate AS hop decides to change its policy, then the path will likely change and the prevalence of that path will decrease.

Most and least prevalent: Next, we focus on the destination ASes that used high (≥ 0.95) and low (≤ 0.25) prevalence routes. For this analysis, we picked the top-50 most

(a) High

Ark	2907, 786, 1213, 680, 553
iPlane	17908, 21565, 786, 7066, 12271

(b) Low

Ark	7922, 31377, 6621, 9829, 8167
iPlane	21433, 20940, 16586, 17974, 6983

Table 5: Top-5 destination ASes associated with the highest/lowest prevalence values.

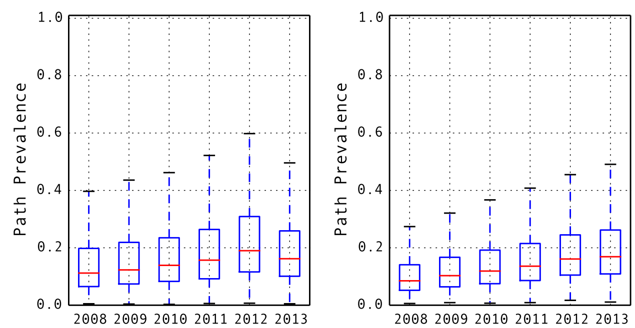

(a) Ark (b) iPlane

Figure 8: Route prevalence between city-city pairs is gradually increasing. The absolute numbers are smaller than that at the AS-granularity.

frequently occurring destination ASes[5] in each set (high and low prevalence) for each year and took the intersection across six years. Table 5 presents the top-5 destination ASes consistently associated with high and low prevalence routes in both datasets.

We found that, in both datasets, ASes associated with high prevalence routes have low in-degree and vice versa. In particular, most high-prevalence ASes in Table 5 have in-degrees around 5, whereas the ASes associated with low prevalence routes have in-degrees around 30. As the in-degree increases, nodes may have larger number of options for load balancing, leading to lower path prevalence. Finally, we note that high- or low-prevalence ASes are not restricted to any particular geolocation.

5.1.2 City-granularity

Figure 8 shows that the median prevalence at the *city-level* granularity is less than 0.20. This is in sharp contrast to Paxson's observations that suggested very high prevalence (≥ 0.97) at the city granularity. (Again, we believe that Paxson's study found very high prevalence at city-level because it mostly used academic sites with a small number of tightly interconnected cities in the academic backbone. In contrast, our dataset covers much more of the "commodity" Internet,

[5] Because the number of source ASes were much smaller than that of destination ASes, many source ASes appear in both high and low prevalence sets; hence, we focus only on destination ASes.

(a) High

Ark	2907, 1213, 12271, 3320, 4713
iPlane	21565, 7080, 5786, 11351, 7066

(b) Low

Ark	7922, 22394, 6621, 8167, 9829
iPlane	16586, 9829, 21433, 6983, 8151

Table 6: Top-5 destination ASes associated with the highest/lowest persistence.

and thus exhibits more flux.) Similar to Figure 7, we see that prevalence at city-granularity is also increasing over time.

Following the same methodology as before, we calculated destination cities that used the most (≥ 0.75) and least (≤ 0.05) prevalent routes. While high prevalence routes end in many cities across the world, we found many low prevalence routes to be destined toward the west coast of United States and the north-east coast of China.

5.2 Persistence analysis

We now analyze the route persistence using our Ark and iPlane datasets. We start our analysis at the AS granularity, followed by analysis at the city granularity.

5.2.1 *AS-granularity*

Figure 9 shows the distribution of persistence over the years for Ark and iPlane datasets at the AS granularity. We found route persistence to be highly correlated with route prevalence in the previous subsection (Pearson's r values between 0.75 and 0.8). Consequently, most of the observations about prevalence qualitatively hold true for route persistence as well. For example, persistence is increasing over the years in both datasets – median persistence improved from 0.76 to 0.88 in Ark and from 0.58 to 0.78 in iPlane between 2008 to 2013. Similar to prevalence, the average number of AS hops and average end-to-end latency are also negatively correlated with persistence.

