Consolidated Review of
Mapping the Expansion of Google's Serving Infrastructure

1. Strengths

The paper presents the first study on the growth of a CDN provider.

Another clever use of spoofing to gain insight into the infrastructure not otherwise practically feasible to obtain by researchers. The paper presents a nice and cute technique for mapping out a CDN provider's infrastructure. Although the techniques provided are specific to Google's CDN, the authors show how their technique can be used to quantify the coverage and accuracy of prior techniques. The presented techniques can also be used to enrich bittorrent client design, geolocation and other domains in networking.

New measurement ideas: rely on EDNS New ways of improving geolocation. Example analysis of Google.

Nice related work section.

The results in the paper are sort of interesting, but the bigger value in this work is likely sharing the data with others to answer different questions about the Internet. Insights could be valuable to the community at large

2. Weaknesses

Techniques not entirely novel: Not the first study to use EDNS.

Focus only on one infrastructure, namely Google. Unclear how complete the mapping really is.

The authors fail to provide intuition as to why many of the techniques were used.

- For example, it is not clear why they chose to use "the smallest RTT of 8 consecutive pings".
- Or why Manhattan distance is the best choice for a distance function.
- Or why the MaxMind geolocation database is used over other geolocation databases.
- The techniques presented are not thoroughly evaluated. The lack of thorough evaluation is highlighted by several explainable results.

Not clear if one can believe the validation results. While they have a long section of validation, it was partial validation and not clearly stating what they haven't, or can't, validate. I found myself wanting quantitative explanations or assessments when they weren't being provided. For example, in the "putting it all together" section I wondered why they authors weren't given numeric/percentage values for how many prefixes were filtered out at each step, how many were excluded at each step, how many were more than a standard deviation beyond the mean. These numbers were no doubt elsewhere in the paper, but I would have liked to see a summary somewhere. I give a few other examples below of places where I thought a numeric value would be more appropriate.

3. Comments

This paper is on a timely topic. The paper has an ambitious goal of mapping Google's serving infrastructure. The paper makes some headway toward that, but doesn't quite achieve that goal. This paper is more aptly named "Toward Mapping...", and must be rewritten accordingly if accepted.

The reason is that the mapping results, while interesting in their own right, do not seem complete: in particular, nothing in the paper gives the reader any sense of the coverage of the mapping effort relative to all of Google's infrastructure. There is no commentary on the impact of potential partial coverage on the overall conclusions the paper is attempting to draw, e.g., observations that many new deployments are in third party ASes -- this is only true for the servers the authors are able to map. IN essence, the overall impact and importance of the mapping effort is unknown.

The techniques used to get there can essentially be described as a mix of known tools and ideas; and, in some cases, the techniques employed have no clear motivation in that it is unclear what the end goal is. The EDNS trick was nice, but not too novel (the use of EDNS has been explored before - e.g., Otto et al's CDN paper (from Northwestern Univ) in IMC'12); the geo-location scheme is fairly straightforward; and the server clustering scheme while interesting has not been mined sufficiently well - the paper looks at number of clusters, but what does simply studying how many clusters deployed in the same area really tell us that's useful in practice? Instead the paper could dig into whether new clusters are deployed in other ISPs, and the reasons for the same. (Also, perhaps there is anecdotal evidence of Google using multiple physically distinct server clusters in the same geo area, which the authors must cite.)

Validation results are unimpressive, if not unconvincing. E.g., the server-clustering scheme is validated by applying to PlanetLab server data. This is a reasonable thing only if we have evidence to believe that Google's deployment pattern reflects PlanetLab's. Other validation results have similar issues. This casts a doubt on the overall mapping analysis. On the whole, this paper is really borderline. In my view, it needs a lot more work to be above the bar.

Overall, this paper is a really nice case and shows how one can introduce new methodology and at the same time use it for a use case, here how to uncover the Google infrastructure evolution.

However, this reviewer finds the term server misleading, as you do not really enumerate servers but server IPs. There can be tens to thousands of servers behind a single IP. The case of Google is one of them as connections can be terminated in frontend servers and then proprietary protocols are used to communicate to the data centers. Datacenters can also host thousands of servers without having publicly addressable IPs. In some places in the text you refer to frontend servers. This is a better term but can you please provide a reference that states that Google frontend servers are assigned one IP per server? It can be also the case that many frontend servers share the same IP (anycast). Later in the text (section 4) you elaborate more on what exactly you are measuring but it would be useful for the reader to mention it up front in the introduction.

This reviewer likes the approach you take to geolocate the Google servers; it is well known that all the publicly available geolocation tools position most of the Google AS15169 in Mountain View. It would be very useful to release your geolocation of Google IPs. In

the case that Google servers are located in other ASes (e.g., GGC servers) how good is the accuracy of Maxmind geolocation, at the country level? My conjecture is that it should be quite acceptable.

