Growth Analysis of a Large ISP

Andrew D. Ferguson
Brown University
adf@cs.brown.edu

Jordan Place
Brown University
jplace@cs.brown.edu

Rodrigo Fonseca
Brown University
rfonseca@cs.brown.edu

ABSTRACT

We present a time-series analysis of Cogent's inter-continental network. The analysis is based on descriptions of Cogent's routers and their interfaces, collected each week for more than one year. These descriptions are collected from public reverse DNS records, which we cross-validate using `iffinder`, a full Internet scan, and limited ground truth data provided by Cogent. For example, our dataset, which we make available to the research community, shows that while the number of Cogent routers grew by approximately 11.3 each week, the average number of interfaces per router, and the effective diameter of the inferred network remained stable over the same period. Our collected dataset includes information about interface types, port identifications, router locations, peer and customer attachments, and more.

Categories and Subject Descriptors

C.2.1 [**Computer-Communication Networks**]: Network Architecture and Design—*Network topology*; C.4 [**Performance of Systems**]: Measurement techniques

Keywords

Reverse DNS; Alias resolution

1. INTRODUCTION

Broadly speaking, IP addresses on the Internet fall into one of two categories – they either represent hosts or other endpoints, or router interfaces. Mapping the connections between, and properties of, the router interfaces (*e.g.*, ownership, geographic location, logical location, etc.) is important for understanding the Internet's topological structure, and has been a topic of extensive research (cf., §2). Typical mapping approaches to reconstruct the router-level topology of the Internet include data from `traceroute`-like probes [25], multicast advertisements [20], IP options probing [9, 21], and DNS records [24].

As shown below, DNS records can be a rich source of information; yet, they are potentially problematic [27], and are not used in large-scale Internet topology mapping performed by CAIDA [16].

IMC'13, October 23–25, 2013, Barcelona, Spain.

ACM 978-1-4503-1953-9/13/10 ...$15.00.
http://dx.doi.org/10.1145/2504730.2504769.

For example, there are no standards for naming interfaces (resulting in idiosyncratic rules for each ISP), and there is no requirement that an interface's DNS record remain up-to-date as interfaces are added, reconfigured, or removed. Within a single ISP, however, such difficulties may be mitigated, as organizations presumably strive for consistent administrative procedures and best practices.

We find that Cogent, one of the most connected Internet service providers [1], with large networks in both Europe and North America, is such an ISP. Cogent provides reverse DNS records for more than 99% of the 51,000+ interfaces on Cogent-owned routers we could identify. For example,

```
te2-1.ccr01.jfk01.atlas.cogentco.com
fa0-2.na01.b003070-1.sfo04.atlas.cogentco.com
```

are, respectively, the reverse DNS records for the Cogent-owned IP addresses `154.54.80.85` and `38.112.5.17` during the week of March 10, 2013.

We assume that all records under the `*.atlas.cogentco.com` DNS hierarchy are part of Cogent's infrastructure.[1] These records include four pieces of information. First, a router location (*e.g.*, `jfk01` and `sfo04`) – we find Cogent has 460 router locations, almost all of which are coded with three-letter airport codes. Second, the router within a location (*e.g.*, `ccr01` and `na01.b003070-1`) – we estimate Cogent had 4,469 routers the week of March 10, 2013. Third, the type of interface, which we infer based on Cisco naming conventions (*e.g.*, `te` for 10 Gbps Ethernet, and `fa` for 100 Mbps Ethernet). And fourth, the interface's position within the router (*e.g.*, `2-1` and `0-2`, which are, respectively, the first and second ports on their line cards).

Ideally, with such structured records, we could determine the existence of a 10 Gbps interface at position `2-2` on the `ccr01.jfk01` router with a simple DNS query. Unfortunately, Cogent only provides *reverse* DNS records. However, by issuing reverse DNS queries for *all* Cogent-owned IPv4 addresses, we find that the IP address at position `2-2` on that router is `154.54.25.17`. In addition, because of the exhaustive lookup, we find that this router appears to have 18 such 10 Gbps interfaces configured across five line cards.

Other Cogent DNS records include information about related business entities. For example, `Tetratech.demarc.cogentco.com` is the reverse DNS record for `38.112.5.18`; hence, we infer that Tetratech is connected to Cogent with up to 100 Mbps of available bandwidth at a router near San Francisco, based on the information about `38.112.5.17` (the other usable address in the `38.112.5.16/30` subnet) inferred above.