As earlier, we identified the destination ASes that were consistently involved in high (≥ 0.95) and low (≤ 0.25) persistence routes (Table 6). While high-persistence destinations had nothing in common across datasets, AS5089, AS6983, AS8151, AS8167, AS9829, and AS20940 appeared in low-persistence pairs in both datasets (some not among the top-5). Reenforcing the strong correlation between prevalence and persistence, we observe significant similarities between Table 5 and Table 6 (especially for low prevalence and low persistence).

5.2.2 *City-granularity*

Similar to the case of route persistence at the AS-granularity, route persistence at the city granularity (Figure 10) is strongly correlated with route prevalence at the city granularity (Pearson's r values over 0.8). We also calculated destination cities with the most (≥ 0.75) and the least (≤ 0.05) persistent routes (not shown for brevity). However, unlike prevalence, there is no concentration of low-persistence cities in the coasts of United States or China.

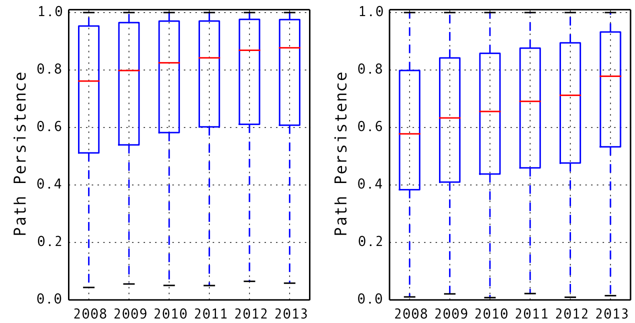

(a) Ark (b) iPlane

Figure 9: Route persistence between AS-AS pairs shows a upward trend in both datasets.

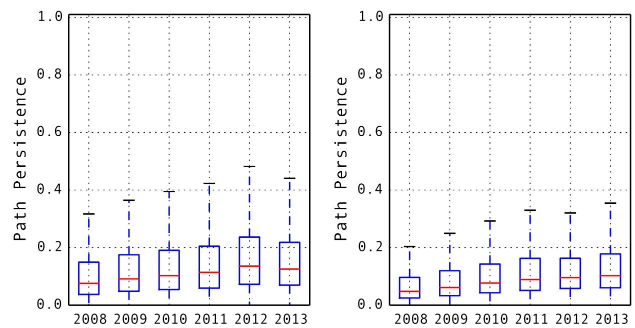

(a) Ark (b) iPlane

Figure 10: Route persistence at the city granularity is improving over time.

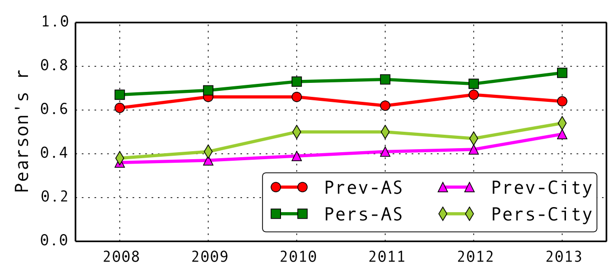

Figure 11: Pearson's correlation coefficients (r) between Ark and iPlane datasets for prevalence and persistence over the years at the AS- and city-granularity.

5.3 Differences across datasets

Our analysis so far was restricted to analyzing trends within each dataset individually. Note from Figure 7 that the absolute values of route prevalence differ significantly across datasets. We found the fraction of source-destination pairs with low prevalence routes to be similar in both datasets (less than 0.20), the routes in the Ark dataset have better prevalence with respect to the median- and the high-prevalence routes when compared to the iPlane dataset.

To understand the difference between the two datasets more formally, we compute the Pearson's correlation coefficients for the common SrcDst pairs in each year (Figure 11). We observe that both route prevalence and persistence across datasets at the AS-granularity are strongly correlated. Furthermore, the agreement between Ark and iPlane is slowly increasing. At the city-granularity, Ark and iPlane appear to have a weaker correlation but their agreement is increasing at a faster rate.

