*Computer Science*
*Technical Report*

Colorado State University

# Characterizing International BGP Detours

Anant Shah, Christos Papadopoulos
{akshah,christos}@colostate.edu

November 18, 2015

Colorado State University Technical Report CS-15-104

Computer Science Department
Colorado State University
Fort Collins, CO 80523-1873

Phone: (970) 491-5792 Fax: (970) 491-2466
WWW: http://www.cs.colostate.edu

**Abstract**

There are currently no requirements (technical or otherwise) that BGP paths must be contained within national boundaries. Indeed, some paths experience *international detours*, i.e., originate in one country, cross international boundaries and return to the same country. In most cases these are sensible traffic engineering or peering decisions at ISPs that serve multiple countries. In some cases such detours may be suspicious. Characterizing international detours is useful to a number of players: (a) network engineers trying to diagnose persistent problems, (b) policy makers aiming at adhering to certain national communication policies, (c) entrepreneurs looking for opportunities to deploy new networks, or (d) privacy-conscious states trying to minimize the amount of internal communication traversing different jurisdictions.

In this paper we characterize international detours in the Internet during the months of August 2014 and April 2015. To detect detours we sample BGP RIBs every 6 hours from 172 RouteViews peers spanning 22 countries; geolocate visible ASes by geolocating each BGP prefix announced by each AS; and analyze each global BGP RIB entry looking for detours. Our analysis shows more than 15K - 19K unique BGP prefixes experienced a detour. A few ASes cause most detours and a small fraction of prefixes were affected the most. We observe about 2 million detours in RIBs of each month. Detours either last for a few days or persist the entire month. Out of all the detours, more than 50% - 60% were *transient* detours that lasted for 72 hours or less.

# 1 Introduction

We define an international detour (detour for short) as a BGP path that originates in an AS located in one country, traverses an AS located in a different country and returns to an AS in the original country. Detours have been observed in the Internet, for example, cities located in the African continent communicating via an external exchange point in Europe [7]. Many autonomous systems are also multinational, which means that routes traversing the AS may cross international boundaries. There have also been suspicious cases of detours. In November, 2013, the Internet intelligence company Renesys (now owned by Dyn) published an online article detailing an attack they called Targeted Internet Traffic Misdirection [6]. Using *Traceroute* data they discovered three paths that suffered a man-in-the-middle (MITM) attack. One path originated from and was destined to organizations in Denver, CO, after passing through Iceland, prompting concern and uncomfortable discussions with ISP customers.

Each of these anecdotes, while interesting in its own right, does not address the broader question about how prevalent such detours are, their dynamics and impact. Characterizing detours is important to several players: (a) as a diagnostic tool for network engineers trying to diagnose problems; (b) policy makers aiming at adhering to potential national communication policies mandating that all intra-country communication be confined within national boundaries, (c) entrepreneurs looking for opportunities to deploy new infrastructure in sparsely

covered geographical areas such as Africa, or (d) privacy-conscious states trying to minimize the amount of internal communication traversing different jurisdictions. The methodology developed to detect detours can also be incorporated in a tool to monitor the Internet routing system in near real-time and produce alerts. Network operators can not only appear informed about the incident, but also may be able to take action in peer selection in response to the alerts. Finally, longitudinal analysis of detours can give us insight into how routing and network infrastructure evolve over time.

In this paper we first develop methodology to detect detours and then use it to characterize them in a global scale. We study detours for two months, August 2014 and April 2015 using data from RouteViews. We first assign a country location to all advertised BGP prefixes and then to each corresponding AS using country-level information from MaxMind. Finally, we examine routing paths during those months, flagging detours.

The rest of this paper is organized as follows. In Section 2 we describe our datasets. Section 3 details the methodology used to quantify detours, reasoning for the choice of our datasets and analysis of our geolocation results. In Section 4 we characterize detours in August 2014 and April 2015. First we present aggregate analysis of entire dataset, then classify detours into different categories and finally focus on *transient detours* in Sections 4.1, 4.2 and 4.3 respectively. In Sections 5 and 6 we discuss value additions of our work and present related work. We summarize and present future work in Section 7.