To capitalize on this wealth of information, we have issued reverse DNS queries each week for more than 17 million Cogent-owned IPv4 addresses (now more than 20 million, see §4.3), start-

[1]Excluding approximately a dozen mis-named addresses such as `fixme_please.atlas.cogentco.com`.

ing the week of January 22, 2012. These weekly snapshots allow us to analyze the growth of – and change in – Cogent's network at the router-level.

This dataset, which we release to the research community, has several interesting features:

- Extensive records of the evolution of a large ISP, providing a platform for future network research – either directly using the dataset, or by offering improvements to existing topology generators [26]
- Novel information about interface types and positions
- Novel estimates of router, module, and interface growth rates for a large ISP
- A new dataset for evaluating methods to de-alias router interfaces, which are used when discovering Internet topologies using `traceroute`-like measurements (§2)
- Results from more than one billion DNS queries issued from approximately 100 globally-distributed vantage points, including more than 100,000 anomalies (*e.g.*, ID mismatches, replies from unexpected sources, corrupt responses, etc.)
- Validity and coverage-rate established through comparison with `iffinder`, public ground truth, and a complete set of IPv4 reverse DNS records

The rest of this paper proceeds as follows: first, we provide context for our study in the space of related work (§2); second, we describe and validate the approach used to build this dataset (§3, §4); next, we conduct an initial time-series analysis of this data (§5); and finally, conclude with potential avenues for further research (§6).

2. RELATED WORK

Constructing router-level Internet topologies has been a project of the networking research community for more than 10 years [14]. A common approach is to use `traceroute`-like probes with successively increasing TTL values. This approach was made scalable by targeting probes more efficiently with Rocketfuel [25], more transparent to network middleboxes through the use of TCP Sidecar [23], and has been augmented with information gathered by the IP Record Route option [22]. At their core, these `traceroute`-like approaches provide a list of router IP addresses which can be viewed as *aliases* of a single router, to be determined [14].

Therefore, numerous solutions to the *alias resolution* problem have been developed including: examining IP ID values [9, 17, 25], using prespecified IP timestamps [21], detecting variants in the source address of a probe's response [14], applying graph analysis techniques [15], and, as in this work, by examining reverse DNS lookups [24].

In this work, however, reverse DNS lookups are the *only* method for discovering router interfaces – we do not send any TTL-limited active probes. Other techniques for topological discovery without `traceroute`-like probes include inference from passive measurements [13], and gathering data with MPLS ICMP extensions [23] and IGMP messages [20].

Using reverse DNS records for this application is not without risks, however. Previous work has shown that using out-of-date and other incorrect reverse DNS records can have a disproportionately negative effect on POP-level path reconstruction [27]. In this work, by contrast, we primarily focus on inferring Cogent's router organization and interface properties, and their change over time. To

(a) Breakdown of all DNS probes we sent.

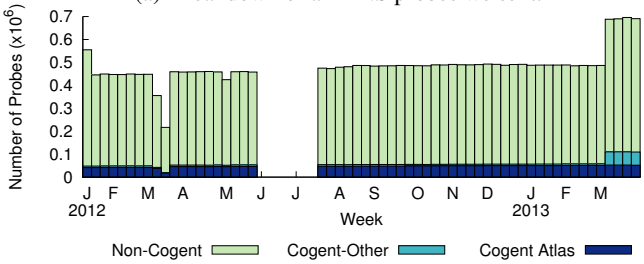

(b) Breakdown of successful DNS probes

Figure 1: Responses to weekly DNS probes.

the best of our knowledge, this information is only publicly available from DNS records; therefore, such inference requires careful validation of our DNS-based mapping, as we show in §4.

Perhaps the most closely related work is CAIDA's "IPv4 Routed /24 DNS Names Dataset" [7]. This dataset consists of reverse DNS records for router interfaces discovered by `traceroute`-like active probes. CAIDA's probes are sent weekly to a random IP address in every Internet-routed `/24` subnet. This measurement naturally aims for *breadth* of coverage. Our work, by contrast, aims to achieve *depth* – complete coverage of a single ISP, as we show in our validation (§4).

3. METHODOLOGY

This paper analyzes weekly snapshots of Cogent's router organization, built from publicly available reverse DNS records. As described above, we systematically issue reverse DNS queries for each Cogent-owned IP address, allowing us to discover all of the corresponding interface names. This initially required more than 17 million weekly DNS queries; we now issue more than 20 million each week.