5.4 Summary of key observations

In this section, we analyzed the route stability in Ark and iPlane datasets at the AS and the city granularity, with a focus on how stability has changed over the years and how it varies across datasets. Our key observations are:

- **Overall trend:** We find that the route prevalence and persistence at both the AS- and the city-granularity is increasing over the past six years, with improvements as high as 20% over the six year period.
- **Absolute numbers:** In a sharp contrast to observations made in prior work [36], we find the route prevalence and persistence to be significantly lower at the AS level (0.75 in our analysis versus 1 in [36]) and even lower at the city level (0.20 in our analysis versus 0.97 in [36]). We attribute this difference to the fact that prior work analyzed routes between university sites while our analysis is done on the "commodity" Internet.
- **Prevalence and persistence vs. latency and multihoming:** We find route prevalence and persistence to be weakly inversely correlated with end-to-end latency and number of AS- or city-hops. Furthermore, we find the in-degrees of destination ASes to be inversely correlated with prevalence and persistence.
- **Cross-dataset study:** Similar to Section 4, we find the trends in one dataset to hold well across datasets. However, absolute numbers from any dataset in isolation do not seem to provide a complete picture (although they strongly correlate), suggesting the need for validation using a cross-dataset analysis.

6 Routing Graph Structure Analysis

In this section, we analyze the *structure* of the routing graph to analyze potential root causes for the observations made in Section 4 and in Section 5. To do so, we map the `traceroute` dataset from Ark and iPlane into a graph (for each individual year) and analyze the evolution of these graphs in terms of key structural properties. We do acknowledge that Internet routing is a complex multi-faceted phenomenon; latency and path stability depend on several other hidden variables (e.g., protocol convergence, business relationships, node/link failures) and routing/graph structure is only one aspect. Nevertheless, we believe this structural analysis is useful as a tool to shed light on aggregate trends in latency and stability.

Methodology: Our goal is to identify the potential causes for *changes* observed in end-to-end latency and path stability over the years. Since a single-year snapshot (such as the one from LargeCDN) does not serve our purposes, we focus on the Ark and iPlane datasets. We map the yearly datasets from Ark and iPlane to an AS-granularity and a city-granularity graph — ASes and cities in the traces represent the nodes in these graphs and the pairwise bigrams (i.e., direct links) in the traces represent corresponding links. Moreover, we do not sample nodes and links since the intermediate nodes and links contribute to the end-to-end latency and path stability. For a given dataset, we make sure that the set of nodes (ASes or cities) remain the same in each year, while the set of links between them can change.

For brevity, we present the results obtained only for Ark throughout the rest of this section. We choose this because Ark has bigger AS- and city-granularity graphs than iPlane (Figure 1). Furthermore, in some plots, we present results only for 2009, 2011, and 2013 for ease of exposition.

Metrics: We consider two primary metrics for our analysis. First, the *density* or the *average degree* of the graph defined as the ratio of the number of links to the number of nodes in the graph. We complement this metric with a longitudinal analysis of degree distributions over the years and with a study of how degrees of "popular" nodes in the graph is changing over the years.

Second, we use *clustering coefficient* of nodes, a fundamental measure that quantifies how tightly-knit the community is around a particular node. Specifically, the clustering coefficient of a node u in the graph is the number of triangles u belongs to normalized by the maximum possible number of such triangles [40]. That is, for an undirected graph $G = (V, E)$ and a node u, the clustering coefficient $c(u)$ is measured as:

$$c(u) = \frac{2T_u}{N(u)(N(u)-1)}$$

where T_u denotes the number of triangles u belongs in and $N(u)$ is the number of neighbors of u. The average clustering coefficient of G is defined as the sum of the individual clustering coefficients of each node divided by the number of nodes.

As discussed earlier, these metrics may neither be sufficient nor be perfectly accurate to provide a complete picture for the observations made in the previous sections. Nevertheless, we use them to gain interesting insights into the evolution of the underlying topologies.