## 2 Datasets

We use RouteViews [5] data for our analysis with sampling rate of four RIBs per day (one every six hours) for a total of 38,186 RIBs over 2 months from 172 peers. Selected peers include those in large exchange points. We use data from 9 RouteViews collectors spanning 22 countries, which amounts to 30G per month. Henceforth, we refer to our datasets as DS-Aug-2014 (August 2014) and DS-Apr-2015 (April 2015). For geolocation we use the freely available MaxMind GeoLite City database [3] from the second week of August, 2014. Finally, we use CAIDA AS Relationship datasets [2] from corresponding two months of our BGP data to evaluate false positives of detours detected.

## 3 Detecting Detours

One way to detect detours is to use *traceroute*, analyze reported hops and use latency as an indication of a detour. This approach was followed by [7] that studied peering relationships in Africa. Collecting data plane information at an Internet scale, however, is very difficult. The second approach, which we adopt, is to use control plane (BGP) information, which provides a direct insight into routing paths. However, detecting detours requires mapping every AS to a country, which is very challenging since an AS may have presence in more than

one country and it is very hard to determine the exact path without data plane information. We reject a third possible approach, namely using location from the RIRs, since (a) that information can be outdated, and (b) the location an AS was registered is a poor indication of the location of its routing infrastructure.

We considered yet another possible source of AS geolocation from CAIDA's Archipelago (Ark) dataset [4]. Ark provides Internet topology data derived from Traceroute-based measurements by collectors spread across the world. Their Internet Topology Data Kit provides router to organization mapping based on the routers' estimated geographical location. One could conceivably combine router-to-ASN and router-to-geo data to obtain ASN-to-geo. However, this method returns the router geolocation rather than the IP prefix geolocation. We create the AS-to-country mapping using MaxMind. We geolocate an AS by first geolocating all its advertised BGP prefixes. To map a BGP prefix to a country we break the prefix down into its /24 blocks and map all addresses from each block to a country using MaxMind. We believe country geolocation is accurate; in [8] and [13], authors determined that MaxMind is 99.8% accurate at country level geolocation. We decided to create the AS-to-country mapping using MaxMind [3]. We geolocate an AS by first geolocating all its advertised BGP prefixes. To map a BGP prefix to a country we break the prefix down into its /24 blocks and map all addresses from each block to a country using MaxMind. We believe country geolocation in MaxMind is accurate; in [8] and [13], authors determined that MaxMind is 99.8% accurate at country level geolocation.

Our methodology to detect detours is divided in three steps. First, we map all BGP prefixes in the global BGP table to a country as described above. Second, we map each AS to a country by extracting all prefixes originated by an AS and assembling a list of countries from those prefixes. The latter step results in each AS mapping to one or more countries. Finally, we examine the RIB of each peer looking for detours. We expand on these steps next.

### 3.1 Prefix Geolocation

In this section we describe our approach for prefix geolocation using MaxMind, depicted in Figure 1. To geolocate a BGP prefix, in this case a /22, we first break it down to its /24 sub-prefixes. Then we look up each IP address in every /24 block with MaxMind to get country location. This results to a list of countries for that BGP prefix. Plotting the results for the DS-Aug-2014 dataset, we see that about 99.98% of all /24 blocks and about 99.5% of all BGP prefixes geolocate to a single country. Figure 2 shows the distribution of countries in BGP prefixes. When BGP prefixes map to more than one country, the average size of the set was 2.66.

### 3.2 AS Geolocation

Next, to geolocate an AS we aggregate geolocation of all the BGP prefixes that originate from this AS. This methodology is shown in Figure 3.

The distribution of AS geolocation is shown in Figure 2. Perhaps surprisingly, only about 7% ASes out of more than 48,000 geolocated to multiple coun-

