

# Towards Characterizing International Routing Detours

Anant Shah

Colorado State University  
akshah@cs.colostate.edu

Romain Fontugne

IIJ Research Lab  
romain@iiij.ad.jp

Christos Papadopoulos

Colorado State University  
christos@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 month of January 2016. To detect detours we sample BGP RIBs every 8 hours from 461 RouteViews and RIPE RIS peers spanning 30 countries. Then geolocate visible ASes by geolocating each BGP prefix announced by each AS, mapping its presence at IXPs and geolocation infrastructure IPs. Finally, analyze each global BGP RIB entry looking for detours. Our analysis shows more than 5K 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 544K detours. Detours either last for a few days or persist the entire month. Out of all the detours, more than 90% were *transient* detours that lasted for 72 hours or less. We also show different countries experience different characteristics of detours.

## Keywords

AS Geolocation, Routing Detours, MITM

## 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 [11]. Many au-

tonomous 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 [10]. 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 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. Using the methodology developed to detect detours we also present a tool, *Netra*<sup>1</sup>, 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, validate it on live traffic using our tool *Netra* and then use it to characterize them at a global scale on historical BGP data of January 2016 from RouteViews and RIPE RIS.

The rest of this paper is organized as follows. In Section 2 we present related work and highlight previous efforts in direction similar to ours and point out key

<sup>1</sup><https://github.com/akshah/netra>

areas where our work differs from them. In Section 3 we describe our datasets, corresponding usage and reasoning for the choice of our datasets. Section 4 details the methodology used to perform AS geolocation and analysis of our geolocation results. Section 5 explains in detail detour detection process, corresponding terminologies used throughout the paper. In Section 6 we explain our data plane measurements and present validation results. In Section 7 we characterize detours seen in January 2016. First we present aggregate analysis of entire dataset, then classify detours into different categories and finally focus on *transient detours* in Sections 7.1, 7.2 and 7.3 respectively. In Sections 8 and 9 we discuss value additions of our work, summarize and present future work respectively.

## 2. RELATED WORK

### Detour detection:

In November 2013 Renesys reported a few suspicious paths [10]. 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 [11] the authors focus on ISP inter-connectivity in the continent of Africa. They 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* [19], 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. This work differs from ours in a number of ways: we characterize detours not just for one but 30 countries using control plane information rather than data plane. We use data plane measurement only for validation purposes. Our goal is to not only detect detours but show characteristics about them which previous work does not present.

### Data plane vs Control plane Incongruities:

In [9] authors focus on routing policies and point out cases where routing decisions taken by ASes do not conform to expected behavior. There are complex AS relationships, such as, hybrid or partial transit which impact routing. Such relationships may lead to false positives in our results. However, the paper points out that

most violations of expected routing behavior caused by complex AS relationships are very few and most violations were caused by major content providers. Our work identifies detours for variety of ASes, including both large content providers and small institutions. Moreover, in [24] authors argue that such incongruities are caused due to incorrect IP to AS mappings. About 60% of mismatches occur due to IP sharing between adjacent ASes. Authors here show that 63% to 88% of paths observed in control plane are valid in data plane as well. The work in [14] 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.

### Malicious AS detection:

In [16] 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 reachability. 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 [16] point out malicious ASes rewire their routes more frequently than legitimate ones, transient detours might be particularly useful to improve detection capability of ASwatch.

### Geolocation Accuracy:

In context of MaxMind geolocation accuracy, [12] and [20] have shown MaxMind country geolocation to be 99.8% in consensus with other geolocation DBs. In [21] 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 [21].

### Control-plane-only for detection:

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 [11] that studied peering relationships in Africa; we too use this approach to validate our results on live data. However, we detect detours using only control plane data. This has a number of advantages: 1) Collecting data plane information at an Internet scale is hard. It needs infrastructure and

visibility provided by Atlas probes or Ark monitors is limited. Moreover, running too many traceroutes from own network to others might lead to blacklisting. 2) Small footprint of our methodology makes it easily reproducible. Any network operator can pull a RIB dump from his/her border router and run *Netra* to detect detours for prefixes they own.

