Implications of the Topological Properties of Internet Traffic on Traffic Engineering

Steve Uhlig[*], **Olivier Bonaventure**
Computer Science and Engineering Department
Université catholique de Louvain, Belgium
suh@info.ucl.ac.be

Vincent Magnin
Laboratory for Computer Communications and Applications
EPFL, Switzerland
vincent.magnin@a3.epfl.ch

Chris Rapier
Pittsburgh Supercomputing Center, USA
rapier@psc.edu

Luca Deri
Centro Serra
University of Pisa, Italy
deri@ntop.org

ABSTRACT

In this paper we study the behavior of Internet traffic on the AS-level topology and discuss its implications on interdomain traffic engineering. We rely on two notable interdomain traffic traces, the first is one month long and the other is one day long. This study shows that interdomain paths are stable for a large majority of the traffic from a routing viewpoint. We show that the aggregation of the traffic occurring on the AS-level graph is essentially limited to direct peers, with almost no aggregation occurring at larger AS hop distances. Furthermore, only part of the AS paths of the AS-level topology that see a lot of traffic are stable, when considering their presence among the largest AS paths on a hourly basis. Relying on the largest AS paths in traffic over a time window to capture the traffic over the next time interval discloses the important variability of the traffic seen by the largest AS paths in traffic. Interdomain traffic engineering is hence due to be difficult because of the limited traffic aggregation on the AS-level topology and the important topological variability of the traffic for a significant percentage of the total traffic.

Categories and Subject Descriptors

C.2.5 [**Computer-communication Networks**]: Local and Wide-Area Networks—*Internet traffic*; C.2.1 [**Network Architecture and Design**]: Network topology

General Terms

Measurement

Keywords

Internet traffic, topological traffic properties, traffic engineering

1. INTRODUCTION

The global Internet is made of two types of ASes: transits and stubs. Transit ASes provide Internet connectivity to stub ASes, by forwarding IP packets across their network. Stub ASes on the other hand only produce or receive IP packets, they do not allow IP packets to transit their network.

The topological distribution of the Internet traffic has been rarely studied in the literature. [12] is among the early papers that analyzed the traffic distribution on the ARPANET in 1974. An important finding from this paper was that a few sites were responsible for the majority of the traffic. Similar results were obtained in 1993 on the NSFNet backbone [7] on the basis of packet-level traces and SNMP statistics from backbone routers. More recently, [8] analyzed several one hour packet level traces from universities and a commercial backbone to evaluate the impact of aggregating flows at the Autonomous System level. Similar results are reported for a large transit ISP in [9] and stub ASes in [22, 15].

The common thread in all of these papers is that a limited number of ASes are responsible for a large fraction of the interdomain traffic. The limitation of these studies concerns the fact that they all consider the traffic distribution for a given time period, assuming that the timescale over which the studies are carried is representative of the behavior of the interdomain traffic, as if the way traffic crosses the AS-level topology does not vary with time. However, the recent proliferation of content distribution networks and peer-to-peer systems may be responsible for large variations of the topological characteristics of interdomain traffic.

More and more ISPs are concerned with how to control the traffic at multiple interdomain access points. Recent studies [20, 2] have shown that a large number of today's ASes are multi-homed. More than 60 % of all ASes have at least two BGP peers [20]. An increasing trend towards multi-homing was shown in [2] and is likely to continue as ISPs become aware of interdomain traffic engineering tools and techniques [15, 5, 3]. Examples of traffic engineering tools include the solutions commonly referred as "route optimization" techniques [3]. These route optimization techniques work on relatively short timescales and target multi-homed ASes. Their principle is to find, based on active and passive measurements, the "best" route to attain a destination or/and to be attained by external hosts. These techniques work on timescales in the order of a round-trip time or more and adapt the traffic flow to the current conditions of the network, i.e. they try to find the best upstream provider to send or receive traffic. Their objective is often to optimize the "instantaneous" QoS experienced by the traffic, by measuring the

---

[*] Corresponding author (URL: http://www.info.ucl.ac.be/∼suh/). This work was carried while Vincent Magnin was a visiting student at UCL.

SAC '04, March 14-17, 2004, Nicosia, Cyprus

"quality" of the routes available from the upstream providers and trying to choose the best upstream in real time. From our knowledge, no paper in the literature has evaluated nor described how these techniques exactly achieve their goals. We are only aware of [23] that proposes a method to engineer the outbound interdomain traffic for stub-ASes by tweaking the `local-pref` attribute of BGP routes and [9] that discusses the predictability of the BGP routes for outbound traffic engineering. Interdomain traffic engineering is still in its infancy. Understanding the dynamics of the interdomain traffic and the relationships between the topology of the traffic and BGP peering relationships is absolutely necessary to do it properly [4].

The control of the outgoing traffic with BGP is based on the selection of the best route among the available ones. This selection can be performed on the basis of various parameters, but it is limited by the diversity of the routes received from upstream providers which depends on the connectivity and the policy of these ASes. The control of the incoming traffic is based on a careful tuning of the advertisements sent by an AS. This tuning has however several drawbacks. First, an AS that advertises more specific prefixes or divides its address space in distinct prefixes to announce them selectively will advertise a number of prefixes larger than required. All these prefixes will be propagated throughout the global Internet and will increase the size of the BGP routing tables of potentially all ASes in the Internet. [6] reports that more specific routes constitute more than half of the entries in a BGP table. Faced with this increase of their BGP routing tables, several large ISPs have started to install filters to ignore BGP advertisements corresponding to more specific prefixes. The deployment of those filters implies that the more specific prefixes will not be announced by those large ISPs and thus the technique will become much less effective. Details concerning interdomain traffic engineering with BGP can be found in [15].