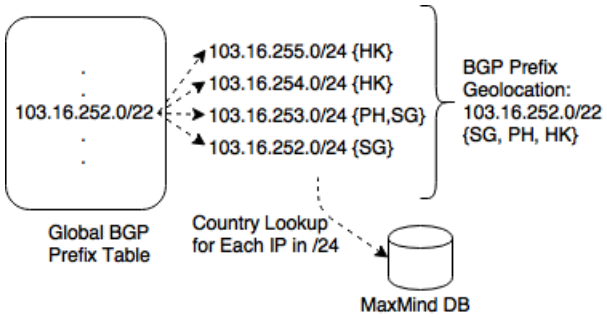


**Figure 1:** Example showing geolocation of a prefix seen in RIB

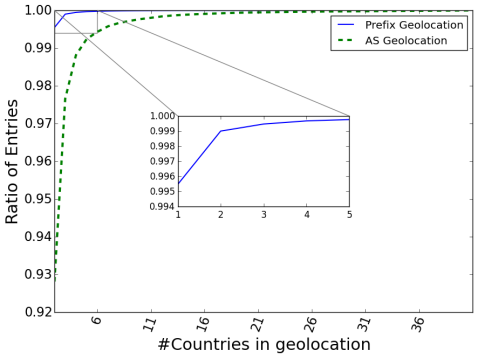


**Figure 2:** CDF: Number of countries in AS and Prefix geolocation

tries. We believe that this is the result of a practice where most organizations use a different AS number in different countries. If an AS does geolocate to multiple countries we use the set of all countries in our analysis.

## 3.3 AS Path analysis

In this step we search all paths in each routing table to find detours. We define a path as having a detour if the origin and destination is country *'A'* but the path unambiguously includes some other country *'B'*. Note that this approach examines paths where the prefix origin AS and the AS where the peer is located are in the same country. To analyze the AS path, we provide the following definitions:

- **Prefix Origin**: The AS that announces the BGP prefix.
- **Detour Origin AS**: The AS that starts a detour in country *'A'* and diverts the path to foreign country *'B'*.
- **Detour Origin Country**: The country where we approximate location of Detour Origin AS.
- **Detour Destination AS**: The AS in foreign country *'B'*.

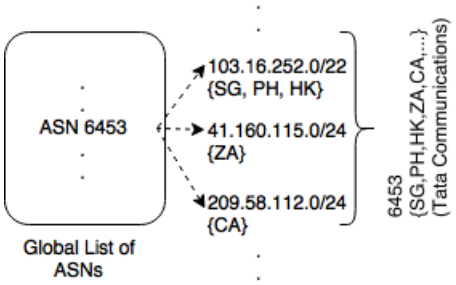


**Figure 3:** Example showing geolocation of an AS

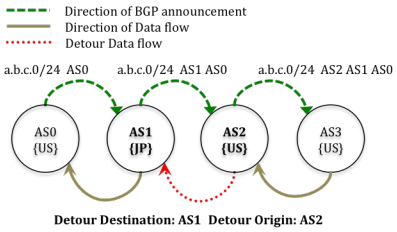


**Figure 4:** Example showing direction of BGP announcement and direction of observed detour

- **Detour Return AS**: The AS where detour returns back in country *'A'*.

Figure 4 illustrates detours. *AS0* announces prefix *a.b.c.0/24* to AS1, AS2 and AS3. AS1 geolocates to JP whereas AS3, AS2 and AS0 are in the US. In this case, data traversing from AS3 to AS0 will contain a detour from AS2 (Detour Origin) to AS1 (Detour Destination). We do not include sub-paths in our analysis; other portions of the path that may experience a detour. For example, in path AS1{US}-AS2{IN}-AS3{CN}-AS4{IN}-AS5{US}, we only count the detour US-IN-US. We do not count the detour IN-CN-IN.

There are some cases where we need to approximate detour origin and country. In a path such as AS1{US}-AS2{US,BR}-AS3{CN}-AS4{US}. We resolve the uncertainty of the detour origin by assuming that it starts in AS2, since there is a likely path to AS2 from AS1 through the US and AS2 starts the detour from US, not BR.

It is possible that we find an AS with no known geographic locations because its BGP prefixes could not be geolocated. 2.2% ASes in our dataset could not be geolocated. We exclude any subpath containing such an AS from our analysis. Finally, since some ASes map to multiple countries it is not always possible to deduce a detour. Thus we characterize some detours as ***possible*** detours. For example, a path that geolocates to {US}-{US,IN}-{US} may in fact stay within the US and never visit India. In this work we focus on paths that contain ***definite*** detours, such as {US}-{IN}-{US}.

## 3.4 Filtering peered AS paths

It is possible that the detour origin and the detour return ASes have a peering relationship and reality traffic was not detoured at all. This, however, is hard to determine with certainty since peering relations and policies are not public. What we can do is provide an upper bound on how many detours may be eliminated due to peering. To detect such cases we use CAIDA's AS relationship dataset [2]. This dataset provides information of provider to provider (p2p) and provider to customer (p2c) relationship between ASes. We count cases

where p2p link might be used, i.e., data originates from the peer itself or from a downstream customer. In case of p2c link we assume this link is always chosen. While it is important to understand total number of detours that may be negated by peering, we simply note that number but do not remove them from our results. We revisit this issue in the next section and summarize the peering relationships in Table 2.