To minimize our load on the DNS infrastructure, we globally distributed our queries at 90-100 sites using PlanetLab [12]. Worker processes at each site issue queries for all IP addresses in a block no larger than a `/19` subnet (8,192 addresses), at a rate of approximately one query per second using the Linux `host` command. Between assignments, workers pause for a random sleep of between 10 to 20 minutes. Workers notify our master server every 256 queries (allowing us to detect and re-assign incomplete blocks), and upload a log when finished. With this approach, we can successfully query millions of DNS records each week.

Figure 1 provides an overview of our DNS queries' responses. During the weeks of March 11 and 18, 2012, our master server experienced data loss. In addition, our measurement infrastructure was unavailable for seven weeks (from the week of June 3, 2012

Code	Type	Class
et	10 Mbps Ethernet	Physical
fa	100 Mbps Ethernet	Physical
gi	1 Gbps Ethernet	Physical
te	10 Gbps Ethernet	Physical
se	Serial link	Physical
pos	Packet-over-SONET	Physical
ism	Integrated Services Module	Physical
lo	Loopback	Virtual
mul	Multilink	Virtual
tu	Tunnel	Virtual
vl	VLAN	Virtual

Table 1: Interface types observed in Cogent's network.

until the week of July 22, 2012, and partially during the week of May 6, 2012) due to hardware failures affecting the master server.

We also run `iffinder` [5] each week on the list of interface IP addresses discovered by the previous week's DNS queries – that is, all of the IP addresses with reverse DNS records under the `*.atlas.cogentco.com` hierarchy. In the first week, this was 42,100 interfaces (62,906 probes); as of the week of April 28, 2013, it is now 53,457 interfaces (81,736 probes). These probes are not distributed, and currently complete in approx. 38 hours (up from approx. 29 hours in the first weeks). The weekly `iffinder` data is used to validate our grouping of interfaces into routers (§4.1).

We group interfaces into routers based on the DNS records; generalizing from the examples in §1, we consider Cogent's reverse DNS records to consist of three fields:
`(interface).(router).(location).atlas.cogentco.com`, where the router field may include one or more levels of the DNS hierarchy. Additionally, we infer interface types using the Cisco naming conventions in Table 1.

Finally, to account for transient DNS failures, we smooth our weekly dataset based on surrounding weeks according to the following rule: if an interface is missing in Week N, yet present in Weeks $(N-1)$ and $(N+1)$ with the same corresponding IP address, we consider its absence in week N to be accidental. Excluding the weeks of March 11 and 18, 2012, this smoothing process increased the number of interfaces each week by 0.18% on average. By this process, an interface not present for two consecutive weeks is considered to be removed.

All of our code for measurement, parsing, smoothing, graphing, and analysis is publicly available.[2]

4. VALIDATION

Internet topographers are faced with two challenges in their work. First, because most Internet maps are built by inference, they must asses a map's validity, comparing it with maps built with other techniques, or by comparing with the limited ground-truth data available. Second, they must determine the completeness of their coverage. In this section, we address these challenges through comparison with three additional datasets: router aliases determined with `iffinder` [5], public information provided by Cogent, and a complete set of reverse DNS records collected by an anonymous "Internet Census" project [6].

4.1 Validity of router de-aliasing

As previously discussed, ISPs are not required to maintain consistency between a router interface's true location and its reverse DNS record. This makes any effort to group interfaces into routers by using DNS records potentially problematic. Therefore, to double-check our efforts, we ran `iffinder` [5], a traditional solution to the alias resolution problem, each week on the interfaces discovered in the previous week.

The `iffinder` tool works in the following manner: for each interface, it sends one or more UDP packets to high-numbered ports; these packets are designed to elicit ICMP Port Unreachable messages in response. If received, the source address of any ICMP message is assumed to also be an interface on the same router, assuming this source address is different from the destination address of the original probe [14]. This technique produces a list of IP address pairs which `iffinder` infers are aliases of the same router. By taking the transitive closure of these pairs, we produce a candidate set of router aliases (*e.g.*, if we see pairs $(A,B),(B,C),(D,C)$, we infer a single router with interfaces $\{A,B,C,D\}$).

