

Topological Trends of Internet Content Providers

Yuval Shavitt
School of Electrical Engineering
Tel-Aviv University, Israel

Udi Weinsberg
Technicolor Research
Palo Alto

ABSTRACT

The Internet is constantly changing, and its hierarchy was recently shown to become flatter. Recent studies of inter-domain traffic showed that large content providers drive this change by bypassing tier-1 networks and reaching closer to their users, enabling them to save transit costs and reduce reliance of transit networks as new services are being deployed, and traffic shaping is becoming increasingly popular.

In this paper we take a first look at the evolving connectivity of large content provider networks, from a topological point of view of the autonomous systems (AS) graph. We perform a 5-year longitudinal study of the topological trends of large content providers, by analyzing several large content providers and comparing these trends to those observed for large tier-1 networks. We study trends in the connectivity of the networks, neighbor diversity and geographical spread, their hierarchy, the adoption of IXPs as a convenient method for peering, and their centrality. Our observations indicate that content providers gradually increase and diversify their connectivity, enabling them to improve their centrality in the Internet, while tier-1 networks lose dominance over time.

1. INTRODUCTION

The Internet is a constantly evolving network, quickly adapting to customer needs and financial forces. Up until recently it was common to picture an hierarchical Internet [12, 9, 5], in which networks are either tier-1, large networks that provide global transit functionalities, tier-2, smaller Internet Service Providers (ISPs) that provide Internet connectivity to their customers, or stub networks that produce and consume content [26].

However, in recent years the Internet is changing. The appearance and rapid growth of large content providers, such as Google, Yahoo! and others, is gradually changing the roles of key Internet players to accommodate their needs. First, large content providers produce huge amount of content that is consumed by users around the globe, inducing heavy traffic on transit networks. Although wholesale tran-

sit prices are decreasing by roughly 30% each providers offer various wholesale pricing plans to accommodate these needs, such as tiered prices [26], content providers still seek ways to significantly cut transit costs. Furthermore, as content providers deploy an increasing number of Software as a Service (SaaS), such as elastic computing, collaboration tools, storage and even complete content delivery networks (CDNs), they seek to reduce reliance on transit providers that were reported to perform traffic shaping.

As a result of these trends, the Internet was reported to be forming a flatter and denser network [13, 16, 8], mostly using observations of traffic flows. In this paper we take a first look at the changing connectivity of large content provider networks, from a topological point of view. Unlike previous work that studied various traffic characteristics [10, 20, 16], we consider the trends observed in the connectivity of the networks in the Autonomous Systems (AS) graph.

We achieve this using a 5-year longitudinal study of the AS graph, focusing on 5 major content providers: Google, Yahoo!, MSN, Amazon, and Facebook. The first three are well established, large content providers, that have been around before the beginning of our study, in 2006. Amazon provides a unique opportunity to study a content provider that changed its scope (from an online store to a cloud service host) and Facebook reveals the high-paced growth of an extremely popular content provider. Using a comparative approach, we examine 5 major transit providers, namely AT&T, Qwest, Level3, Sprint, and Global Crossing (Glbx). All of these are large, tier-1 transit networks, that have been used by content providers for transit over the years.

In this paper we create a snapshot of the AS-level graph every 3 months, using a month of active traceroutes, from late 2006 till early 2011. We then study the connectivity trends, meaning, how content and transit providers are connected and evolve over time. Understanding the evolving trends of AS connectivity has implications on different aspects of the Internet ecosystem. The decreasing dominance of large transit providers we observe indicates a change in the way traffic flows and networks interconnect. These in turn have direct implications on the operational decisions that drive ASes, their connectivity and profitability. Additionally, understanding these trends can help improve Internet research, such as growth models and traffic flow analysis.

2. RELATED WORK

Several recent papers study the emerging change in the Internet ecosystem, which is driven mostly by large scale content providers. One of the early observations of this change

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SIMPLEX '12 April 17 2012, Lyon, France

Copyright 2012 ACM 978-1-4503-1238-7/12/04 ...\$10.00.

was made by Gill *et al.* [13]. The authors showed that large content providers bypass many tier-1 ISPs by pushing their networks closer to the users, and suggested that such a trend can possibly flatten the Internet.