# 4 Results

In this section we quantify detours detected in our dataset. Since we sample 4 RIBs per day the sampling rate is six hours. First, in Section 4.1 we present an overview of all the detours detected in our datasets. In Section 4.2 we define metrics and classify detours based on their stability and availability. In Section 4.3 we focus on transient detours.

## 4.1 Aggregate Results

We begin by characterizing aggregate results, namely all detours seen by all peers. For these results we count all detour incidents; in other words, we count an incident every time a path appears in a RIB of any peer that contains a detour. Many of these incidents are duplicates. Therefore in addition to the total we also present the number of unique detours.

We observe that detours are not generally common. On an average we find about 70K - 80K detoured paths out of a total of 240M examined entries per day. Also, not all peers see a detour. Only 65 peers in August 2014 and 73 in April 2015 saw one or more detours. Figures 5 and 6 show the number of detours for each day in DS-Aug-2014 and DS-Apr-2015. Table 1 details the number of detours seen. We analyzed about 7 billion RIB entries in each month and only about 2 million entries showed a detour; out of the 2 million detours only 60K were unique.

Next we examine the visibility of detours, where we observe an uneven distribution among ASes. Just 6 ASes originate more than 50% of the detours. Similarly, some prefixes experience detours more than others. About 800 prefixes in DS-Aug-2014 and about 1200 prefixes in DS-Apr-2015 experienced more than 50% of the total detours. Looking at the average length of a detour, we

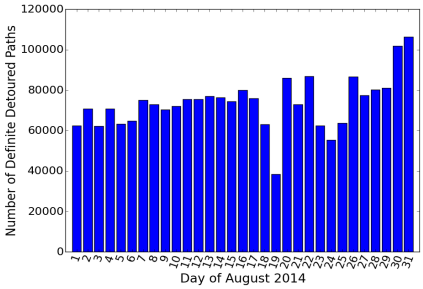


**Figure 5:** Total number of definite detours per day in DS-Aug-2014

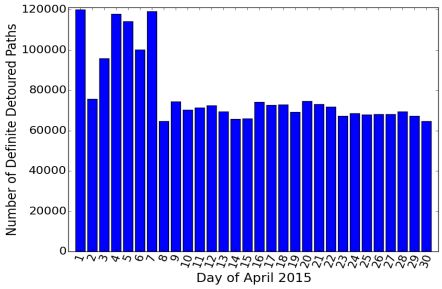


**Figure 6:** Total number of definite detours per day in DS-Apr-2015

**Table 1:** Aggregate number of detours detected

| | #Total RIB Entries | #Detoured Entries | #Unique Detours |
|---|---|---|---|
| DS-Aug-2014 | 7,105,389,810 | 2,282,762 | 61,845 |
| DS-Apr-2015 | 7,247,881,444 | 2,347,600 | 62,883 |

see that a detour visits 1 to 2 foreign ASes before returning to its origin country. The maximum length of a detour we observed was 4 - 5 hops. We do not characterize the geographical distance of detours, a task not supported by our data.

We now estimate the effect of peering links on detours. Specifically, we are interested in cases where a peering relationship exists between the *Detour Origin AS* and the *Detour Return AS* as described in Section 3.4. If such a link exists, it is possible that traffic traverses that link instead of the detour. Table 2 shows the number of detours between ASes that also have peering relations compared to total number of detours. We find that less than 15% of the detours could be false. However, it is not possible to be certain about when traffic actually follows a peering link as opposed to the detour. Also, data from a downstream AS which is not a customer may still experience the detour. Because of this uncertainty we do not alter our original estimate of detours but acknowledge that we may be over-counting by as much as 15%.

**Table 2:** Routes that may have peering relations

| | #Total Detours | #Detours with possible peering | % |
|---|---|---|---|
| DS-Aug-2014 | 2,282,762 | 315,323 | 13.81% |
| DS-Apr-2015 | 2,347,600 | 286,555 | 12.2% |

## 4.2 Characterizing Detours

In this section we define metrics and characterize the detours we observed. We focus on two metrics:

1. **Detour Dynamics**
   (a) **Flap Rate**: Measure of stability of a detour; how many times a detour disappeared and reappeared.
   (b) **Duty Cycle**: Measure of uptime of a detour throughout the month measurement period.
2. **Persistence**: Total number of continuous hours a prefix was seen detoured.