From an interdomain traffic perspective, not much is known about how the traffic that leaves or enters a stub AS, nor how the topological distribution of this traffic varies with time. The implications of traffic engineering on the traffic pattern on the Internet topology are largely unknown, mostly because the topological properties of the traffic pattern are themselves not known. In this paper, we study the aggregation of the traffic on the AS-level topology and its dynamics, to understand its implications on the feasibility of interdomain traffic engineering. Contrary to [16, 18, 17, 13] that focus on a subset of the applications in the whole traffic, our goal is to study the dynamics of all traffic sent by stub ASes, independently of the underlying applications. This paper focuses on stub ASes because they constitute the vast majority of all ASes [20, 2].

The structure of the paper is the following. In section 2 we explain the context of the paper. In section 3 we present the traffic traces on which we rely. Section 4 then studies the stability of the BGP routing for the interdomain traffic. Section 5 analyzes the topological properties of the interdomain traffic over the whole traces while section 6 looks at the hourly dynamics of the traffic on the interdomain topology. Finally, section 7 discusses the implications of our study on traffic engineering.

2. STUDYING THE INTERDOMAIN TOPOLOGY

In order to understand the topological distribution of interdomain traffic we rely on an AS-level graph. This graph is built from the AS path information contained in the BGP advertisements received by the studied stub ASes. More precisely, each AS is one node of the interdomain graph and an edge correspond to the peerings between two ASes. A node of our AS-level graph may correspond to several distinct prefixes. For example, the node corresponding to a large provider will correspond to the domain of this provider and to several of its smaller clients that do not have an AS number. Furthermore, an edge in the interdomain graph may correspond to several distinct physical links or peerings. Since BGP only advertises one path to each destination, the considered interdomain graph does not show all interconnections between domains.

On this interdomain graph, we are interested in the paths followed by IP packets sent by the monitored stub AS to the global Internet. Those paths are the one selected by the BGP decision process [19] based on the advertisements received from each of its BGP peers. Since BGP is a dynamic routing protocol, these paths may change with time. Furthermore, since an AS may advertise several prefixes and the AS Path information is associated with one prefix, there might be several distinct paths between the studied stub AS and a given destination AS.

In this paper, we call an "edge" an AS pair appearing as two consecutive and distinct ASes in the AS path of the *best BGP route* for some traffic destination. With the BGP routes, we know how the interdomain traffic gets forwarded across the AS-level topology without respect to physical links or multiple peering points between two ASes. With the AS path information of the BGP routes, only "edges" of the AS-level topology are known. An example will better illustrate this notion of an "edge". In the left part of Figure 1, the stub AS (AS A) has two BGP peers (AS B and AS C) and each BGP peer announces one route towards each of the two destination ASes (AS F and AS G). In this topology, the stub AS receives two routes towards destination AS F from its two peers. The route advertised by peer 1 has AS path *B-F* while the route advertised by peer 2 has AS path *C-B-F*. Assuming that the stub chooses as its best route towards AS F the route advertised by AS B having the shortest AS path, we have two edges in the interdomain topology for this AS path, namely edge <A,B> and edge <B,F>. Our stub AS receives two routes for destination AS 2 having the same AS path length of 3 from its two peers, *B-D-G* for peer 1 and *C-E-G* from peer 2. If we also assume that our stub chooses the route advertised by peer 2 we have three new edges; <A,C>, <C,E> and <E,G>.

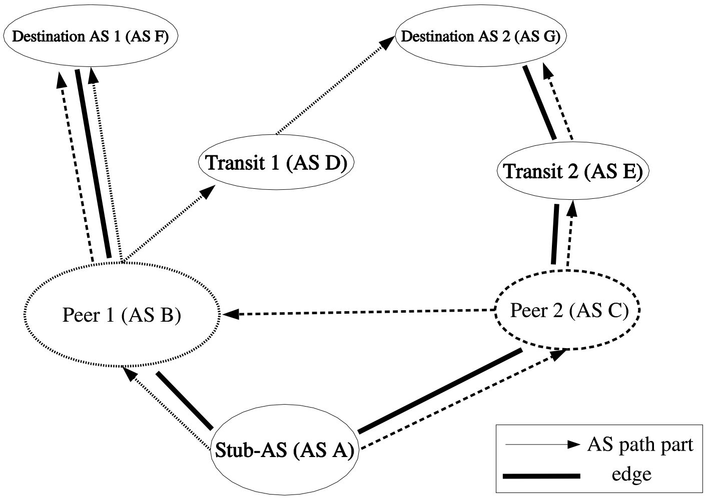

Figure 1: Example AS-level topology.

In the AS-level graph, we are interested in studying and understanding the topological distribution of the interdomain traffic. For this, we need to associate each "edge" with the amount of traffic it carries. We thus count all network traffic that transits between two ASes connected by that edge. An edge does not represent a physical link or a single peering between two ASes, but all peerings between two ASes. From a topological characterization perspective, “edges” represent an aggregate of links or peerings. From a traffic engineering perspective on the other hand, “edges” represent a meaningful traffic aggregate that could be controlled through BGP tweaking like AS path prepending, MED, or redistribution communities [15].

3. DATA

In this paper, we rely on two traffic traces along with the associated BGP routing information. The first trace is a one entire month trace of the outbound traffic from the University of Louvain. The BGP data was gathered from a session established with the associated network provider, BELNET. A literature review indicates that the length of this trace and its continuous nature make it notable. The second trace, mainly used for validation purposes, was gathered from the Pittsburgh Supercomputing Center and covers 24 hours of the outbound commodity traffic. Its BGP information was gathered from an on-site route server.