Recent works [4, 15, 14] on the impact of IXPs on the AS-graph report significantly higher number of peering relationships discovered among ASes that are IXP participants than among ASes that are not connecting via an IXP.

Dhamdhere and Dovrolis [8] presented a new Internet model that captures the Internet transition from a hierarchy of transit providers to a flatter interconnection of peers. Most recently, Labovitz *et al.* [16] performed a large-scale two-year study of the inter-domain traffic, showing that the amount of traffic originated from content providers is rising, and most of it is routed outside of the traditional Internet core. Specifically, more than 5% of all inter-domain traffic in July 2009 originated from Google’s networks. Additionally, they showed that content providers often use their own networks for saving transit costs.

This paper takes a different approach by studying the trends observed in the topology of the Internet using complex networks metrics. These topology trends reveal how the different networks interconnect and change the structure of the Internet.

3. METHODOLOGY

For the purpose of this study we build the Internet AS graph every three months since January 2006 until July 2010. The AS graph is built by traversing IP-level traceroutes and resolving each IP into its corresponding AS. We resolve IP to AS using the published iPlane nano [18] database, which is a set of mappings between IP prefixes into ASes, collected from all major BGP monitors. This set contains 326,102 prefixes mapped to 30,779 ASes. IPs that were mapped into AS sets or multi-origin ASes (MOAS) were treated as unresolved.

Additionally, we used Internet Exchange (IXP) prefix mapping provided from [4]. The IXP list is comprised of prefix lists collected from Packet Clearing House (PCH) [1], PeeringDB [2] and additional manually collected sources. This list provides us with 393 prefixes belonging to 278 IXPs worldwide. We use the same IXP list, obtained in late 2009 for all years, assuming that the assignment of prefixes to IXPs did not decrease overtime, i.e., even if an IXP becomes defunct [1], there are only a few or no cases that a prefix, which was assigned to an IXP in 2009, was assigned to a real AS in other times.

Each AS graph is built by traversing AS traces of a single month, creating a link between two ASes that follow each other in the AS trace, or have an IP that belongs to an IXP prefix between them (the latter follows the technique described in [14] and extended in [4]).

For the IP-level traces we use two datasets, DIMES [24] and iPlane [17]. DIMES is a community-based Internet mapping effort, measurement from thousands of vantage points, located at user homes and since 2010 also in PlanetLab [6]. iPlane uses PlanetLab nodes and traceroute servers, measuring from a relatively stable set of 300 servers. Although DIMES, due to its diverse distribution, uncovers more links than iPlane [25], it suffers from vantage point churn, making its observed topology more “noisy” and susceptible to measurement artifacts. iPlane on the other hand, is more stable, both in the number of measuring vantage points and

the target IPs, however its topology is smaller.

In this paper we are interested in global trends observed in the ecosystem of large Internet players. Since these networks are well observed by both DIMES and iPlane, we expect both platforms to capture similar trends, even if the exact numbers are somewhat different.

Therefore, when analyzing trends of the entire AS-level graph, we use DIMES data, which brings a more accurate view of the topology, while the noise gets smoothed due to the large amount of data. When analyzing specific ASes, we use iPlane’s data, since it is more stable. Indeed, in most cases, both infrastructures resulted in the same overall trends, assuring that the observations we make are indeed due to real topology and routing characteristics and are not the results of some measurement bias. Whenever the two datasets do not agree, we present both results and discuss the differences and their causes.

It is important to note that since we are interested in trends, the exact numbers we obtain (for number of connections, clustering coefficient, etc.) are not important. We are interested in their scale and especially in their evolution over time. Thus, the effect of measurement errors [3], errors in the IP to AS translation [21, 22], and similar inaccuracies, which are discussed in the measurement literature, are not significant for this study since they can not affect trends.