### 3. DATA SOURCES

We use variety of data sources to perform AS geolocation, BGP RIBs for detour detection and Traceroutes from RIPE Atlas for detour validation. In Table 1 we list different datasets with their usage and relevant information about each. Our sampling rate is 3 RIBs per day (one every eight hours, as provided by RIPE RIS) for a total of 38,688 RIBs from 416 peers. This spans 30 countries, which amounts to about 55GB of compressed MRT data. We acknowledge that 30 countries do not necessarily represent global scale, but our scope is limited by placement of peers that provide BGP feeds. We used all v4 peers in our analysis.

For geolocation of IP addresses we use MaxMind GeoLite City DB [17]. We treat end user IPs and infrastructure IPs differently since MaxMind is known to be more accurate for eye-ball networks only. To gather the list of infrastructure IPs we used list of routers from CAIDA Ark traceroutes [1], OpenIPMap [8], iPlane [5] and RIPE Atlas built-in measurements and the anchoring measurements. The built-in measurements use all the RIPE Atlas probes and the destinations are root servers. The anchoring measurements are from 400 Atlas probes to other 189 Atlas anchors. These infrastructure IPs are then mapped to AS using IP to AS mappings from CAIDA ITDK [4], iPlane or longest prefix match.

In addition to BGP sources, we use AS-to-IXP mapping to estimate presence of an AS in a country. We gather AS to IXP mappings from Packet Clearing House (PCH) [6], PeeringDB [7] and by crawling 368 IXP websites that make their participant list public. Finally, we use CAIDA AS Relationship datasets [3] to eliminate false positives from detours detected. In Section 4 we provide more details on how these datasets are used in AS geolocation along with a flowchart (Figure 1).

## 4. AS GEOLOCATION

To detect detours we are interested in country level geolocation. We define AS geolocation as presence of an AS in a country. An AS can have presence in multiple countries, especially ASes that belong to large providers. We detect the presence of an AS in country *A* if it :

1. Announces a prefix that geolocates to *A* or
2. Has infrastructure IPs that geolocate to *A* or

3. Has a presence at an IXP in *A*.

In Figure 1 we show a flowchart detailing AS geolocation processes. There are 3 main steps as described above. In next sections we elaborate on each.

### 4.1 Prefix Geolocation

We begin by geolocating all advertised BGP prefixes by an AS. It is possible that during our analysis in January 2016 some AS erroneously announced prefixes that it did not own. Therefore we perform a simple filtering; we trust an AS to be owner of a prefix if it announced the prefix for at least 15 days in our dataset. We assume most mistakes or hijacks will be less than this duration. Next, to map a BGP prefix to a country we geolocate each IP in the prefix using MaxMind-free. MaxMind could not geolocate 3.8M IPs. We could successfully geolocate 1.48M of these IPs with MaxMind-paid, we could not use remaining 2.32M IPs. Now we use the union of IP geolocation sets to get the BGP prefix geolocation. Due to the 2.32M IPs not geolocating even with paid version of MaxMind, 614 BGP prefixes could not be geolocated. For the remaining 610,722 BGP prefixes<sup>2</sup> which were geolocated we observe that more than 99% geolocated to single country. We note that 328,398 BGP prefixes were /24s. When BGP prefixes map to more than one country, the average size of the set was 2.9 countries. Finally, we perform union of geolocation sets of all BGP prefixes that an AS announces to create 1<sup>st</sup> AS to country set.