Given that the aim of this paper is to study macroscopic properties of interdomain traffic the decision was made to use a granularity of one hour for the timescale. This also has a practical consideration in that for a trace of one month duration a finer granularity would prove to be unworkably cumbersome.

3.1 The UCL trace

The primary data set used is a month long trace of all the outbound traffic from the University of Louvain, Belgium, between March 19^{th} 2003 0:00 and April 18^{th} 2003 23:59 CET. During this period 8.4 terabytes of traffic were observed with an average bit rate of 25.1 Mbps. The University’s connectivity is a heavily loaded full duplex fast Ethernet link to the network service provider BELNET. We are not aware of analyses of such long and detailed traffic traces in the literature.

The majority of the University’s user base consists of about 25,000 students and 5000 staff and faculty members using the internal University network of 10 Mbps and 100 Mbps Ethernet links. The University also provides ADSL and cable modem access to some students and staff. The University’s Internet connectivity is provided by the Belgian research network, BELNET which is connected to two tier-1 providers and the European research network GEANT, in addition to more than 100 peers at several interconnection points (BNIX, AMS-IX, SFINX). The internal network of BELNET consists of all Belgian universities and research institutions, linked by a 2.5 Gbps star configuration between its main routers in Brussels and at each University. BELNET did not perform any kind of traffic engineering on its commercial providers but relied on the shortest AS path towards each destination during the time period of the trace. All academic traffic was sent and received through the GEANT network. A limit of 45 Mbps was enforced by BELNET for commodity traffic in the incoming direction only, and no traffic limitation was applied for traffic among the BELNET sites nor for traffic exchanged through the GEANT network.

3.2 The PSC trace

The Pittsburgh Supercomputing Center (PSC) is a regional aggregation point of presence located in western Pennsylvania, USA. PSC provides commodity and Internet2 access to local universities and organizations including Carnegie Mellon University, The University of Pittsburgh, Pennsylvania State University, West Virginia University, The City of Pittsburgh, etc. Currently PSC has a maximum capacity of 395 Mbps of commodity access though AT&T at 145 Mbps and Verio with 250 Mbps via an OC-12. Furthermore, PSC has a full OC-48 of Internet2 connectivity through the Abilene network. The user community consists of approximately 100,000 students and an additional 25,000 faculty and staff members.

The PSC trace used in this study is composed of all outbound commodity traffic from 8:47 AM 18 April 2003 to 8:45 AM 19 April 2003. During this period 1.7 terabytes of traffic were observed for an average rate of 164 Mbps.

4. STABILITY OF THE INTERDOMAIN PATHS

Any study of the topological properties of Internet traffic depends on the premise that the interdomain topology advertised by BGP itself be stable. Therefore we first need to explore the stability of the AS paths of the BGP routes. To do this we rely exclusively on the traces and BGP data collected from UCL. The month long trace provides the necessary depth of information which the PSC trace of 24 hours simply cannot.

First we look at how long AS paths were seen in the BGP table. The top part of Figure 2 compares the number of days distinct AS paths were seen, for all AS paths as well as AS paths that saw at least one byte of data over that month. During this period 103,853 distinct AS paths were seen in the BGP advertisements but only 31,151 carried any traffic at all. Due to BGP dynamics more than fifty percent of the distinct AS paths were present in the BGP routing table for less than nine minutes. On the other hand, more than 42 percent of the AS paths over which traffic was sent were found in the BGP routing table for 99 percent of the duration of the month. These stable AS paths having traffic represented slightly more than 13 percent of all AS paths seen over the month. These results are similar to [16] where it was shown that most BGP update events occur for a small fraction of the prefixes. [24] also showed that the primary AS path for a DNS A root server lasted for long time periods.

The top part of Figure 2 provides an overview of the stability of the AS paths but without respect to the traffic they see. When studying the traffic, we are interested in whether the AS paths which carry the majority of the traffic are stable or not, AS paths that do not carry traffic are irrelevant. The bottom part of Figure 2 compares the longevity in presence of the AS paths and their associated traffic over time. We define “presence” as the number of seconds during which an AS path was found among the BGP best routes over the one month of the trace. The curve labeled “AS paths” on the right part of Figure 2 gives the percentage of the total time AS paths were present in the BGP routing table. The curve labeled “traffic” shows the percentage of traffic for AS paths present for more than some percentage of the one month trace period. The latter curve shows that more than 95 percent of the total traffic was carried by AS paths that were present in the BGP routing table more than 99 percent of the time. In contrast, only 42 percent of the AS paths having traffic were present for more than 99 percent of the time in the BGP routing table. The prefixes for the largest fraction of the traffic of our stub AS have a single AS path that lasts for more than 99 percent of the one month period of our trace. Our results are similar to those of [16] that showed that the prefixes associated with popular web sites for AT&T traffic had stable BGP routes. Our results however are more precise and detailed than [16] since we consider the routing stability of all the interdomain traffic, not a subset of the prefixes.