6.1 Density and degree distribution

Figure 13 shows the density of the AS-granularity and city-granularity graphs for the Ark dataset. Note that the city-granularity graph is denser than that at the AS-granularity, because ASes cover large geographical regions and connect many cities using intra-AS links. The table shows that the density of each of the graphs has gradually increased over the years (for ASes/cities that appear across all years). Each

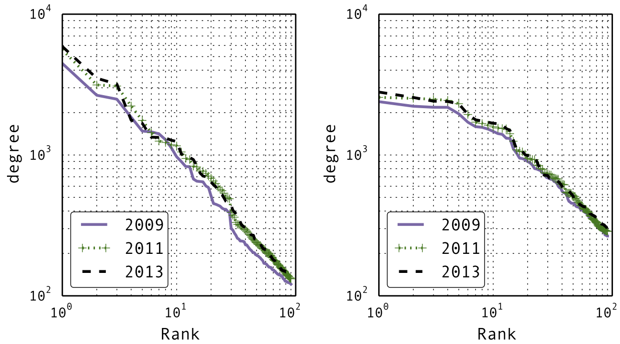

(a) AS (b) City

Figure 12: Degree distributions of the top-100 ASes and cities in the Ark dataset. High-degree entities are connecting to more over time.

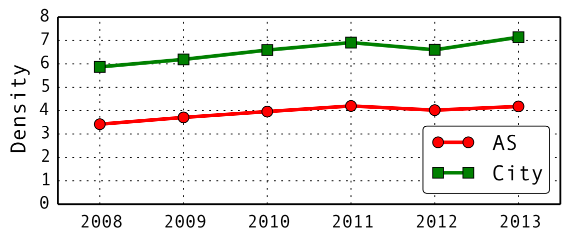

Figure 13: AS- and city-granularity graphs in the Ark dataset are getting denser over time.

node in the graph connecting to more nodes suggests that, ignoring routing policies, the average number of hops between node pairs reduce over the years, leading to potential reductions in latency.

To dig deeper, we computed the change in degrees of the "popular" (high-degree) nodes over years. Figure 12 zooms into the 100 highest-degree ASes and cities. We observe that while the average degree for the entire graph has changed by at most 2.5, the degree of popular nodes have increased by a much larger number over the years; for instance, the average degree in the AS-granularity graph increased by less than 1 but the degree of the most popular AS increased by almost 1500. A closer look at the high-degree ASes reveal, as expected, that these are all Tier 1 ASes: in both datasets the highest-degree AS in each year was AS3356; AS7018 and AS174 were the second- and third-highest from 2008 to 2010, and they reversed the order from 2011 to 2013. We speculate that by way of directly connecting to (AS/city corresponding to) the core directly, nodes in the graph observe reduced end-to-end latency and increased path stability. Finally, we found New York, Chicago, London, and Frankfurt consistently among the best-connected cities in the world.

6.2 Clustering coefficient

The increase in degree of popular and Tier-1 ASes provides one possible root cause for reduction in end-to-end latency

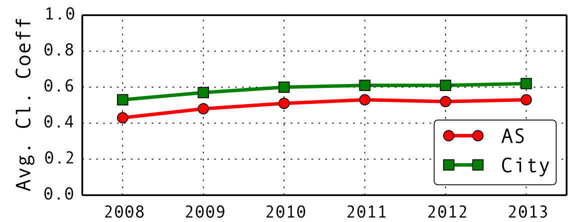

Figure 14: Average clustering coefficients for Ark at the AS- and city-granularity are increasing.

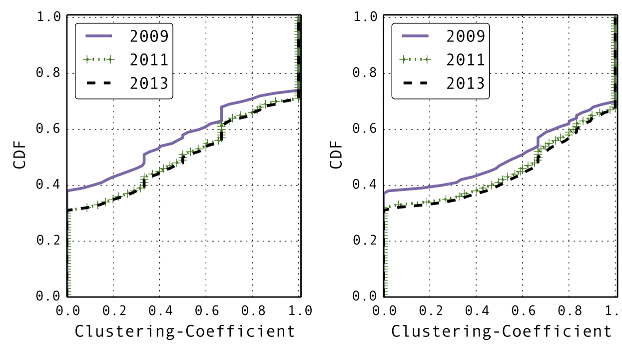

(a) AS (b) City

Figure 15: CDFs of clustering coefficients for AS- and city-granularity graphs for Ark.

and increase in path prevalence. In this subsection, we analyze changes, over the years, in another structural property of the graph — clustering coefficient. As discussed earlier, a higher clustering coefficient suggests more closely knit neighborhoods. For instance, at two extremes are the trees (with each node having clustering coefficient of 0) and complete graphs (with each node having clustering coefficient of 1), with the latter suggesting paths of fewer hops between node pairs.