### 4.2 Infrastructure IP Geolocation

As mentioned previously, we treat infrastructure IP addresses separately. Router geolocation is known to be inaccurate [15]. Therefore for these IPs we want to create country geolocation set as large as possible. We populate list of router IPs from CAIDA Ark Traceroutes, iPlane IP to PoP mappings, OpenIPMap and RIPE built in measurements. Our list included 3M router IPs. This is the ‘Read Infrastructure IPs’ step shown in flowchart Figure 1. To geolocate each router IP we look at country location provided by iPlane, OpenIPMap<sup>3</sup> and Maxmind-paid and perform a union to give a set of countries. Next step is to map these routers to ASes. IP to AS is a challenging problem and active area of research. We use the best datasets available to create these mappings. Both CAIDA ITDK and iPlane datasets provide IP to AS mappings using the methodology described in [13]. For cases where either of these datasets fail to provide IP to AS mapping, we perform

<sup>2</sup>We use BGP prefixes ‘as is’ from the RIBs and do not perform any prefix aggregation. For example, if both /8 and /9 blocks of a prefix were seen in RIBs of the same or different peers, they are treated as 2 separate prefixes.

<sup>3</sup>OpenIPMap is crowdsourced and may not be very accurate. We use cases where confidence level for router geolocation is higher than 90%.

**Table 1: Dataset Description**

Name	Usage	Date	Sources	Info
BGP	AS Geolocation; Detour Detection	2016-01	RouteViews, RIPE RIS	38,688 RIBS, 416 peers, 30 countries, 55GB
Infrastructure IP List	AS Geolocation	2016-01 to 2016-03	CAIDA Ark, iPlane, OpenIPMap, RIPE Atlas Measurements	3M Router IPs
Infrastructure IPs to AS Mapping	Infrastructure IP geolocation	2015-08	CAIDA ITDK, iPlane	6.6M IP to AS mappings
AS to IXP Mapping	AS Geolocation	2016-01 to 2016-03	IXP websites, PeeringDB, PCH	368 IXP websites crawled
AS Relationship	Filtering peered paths from detection	2016-01	CAIDA AS Relationship	482,657 distinct relationships
Traceroute	Detour Validation	2016-05-01	RIPE Atlas	Used by <i>Netra</i> , 163 traceroutes
MaxMind	Prefix Geolocation; Detour Validation	2016-01, 2016-03	MaxMind GeoLite City (free and paid)	Paid version used only for geolocating infrastructure IPs and detour validation


longest prefix match on the global routing table and map the IP to the AS announcing the longest matching prefix. Lastly, we combine IP to Country and IP to AS mappings to give 2<sup>nd</sup> AS to country set.

### 4.3 IXP Presence of an AS

We extract presence of ASes at different IXPs and add the geolocation of IXP to the AS geolocation. As shown in ‘Read IXP data’ step in Figure 1, we use 3 sources of AS to IXP mappings. First, we crawl 368 IXP websites and extract their corresponding participants. Next, we use PeeringDB 2.0 API [7] and lastly, we use dataset from Packet Clearing House (PCH) that lists participants at IXPs that PCH is also a part of. We then combine geolocation obtained from these IXP sources to obtain 3<sup>rd</sup> AS to country set. We acknowledge that IXP mappings from websites, PCH and PeeringDB might not be updated regularly and hence lead to mapping of an AS to a country that it does not have a presence in. Note that this will lead to false negative (not false positives) in detour detection, a trade-off we make to error on safe side.

### 4.4 AS to Country Set

Finally, we map an AS to a set of countries by taking a union of all the 3 steps above. This is the merge step in Figure 1. The distribution of AS geolocation is shown in Figure 2. Perhaps surprisingly, only about 11.6% ASes out of a total of 52,984 geolocated to multiple countries. 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. We could not geolocate 24 ASes because none of their BGP prefixes could be geolocated, no infrastructure IP from our set mapped to it nor did we find its IXP presence in public datasets. These ASes on an average announced only 2 to 3 BGP prefixes.