BGP hence does not significantly interfere with the traffic dynamics. The interdomain topology used to carry outbound traffic

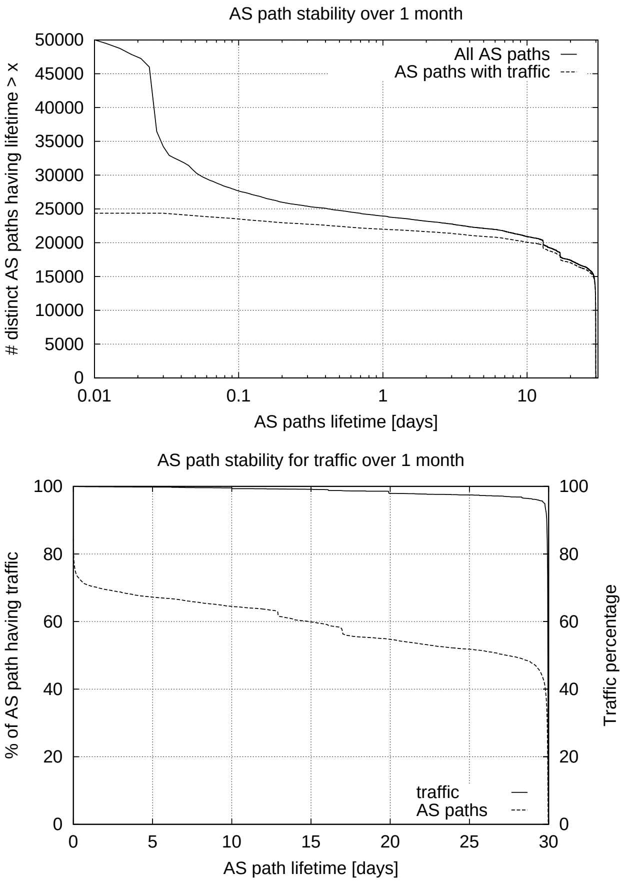

Figure 2: Stability of AS paths: without traffic (top) and with traffic (bottom).

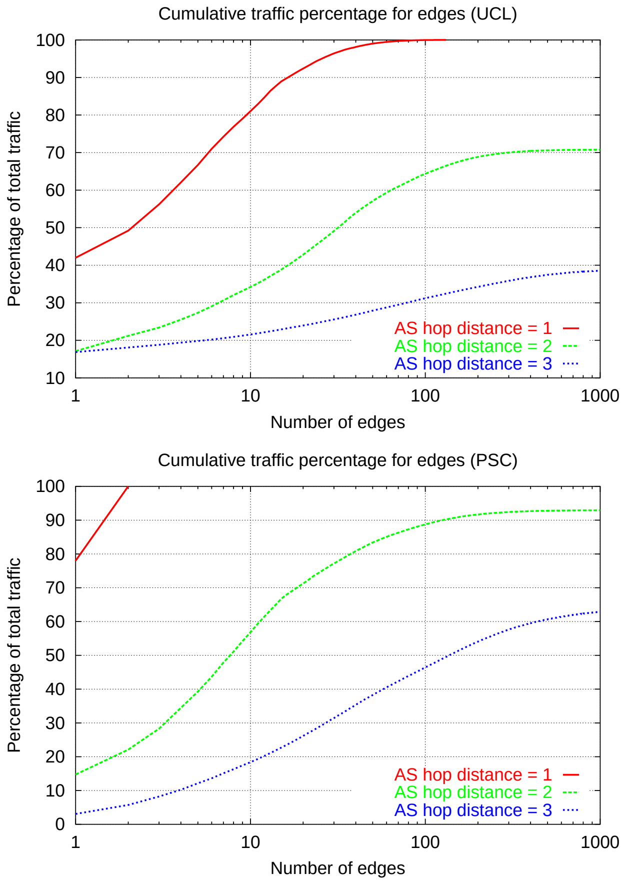

Figure 3: Cumulative traffic captured by edges : UCL (top) PSC (bottom)).

can thus be considered as stable even over a long time period as one month.

5. TOPOLOGICAL TRAFFIC AGGREGATION

Before looking at the dynamics of the traffic on the AS-level topology, it is good to build an understanding of how the traffic gets aggregated on the AS-level topology, given that the whole AS-level graph visible from BGP routing tables contains about 15,000 ASes and more than 30,000 edges. The definition of an "edge" serves this particular purpose. In interdomain traffic engineering, a primary concern is the size of the interdomain topology for the majority of the traffic. Note that the AS-level topology where only edges that saw traffic are considered can be seen as a tree rooted at the traffic capture point and whose leaves are the destination ASes of the traffic.

We begin with how much traffic edges see for each AS hop distance. This viewpoint will show us how many edges we need at each AS hop distance to influence some percentage of the total traffic. Figure 3 provides the cumulative traffic percentage carried by edges at each AS hop distance. So we take all edges at some AS hop distance and count how many bytes have been carried by these edges. We then sort the edges at this AS hop distance by decreasing byte count and plot the cumulative percentage of traffic carried by the largest n edges at this AS hop distance from the local AS. The curves for an AS hop distance of one on Figure 3 illustrate the main difference between UCL and PSC. The largest direct peer of UCL's provider sees only a little more than forty percent of the total traffic, while the largest peer of PSC sees a little less than eighty percent of the total traffic. UCL's provider has many direct peers with whom it exchanges traffic at local interconnection points. The PSC trace on the other hand shows that the two largest direct peers see all the traffic, simply because the PSC trace contains only commodity traffic that is sent through its commercial providers.