The `iffinder` approach naturally leads to a higher rate of false negatives (failing to infer two interfaces are on the same router) than false positives (inaccurately inferring two interfaces are on the same router). Therefore, to validate our DNS-based grouping, we search for instances in which our DNS-based approach infers two or more routers for a set of interfaces, yet `iffinder` infers only one. If such instances occurred frequently and persistently, it would suggest that Cogent fails to keep their interface's reverse DNS records consistent with reality.

Our analysis implies that Cogent maintains their interface's DNS records consistently. In each week, less than 1% of the candidate routers inferred by `iffinder` contained IP address with reverse DNS records suggesting multiple routers. Furthermore, 95.8% of these discrepancies last for one week or less, and could be due to the delay between our platform's DNS queries and `iffinder` probes. We find that only 1.7% of discrepancies persist for more than two weeks, affecting just 28 interfaces (0.053% of all interfaces).

Finally, we note that `iffinder`'s technique also discovers new router interfaces, not in our DNS-based dataset. On average, these interfaces, which lack reverse DNS records, accounted for 0.86% of the total number of interfaces discovered by either approach.

4.2 Comparison with public information

Cogent provides a graphical network overview on their public website [4]. This map is very coarse, unfortunately, and only contains a single point for each of the 192 cities in which Cogent maintains a presence. Furthermore, this public graph only contains edges between geographically neighboring cities. Nevertheless, this graph provides a public source of ground truth about Cogent's network.

To compare, we extracted the list of 187 airport-like codes discovered in our dataset (*e.g.*, `jfk`, `sfo`, `fco`), and plotted each inferred location on a map. Through manual inspection, we determined that our dataset is missing only nine cities (4.7%) – three in the US, and six in Europe. In addition, our dataset contains airport codes for four cities not shown on the published map – one in the US, and three in Europe. `traceroute` probes to interfaces we believe are located in these four cities revealed IP paths consistent with the new cities' presumed locations.

While not conclusive, we believe the close similarity between our map and Cogent's suggests our DNS record-based dataset covers around 95% of the cities in Cogent's network.

In addition, Cogent's 10-K Annual Report provides public declarations about their network [2]. This legal document describes the services Cogent provides to its customers, which are consistent with the list of inferred interface types given in Table 1. The annual report also states that their primary service is provided at 100 Mbps,

[2] `http://github.com/brownsys/pl-mapping/`

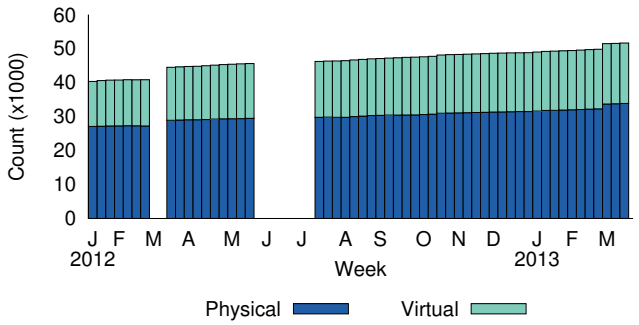

Figure 2: Count of physical and virtual interfaces in Cogent's network (stacked to show total).

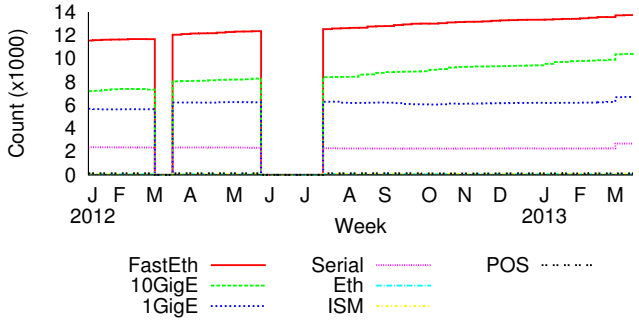

Figure 3: Weekly count of physical interfaces by type.

which matches the composition of physical interfaces we observed, as shown in Figure 3.

4.3 Reverse DNS coverage

The initial list of IP addresses for which we issue reverse DNS queries each week was chosen based on IP addresses observed in select `traceroute` probes. With these addresses, we identified the corresponding Cogent-owned CIDR address blocks using the ARIN and RIPE registry databases. This list contained 17,731,584 IPv4 addresses, and was used through the week of March 3, 2013.