The main contributor to the number of server (server IPs) is the ubiquitous deployment of Google Global Cache (GGC) in IXPs and inside a number of ISPs. Can you assess the growth of the overall capacity of Google? My conjecture is that despite the fact that the number of server (server IPs) is doubled mainly due the GGC initiative, the capacity of Google grows (how much is an open question) due to the increase of capacity in the (possible) already assigned IPs to Google. Do you consider the use of Ono or Planet necessary and/or sufficient condition for periodically mapping the Google's infrastructure and geolocate the Google server IPs?

Do you think that you can apply your methodology for CDNs/content providers other than Google?

Were there any frontends your technique was not able to place?

Using airport codes is not exactly ground truth unless someone from Google told you these are always accurate.

"First, we exclude locations that are clearly wrong." -- How is that defined?

Presentation issues:

- The order of figure 4 to 9 is very confusing. Please make sure that the plots appear in incremental number order because it is very difficult for the reader to follow the text.
- You may consider swapping section 7 (Related Work) with section 8 (Using our mapping).
- Section on Accuracy of Frontend Clustering. So what percent of total nodes are in your ground truth? Some quantitative assessment seems glossed over here.
- Figure 7 was hard to see what you're trying to show; maybe additional colors to differentiate new vs. old would help. (And not having one symbol be the same color as the country outlines?)
- "We restrict the clients to those in out BT eyeball dataset and geolocate all client locations using MaxMind" -- could you remind me the numbers/percentages here?
- "in mid-April 2013, many clients are substantially closer" -- how many? Can you quantify "substantially"?
- Section 4 clearly states what is being done but leaves out why the specifics are better than their alternatives. A little bit of intuition would go a long way. In the evaluation of the front-end clustering, it would be nice if the results were further broken down to display the accuracy for clustering different serving sites.
- For the CA experiments, it is not clear if 97% correctly classified belong to different geolocations and the 3% incorrectly classified pertain to serving sites.
- Similarly, it would be nice if figures 3 & 4 were on the same data set. This would help facilitate comparisons.
- The concerns about the accuracy of your technique on clustering geographically close serving sites are further highlighted in section 6.2 when you discuss the problem with Mexico.
- I often wanted to see absolute numbers, sometimes in addition to percentages. "at least 29% more Google front-end addresses". Or vice-versa "we use the subset of Google frontends with hostnames that contain airport codes" -- 550 out of how many? I understand it's in the paper somewhere but it made it hard to read to have to keep all these numbers straight.

4. Summary from PC Discussion

The PC felt the paper attacked an interesting problem space and provided interesting results. The PC found fault with the validation of the proposed techniques and the analysis of the results. Much of the analysis was qualitative rather than quantitative. The PC felt that with some shepherding the analysis could be made quantitative and the validation can be improved. The PC decided to accept with shepherding.

5. Authors' Response

We are grateful for the detailed reviews, which have helped significantly improve the manuscript. Beyond fixing all of the presentation issues that the reviewers have identified, we have sharpened the description of our goals, which include not just the mapping of Google's serving infrastructure, but also an understanding of its evolution. We have also clarified that estimating the capacity increase in Google's infrastructure was an explicit non-goal; in placing front-ends in ISPs around the world, Google's expansion presumably focused on improving the *latency* of Web accesses through split-TCP connections, so proximity of front-ends to clients, and good path performance between clients and front-ends was more important than capacity increases. We capture the proximity in this paper and will explore performance in future work. To achieve these goals, we use a novel combination of mechanisms: the DNS client subnet prefix extension to enumerate client to front-end mappings, a new frontend geolocation scheme for front-ends that uses a large number of possibly inaccurate client locations together with aggressive noise filtering, and a new clustering algorithm that identifies co-located frontends. We believe we are the first to expose this level of detail on the infrastructure, enumerating the front-end IPs, showing where they are and how they are clustered, and mapping which client prefixes are directed to which front-ends.

We have added more quantification to our validation since the submission, within the limits of the ground truth data available. In particular, we are able to estimate the completeness of our approach, by using the property that returned EDNS responses also specify the address prefix for which the response is valid. Using this mechanism, we issued EDNS responses that cover half of the IPv4 address space, and those brute force measurements did not uncover any new frontend IP addresses, suggesting that the /24 prefix approach in our paper likely provides complete coverage of Google's entire frontend serving infrastructure. Furthermore, while the submitted version of the paper validated clustering not just with PlanetLab servers, but also selected Google clusters which had airport codes in them, our revisions incorporate a more careful analysis of the clustering performance using internal consistency checks on the data. Finally, we have carefully explained our choice of the many different techniques we use and how they fit into the broader picture.

Most important, our revised manuscript includes the latest numbers from our continuous measurement of the Google infrastructure, which shows that it has grown by a factor of 6-8x the last eight months or so. It is somewhat unusual to have a ringside seat during the growth of such an important piece of the Internet infrastructure, and we hope readers benefit from the measurement results (publicly available, link in the paper) we have presented.