Figure 14 shows that, over the past six years, the average clustering coefficient has increased by more than 23% for the AS-granularity and by almost 17% for the city-granularity graph for the Ark dataset. Increase in average clustering coefficient suggests that the graph has become more tightly knit; as a result, the average number of hops between node pairs has reduced. As discussed earlier (§4.1), this positively correlates with reduction in latency observed in Section 4.

Figure 15 shows the CDF for node clustering coefficients. We make two observations. First, the median clustering coefficient has increased by more than 50% between 2009 and 2013 for both AS-level and city-level graphs. This is a significant increase, suggesting a quick departure from a tree-like topology to a much flatter topology. Indeed, this provides a possible second reason for reduction in average number of hops and in turn, reduction in end-to-end latency between node pairs.

The second observation from Figure 15 follows from the very definition of clustering coefficient. In particular, consider a node with degree 5 and suppose it belongs to 5 triangles; then, the node has clustering coefficient of 0.5. Now suppose the degree of the node increases by $2\times$ to 10 and the number of triangles the node belongs to increases by $4\times$ to 20; then, the new clustering coefficient of the node is $20/45 < 0.5$. That is, nodes whose degree increases over the years must be contained in many more triangles for an increase in their clustering coefficients. Comparing it with the observations made in the previous subsection, it is highly unlikely that the clustering coefficient of popular nodes is increasing over the years (since their degree is increasing significantly). Indeed, we confirmed this intuition — the clustering coefficient of all the popular nodes decreased between 2009 and 2013.

The above indicates that the clustering coefficient of nodes that have small degree is increasing much more rapidly over the years. As a result, these nodes are conforming to a significantly more clustered topology (forming denser communities), which assumably leads to paths of much fewer hops to other non-core nodes in the network.

7 Related Work

Measurement and analysis tools: Understanding Internet path properties requires developing tools for collecting, de-noising, and extracting interesting observations from data. A number of tools have been developed in the past (e.g., [1, 7, 23, 30, 32]). While these tools differ in terms of underlying mechanism to collect data, source and destination selection, coverage etc., they have one commonality — they generate enormous amounts of data. Even after sampling a part of data, extracting meaningful information remains a time consuming process. Our work complements the above measurement tools. Using a state-of-the-art data analytics platform, we developed an analysis suite (§2) that will help future studies to quickly extract information from these large datasets.

Studies of Internet path properties: Paxson [36] performed the first-of-its-kind study of Internet path properties using 40000 measurements between 37 Internet sites. The study provided insights into routing pathologies, routing stability, and routing asymmetry. This motivated many follow-up studies on analyzing network latencies [19, 21, 32], path stability [11, 12, 38], path availability [16, 18, 39, 43] and on understanding the structure of the network [6,20,25,27]. Our work follows in the spirit of this rich body of measurement work, enabled by the availability of much larger-scale measurements from multiple datasets over a longer timeframe.

Topology studies: There is a rich body of work in understanding Internet topology, peering relationships, and their evolution over the years. This includes work on showing power-law relationships [17], explaining the rise of such power-law relationships [29], their use in topology generators [33], and techniques for inferring hidden AS-links [34]. More recent studies have shown emerging trends such as "flattening" [20, 27] and the rise of Internet exchange points (IXPs) [5, 8] that have made the AS-level graph denser. Our analysis of routing graph structure is naturally related to such topology measurements. While our observations are largely in line with this prior work, we use this structural analysis more specifically to explain the latency and stability trends we see.