4. AS-LEVEL CONNECTIVITY

The simplest way to measure the connectivity of an AS is the number of neighboring ASes it is connected to, i.e., its degree in the AS graph. Fig. 1a and Fig. 1c show the degree of the content and transit networks over time, exhibiting a large difference between their degrees. However, while Fig. 1a shows that content providers are increasing their connectivity over time, the connectivity of transit networks depicted in Fig. 1c exhibits a slow decrease (except Qwest).

Fig. 1b shows that the average degree of the neighboring ASes of the three ‘veteran’ content providers slowly decreases over time, while the new ones, Amazon and Facebook, start with a high average neighbor degree, which is quickly reduced to match the other three. This indicates that the content providers start by connecting mostly to tier-1 providers, but as they expand they connect to additional providers which are mostly not tier-1 (see additional discussion below). Fig. 1d shows that the average neighbor degree of the transit networks increase. This can be an indication that small customers are disconnecting from the transit networks, or alternatively, that the tier-1 neighbors are becoming better connected, hence increase their degree. Both reasons imply that transit networks are losing some of their dominance in Internet connectivity.

To further understand the reasons behind these trends, we classify the neighbors of each AS. We use the classification provided by Dhamdhere and Dovrolis [7], in which an AS is classified as an enterprise customer (EC), a small or large transit provider (STP and LTP), or a content, access and hosting provider (CAHP). The authors base their classification on the average customer and peer degrees of the AS over its entire lifetime within a 10 year longitudinal study, and claim to reach over 80% accuracy.

Analysis of the AS types of the neighbors of the content providers revealed that the number of LTP neighbors did not significantly change. However, since the number of neighbors increased over time for all content providers, the