An AS may contain multiple RouteViews peers with similar views, which contributes potentially duplicate detours to our dataset. If we do not account for such duplicates our results will be skewed towards ASes with multiple peers. We follow a simple approach to deal with this problem: if an AS contains more

than one peer we select the peer that saw the most detours as the representative of that AS. This may potentially undercount detours since some peers in same AS may see different detours. After selecting a representative we are left with 56 and 57 peers in DS-Aug-2014 and DS-Apr-2015 respectively.

We now continue our characterization of detours by looking at **detour dynamics**. Specifically we focus on flap rate and duty cycle, defined as follows:

$$FlapRate = \frac{TotalTransitions}{TotalTime} \times 100 \qquad DutyCycle = \frac{TotalUptime}{TotalTime} \times 100$$

Figures 7 and 8 show a scatter plot of flap rate vs. duty cycle for all detours in our datasets. We see a triangular pattern with some outliers. Next we drill into country specific detours. woFigures 9 and 10 show a scatter plot of flap rate vs. duty cycle for various detours in US, Brazil and Russia. We selected these three countries because they show the most detours in our datasets; they account for 92% and 93% of all detours in DS-Aug-2014 and DS-Apr-2015 respectively. We divide each figure into 4 quadrants based on average flap rate and average duty cycle of all detours. In DS-Aug-2014, US detoured paths appear more stable (lower flap rate and higher duty cycle). On the other hand, Russian and Brazilian detoured paths fall mostly in the *Ist* and *IIIrd* quadrant. DS-Apr-2015 shows a similar pattern. We also studied a similar scatter plot for all the non US, BR and RU detours. In this case we observed detours mostly in extreme ends on *Ist* and *IIIrd* quadrant indicating two categories of detours, either frequently reoccurring or very rare events. Lastly, we looked at behavior of detours in South Africa and Kenya where lack of peering has been previously reported. Here, detours mostly lie in *IInd* quadrant indicating very stable detours.

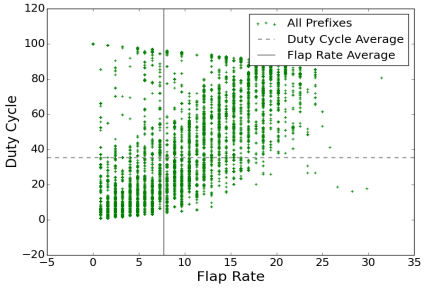


**Figure 7:** Flap Rate vs DC for all prefixes in DS-Aug-2014

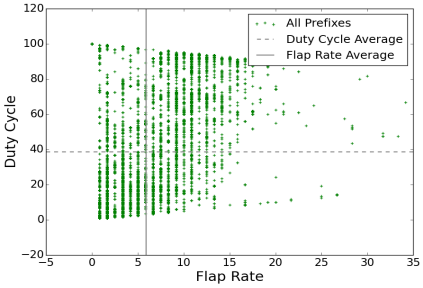


**Figure 8:** Flap Rate vs DC for all prefixes in DS-Apr-2015

Next, we examine the **persistence** of detours. Figures 11 and 12 show the number of consecutive days a detour was visible by any peer. We see a U-shaped pattern in the figures, meaning that many detours are either short lived (one day) or they persist for entire month. We take a different view at persistence in Figure 13 by plotting CDF of duration in hours. We see that most detours are short-lived, with about 50% to 60% lasting less than 72 hours, defined as *transient* detours. The selection of the 72-hour threshold is somewhat arbitrary