The curves for an AS hop distance of two and three provide us with the interesting information about traffic aggregation on the AS-level topology. Obviously, traffic aggregation occurs on the edges with the direct peers. However, traffic aggregation on edges is desirable farther in the topology because traffic engineering techniques are due to perform finer operations than moving all the traffic from one direct peer to another. Traffic engineering will have to work on traffic aggregates to allow an AS to control the amount of traffic that crosses the links with its direct peers. This is why it is important to know how traffic is split a few AS hops away in the AS-level topology. Figure 3 tells us that if we were to influence the largest ten edges at two AS hops, then we would be able to control about 35 percent of the total traffic for UCL and a little less than 60 percent for PSC. This information is relevant for traffic engineering since we would be able to influence a large fraction of the total traffic by tweaking a few edges. If one wants to perform fine-grained inbound traffic engineering, relying on edges at two AS hops or more will be necessary to control sufficiently small traffic aggregates because traffic aggregation is too important on the links with the direct peers. This means that we would then need to tweak a larger number of edges due to the lack of traffic aggregation beyond direct peers. Although a few edges at two or three AS hops may allow to influence a large fraction of the traffic, we know that inbound traffic engineering will not always work due to local policies applied by other ASes. The topological analysis of the traffic hence provides an overly optimistic view of the feasibility of interdomain traffic engineering.

6. SHORT-TERM STABILITY OF INTERDOMAIN TRAFFIC

The previous sections examined the properties of the traffic-aware interdomain topology over the entire temporal length of the measurements. Another important issue is the stability of the traffic over smaller timescales. If its distribution is stable, then it should be possible to engineer it. Otherwise, this will be more difficult. Therefore in this section we outline the connection between the large timescale features of the traffic-aware interdomain topology previously discussed and a smaller timescale of one hour granularity. The analysis in this section is based solely on the UCL dataset while using the PSC dataset for confirmation of our results. In addition, we rely solely on traffic carried by AS paths in the remainder of the paper because of the limited traffic aggregation occurring for the edges of the AS-level topology. Given that few distinct AS paths get aggregated on a given edge, the dynamics of the AS paths is similar to the one of the edges (we do not show the corresponding figures due to space limitations).

6.1 Short-term evolution of top AS paths

In this section, we analyze the evolution with time of what we call the top AS paths. The “top AS paths” are nothing more than the AS paths weighted by the traffic they carry and sorted by decreasing amount of traffic over the considered time period (month or hour in our case). For instance, the top fifty percent AS paths for some hour are the top AS paths that capture fifty percent of the total traffic over that particular hour.

Figure 4 shows the hourly evolution of the number of AS paths present in the top x percent AS paths for 4 different values of x. To create this graph we computed the amount of traffic seen by each AS path for each hour, sorted them by decreasing amount of the total traffic, computed the cumulative traffic distribution and finally plotted the hourly evolution of the number of AS paths for various values of the traffic percentage every hour. Each point of the curves shows how many AS paths are necessary to carry an arbitrary percentage of the total traffic during a given hour. The hourly evolution of the number of AS paths for the traffic percentage follows a daily periodicity similar to the one of the total traffic evolution. The larger number of AS paths seen during the busiest hours of the day is likely to be due to a larger number of IP addresses. Note that for fifty percent of the hourly traffic, this daily periodicity is less evident because the number of AS paths is small.

The use of a logarithmic scale for the y-axis on Figure 4 reduces the visual importance of the dramatic increase in the number of AS paths needed as the percentage of traffic increases. Capturing fifty percent of the traffic every hour requires between five and ten AS paths, while ninety percent requires one hundred AS paths, and capturing all the traffic demands as many as five thousand AS paths. These numbers are to be compared with the number of AS paths that capture the same percentages of traffic over the whole month: 17 for fifty percent, 332 for ninety percent, 1923 for ninety-nine percent, and 31,150 for one hundred percent. The largest hourly AS paths thus are fewer in number than monthly ones, indicating that it should be easier to engineer some percentage of the total traffic on a hourly basis. This statement however does not take into account whether the important AS paths change a lot from one hour to another. If the important AS paths are stable in time, then performing traffic engineering on small timescales is feasible. If they are not stable, traffic engineering can be difficult.

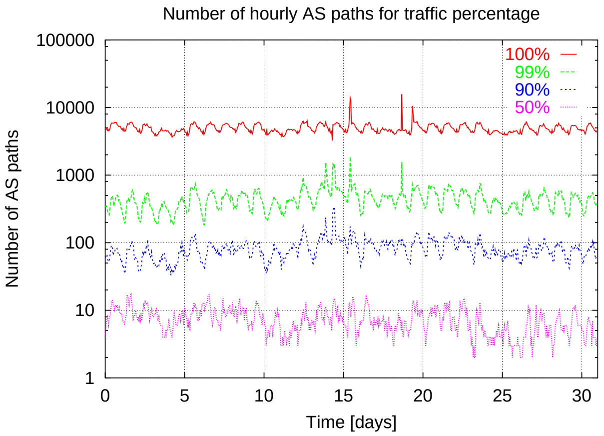

Figure 4: Hourly evolution of the number of top AS paths.

6.2 Presence of short-term top AS paths

This section studies for how much time most AS paths are among the hourly top AS paths. For this, we compute for each hour the top ninety percent AS paths and count how many times the AS paths reappear in the top ninety percent over the length of the trace.

Figure 5 counts the number of hours over the month that each AS path appears in the hourly top 90 % (top of Figure 5) and top 50 % AS paths (bottom of Figure 5). We also compute for each of these AS paths the amount of traffic it captures during the hours it is among these top AS paths. The curve labeled “AS paths percentage” gives the cumulative percentage of the AS paths which have ever appeared among the top AS paths in terms of the number of hours it was present, hence the name “presence”. The other curve labeled “traffic percentage”, gives the corresponding cumulative percentage of the total traffic carried by the previous AS paths during a given number of hours over the month.