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

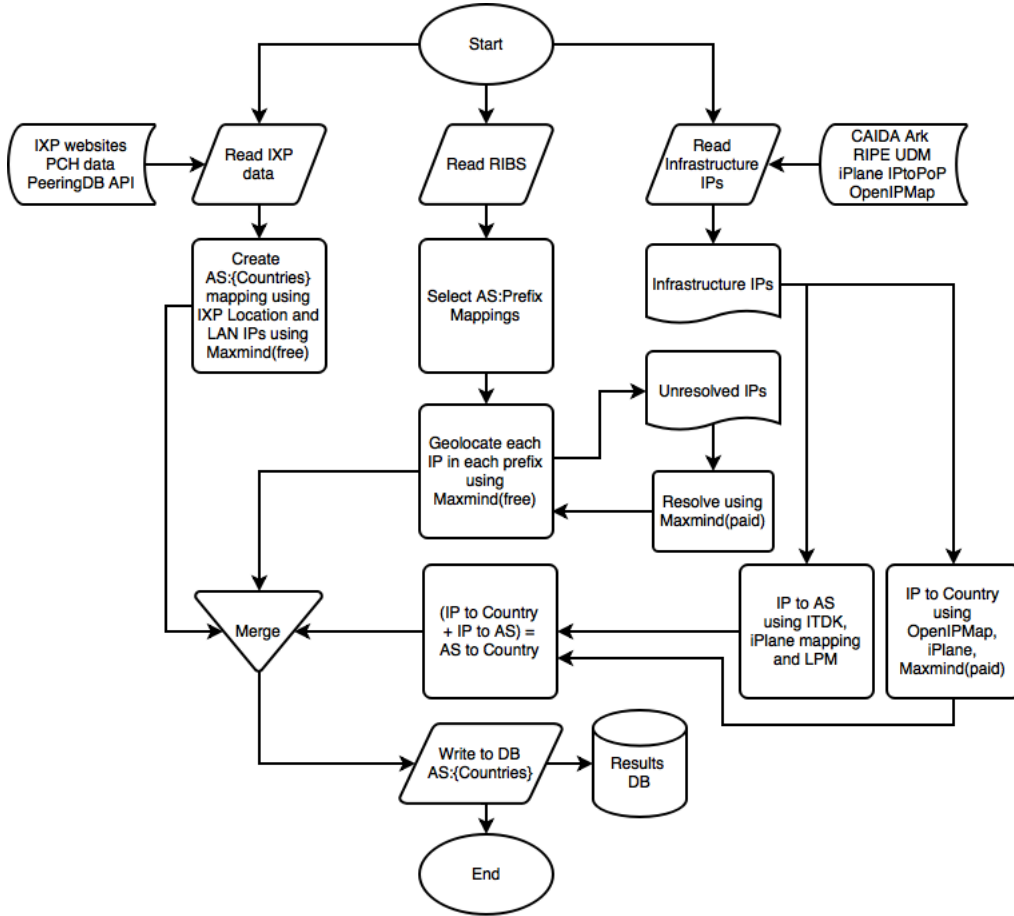
### Comparison with CAIDA’s AS Rank:

Although our end goal is to detect detours, these geolocation results provide interesting insights. To understand more about which ASes geolocate to more than one country we use CAIDA’s AS Rank [2]. This dataset gives higher ranks to ASes that have large customer cones. Intuitively, ASes with higher rank should resolve to many countries due to their wider presence. Table 2 shows ASes with their CAIDA AS rank and corresponding number of countries the AS geolocated to for top 3 and bottom 3 in the first 1000 ranked ASes. As expected, we see that ASes which have large presence with many customers across the world geolocate to large number of countries and low rank ASes with smaller customer cones geolocate to fewer countries.

## 5. DETOUR DETECTION

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 de-



**Figure 1:** Flowchart: AS-to-Country mapping creation

tour in country ‘A’ and diverts the path to foreign country ‘B’.

- **Detour Origin Country:** The country where we approximate location of Detour Origin AS, country ‘A’.
- **Detour Destination AS:** The AS in foreign country ‘B’.
- **Detour Return AS:** The AS where detour returns back in country ‘A’.



