Consolidated Review of

On the Benefits of Using a Large IXP as an Internet Vantage Point

1. Strengths:

Great dataset and an interesting read. I liked the overall contribution of the paper. The notion that in fact there is a vantage point that can show some of the large scale behavior we know to be happening is cool. And, this is done in a concrete and empirical way such that we get a firm understanding rather than just anecdotal "well, we know Akamai is widely distributed" sort of notions.

A good assessment of what Internet traffic at one particular point looks like today. Some neat observations, like the bit about the hurricane. The paper confirms through a large and meticulous set of measurement data information about the evolution of the Internet -- and about what can be observed through traffic measurements at a large IXP -- using a uniquely powerful vantage point.

2. Weaknesses

As good as the dataset is, it misses a lot of traffic. Further, trends seem to suggest that it will miss more and more. The analysis all seems quite straightforward, with little in the way of new techniques or surprising findings.

The presented data will not surprise anyone (but it does offer quantitative evidence of known trends, which is always a good thing). Most of the findings in this paper are either intuitive or well known in the networking community. So, the work seems of limited impact (beyond usefully putting data behind expectations) unless the researchers make the data available or continue the analysis on an ongoing basis indefinitely, to capture longer term trends. Ideally, both of these will happen, but the paper doesn't seem to suggest either will. I worry that some of the "global role" and coverage arguments are overstated, which distracts from otherwise interesting results. I can see where a lot of people will find the paper somewhat obvious. I.e., if you want to see broad range of Internet traffic, watch at a place where everyone knows broad range of Internet traffic is swapped!

It is nice to have a confirmation through a large set of traffic measurements at a unique vantage point, however the paper is too verbose and 14 pages look too much for content that can be easily summarized (e.g., "summary" subsections of sections 3 and 4 and few sentences in section 6). The initial claims about how the view from this IXP is representative of what can be viewed from other large IXP should be substantiated by measurements. One of the biggest issues of the paper is the claim around the placement of third-party nodes inside of ASes from Akamai, Netflix, etc. The authors claim that this makes ASes (and links) heterogeneous because the AS announcing the address space are not the organizations using it. This does not alter the fact that an AS is still a unit of routing policy as it traditionally has been. Why is it that a server being hosted inside a third-party network makes the AS boundary less defined? The AS is still responsible for carrying that traffic and while it might ultimately belong to a "complex" customer does not mean that the AS did not carry that traffic. If this traffic is carried between ASes then it will still be counted in commercial negotiations between ASes. The ability to attribute traffic to the "correct" party is orthogonal to what can be learned through AS-level modeling. Which brings us back to page 2: "these finding argue that any future attempts at accurately and meaningfully studying the relationships between the different Internet constituents have to move beyond the traditional AS-level view of the Internet and must account for the complexities that result from today's Internet realities". The sentence is puzzling and ambiguous. The context is attributing traffic (previous sentence). Are c2p/p2p/s2s relationships insufficient? What is their suggested alternative? How does a third-party placing servers in a network change any AS relationships in the graph, or impair our ability to model AS-level traffic flow?

I think the paper should note that the observations are largely structural in nature. I.e., the paper does show that there is hosting going on within ASes that is not under the control of that given AS. That pertains to how the network is setup and structured. But, the data they have is in fact skewed in terms of things like determining popularity. And, the sampling makes it less useful for piecing together web sessions. Etc.

3. Comments

A few things in the abstract seemed overstated: Has it really been a long-standing question about whether a single representative vantage point existed? That question never occurred to me. I would agree that an important question is whether we can get access to representative datasets, but that seems like the important part, rather than the question of one VP vs. a number that are accessible and combine well. Also, it's hyperbolic to state that you showed that network heterogenization "impacts each and every aspect." First, you didn't show anything like that. Second, there are plenty of aspects that aren't impacted. Just state that you show that it should affect how we think about the traditional view of the Internet.