Using data collected by a recent "Internet Census" [6], which issued multiple reverse DNS queries for the entire IPv4 address space, we identified an additional 2,426,624 Cogent-owned IP addresses which could potentially contain `*.cogentco.com` reverse DNS records. While our original query list contained just 88% of the expanded list, we found the number of `*.atlas.cogentco.com` records only increased from 50,981 to 52,643 (a 3.3% increase, cf. Figure 1(b)). Therefore, as with the map comparison described above, we believe this comparison suggests our targeted DNS-based mapping well covers Cogent's network.

5. EXPLORATION

Having established the validity and coverage of our data, we now conduct an initial time-series exploration of Cogent's network. Our analysis proceeds at three levels: first, considering the interfaces, independent of the routers; second, considering the routers, independent of their connections; and third, considering the inferred network graph.

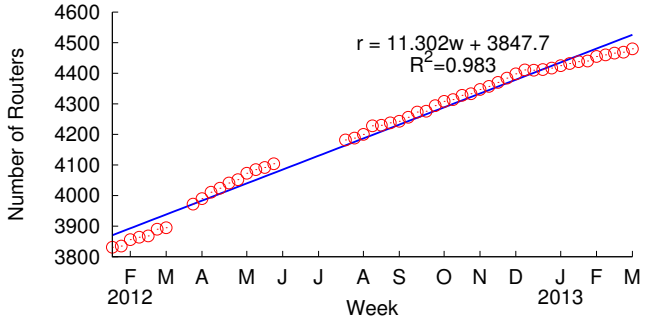

Figure 4: Number of routers over time.

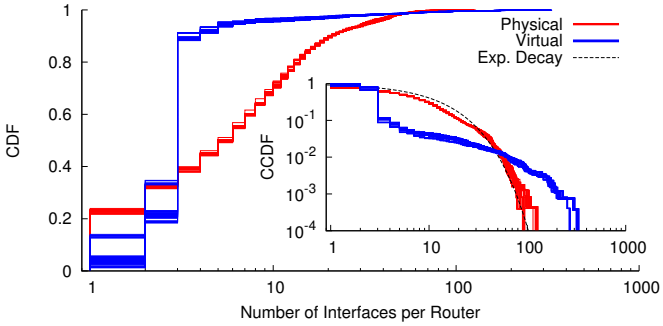

Figure 5: CDF and complementary-CDF of the number of physical and virtual interfaces per router for all weeks.

5.1 Interface evolution

First, we consider Cogent's interfaces independently of their associated routers. In Figure 2, we see that a majority of interfaces are physical interfaces, which we further breakdown in Figure 3. The virtual interfaces are predominantly VLAN interfaces, plus loopback interfaces for each router. Based on our dataset, we infer that the three most common types of physical interfaces in Cogent's network are 100 Mbps Ethernet, 10 Gbps Ethernet, and 1 Gbps Ethernet, respectively.

We find that the number of interfaces grew linearly over the period of our data collection. We infer that Cogent adds an average of approximately 153 new interfaces each week ($R^2 = 0.92$). We also infer that this growth is not proportional to the existing distribution of interface types. Figure 3 shows that the rate of growth is highest for 10 Gbps Ethernet interfaces, followed by that for 100 Mbps Ethernet.

5.2 Router evolution

Next, we consider Cogent's routers, independent of the connections between them. In Figure 4, we plot the number of inferred Cogent routers over time, and find that Cogent adds an average of approximately 11 new routers per week. We note that the apparent recent dip in growth rate occurs at the start of 2013.

We find that the distribution of the number of interfaces per router remains relatively stable over the course of our data collection. The average number of interfaces only rises from 10.7 per router to 11.4, while the median remains six. In Figure 5, we separate the interfaces into physical and virtual interfaces (see Table 1), and find that the weekly distributions remain stable throughout the study. The number of physical interfaces decays as an exponential

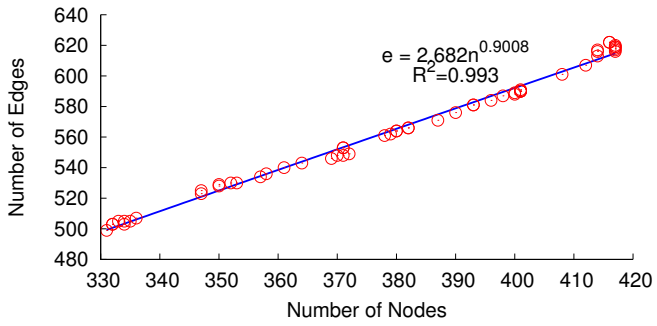

Figure 7: Number of nodes versus the number of edges of the inferred network topology with one point per week.