**Figure 3:** Example showing direction of BGP announcement and direction of observed detour

Figure 3 illustrates detours.  $AS_0$  announces prefix  $a.b.c.0/24$  to  $AS_1$ ,  $AS_2$  and  $AS_3$ .  $AS_1$  geolocates to JP

whereas  $AS_3$ ,  $AS_2$  and  $AS_0$  are in the US. In this case, data traversing from  $AS_3$  to  $AS_0$  will contain a detour from  $AS_2$  (Detour Origin) to  $AS_1$  (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  $AS_1\{US\}$ - $AS_2\{IN\}$ - $AS_3\{CN\}$ - $AS_4\{IN\}$ - $AS_5\{US\}$ , 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  $AS_1\{US\}$ - $AS_2\{US, BR\}$ - $AS_3\{CN\}$ - $AS_4\{US\}$ . We resolve the uncertainty of the detour origin by assuming that it starts in  $AS_2$ , since there is a likely path to  $AS_2$  from  $AS_1$  through the US and  $AS_2$  starts the detour from US, not BR. We do not characterize *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 only focus on paths that contain *definite* detours, such as  $\{US\}$ - $\{IN\}$ - $\{US\}$  or  $\{US\}$ - $\{IN, CN\}$ - $\{US\}$ . Again, we re-emphasize that in this work we only look at paths that confidently start and end in the same country; paths like  $\{US, BR\}$ - $\{IN\}$ - $\{US\}$  or  $\{US\}$ - $\{IN\}$ - $\{BR\}$  are not considered. We discard paths where we see an AS whose geolocation is unknown and a detour is not certain. For example, paths like  $AS_1\{US\}$ - $AS_2\{?$ -

**Table 2:** Comparison of CAIDA’s AS Rank with number of countries in AS Geolocation

AS Rank	ASN	Customer Cone Size	AS Name	#Countries
1	3356	24,553	Level 3 Communications	63
2	174	17,891	Cogent Communications	58
3	3257	16,963	Tinet Spa	34
998	25394	18	MK Netzdienste GmbH Co. KG	2
999	6724	18	Strato AG	4
1000	52925	18	Ascenty DataCenters Locacao e Servicos LTDA	2

AS3{US} are discarded. However, if we see the detour occurring before the AS that could not be geolocated we do count it as a valid detour i.e., in AS1{US}-AS2{BR}-AS3{US}-AS4{}-AS5{US}, AS4 does not have geolocation information but the US-BR-US detour occurred earlier. We treat this path as definite detour. We note that in addition to geolocation accuracy there is also some ambiguity about exact country boundaries. Some territories and relationships are currently disputed between multiple authorities and no worldwide consensus exists. For example, Hong-Kong and the People’s Republic of China could be considered one or two entities. Hong-Kong is affiliated with China but it is a charter city and has its own independent constitution and judiciary system. For our analysis, we left the resolution of boundaries and countries to the MaxMind database. With this particular example, Hong-Kong and China are treated as two separate entities. MaxMind follows ISO 3166 country codes. In some cases the geolocation from MaxMind is ambiguous: ‘A1: Anonymous Proxy’, ‘A2: Satellite Provider’, ‘O1: Other Country’, ‘EU: Europe’, ‘AP: Asia/Pacific’. We discard detours caused by these ambiguous codes, such as {DE}-{EU}-{DE}.

#### Filtering peered AS paths:

It is possible that the detour origin and the detour return ASes have a peering relationship and in 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 [3]. 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. We eliminate such paths from our analysis and revisit this issue in the next section summarizing the peering relationships in Table 4.

#### Multi-Origin Prefixes:

Some prefixes are announced by more than one ASes. We do not eliminate such cases. So, if a prefix  $a.b.c.0/24$

is seen in RIBs of 2 peers with AS paths ‘X Y Z’ and ‘P Q R’ then we treat each path as independent and detect detour if it fits above mentioned criteria of starting and ending in the same country. In our geolocation dataset we observed 7,579 prefixes of multi-origin (7,247 originated from 2 ASes). Out of these 6,104 suffered a detour. Motivation to not eliminate these prefixes is as follows: Network operators of such prefixes might want to re-evaluate their decisions especially if the ASes originating the prefix are in different countries. This might be a cause of high latency.