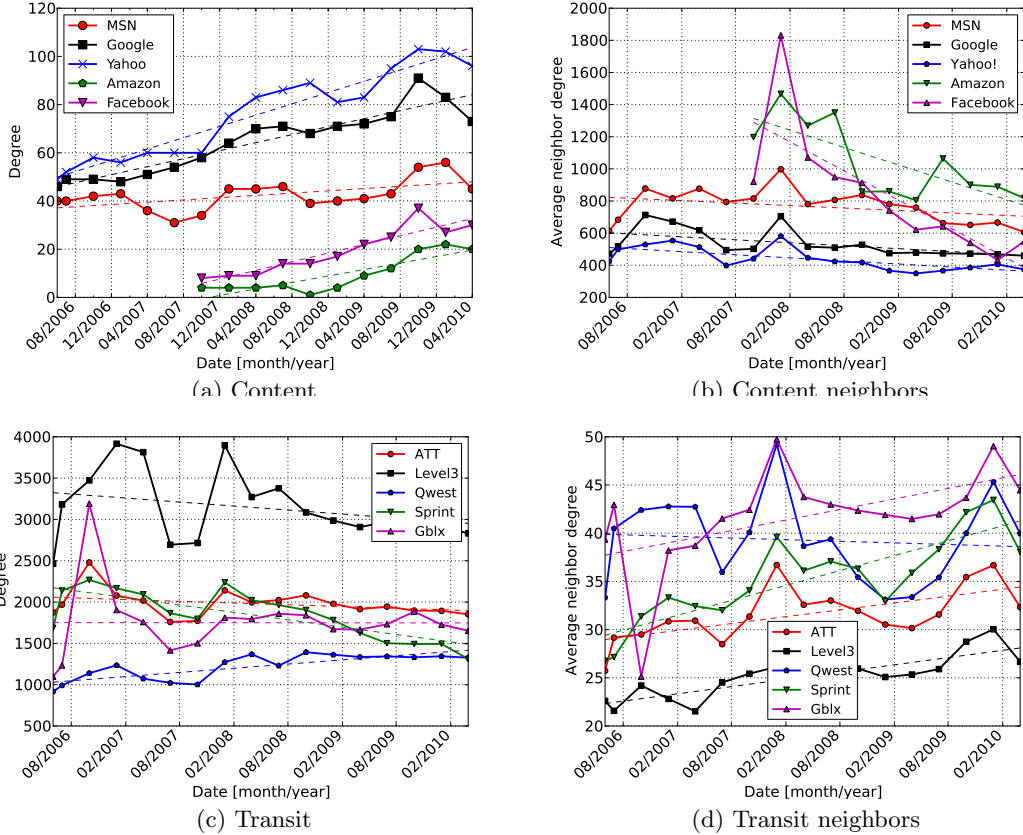


Figure 1: The degree and average neighbor degree of transit and content networks

overall LTP percentage decreased. The most drastic change is observed for Amazon and Facebook, with a 2/3 reduction. Amazon and Facebook greatly increased the number of CAHP and STP neighbors. Similar analysis on transit providers revealed that they mostly interconnect with EC and CAHPs. These neighbors exhibit the largest reduction, indicating a market loss in their core business, probably impacting their dominance in the Internet connectivity.

Finally, we measure the assortativity coefficient values [23] of the complete AS graph, using data from DIMES. These coefficients are positive when vertices tend to connect with similar-degree vertices and negative otherwise. The AS-graph was shown to be disassortative [23, 19], i.e., comprise mostly of radial links, connecting ASes towards the tier-1 ASes. Analysis of the trend of the assortativity coefficient indicates that the AS-graph becomes slightly less disassortative, partly indicating that ASes increase their connections with other similar-degree ASes.

Overall, these findings show a trend of large content providers that shift from relying mostly on large transit providers towards a flatter topology, interconnecting with smaller networks for gaining transit and direct access to last-mile customers. Content providers that have been around longer, such as Google, Yahoo!, and MSN do not increase the number of LTPs, mainly since they already exploit the benefit of connecting with them. The “younger” content providers follow this trend and mostly use STPs and CAHP. This

trend reveals an overall decrease in the dominance of tier-1 networks in the Internet ecosystem.

4.1 Density and Clustering

The connectivity of the ASes in the graph directly affect its density and clustering. Given a graph $G = (V, E)$, the density D of the graph is defined as the number of existing links out of the number of potential links, i.e., $D = \frac{2|E|}{|V|(|V|-1)}$.

Fig. 2 shows that the density of the AS-graph decreases over time, mainly since new ASes that join the Internet significantly increase the potential of links, however, they connect only to a small portion of the already existing ASes. This is expected since extremely large degrees are observed only in a relatively small number of ASes, and even these connect to a few thousands ASes.

In order to better understand the local connectivity of the ASes, we look at their clustering coefficient (CC), which is a measure of the local density of an AS based on its neighbors. More formally, the CC of a vertex in a graph is the number of triangles it forms with its immediate neighbors out of the potential number of triangles.

Fig. 3 depicts the CC of content and transit providers. Fig. 3a shows that the CC of content providers is mostly decreasing, a result of the increasing number of neighbors having few or no links amongst themselves. On the other hand, transit networks exhibit an increasing CC, meaning

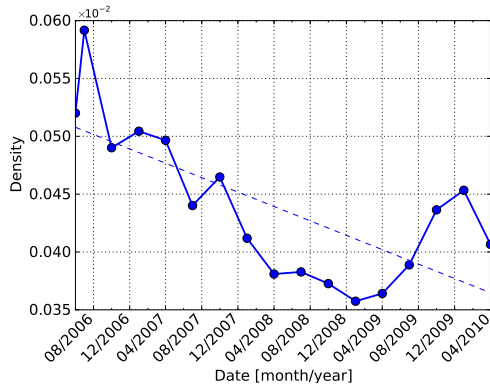


Figure 2: AS graph density

that they are losing neighbors with low connectivity and maybe hints at an increase in neighbor interconnections.

These two observations strongly indicate that content providers are increasing connectivity with access providers that are not in the core and are thus not interconnected. On the other hand, transit networks lose small customers and the remaining customers increase their interconnections that are used to bypass the core [16], which causes the transit networks to gradually lose their centrality.