Applications: `traceroute` measurements have been used in the past to predict Internet path performance [19, 31, 32], to track and predict Internet path changes [13], to debug failures [16], to identify Internet black holes [22], to identify long term changes in underlying network routing topology [20, 27], and to select CDN servers among the available pool of servers [26]. We add a new dimension to be considered in future studies using existing datasets — identifying and de-touring the blind spots and biases in one dataset by performing a cross-dataset analysis. Indeed, studies similar to above can also benefit by our observations about the gradual changes in the Internet path properties across years.

8 Lessons and Implications

We conclude with some key lessons learned and implications for future research both inside and outside the measurement community.

- *Blind spots:* Our country-level analysis revealed that even the largest measurement datasets can have key weaknesses. For example, Ark relies on random sampling of prefixes (/24 or routable prefixes) to select destination IP addresses, creating a large number of blind spots precisely in Africa where Internet traffic and adoption is likely to grow exponentially in the future. Even though African countries were among the most improved (§4), we could not get enough information on many of them, because random sampling of prefixes often missed 1.55% of IP addresses located in Africa.[6] While iPlane uses a fixed set of prefixes, it still has a large number of blind spots. An ideal solution, for the lack of a "complete" coverage, would be to use a careful combination of both random and selective destination selection strategies.

  Furthermore, even with these larger scale measurement datasets, we still get a very Euro- and US-centric view of the broader Internet landscape as most of the vantage points (at least the most stable ones) are located in Europe and US. Consequently, we may not have good visibility into localized and regional phenomena; e.g., local IXPs or country-level Internet structures. This is especially relevant as content providers increasingly seek to localize the bulk of their data transfers [26].
- *Need for cross-dataset validation:* The anecdote from Section 4.3 regarding Ark once again reiterates a well-known but nevertheless important cautionary tale for Internet measurement research—the need for cross-dataset

[6] Calculated from our IP-to-AS/Geolocation conversion table.

validation of findings. If we had relied solely on Ark, we may have incorrectly concluded that several Eastern European nations are seeing degraded connectivity. But in reality, that is an artifact of the specific view that the Ark "lens" gives us as we do not find similar trends using iPlane and LargeCDN.

- *Path toward potential improvements:* While Internet connectivity is improving, we do observe that there is still significant room for improvement in many geographical regions (e.g., Africa). Some of these regions can also benefit from understanding the roadmaps of how/why some of the biggest gainers have improved their latency: for example, AS58587 (Fiber@Home Limited, Bangladesh) became a prominent last hop in AS routes to Bangladesh as Bangladesh's average latency improved. The rise of new last-hop ASes was common for several other most-improved countries (§4.2). Combining this with the increasing stability of AS-granularity paths (§5), we speculate that last-mile upgrade is quite possibly the best way to improve Internet connectivity.
- *Impact on applications:* The generally increasing trends in long-term path stability (§5) as well as graph density and clustering (§6) bode favorably for applications that rely on path/latency prediction [14, 30]. This also has favorable implications for content providers and CDNs as they naturally benefit from lower latency and increased routing stability [26, 37, 41]. That being said, we also note that given the trends toward "localizing" bulk transfers [10], content providers and the measurement systems that they rely on (e.g., [3]), will also need to deploy deeper lenses to avoid the aforementioned blind spots.