## 6. DETOUR VALIDATION

In this section we validate detours in near real time using *traceroutes* from RIPE Atlas probes. Our validation comprises of four steps:

1. Run **Netra** with live BGP feeds from 416 peers to detect detours.
2. When a detour is detected, run corresponding traceroutes (from same country and same AS) using RIPE Atlas.
3. Check if the traceroute and detour see similar AS path.
4. Validate using traceroute IP hops and RTT.

### 6.1 Data-plane Measurements

We ran **Netra** from May 2<sup>nd</sup> 2016 noon to midnight (using BGP feeds from 416 peers). When a detour was detected in control plane we selected RIPE Atlas probes in the same country and same AS which we detected detour from and ran traceroute (ICMP Paris-traceroute [18]) to IP addresses in the detoured prefix. The methodology to run data plane measurements is shown in Figure 4. There are a few cases where more than two Atlas probes are present in selected AS; in this case we selected 2 probes that are geographically farthest from each other. By doing this we aimed to account for cases where routes seen from geographically distant vantage points within the same AS are different. To select target IPs from detoured prefix, we break the prefix into its constituent /24s and randomly select an IP from each /24. For example, in a /23 prefix we select 2 IPs belonging to different /24s. By doing this we



**Figure 4:** Data plane measurements: Example showing selection of RIPE Atlas probes and target IPs

account for cases where a large prefix, even though in the same country, has different connectivity via different upstream provider. During this live run we detected 6,175 detours. Out of these 5,787 were unique detours ( $\{\text{peer}, \text{prefix}, \text{aspath}\}$  tuple).

## 6.2 Selecting Congruent Paths

Only 72 peers saw the 6,175 detours and the 72 peers belong to only 63 ASes. From these 63 ASes we then select ASes that also have active RIPE Atlas probes; there were only 10 ASes that both saw a detour and host a RIPE Atlas probe. 169 detours were seen from these 10 ASes corresponding to 6 countries: {Brazil, Italy, Norway, Russia, United States, South Africa}. From the 169 traceroutes we initiated to detoured prefixes, we discard 6 traceroutes where less than 3 hops responded since drawing detour conclusion from these is not possible. Finally, we are left with 163 traceroutes that can be used for validation. We acknowledge that 163 is not a very large number for validation purposes. However, running *Netra* for more hours does not necessarily increase the number of usable traceroutes for validation by a lot, we are limited by the number of ASes that have RIPE Atlas probes which also see a detour and detour-origin and detour-destination have no peering.

In total we detected 85 prefixes (corresponding to the 163 traceroutes) that suffered a detour that was visible from an AS which has RIPE Atlas probes. Note that some detoured prefixes were larger than /24, so we traceroute multiple IPs within it as explained in Section 6.1. The validation methodology is stated in Algorithm 1. As previous work [22] has pointed out, we found many cases where AS path seen in control plane and AS path seen in data plane do not match. However, these paths can still show detour if the detour origin AS and the detour destination AS are still present in the traceroute observed AS path. We call such AS paths *congruent*. More specifically, we consider the detoured AS path congruent only if detour origin AS and detour return AS both are present in the traceroute-observed AS path in the same order (detour origin first). For example, if an AS path ‘A B C D E’ in control plane changed to ‘A X B C E’ in data plane where ‘B’ was detour origin and ‘C’ was detour destination, we consider

it as a congruent path. To resolve traceroute path to AS path we used CAIDA ITDK and iPlane IP to AS mappings and in cases where no match was found we use longest prefix match on the global routing table for the hop IP. Then we map the longest prefix match to the AS that originated it. Out of all the IPs we saw in 163 traceroutes, only 44 could be mapped to an AS using the IP to AS datasets. All other IPs were mapped using longest prefix match.

We observed 113 congruent AS paths. This includes 3 cases, insertions, deletions and mix of both. We borrow nomenclature of these paths from [22]. We saw 73 deletions, 29 insertions, 4 mix of insertion and deletions. The remaining 7 AS paths were exact matches. Note that these insertions and deletions occurred only for ASes that were not involved in the detour.