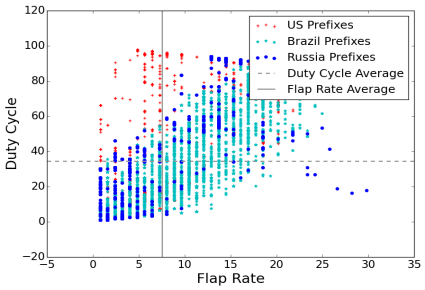


**Figure 9:** Flap Rate vs DC for US, RU and BR prefixes in DS-Aug-2014

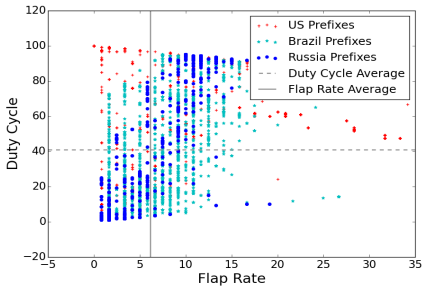


**Figure 10:** Flap Rate vs DC for US, RU and BR prefixes in DS-Apr-2015

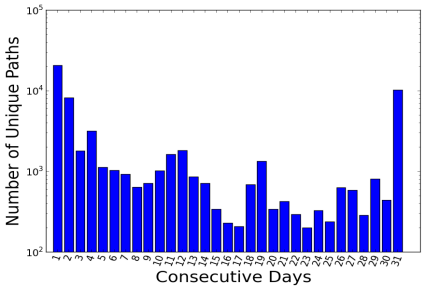


**Figure 11:** Persistence of definite detoured paths as seen by all peers in DS-Aug-2014

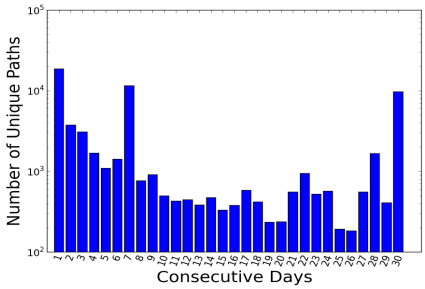


**Figure 12:** Persistence of definite detoured paths as seen by all peers in DS-Apr-2015

and we justify it as follows: this is the time between a misconfiguration, which is accidentally made on a Friday, and when it is fixed on Monday morning. We do realize that most networks will fix problems sooner than that, but there is a wide range in how networks are managed. Finally, we also examine a specific case of a transient detour, namely *flash detours* which appeared only once and never appeared again during the month.

In the following section we focus on transient and flash detours. Due to space limitations we do not characterize persistent detours further. We do note, however, that characterizing persistent detours is important for at least some of the reasons we enumerated earlier. We chose to focus on transient detours as they shed light on misconfigurations or even malicious activities, both aspects of routing we understand less.

## 4.3 Transient and Flash Detours

We first present an understanding of the transient detours on per-country basis. Since there are more than one peers in some countries and different peers see varying number of transient detours, we calculate an average number of transient detours per country by dividing total number of transient detours in a country by number of peers in the given country. This average value per country is presented in Figure 14. We detected detours in 16 different countries and some

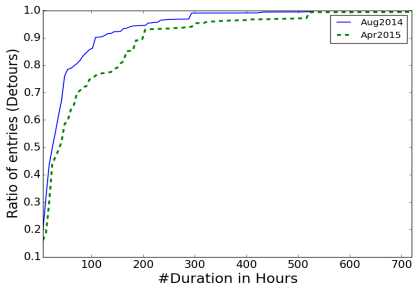


**Figure 13:** Distribution of detour duration observed in DS-Aug-2014 and DS-Apr-2015

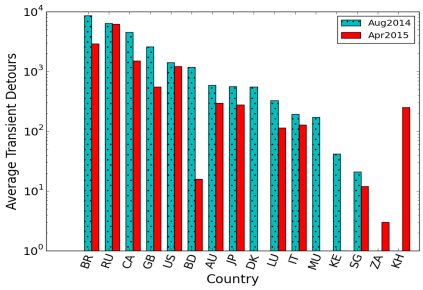


**Figure 14:** Average number of transient detours in DS-Aug-2014 and DS-Apr-2015

countries showed detours only in one dataset. Moreover, some countries showed very similar average number of transient detours in both datasets. Figures 15 and 16 show a distribution of ASes that initiate detours and prefixes affected by detours. We observe that less than 5 ASes originate 50% of the transient detours and about 1500 prefixes account for 50% of the transient detours. Finally, we

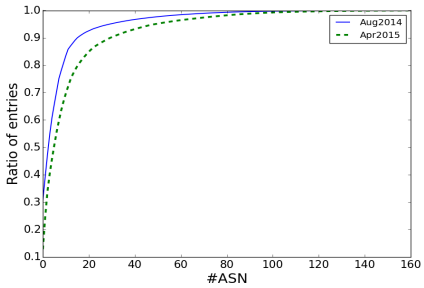


**Figure 15:** Distribution of ASes that originated a transient detour. The top 5 Detour Origin ASes account for 50% of all transient detours