9 References

[1] CAIDA Ark. http://www.caida.org/projects/ark.
[2] For impatient web users, an eye blink is just too long to wait. New York Times, http://goo.gl/NNhmw.
[3] Keynote systems. http://www.keynote.com/.
[4] MaxMind. http://dev.maxmind.com/geoip/legacy/.
[5] B. Ager et al. Anatomy of a large European IXP. In *SIGCOMM*. 2012.
[6] D. G. Andersen et al. Topology inference from bgp routing dynamics. In *IMW*. 2002.
[7] B. Augustin et al. Avoiding traceroute anomalies with Paris traceroute. In *IMC*. 2006.
[8] B. Augustin et al. IXPs: Mapped? In *IMC*. 2009.
[9] D. Borthakur. The Hadoop distributed file system: Architecture and design. Hadoop Project Website, 2007.
[10] M. Calder et al. Mapping the expansion of Google's serving infrastructure. In *IMC*. 2013.
[11] G. Comarela et al. Studying interdomain routing over long timescales. In *IMC*. 2013.
[12] Í. Cunha et al. Measuring and characterizing end-to-end route dynamics in the presence of load balancing. In *PAM*. 2011.
[13] I. Cunha et al. Predicting and tracking Internet path changes. In *SIGCOMM*. 2011.
[14] F. Dabek et al. Vivaldi: A decentralized network coordinate system. In *SIGCOMM*. 2004.
[15] J. Dean et al. MapReduce: Simplified data processing on large clusters. In *OSDI*. 2004.
[16] A. Dhamdhere et al. Netdiagnoser: Troubleshooting network unreachabilities using end-to-end probes and routing data. In *CoNEXT*. 2007.
[17] M. Faloutsos et al. On power-law relationships of the Internet topology. In *SIGCOMM*. 1999.
[18] N. Feamster et al. Measuring the effects of Internet path faults on reactive routing. In *SIGMETRICS Performance Evaluation Review*, volume 31, pages 126–137. 2003.
[19] P. Francis et al. Idmaps: A global Internet host distance estimation service. *IEEE/ACM ToN*, 9(5):525–540, 2001.
[20] P. Gill et al. The flattening Internet topology: Natural evolution, unsightly barnacles or contrived collapse? In *PAM*. 2008.
[21] A. Gupta et al. Peering at the Internet's frontier: A first look at ISP interconnectivity in Africa. In *PAM*. 2014.
[22] E. Katz-Bassett et al. Studying black holes in the Internet with Hubble. In *NSDI*. 2008.
[23] E. Katz-Bassett et al. Reverse traceroute. In *NSDI*. 2010.
[24] K. Keys. Internet-scale IP alias resolution techniques. *SIGCOMM CCR*, 40, 2010.
[25] A. Khan et al. AS-level topology collection through looking glass servers. In *IMC*. 2013.
[26] R. Krishnan et al. Moving beyond end-to-end path information to optimize CDN performance. In *IMC*. 2009.
[27] C. Labovitz et al. Internet inter-domain traffic. In *SIGCOMM*. 2011.
[28] A. Lakhina et al. Sampling biases in IP topology measurements. In *INFOCOM*. 2003.
[29] L. Li et al. A first-principles approach to understanding the Internet's router-level topology. In *SIGCOMM*. 2004.
[30] H. V. Madhyastha et al. iPlane: An information plane for distributed services. In *OSDI*. 2006.
[31] H. V. Madhyastha et al. A structural approach to latency prediction. In *IMC*. 2006.
[32] H. V. Madhyastha et al. iPlane Nano: Path prediction for peer-to-peer applications. In *NSDI*. 2009.
[33] P. Mahadevan et al. Systematic topology analysis and generation using degree correlations. In *SIGCOMM*. 2006.
[34] R. V. Oliveira et al. Observing the evolution of Internet AS topology. In *SIGCOMM*. 2007.
[35] V. Paxson. Reflections on measurement research: Crooked lines, straight lines, and moneyshots. SIGCOMM 2011 Keynote Speech.
[36] V. Paxson. End-to-end routing behavior in the Internet. *IEEE/ACM ToN*, 5(5):601–615, 1997.
[37] L. Qiu et al. On the placement of web server replicas. In *INFOCOM*. 2001.
[38] J. Rexford et al. BGP routing stability of popular destinations. In *IMW*. 2002.
[39] R. Teixeira et al. Bgp routing changes: Merging views from two isps. *ACM SIGCOMM Computer Communication Review*, 35(5):79–82, 2005.
[40] D. J. Watts et al. Collective dynamics of 'small-world' networks. *Nature*, 393(6684):440–442, 1998.
[41] M. Yu et al. Tradeoffs in cdn designs for throughput oriented traffic. In *CoNEXT*. 2012.
[42] M. Zaharia et al. Resilient distributed datasets: A fault-tolerant abstraction for in-memory cluster computing. In *NSDI*. 2012.
[43] Y. Zhang et al. The stationarity of Internet path properties: Routing, loss, and throughput. Technical report, ACIRI Technical report, 2000.