The authors report that traffic is visible from all routed prefixes at the IXP (a global view); it would be interesting to see a distribution of traffic per prefix and per country. Figure 3 shows the fraction of IPs geolocated to a country, which are seen at the IXP, only a few countries have more than 5% visible, which seems underwhelming. Perhaps this is due to blind spots (private peering) at the IXP. This seems like something that the IXP could help them with to quantify: i.e. an overview of how many cables are directly between IXP member routers, and the characteristics of the ASes involved. But the claim the IXP provides a global view seems completely insupportable: does the IXP really provide a global view of (e.g.) Chinese traffic? For example, is the traffic to/from these prefixes merely scanning traffic "leaking" to the IXP?

I hope you'll make the data available and also rerun this analysis regularly. Your paper correctly points out that some questions are hard to answer without access to hard-to-get datasets, and your work falls into that category. It would be a great service to the community if you could make both the raw (preprocessed as necessary) and processed (towards the results in your paper) data available regularly, and it would magnify the impact of your

work. I wish there was some capturing of the traffic you miss for the following reasons:

Private peerings. I have no idea how common these are or what trends look like. Is the public fabric used to connect customers to providers, or just peer-to-peer? I worry that it might not be used for customer links (or for peer links with strict traffic ratios) because of the need for more rigorous accounting. Can you at least characterize what types of peerings tend to use private? I assume that the majority of IXP links on public fabric are via open peerings at route servers, but I'm less clear on what happens to other links. Have you asked the IXP for a rough estimate of private vs. public peering? Didn't the earlier SIGCOMM paper on this IXP dataset find that tier 1 traffic was on private peerings? That seems like a lot to miss.

The amount of traffic that stays in network. The 3rd-party serving trend you talk about means that more traffic will stay in network, and you'll see (at most) the cache fill. Does this mean that the IXP will become less useful as a VP over time, as the trend continues? I would think you could capture some of this by comparing the client prefixes you see traffic for non-3rd party servers (like ones that are only served from a single location, say) to the ones you see traffic for Google/Akamai. Or, you could use EDNS client-subnet and open resolvers to see where clients are directed for popular services. Do the trends towards more 3rd party deployments suggest that your IXP will see less and less of the traffic over time? For example, in Figure 6(a), the IXP misses the majority of links (and presumably misses the "priority" ones that Google and Akamai will attempt to serve from as much as possible).

2.2.1: Is the TCP vs. UDP split by packet count or by volume? By volume this is quite a bit different than I have seen recently.

2.2.2: Given the sampling, don't you miss a ton of possible HTTP traffic? And, even that which you do see might not include the HTTP headers, right? I.e., it might be some packet from the middle of a jpeg transfer or something. I assume the idea is that you overcome such things with shear volume. I.e., all it takes is catching one good header to ID a server. A little more explanation here would be useful.

2.2.2: I liked the active probing step. Makes the passive observations more sound. Well done.

In general it is a plus that the paper uses various bits and pieces of ancillary data to confirm the findings from the main passive trace data. This makes the findings sounder. The paper does use subjective assessment too much in some places. E.g., right at the end of section 2 we find the words "only so slightly". It'd be better if the change were quantified. E.g., "decrease by 0.1%". I really have no idea if "every so slightly" is a fraction of a percent or a few percent or what. This is an exemplar. I would suggest the authors re-read the paper and scrub all these sorts of things from the paper---especially when they are the only assessment given (e.g., something like "the reduction is small (1.2%)" would be OK).

3.1: Can you characterize the prefixes/ASes you don't see (out of all that exist)? Are there regional / tier / etc. trends? Similarly, are the 45K server IPs seen by the large ISP but not the IXP from orgs that you already see server IPs from, or from different ones? One thing I'm curious about: how many of the links do you need to see as many prefixes as you see? Are there a few links that are absolutely critical for the coverage you see?