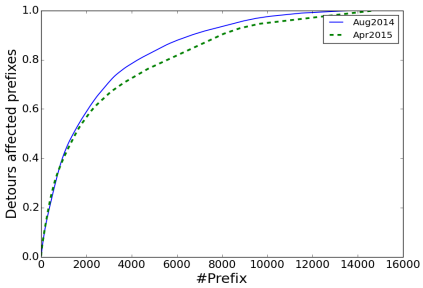


**Figure 16:** Distribution of prefixes that experienced a transient detour. About 1500 prefixes account for 50% of all transient detours

look at *flash* detours. These are detours that appeared only once and were observed in one RIB of at least one peer. While flash detours account for only 0.34% and 0.36% of the total number of detours in August 2014 and April 2015 respectively, 43% of the prefixes in August 2014 and 35% of the prefixes in April 2015 experienced at least one flash detour.

# 5 Discussion

In this paper we present a first attempt to characterize detours in the Internet. We find many detours to be persistent and most likely the result of traffic engineering decisions. We also find a significant portion, however, that are

transient, appearing only for a few hours or days.

Characterizing detours in the Internet is very useful. Customers gain more insight into how their providers route traffic. There is perhaps an expectation from users that if they send traffic to other users in the same country the packets will not step outside national borders; our work provides evidence to the contrary. Network operators can use our methodology and results for diagnostic purposes. A sudden change in RTT may be traced to a detour, or keeping track of what the routing system does. The latter is important to assure customers that their traffic is not subject to monitoring by other governments.

Our work is useful to regulators and state officials responsible for network infrastructure, since our work quantifies information about a practice that may run afoul of state policy. State officials can use such information to assure citizens that their traffic stays within national borders or direct ISPs to alter their practices. State agencies that transmit sensitive information may monitor detours to alert for potential MITM attacks. For example, we did observe cases where prefixes belonging to US DoD were detoured through Japan.

Finally, entrepreneurs may use our results when deciding where to establish new Internet Exchange points (IXP) or deploy infrastructure in developing countries.

Our results on AS and prefix geolocation are available using RESTful API at http://geoinfo.bgpmon.io.

# 6 Related Work

In November 2013 Renesys reported a few suspicious paths [6]. One went from Guadalajara, Mexico to Washington, D.C. via Belarus; another went from Denver, CO through Reykjavik, Iceland, back to Denver. They used mostly data plane information from traceroute for their analysis. In [7] the authors focus on ISP inter-connectivity in the continent of Africa. They too, searched for paths that leave Africa only to return back. The goal, however, was to investigate large latencies in Africa and ways to reduce it. The premise was that if a route crosses international boundaries it would exhibit high latency. The work pointed to cases where local ISPs are not present at regional IXPs and IXP participants don't peer with each other. Similar to Renesys, they also use traceroute measurements, this time from the BISmark infrastructure (a deployment of home routers with custom firmware) in South Africa. Our study extends beyond Africa and investigates transient in addition to long-lasting detours. In *Boomerang* [12], the authors again use traceroute to identify routes from Canada to Canada that detour through the US. In this work the motivation was concerns about potential surveillance by the NSA. The work in [10] also analyzes the control plane (RIBs and AS paths) to construct a network topology and then uses traceroute to construct country-level paths. The goal of this work was to understand the role of different countries that act as hubs in cross-country Internet paths. Their results show that western countries are important players in country level internet connectivity. In [11] authors present

*ASwatch*, an AS reputation system to detect bulletproof hosting ASes. Similar to our work ASwatch relies on control plane information to detect malicious ASes (that may host botnet C&C servers, phishing sites, etc). The motivation of this work is different than ours. ASwatch attempts to detect malicious ASes by mining their link stability, IP space fragmentation and prefix rechaebility. ASwatch will not detect ASes that cause detours. The detour origin ASes that our work detects could complement features that ASwatch uses. As authors in [11] point out malicious ASes rewire their routes more frequently than legitimate ones, transient detours might be particularly useful to improve detection capability of ASwatch.

In context of MaxMind geolocation accuracy, [8] and [13] have shown MaxMind country geolocation to be 99.8% accurate. In [14] authors use data from Routing Information Registries (RIRs), RIPE DB and Team Cymru to determine all IP blocks and ASes that geolocate to Germany. To validate their geolocation accuracy, authors query the MaxMind database which allows mapping IP addresses to their country of presence. We adopt a more exhaustive strategy than [14].