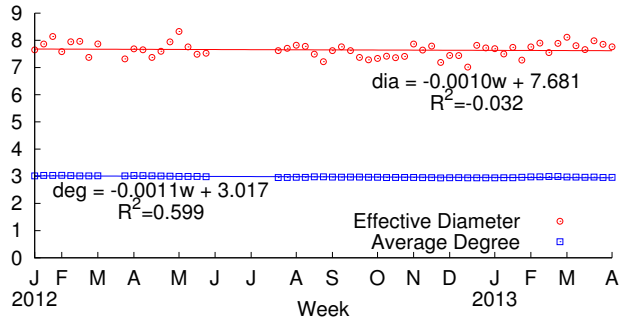

Figure 8: Evolution of the effective diameter and average degree of the inferred network topology. Both quantities remained fairly constant in the measurement period.

distribution, while the number of virtual interfaces is heavily concentrated around 3, and has a heavier tail.

Cogent's physical interface names also contain line card and port positions. Examining data for the week of April 14, 2013, we find that routers have 1.56 line cards on average, with an average port density per line card of 4.21. The maximum number of physical ports we observe on a single router is 90, and three routers appear to have nine line cards configured (the maximum we observed). The maximum number of configured ports we observe on a single line card is 43, with 1 Gbps Ethernet interfaces; the highest port number we record on such line cards is 48. For 10 Gbps Ethernet interfaces, the highest port number we record is eight.

5.3 Network evolution

Finally, we analyze our inferred graph of Cogent's network. While passively collecting Cogent's reverse DNS records each week provides us with a detailed list of interfaces, the records provide no explicit information about connections between routers. Therefore, we use the following procedure to infer a graph of their network. First, we observe that connected routers must each have an interface on a shared subnet. Thus, as in Sidecar [23], we collect candidate pairs of interfaces whose IP addresses appear to share a `/30` subnet (addresses must be off-by-one, and the inferred network and broadcast addresses must be valid and not assigned to other interfaces). Second, we conservatively consider only candidate pairs for which both are physical interfaces, and of the same type. Finally, we consult Cogent's BGP looking glass server [3], and discard any pairs where the corresponding prefix is larger than `/30`.

Figure 6 shows a visualization of the inferred network for the week of April 7, 2013, where we have grouped together all routers belonging to the same site. The graph, produced using Gephi [8], is colored according to a community detection algorithm [10], and the size of the nodes is proportional to their betweenness centrality, *i.e.*, related to the number of shortest paths the node is part of. Even though the node positions resemble geographical positions, they are computed by a force-directed layout algorithm with no geographic information. This is a further indication that the inferred network is related to the real one.

We constructed an equivalent graph for each valid week in the dataset, and now look into the evolution of a few important graph metrics. Leskovec *et al.* [18] examined the evolution of several real networks, such as the AS graph and the arXiv citation network, and found two phenomena: densification, and shrinking diameters. Densification states that the number of edges $e(t)$, and the number of nodes $n(t)$, over time follow the relation $e(t) \propto n(t)^{\alpha}$, with $1 < \alpha < 2$. This implies that the average degree increases over time. They also found that the effective diameter of the network shrinks over time.[3] We found neither of these phenomena to be significant in the evolution of the inferred Cogent graph.

Figure 7 shows a scatterplot of the number of nodes versus the number of edges for the inferred network graph. The relation is very close to linear, but the exponent of the densification law is 0.9. Correspondingly, the bottom curve in Figure 8 shows that the average degree actually decreases, with a small negative slope of -0.0011 in a linear fit. The top curve in Figure 8 shows the evolution of the effective diameter of the graph. The diameter remains nearly constant, with a negative slope of -0.001 in a linear fit.

Since our collection period of slightly more than one year is relatively short compared to the analysis in [18], it remains to be seen if these trends continue as we collect more data.

6. CONCLUSION

As noted by previous researches, DNS queries are a potentially problematic source of topological information; we find this is not the case for Cogent's network. We were lucky that Cogent's interface naming scheme was both rich with data and obvious to parse. Finding other ISPs with similar naming schemes and sensibilities is simply a matter of chance; if even feasible, it would require human intervention to build a database of rules, similar to Rocketfuel's `undns` tool, which attempts to map routers to physical locations based on DNS records [25].