3.1: I'm concerned that your large European Tier-1 ISP and your large European IXP datasets are not nearly as orthogonal as you claim. Wouldn't we expect the fact that they are hosted in the same country to mean that they serve many of the same customers? I realize that the tier 1 has a wide presence, but I would guess that it has a particularly big presence near where it is headquartered. I also didn't understand how you were using the validation: "for the server IPs seen both at the IXP and by the ISP, those we identified as server IPs using the IXP data are confirmed...." At least as phrased, this is circular: if we remove everything about the IXP data, it says that server IPs seen by the ISP are confirmed as server IPs using the ISP data. It doesn't comment on IPs that were not identified as servers by the ISP data (but may have been by the IXP data).

Section 3.1: Could the authors justify their statement that the IXP and ISP datasets are "orthogonal"? Whether this is the case or not depends on the relative location of the IXP and ISP, as well as their business relationship (e.g., does the ISP sends a significant fraction of its traffic through the IXP?)

Sections 3.1 and 3.2: It is not 100% clear what it means that the IXP is "local yet global". Does it mean that it sees traffic from all over the world (hence "global"), yet a significant fraction of this traffic still comes from Europe (hence "local")? Could the authors explain why this "dual role" is important? What kind of conclusions can it help us draw that we would not be able to draw, say, from a global-only dataset?

3.2: Similarly, I'm a bit confused about how you are showing global role in 3.2.

In Table 2, how do you account for the fact that traffic has both a sender and a receiver? My best guess is that you are just looking at sources, but it isn't stated clearly. For the top part of the table (IPs), a flow could count for both source and destination? What about for traffic? When you say, for example, that the IXP is important for US traffic, does this just mean that lots of European clients access US servers (or vice versa), or does it actually mean that there are lots of instances when neither of the endpoints is in the Europe) (which seems surprising)? It seems like you need to account for locations of both endpoints, not just one. This again makes Table 3 confusing: traffic has two endpoints, and (say) one could be in A(L) and one could be in A(M), but the rows sum to 100%, so I have no idea how your accounting assigns that traffic (which seems like it should count in both groups). Maybe you need the table to show source and destination breakdowns separately, and a combined one that assigns each flow to the nearest group between source and destination? The phrase "there is potentially significant overlap between the sets" is hard to make sense of when you defined A(G) as the complement (and hence by definition non-overlapping). When you assign country in Table 2, will it assign all traffic / IPs from Google to US, or will it actually look for the location of the Google server IP (for which you can't use something like Maxmind)? It would be interesting to do more of a classification of traffic by AS type and where sources/destinations are relative to IXP. Some of your results start towards that, but I think it would be interesting to look further.

3.3: How/why would many dynamic or ephemeral hostnames make the list of the top 1000 sites? If they make that list, aren't they by definition important/popular?

Section 3.3: Could the authors explicitly state their conclusion from this section? My conclusion is this: the IXP observes a large number of server IPs, but there is no way of knowing how many more server IPs there are in the world. But I get the sense that I am missing something.

Section 4.1: What are plausible reasons why a significant number of server IPs disappear? Could we somehow use the data collected at the IXP to reason about the "stability" of e-commerce, i.e., the rate at which such sites appear and go out of business? - Could we somehow use the data collected at the IXP to estimate what fraction of Internet traffic is "suspicious"? E.g., try to identify how much traffic could be part of a flooding attack or port/HTTP scanning, how much traffic involves blacklisted domain names, etc.?

Figure 5 is impossible to really read. At least the authors should have used the full column width. But, even so, I think there is just too much on this one.

Likewise 4(b) just includes too many plots to take all that much away. I can see some patterns, but the magnitudes escape me.

While the paper finds the server population is somewhat stable across time, I wondered about the client population. Is that similar? I am having an argument in my head about whether I think it would be or not.

4.2 is just ***full*** of subjective language when surely there are quantitative assessments behind these that could have been better described.

Section 5 (and really throughout) notes that this AS-centric view of the world we have in our head is outdated. I buy the argument the authors are making (even before they made it!). But, the AS-centric view is actually still useful in some aspects. E.g., routing. I am not sure I really bought the "new way is better than the old" argument. Rather, it seems this way of looking at the network advocated by the authors is an *additional* way to view things. In some cases and for some things it is better. But, not for everything. Perhaps a slight re-framing this as an added perspective would be useful.