# 7 Conclusion and Future Work

In this paper, we sampled BGP routing tables from 172 peers around the world over the entire months of August 2014 and April 2015 to investigate international detours. We see about 60K distinct entries in RIBs that show a detour. More than 50% of the detours last less than 72 hours. We also discover that a few ASes cause most of the detours and detours affect a small fraction of prefixes. Some detours appear only once.

In the future, we aim to analyze BGP update messages in real time with the BGPmon data stream [1]. This would give us the finest detail of BGP data and would best describe the nature of transient paths. Additionally, we will investigate a tool that conducts a Traceroute every time a route with detour is found in the AS path of a BGP announcement. This Traceroute would be used to see if the path taken at the data plane is congruent with the path seen in the control plane, similar to the work done by Hyun et al. [9]. In addition to their work, however, this tool would also seek to find if the geographical distance the packets travel is unusual and to what degree these detours affect latency. Doing so would give great insight as to whether transient BGP detours are due to converging BGP updates or malicious activity, and it could also prove to be useful heuristic for BGP alert and defense systems.

**Acknowledgement**

The work in this paper is partially sponsored by the Department of Homeland Security (DHS) Science and Technology Directorate, HSARPA, Cyber Security Division, via SPAWAR Systems Center Pacific under Contract No. N66001-13-C-3001 (all authors), and via BAA 11-01-RIKA and Air Force Research Laboratory, Information Directorate under agreement numbers FA8750-12-2-0344 and FA8750-15-2-0224 (Papadopoulos). The U.S. Government is authorized to make reprints for Governmental purposes notwithstanding any copyright. The views contained herein are those of the authors and do not necessarily represent those of DHS or the U.S. Government.

# References

[1] BGP monitoring system. http://www.bgpmon.io.

[2] CAIDA AS Relationships. http://www.caida.org/data/as-relationships/.

[3] (maxmind geoip country database.

[4] The CAIDA Internet Topology Data Kit. http://www.caida.org/data/internet-topology-data-kit/.

[5] University of Oregon Route Views Project. http://www.routeviews.org/.

[6] JIM COWIE. The New Threat: Targeted Internet Traffic Misdirection, Nov 2013. http://www.renesys.com/2013/11/mitm-internet-hijacking/.

[7] Arpit Gupta, Matt Calder, Nick Feamster, Marshini Chetty, Enrico Calandro, and Ethan Katz-Bassett. Peering at the Internets Frontier: A First Look at ISP Interconnectivity in Africa. In Michalis Faloutsos and Aleksandar Kuzmanovic, editors, *Passive and Active Measurement*, volume 8362 of *Lecture Notes in Computer Science*, pages 204–213. Springer International Publishing, 2014.

[8] B. Huffaker, M. Fomenkov, and k. claffy. Geocompare: a comparison of public and commercial geolocation databases - Technical Report . Technical report, Cooperative Association for Internet Data Analysis (CAIDA), May 2011.

[9] Young Hyun, Andre Broido, and KC Claffy. Traceroute and BGP AS path incongruities. 2003.

[10] Josh Karlin, Stephanie Forrest, and Jennifer Rexford. Nation-State Routing: Censorship, Wiretapping, and BGP. *CoRR*, abs/0903.3218, 2009.

[11] Maria Konte, Roberto Perdisci, and Nick Feamster. ASwatch: An AS Reputation System to Expose Bulletproof Hosting ASes. *SIGCOMM Comput. Commun. Rev.*, 45(5):625–638, August 2015.

[12] Jonathan A. Obar and Andrew Clement. Internet Surveillance and Boomerang Routing: A Call for Canadian Network Sovereignty. *SSRN Electronic Journal*, 2013.

[13] Yuval Shavitt and Noa Zilberman. A Study of Geolocation Databases. *CoRR*, abs/1005.5674, 2010.

[14] Matthias Wählisch, Thomas C. Schmidt, Markus de Brün, and Thomas Häberlen. Exposing a Nation-centric View on the German Internet — a Change in Perspective on AS-Level. In *Proceedings of the 13th International Conference on Passive and Active Measurement*, PAM'12, pages 200–210, Berlin, Heidelberg, 2012. Springer-Verlag.