A total of 2139 AS paths were seen among the hourly top 90 % AS paths. This is much more than the 332 AS paths capturing 90 % of the traffic over the one month, indicating that there is much turnover among the hourly AS paths. 626 of them were seen only once while just six were seen during all time periods. The curve labeled “AS paths percentage” shows that most AS paths were seen for a small fraction of the time. More than 96 percent of these AS paths were seen for less than one third of the time and represent about twenty percent of the total traffic during the whole month. Additionally, 50 % of the total traffic appears for AS paths seen among the hourly top 90 % for about 80 % of the time. The six AS paths that are present during all time periods carry about 36 % of the total traffic which is a significant portion. However, more than 20 % of the traffic appears for AS paths that are seen in less than one third of the time periods during the month. This leaves about 25 % of the traffic being carried by AS paths that are not entirely stable.

The relatively limited presence of the top 90 % AS paths might be due to the fact that most of them carry too limited an amount of traffic. The presence of larger AS paths might be better than what

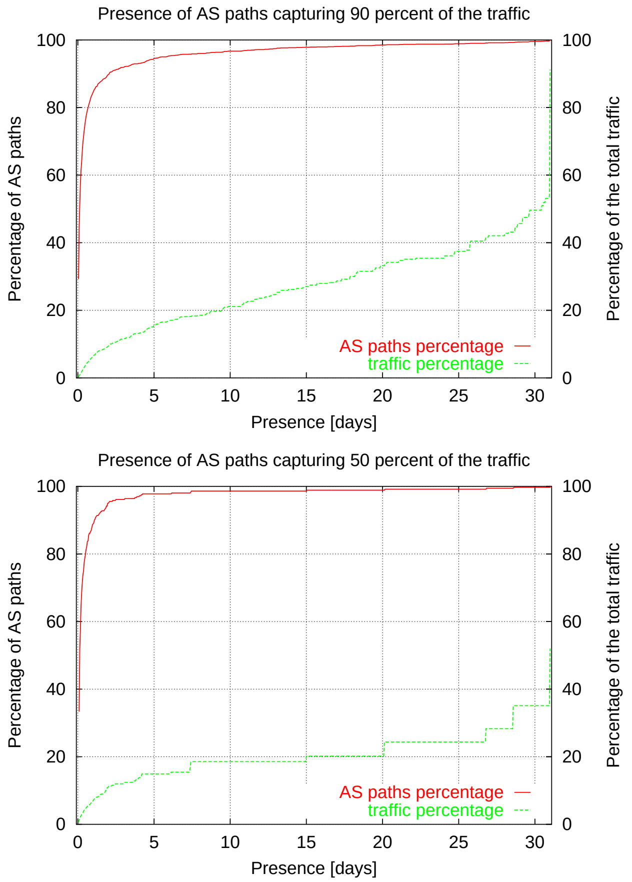

Figure 5: Presence of hourly top AS paths: top 90 % AS paths (top) and top 50 % AS paths (bottom).

the top graph of Figure 5 showed. The bottom graph of Figure 5 shows the same information as the top Figure 5 but for the top 50 % top AS paths. 360 AS paths were seen among the hourly top 50 % AS paths. One third of these AS paths were present during only two hour intervals over the one month among the hourly top 50 % top AS paths. Slightly less than 90 % of these 360 AS paths were present during 24 one hour intervals over the month. Among these 360 AS paths, those that are present for less than one third of the hourly intervals carry a little less than 20 % of the total traffic. AS paths that are present more than half of the hourly intervals hence carry 30 % of the total traffic. This shows that even the largest AS paths in terms of the hourly traffic have a high variability in presence. This variability in presence of the AS paths thus concerns all AS paths, irrespective of the amount of traffic they carry.

6.3 Traffic Capture for top AS paths

To this point we have assumed foreknowledge of which AS paths carry the largest portion of the traffic. Under that assumption it is easy to demonstrate that a significant percentage of the hourly traffic travels AS paths that are prominent in the monthly traffic statistics. However, we have shown in Figure 5 that while some of the top 90 % paths are stable the majority are less stable and do account for a significant portion of the traffic. This means that knowledge of the AS paths that carry large amounts of traffic over short times scales will not necessarily give an indication of those carrying large amounts of traffic over longer time scales.

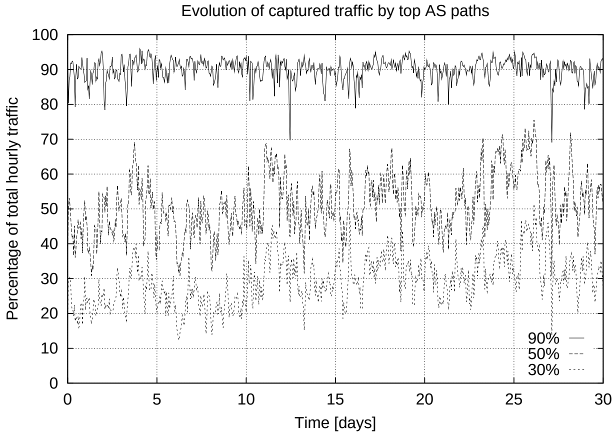

Figure 6: Evolution of percentage of traffic captured by top AS paths.

Figure 6 shows the hourly percentage of traffic captured by the top 30 %, 50 %, and 90 % AS paths over the length of the one month UCL trace. To do this we first determined the AS paths that carried the most traffic over the month. This information is then used to determine how much traffic these paths carried on an hour by hour basis and the results plotted over time. If these paths are stable we would expect to see a relatively flat line with little variation over time. However, we see that the AS paths that carry 30 % of the monthly traffic rarely carry 30 % of the hourly traffic. Instead it displays large variations in the hourly traffic, generally between 20 % and 40 %. We see similar variation as we include more AS paths to increase the relative amount of carried traffic although the absolute level of variation decreases as we include more paths.