---

### Algorithm 1 *Netra* Validation

---

```

1: procedure VALIDATEASPATH
2:   aspath  $\leftarrow$  AS Path from Traceroute
3:   doas  $\leftarrow$  Detour Origin AS from Netra
4:   ddas  $\leftarrow$  Detour Destination AS from Netra
5:   if doas, ddas in aspath then
6:     if doas before ddas in aspath then
7:       Return True
1: procedure VALIDATEIPHOPS
2:   ipHops  $\leftarrow$  IP hops from Traceroute
3:   ipHopCountries  $\leftarrow$  MaxMind-paid
4:   if ipHopCountries show detour then
5:     detourDestTR  $\leftarrow$  Dest. from traceroute
6:     detourDestNetra  $\leftarrow$  Dest. from Netra
7:     if detourDestTR in detourDestNetra then
8:       Return True
1: procedure VALIDATERTTs
2:   hopRTTs  $\leftarrow$  RTTs from Traceroute
3:   if hopRTTs show magnitudeJump then
4:     Return True
1: procedure MAIN
2:   loop: Each Detected Detour
3:   if validateASPath then
4:     validateIPHops
5:     validateRTTs

```

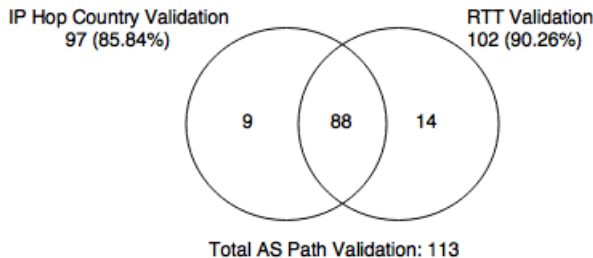
---



### 6.3 Validation

Now we validate detours detected by our methodology by comparing it with detours seen in data plane. For the 113 congruent AS paths, we evaluate if a data plane detour was seen. We chose to perform two tests. First, we resolve IPs observed in the hops of traceroute to country level geolocation using Maxmind-paid. We detect data plane detour if a path traversed foreign country and returned. We make sure that country visited (detour destination country) in data plane is present in the set of destination countries expected for this particular detour by **Netra**. We do this filtering to avoid false positives like: **Netra** detected detour {US}–{GB,DE}–{US} and traceroute detected detour {US}–{IT}–{US}. Although still a detour, since it was not accurately captured we count it as a miss. However, no such case was found. Second, we validate using RTT measurements. We detect RTT based detour if a hop in the traceroute showed increase in RTT by an order of magnitude (at least 10 times increase). The results of this analysis are shown in Figure 5. We observed accuracy of about 85% (97 out of 113) in country-wise method and 90% (102 out of 113) by RTT measurements. The overlap between these two different tests was also large. 88 detours were detected in both (77.8%).

We investigate further the 9 detours that were seen in country-wise method but not in RTT. These detours covered small geographic area; 4 from Italy to France, 2 Norway to Sweden, 2 from Brazil to US and 1 from Russia to Sweden. RTTs between these countries have been previously reported to be low. Next we investigate 14 cases which were captured in RTT measurements but not in country-wise method. All of these do cross international boundaries. For 12 of these cases, due to large number of traceroute hops (especially towards the end of the traceroute) not responding we don’t see the route returning to the origin country, hence not detected by country-wise method. We attribute remaining 2 cases as false positives due to inaccurate AS geolocation.



**Figure 5:** Validation Results: Live traceroutes using RIPE Atlas

In Figure 6 we provide a visualization of the most common detour we observed from Russia. Only visual-

ization is done using OpenIPMap.

### Validation Discussion:


We show that large percentage of detours seen in control plane are accurately reflected in data plane as well. The main challenge is AS paths in both data plane and control plane don’t agree in about 30% cases. We note that this could be an artifact of Atlas probes connected differently than the peers which provide BGP feeds. It is, however, possible to learn common AS insertions and deletions over a period of time and evolve detection capabilities.