Fortunately, Chabarek and Barford are currently developing such a database [11]. Using a combination of machine learning-based clustering and hand-crafted regular expressions, they have partially discovered interface speeds, equipment types, and router functions from the DNS addresses of approximately 26,000 organizations. While their initial study used only the DNS records in the CAIDA dataset previously mentioned [7], the resulting rule database could be used to expand the approach taken in this work to other ISPs.

Unfortunately, the future of performing reverse DNS queries on a large-scale looks hazy. The simple approach taken here of issuing queries at a low rate from hosts at around 100 PlanetLab locations, will not scale as ISP networks transition to IPv6. Therefore, it is important for the research community to measure as many network properties as possible while the scale of IP address blocks is still relatively small.

[3]From [18], the effective diameter is the 90th percentile of the continuous interpolation of the node distance cumulative distribution.

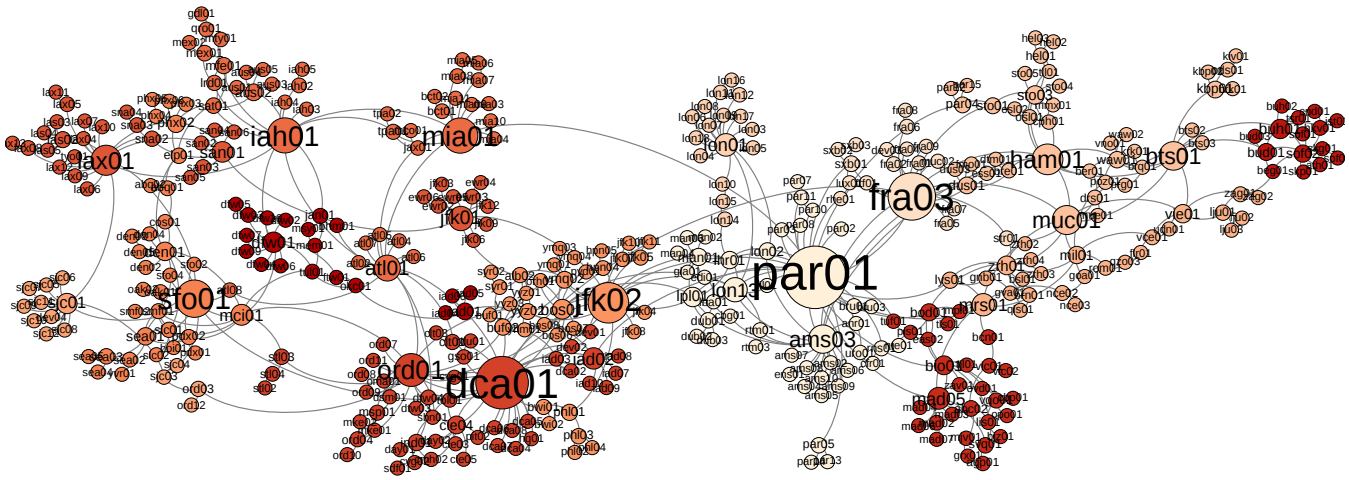

Figure 6: Visualization of paths in Cogent's network based on data from the week of April 7, 2013; nodes represent routers, edges link routers sharing the same IP subnet, and nodes are scaled to represent *betweenness* – larger nodes have a greater number of paths passing through them. The layout is force-directed, with no geographical information.

We hope that releasing the complete record of our weekly probes will be interesting and relevant to the research community. While this paper provides an extensive verification of our measurement approach, and an initial analysis of the dynamics within Cogent's own routers, we suspect additional analysis – such as of Cogent's peers and customers, or of the occasional DNS anomalies experienced by our more than one billion queries – could be possible. In addition, combining this dataset with others, such as those gathered by the iPlane project [19] or CAIDA, could be fruitful avenues of further research.

Acknowledgments

The authors thank David Trejo for plotting Cogent's routers on a physical map, and the staff of the Planet Lab Consortium for technical support. We also wish to thank our reviewers and shepherd, kc claffy, for their many helpful comments. This work was partially supported by NSF grant 1012060. Andrew Ferguson is supported by an NDSEG fellowship.