5: The paper notes that traffic from some non-AS organizations "goes unnoticed". But, its worse than that ... it is noticed and mis-attributed to others.

5.1: You seem to state how you post-process the clusters to assign them but not how you perform the initial clustering. For example, how do you get the clusters that have mostly the same authority (second step)? I was a bit confused by the third step. Aren't Akamai and Google the CDNs most deeply deployed inside ISPs, and they are taken care of in the 1st step? Which are the 3rd step CDNs? What is the cause of the <3% false positives? Do you have false negatives, splitting apart common infrastructures?

5.2: It seems well known that these sorts of serving infrastructures are becoming more common. It would be great if you could carve out space to talk more about what you observe them to look like (the "bewildering array of scenarios").

5.3: Is the CloudFront vs. EC2 explanation just that they only announce most CF prefixes to peers, whereas EC2 has to also go to providers? At least a number of years ago, CF had the reputation for being "cheap" when it came to paying for transit.

Fig 7(b): I don't understand how to read this graph. My understanding is that X axis is, for a given AS, out of all Akamai traffic you observe towards that AS, the % that you observe over the Akamai link. Is the Y axis the fraction of all traffic you observe to that AS that is Akamai traffic? (I'm now convinced that is what it shows, but the explanation should make that clearer - it doesn't say what the fraction is out of) You mention that the traditional assumption would be all points at x=100. However, Akamai's "ideal" (in one sense of extreme deployment to maximize locality) would also be all points at X=100: all traffic on these links is cache fill / backend fetching, and all clients are served from their local ISP (and hence invisible at the IXP). Given the traffic that doesn't cross the IXP, I don't think it makes sense to say, "Akamai sends 11.1% of its traffic not via its peering links." I also didn't really understand the statement that "traffic from more than 15K...is seen at the IXP via non-IXP member links." Is this the same as stating the number of server IPs that aren't hosted inside Akamai?

4. Summary from PC Discussion

The reviewers were mostly positive, because of the rich dataset and the useful observations. There was some concern regarding the writing style of the paper (which the reviewers thought was in places non-scientific), but the consensus was that this could be fixed through shepherding

5. Authors' Response

In response to the reviewer's comments we have made the following changes:

1) We modified Sections 1 and 7 to better articulate one of our main observations; that is, while the traditional AS-level view of the Internet has some value for exploring and understanding various connectivity- and reachability-related questions, it's largely traffic-agnostic nature severely limits its usefulness for accurately and meaningfully studying either the business strategies of the different Internet constituents or the business relationships among them. In particular, we argue that future attempts at studying the AS-level Internet need to move beyond this traditional view and have to account for how traffic and hence money flows in today's Internet.

2) In Section 2, we expanded the description of our methodology for identifying HTTPS servers and also added a discussion of the methodology's limitations. In addition, where possible, we replaced some "subjective language" with more quantitative assessments or actual numbers (we did this also for Section 4).

3) We modified the discussion in Section 3 to better articulate our two key observations. First, the main reason for why IXPs exist in the first place is to keep local traffic local; as such, their role as local players is commonly accepted. Second, at our large European IXP, we see traffic from every part of the world/Internet which in turn highlights that this IXP also plays the role of a global player. In addition, we also explain why the ISP and IXP data sets used in this section are orthogonal.

4) In Section 5, we expanded the discussion of our clustering methodology by including a more detailed description of each of the three steps. We also clarified the discussion concerning Akamai traffic on peering links where its existence would go un-noticed in the traditional AS-level view of the Internet.

5) In Section 6, we included a reference to a case that supports our claim the largest IXPs in Europe are generally open to collaborations with researchers and are supportive of research efforts that make explicit use of their data.

6) In Section 7, we now mention that we expect that the IXP will continue to be valuable vantage points and that our expectation is that as a consequence of more servers being deployed close to the end users, IXPs in the future will "see" less end user-to-server traffic but an increasing amount of server-to-server traffic.