4.2 IXPs

IXPs are a convenient method for ASes to interconnect, since it provides a shared facility and infrastructure [4]. ASes that seek to expand their connectivity have an incentive to use such facilities as it enables them to connect to a wide range of other ASes. As such, content providers can leverage IXP connectivity to gradually increase the number of peering ASes with minimal setup costs. In this analysis we use DIMES as it manages to detect more IXP links than iPlane.

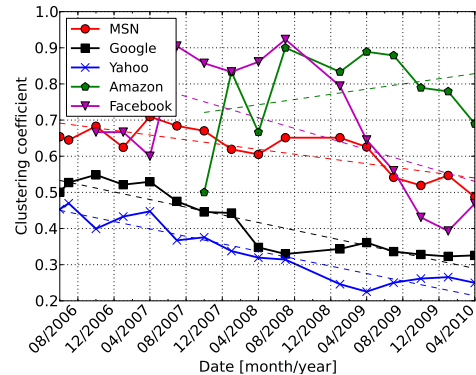
Fig. 4 shows the number of IXPs used by content and transit providers. Fig. 4a shows that content providers gradually increased their adoption of IXPs mainly since early 2009. Most of the content providers increased the number of IXPs they connect through by more than 100% in only a few years, emphasizing the important role that IXPs play in the Internet.

Fig. 4b shows that transit providers use more IXPs, but their percentage out of the overall neighbors is extremely low. Additionally, although transit providers increase their usage of IXPs, the number of IXPs they connect through grows much slower than content providers.

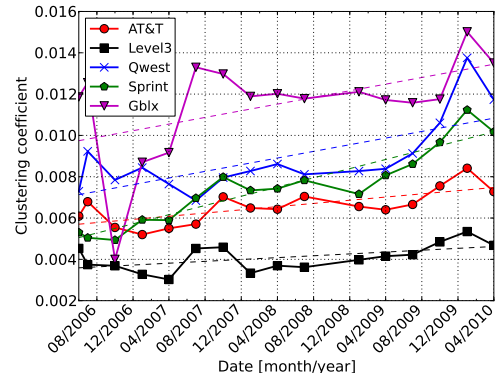
We further looked at the number AS-links that use an IXP, i.e., for each AS, the number of other ASes that it connects through an IXP. Although both transit and content providers exhibit a rising trend, the percentage of IXP links used by content providers is significantly higher, reaching almost 40% of their links, indicating that content providers indeed embrace IXPs as method for increasing connectivity.

4.3 Hierarchy

One of the methods used for studying the hierarchical structure of the Internet is k -pruning [5], which decomposes graphs into *shells*, based on the node connectivity towards the graph center. In the AS graph, ASes in the first shell are those who have only one link leading to the ‘center’ of



(a) Content providers



(b) Transit providers

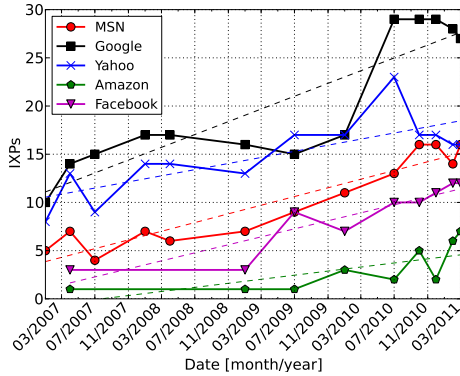
Figure 3: Clustering coefficient

the graph, whereas ASes in the k th shell have k -connectivity towards the center. The *nucleus* (or *core*) is the shell with the highest index, which is considered to contain top level providers, mostly tier-1 transit networks.

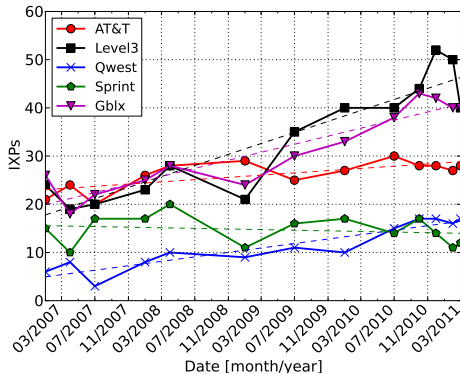
All of the transit networks (not shown) are in the same shell, which is the nucleus (varying between 35 and 41). This is expected, since these ASes are the top-level providers of the Internet, making them extremely central [25].