It is possible that windowing the previous traffic periods might give some insight into the stability of AS paths in subsequent time periods. The results of this inquiry can be seen on Figure 7 where the average percentage of traffic seen by 50 % and 90 % AS paths during time windows of hours and days is given. For each hour we first determined the AS paths that carry 50 % and 90 % of the traffic during the previous time window and computed how much traffic these ASes carry over the subsequent hour. Finally the average percentage of traffic seen for each time window over the length of the trace is determined. The same process was then repeated for time windows of days. The corresponding results can be seen in the top of Figure 7 with the standard deviation found below it.

The top of Figure 7 shows that increasing the time window increases the average percentage of traffic captured, however the gains in accuracy for larger time windows is limited. This is especially true for increasing the time window to periods of multiple days. However, if we compare the standard deviations of the hours window and the days window we do discover a notable difference. The standard deviation for all time windows and the AS paths capturing 50 % over the time window is larger than for the 90 % AS paths. The smaller number of AS paths for the 50 % capture are responsible for this since the smaller set of AS paths account for a comparatively larger portion of the total traffic and thus more vulnerable to random traffic surges. Second, the timescale of days with 90 % capture seems to provide a reduced variability as is evidenced by the plot of the standard deviation.

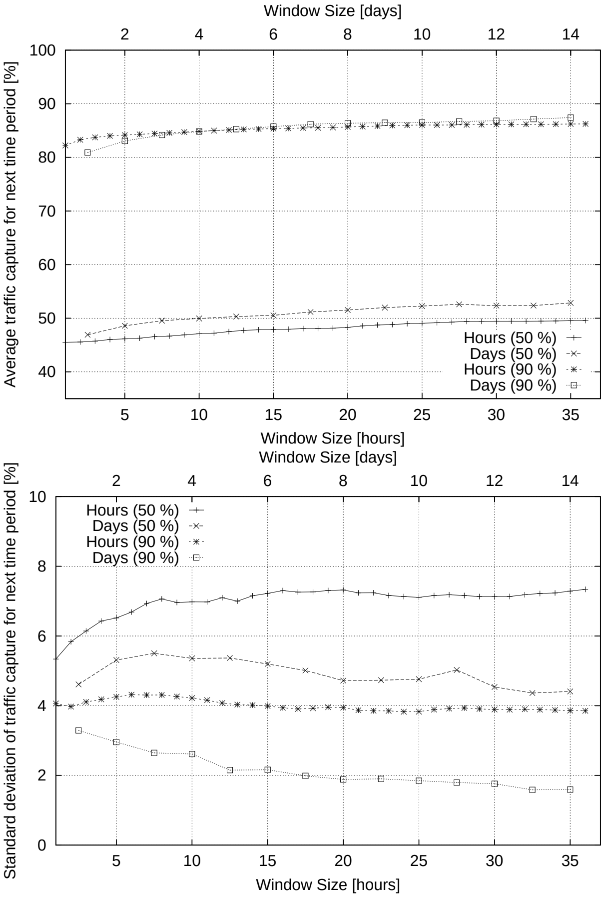

Figure 7: Traffic capture for windowed 50 % and 90 % top AS paths: average (top) and standard deviation (bottom).

7. IMPLICATIONS ON INTERDOMAIN TRAFFIC ENGINEERING

Most studies on interdomain traffic insist that a few sources and destinations generate most of the interdomain traffic. This optimistic news is misleading because it hides an equally important facet of the network topology: influencing most of the interdomain traffic requires that a significant number of ASes be considered – far more than previously believed. For instance, in section 6.1 it was shown that to control 50 % of its traffic UCL would have to tweak about 10 AS paths every hour, or about 100 AS paths every hour to control 90 % of its traffic. While these numbers might not seem particularly large, making the necessary adjustments to BGP on a nearly continuous basis is not a trivial task.

This is not to say it is impossible though. For outbound traffic, tweaking BGP is much easier and does not impact BGP routing outside the local domain of the stub AS. For inbound traffic things are a little more intricate. First, the AS path information contained in the BGP routing table of the local stub does not provide the right information concerning the inbound path followed by IP packets [21]. Second, even if the local stub had the whole map of the AS-level Internet, which is a very strong assumption indeed, there is no guarantee that influencing remote ASes located a few AS hops away would be possible given that policy routing is shaped independently by each AS. Third, even if there existed some manner in which this could be done, it would mean that a stub would need to frequently change its BGP announcements for potentially hundreds of ASes. If a significant percentage of the stub ASes begin to perform that kind of BGP tweaking, it is difficult to predict the global effect on BGP routing, which has already been demonstrated to suffer from convergence problems [4, 14, 11].

Additionally, we showed in section 6 that the variability of the AS paths that carry the largest amounts of traffic is quite important. The uncertainty due to the traffic dynamics implies that controlling the interdomain traffic requires to influence more AS paths than what is believed based on an a priori knowledge of the important AS paths in the traffic. The topological dynamics of the interdomain traffic indicates that the only way to improve the certainty of influencing a large fraction of the traffic is to tweak more AS paths. To do this, one needs to influence more BGP routes, which in turn will potentially affect traffic that spans a larger fraction of the topology of the Internet.

Finally, tweaking BGP will impact on the intradomain traffic distribution of large ISPs [1]. Solutions that rely on intradomain routing to optimize the intradomain traffic of large ISPs [10] can be affected by the tweaking of BGP routing. The interactions between intradomain and interdomain traffic engineering thus need to be studied before deploying interdomain traffic engineering tools.

8. CONCLUSION

In this paper we have studied the properties of Internet traffic on the AS-level topology and discussed its implications on traffic engineering.