7. REFERENCES

[1] AS Rank: AS Ranking, 2013. http://as-rank.caida.org.

[2] Cogent Communications Group, Inc. - SEC Form 10-K Annual Report, 2012. http://www.cogentco.com/files/docs/about_cogent/investor_relations/reports/10k_report_2012.pdf.

[3] Cogent: Looking Glass, 2013. http://www.cogentco.com/en/network/looking-glass.

[4] Cogent: Network Map, 2013. http://cogentco.com/en/network/network-map.

[5] iffinder Alias Resolution Tool, 2012. http://www.caida.org/tools/measurement/iffinder/.

[6] Internet Census 2012. http://internetcensus2012.github.io/InternetCensus2012/.

[7] The CAIDA UCSD IPv4 Routed /24 DNS Names Dataset. http://www.caida.org/data/active/ipv4_dnsnames_dataset.xml.

[8] M. Bastian, S. Heymann, and M. Jacomy. Gephi: An Open Source Software for Exploring and Manipulating Networks. In *International AAAI Conference on Weblogs and Social Media*, 2009.

[9] A. Bender, R. Sherwood, and N. Spring. Fixing Ally's Growing Pains with Velocity Modeling. In *IMC '08*.

[10] V. D. Blondel, J.-L. Guillaume, R. Lambiotte, and E. Lefebvre. Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment*, October 2008.

[11] J. Chabarek and P. Barford. WhatâĂŹs in a Name? Decoding Router Interface Names. In *HotPlanet '13*.

[12] B. Chun, D. Culler, T. Roscoe, A. Bavier, L. Peterson, M. Wawrzoniak, and M. Bowman. PlanetLab: An Overlay Testbed for Broad-Coverage Services. *SIGCOMM CCR*, July 2003.

[13] B. Eriksson, P. Barford, R. Nowak, and M. Crovella. Learning Network Structure from Passive Measurements. In *IMC '07*.

[14] R. Govindan and H. Tangmunarunkit. Heuristics for Internet Map Discovery. In *INFOCOM '00*.

[15] M. Gunes and K. Sarac. Analytical IP Alias Resolution. In *IEEE International Conference on Communications (ICC 2006)*.

[16] K. Keys. Internet-Scale IP Alias Resolution Techniques. *SIGCOMM CCR*, Jan. 2010.

[17] K. Keys, Y. Hyun, M. Luckie, and k. claffy. Internet-Scale IPv4 Alias Resolution with MIDAR. *IEEE/ACM Transactions on Networking*, PP(99), May 2012.

[18] J. Leskovec, J. Kleinberg, and C. Faloutsos. Graph evolution: Densification and shrinking diameters. *ACM Trans. Knowl. Discov. Data*, 1(1), Mar. 2007.

[19] H. V. Madhyastha, T. Isdal, M. Piatek, C. Dixon, T. Anderson, A. Krishnamurthy, and A. Venkataramani. iPlane: An Information Plane for Distributed Services. In *OSDI '06*.

[20] P. Mérindol, V. V. den Schrieck, B. Donnet, O. Bonaventure, and J.-J. Pansiot. Quantifying ASes Multiconnectivity Using Multicast Information. In *IMC '09*.

[21] J. Sherry, E. Katz-Bassett, M. Pimenova, H. V. Madhyastha, T. Anderson, and A. Krishnamurthy. Resolving IP Aliases with Prespecified Timestamps. In *IMC '10*.

[22] R. Sherwood, A. Bender, and N. Spring. DisCarte: A Disjunctive Internet Cartographer. In *SIGCOMM '08*.

[23] R. Sherwood and N. Spring. Touring the Internet in a TCP Sidecar. In *IMC '06*.

[24] N. Spring, M. Dontcheva, M. Rodrig, and D. Wetherall. How to Resolve IP Aliases. Technical Report 04-05-04, UW CSE, 2004.

[25] N. Spring, R. Mahajan, and D. Wetherall. Measuring ISP Topologies with Rocketfuel. In *SIGCOMM '02*.

[26] H. Tangmunarunkit, R. Govindan, S. Jamin, and S. Shenker. Network Topology Generators: Degree-Based vs. Structural. In *SIGCOMM '02*.

[27] M. Zhang, Y. Ruan, V. Pai, and J. Rexford. How DNS Misnaming Distorts Internet Topology Mapping. In *USENIX ATC '06*.