Fig. 5 shows the shell index of content providers. Google, Yahoo!, and MSN have a high shell index, and are either in the nucleus or in a very close shell. Facebook and Amazon exhibit a dramatic increase in their shell index. Facebook enjoyed increasing popularity, which drive it into connecting with high tier networks, improving its connectivity and service levels. Amazon made a market shift, from being an online store that resides in a low shell into a major cloud service provider, hosting numerous virtual servers. This change mandates significantly more connections with tier-1 networks, resulting in the observed climb.

For a content provider, being close to the nucleus has direct capital effect, because the higher the shell a network is located in, the less it needs to be a paying customer of transit networks [11], mainly since it does not need transit services in order to reach most other customer networks. Looking back at the degree of content providers in Fig. 1, their degree is significantly lower than tier-1 transit networks. Since tier-1 networks are expected to comprise the nucleus, it is unexpected that content providers manage to reach high



(a) Content providers



(b) Transit providers

Figure 4: Number of IXPs per AS

shells. However, examining the types of ASes in the nucleus, reveals that unlike previously thought [5], there is a significant portion of STPs (roughly 40%) and even CAHPs (15-25%) in the nucleus. As content providers increase the portion of neighbors that reside in the core, they manage to increase their shell index, thus reduce their dependency on top-tier transit providers. We note that the nucleus index has slightly increased since 2006 (roughly 5% increase in iPlane and 10% in DIMES). An increase in this index indicates a richer interconnection within the Internet core.

5. LOAD APPROXIMATION

Approximating the potential load on an AS is an estimate of its importance, since it indicates that an AS serves many other ASes for routing packets [19]. This approximation is commonly achieved using betweenness centrality (BC). In graph theory, BC measures the number of appearances of a vertex in the shortest paths between all other vertices, relative to the total number of shortest paths. Given a vertex $v \in V$, its bc $B(v)$ is calculated using: $B(v) = \sum_{s \neq v \neq t \in V} \frac{\sigma_{st}(v)}{\sigma_{st}}$, where σ_{st} is the number of shortest paths from s to t and $\sigma_{st}(v)$ is the number of shortest paths from s to t that pass through v . This value is normalized by dividing it by the number of possible pairs, i.e., $n(n-1)$.

BC is commonly applied to the AS-graph for measuring the possible load that an AS sustains. However, since packets traverse the Internet in valley-free paths [11], BC can

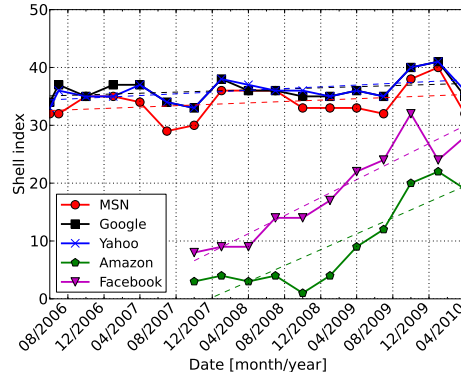


Figure 5: Shell index of content providers

only serve as an approximation for node centrality. We compared the approximated value with the BC of each AS as computed directly from the probed paths of each month. The latter was computed by dividing the number of traceroutes that traverse an AS (in the middle of the trace) by the total number of traceroutes. We found that the latter gives similar results to those calculated using the shortest path on the undirected AS graph, exhibiting slightly lower BC values and significantly more noise. Since we are interested in trends, we use the common method, which is less accurate but provides a clear view of the trends.