We first showed that the interdomain paths that carry the largest fraction of the traffic are stable from a routing viewpoint. Second, traffic aggregation on the AS-level topology is limited to the direct peers, with a limited aggregation occurring at AS hop distances larger than one. Next, while the AS paths carrying the most traffic over the month were present among the largest hourly AS paths most of the time, a significant percentage of the total traffic was carried by AS paths whose presence among the largest hourly AS paths is limited. Finally, we showed the important variability of the traffic captured during the next time period by the largest AS paths over a time window.

This paper showed that the topological aggregation and variability of the traffic is an issue for traffic engineering. The variability of AS paths carrying a significant amount of traffic is an issue for controlling interdomain traffic, because a significant fraction of the traffic is carried by AS paths that are present for a limited amount of time. Furthermore, since traffic aggregation does not occur on the AS-level topology, an AS wishing to control a large fraction of its inbound interdomain traffic will need to influence a number of ASes that can grow quite large.

9. AKNOWLEDGMENTS

The authors wish to thank BELNET for the BGP session they kindly provided. The authors also aknowledge the anonymous reviewers for their valuable comments.

10. REFERENCES

[1] S. Agarwal, C. Chuah, S. Bhattacharyya, and C. Diot. The Impact of BGP Dynamics on Intra-Domain Traffic. Sprint ATL Research Report Nr. RR03-ATL-080677, August 2003.

[2] S. Agarwal, C. Chuah, and R. Katz. OPCA: Robust Interdomain Policy Routing and Traffic Control. In *Proc. of IEEE OPENARCH*, April 2003.

[3] D. Allen. NPN: Multihoming and route optimization: Finding the best way home. *Network Magazine*, February 2002.

[4] R. Atkinson and S. Floyd. IAB Concerns and Recommendations Regarding Internet Research and Evolution. Internet draft, draft-iab-research-funding-01.txt, work in progress, July 2003.

[5] J. Bartlett. Optimizing multi-homed connections. *Business Communications Review*, 32(1):22–27, January 2002.

[6] A. Broido, E. Nemeth, and K. Claffy. Internet expansion, refinement and churn. *European Transactions on Telecommunications*, January 2002.

[7] K. Claffy, H. Braun, and G. Polyzos. Traffic characteristics of the T1 NSFNET backbone. In *INFOCOM93*, 1993.

[8] W. Fang and L. Peterson. Inter-AS traffic patterns and their implications. In *IEEE Global Internet Symposium*, December 1999.

[9] N. Feamster and J. Rexford. Network-Wide BGP Route Prediction for Traffic Engineering. In *Proc. of ITCOM 2002, Boston, MA*, August 2002.

[10] B. Fortz, J. Rexford, and M. Thorup. Traffic engineering with traditional IP routing protocols. *IEEE Communications Magazine*, October 2002.

[11] T. Griffin and G. Wilfong. An analysis of BGP convergence properties. In *Proc. of ACM SIGCOMM'99*, September 1999.

[12] L. Kleinrock and W. Naylor. On measured behavior of the ARPA network. In *AFIS Proceedings, 1974 National Computer Conference*, volume 43, pages 767–780. John Wiley & Sons, 1974.

[13] B. Krishnamurthy, C. Wells, and Y. Zhang. On the use and performance of content distribution networks. In *Proc. of the first ACM SIGCOMM Internet Measurement Workshop*, November 2001.

[14] C. Labovitz, G. Malan, and F. Jahanian. Internet routing instability. In *Proc. of ACM SIGCOMM'97*, September 1997.

[15] B. Quoitin, S. Uhlig, C. Pelsser, L. Swinnen, and O. Bonaventure. Interdomain traffic engineering with BGP. *IEEE Communications Magazine, Internet Technologies Series*, May 2003.

[16] J. Rexford, J. Wang, Z. Xiao, and Y. Zhang. BGP Routing Stability of Popular Destinations. In *Proc. of the second ACM SIGCOMM Internet Measurement Workshop*, pages 197–202, November 2002.

[17] S. Saroiu, P. Gummadi, and S. Gribble. A Measurement Study of Peer-to-Peer File Sharing Systems. In *Proc. of Multimedia Computing and Networking Conference*, January 2002.

[18] S. Sen and J. Wong. Analyzing peer-to-peer traffic across large networks. In *Proc. of the second ACM SIGCOMM Internet Measurement Workshop*, November 2002.

[19] J. Stewart. *BGP4 : interdomain routing in the Internet*. Addison Wesley, 1999.

[20] L. Subramanian, S. Agarwal, J. Rexford, and R. Katz. Characterizing the Internet hierarchy from multiple vantage points. In *IEEE INFOCOM 2002*, June 2002.

[21] H. Tangmunarunkit, R. Govindan, S. Shenker, and D. Estrin. The impact of routing policy on internet paths. In *Proceedings of INFOCOM'01*, pages 736–742, 2001.

[22] S. Uhlig and O. Bonaventure. Implications of Interdomain Traffic Characteristics on Traffic Engineering. *European Transactions on Telecommunications*, January 2002.

[23] S. Uhlig, O. Bonaventure, and B. Quoitin. Interdomain Traffic Engineering with minimal BGP Configurations. In *Proc. of the* 18^{th} *International Teletraffic Congress, Berlin*, September 2003.

[24] L. Wang, X. Zhao, D. Pei, R. Bush, D. Massey, A. Mankin, S. Wu, and L. Zhang. Protecting BGP Routes to Top Level DNS Servers. In *Proc. of 23rd International Conference on Distributed Computing Systems (ICDCS)*, May 2003.