Fig. 6 shows the normalized BC of the selected content networks. The BC values of the tier-1 transit networks (not shown) retain a consistent value and are two orders of magnitude higher than content network, as the latter are usually the last hop of the routes. Interestingly, the figure shows an increasing trend in the BC of Google and Yahoo! (also validated with the DIMES dataset). We followed the AS-level traces of Yahoo! and Google during weeks 39 and 40 of 2010, and found that Yahoo! AS10310 and Google AS15169 appear mid-trace in roughly 22% and 14% of the traces they appear in, respectively. We further looked at the traces themselves, and found that both networks always appear as the hop-before-last, and are siblings of the last hop (owned by either Google or Yahoo!). In Google's case, the majority of traces terminate with YouTube (AS36561). Yahoo! on the other hand is comprised of several different ASes, such as Yahoo! Japan (AS23926), Yahoo! US (AS7233), and Yahoo! Backbone (AS24018).

These two networks exhibit a similar behavior: they both provide transit for sibling ASes as well as content, but only for data belonging to their networks. Yahoo!'s major AS does this because their network is comprised of several regional and probably specialized networks. Google provides transit to companies they purchase, leveraging their major AS's connectivity. In both cases, the end result is the same – content providers, which were once mainly stub networks that terminate routes, are becoming more central by providing transit towards other sibling and peering ASes, enabling network operators to save costs by leveraging their gradually improved connectivity. This observation is backed in previous work [16], that showed a decreasing traffic trend observed from YouTube AS and an increasing traffic trend from Google's AS, which is the result of Google's AS acting as the transit provider for YouTube's traffic.

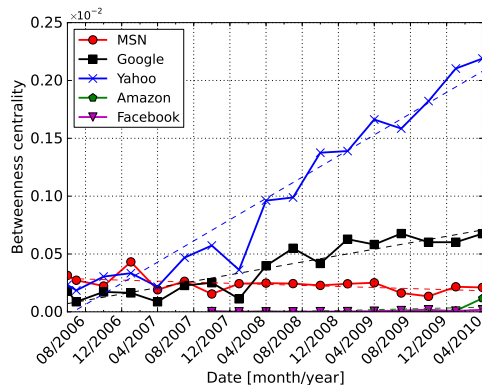


Figure 6: Betweenness centrality of content providers

6. CONCLUSION

The AS-level graph exposes the trend that is changing the Internet – content providers become key players in the Internet. The connectivity trends indicate that content providers increase the number of neighbors and diversify their types and geographical spread. New connections are made with small transit and access providers, enabling better reach to worldwide customers, while reducing transit costs. These changes are also witnessed in the way content providers are gradually climbing towards the core of the Internet.

The centrality of content providers is a strong indication to their increasing dominance, and on the other hand, the changing role of top-tier transit providers. In addition to saving transit costs, content providers increase their offerings to customers by providing Software as a Service (SaaS), such as hosting, cloud services, collaboration tools and others, without worrying for network neutrality violations or other traffic shaping that can interfere with the service grade they provide.

7. REFERENCES

- [1] Packet Clearing House, <http://www.pch.net/>.
- [2] PeeringDB, <http://www.peeringdb.org/>.
- [3] B. Augustin, X. Cuvelier, B. Orgogozo, F. Viger, T. Friedman, M. Latapy, C. Magnien, and R. Teixeira. Avoiding traceroute anomalies with paris traceroute. In *IMC'06*, 2006.
- [4] B. Augustin, B. Krishnamurthy, and W. Willinger. IXPs: mapped? In *IMC*. ACM, 2009.
- [5] S. Carmi, S. Havlin, S. Kirkpatrick, Y. Shavitt, and E. Shir. A model of Internet topology using k -shell decomposition. *Proceedings of the National Academy of Sciences USA (PNAS)*, 104(27), July 2007.
- [6] B. Chun, D. Culler, T. Roscoe, A. Bavier, L. Peterson, M. Wawrzoniak, and M. Bowman. Planetlab: An overlay testbed for broad-coverage services. *ACM SIGCOMM CCR*, 33(3), July 2003.
- [7] A. Dhamdhere and C. Dovrolis. Ten years in the evolution of the Internet ecosystem. In *IMC*, pages 183–196. ACM, 2008.
- [8] A. Dhamdhere and C. Dovrolis. The Internet is flat: modeling the transition from a transit hierarchy to a peering mesh. In *CoNext '10*. ACM, 2010.
- [9] X. Dimitropoulos, D. Krioukov, M. Fomenkov, B. Huffaker, Y. Hyun, kc claffy, and G. Riley. AS relationships: Inference and validation. *ACM SIGCOMM CCR*, 37, 2006.
- [10] C. Fraleigh, S. Moon, B. Lyles, C. Cotton, M. Khan, D. Moll, R. Rockell, T. Seely, and C. Diot. Packet-Level Traffic Measurements from the Sprint IP Backbone. *IEEE Network*, 17(6):6–16, 2003.
- [11] L. Gao. On inferring autonomous system relationships in the Internet. *IEEE/ACM Trans. on Networking*, 9(6), 2001.
- [12] Z. Ge, D. Figueiredo, S. Jaiwal, and L. Gao. On the hierarchical structure of the logical internet graph. In *SPIE ITCOM*, Aug. 2001.
- [13] P. Gill, M. Arlitt, Z. Li, and A. Mahanti. The flattening Internet topology: natural evolution, unsightly barnacles or contrived collapse? In *PAM*, pages 1–10. Springer-Verlag, 2008.
- [14] Y. He, G. Siganos, M. Faloutsos, and S. Krishnamurthy. Lord of the links: A framework for discovering missing links in the Internet topology. *IEEE/ACM Trans. on Networking*, 17(2):391–404, Apr. 2009.
- [15] X. Kuai, D. Zhenhai, Z. Zhi-Li, and C. Jaideep. On properties of internet exchange points and their impact on AS topology and relationship. In *Networking*, 2004.
- [16] C. Labovitz, S. Iekel-Johnson, D. McPherson, J. Oberheide, and F. Jahanian. Internet inter-domain traffic. In *SIGCOMM*, pages 75–86. ACM, 2010.
- [17] H. V. Madhyastha, T. Isdal, M. Piatek, C. Dixon, T. E. Anderson, A. Krishnamurthy, and A. Venkataramani. iPlane: An information plane for distributed services. In *OSDI*, pages 367–380, 2006.
- [18] H. V. Madhyastha, E. Katz-Bassett, T. Anderson, A. Krishnamurthy, and A. Venkataramani. iPlane Nano: path prediction for peer-to-peer applications. In *NSDI*, pages 137–152, 2009.
- [19] P. Mahadevan, D. Krioukov, M. Fomenkov, X. Dimitropoulos, kc claffy, and A. Vahdat. The Internet AS-level topology: three data sources and one definitive metric. *ACM SIGCOMM CCR*, 36(1), 2006.
- [20] G. Maier, A. Feldmann, V. Paxson, and M. Allman. On dominant characteristics of residential broadband Internet traffic. In *IMC*, pages 90–102. ACM, 2009.
- [21] Z. Mao, D. Johnson, J. Rexford, J. Wang, and R. H. Katz. Scalable and accurate identification of AS-level forwarding paths. In *INFOCOM*, 2004.
- [22] Z. M. Mao, L. Qiu, J. Wang, and Y. Zhang. On AS-level path inference. In *SIGMETRICS'05*, 2005.
- [23] M. E. J. Newman. Assortative mixing in networks. *Physical Review Letters*, 89(20), Oct. 2002.
- [24] Y. Shavitt and E. Shir. DIMES: Let the Internet measure itself. *ACM SIGCOMM CCR*, 35(5), 2005.
- [25] Y. Shavitt and U. Weinsberg. Quantifying the importance of vantage points distribution in internet topology measurements. In *IEEE Journal on Selected Areas in Communications*, Oct. 2011.
- [26] V. Valancius, C. Lumezanu, N. Feamster, R. Johari, and V. Vazirani. How many tiers? pricing in the Internet transit market. In *SIGCOMM*, 2011.