## 7. RESULTS

In this section we quantify detours detected in January 2016. First, in Section 7.1 we present an overview of all the detours detected in our dataset. In Section 7.2 we define metrics and classify detours based on their stability and availability. In Section 7.3 we focus on transient detours.

### 7.1 Aggregate Results

We begin by characterizing aggregate results, namely all detours seen by all peers; in other words, we count an incident every time an AS 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. As expected, we observe that detours are not generally common. Also, not all peers see a detour. Only 79 peers, out of 416, saw one or more detours. Table 3 details the number of detours seen. We analyzed about 14 billion RIB entries and about 544K entries showed a detour; out of theses only 18.9K were unique (most detours re-appear during the month). Figure 7 shows the number of detours for each day in January 2016. On an average we find about 17.5K detoured entries per day.

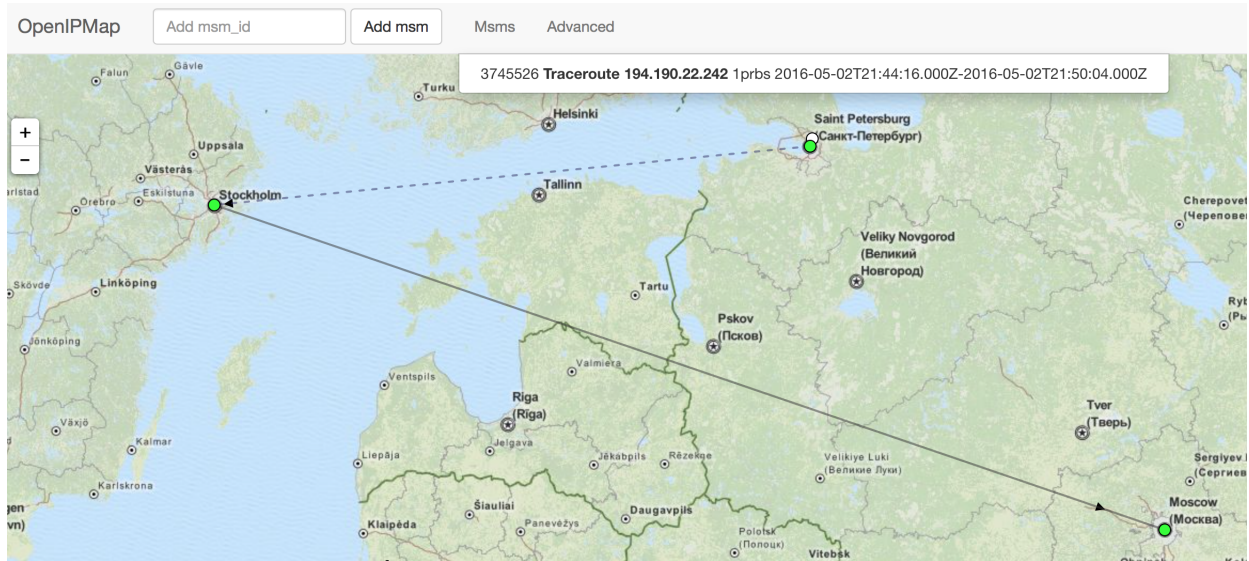


**Figure 7:** Total number of definite detours per day in January 2016

**Table 3:** Aggregate number of detours detected

#Total RIB Entries	#Detoured Entries	#Unique Detours
14,366,653,046	544,484	18,995





**Figure 6:** Top Detour on May 2<sup>nd</sup> 2016: Detected using Netra, visualization using OpenIPMap. Dotted arrow represents multiple hops and solid arrow represents direct hop.

**Table 4:** Routes that may have peering relations

#Total Detours without filtering peered paths	#Detours with possible peering	%
659,569	115,085	17.4%

Next we examine the visibility of detours, where we observe an uneven distribution among ASes. Just 9 ASes originate more than 50% of the detours. Similarly, some prefixes experience detours more than others. 132 prefixes experienced more than 50% of the total detours. Looking at the average length of a detour, we see that a detour visits 1 to 2 foreign ASes before returning to its origin country.

### Impact of Peering:

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 5 using CAIDA AS relationship dataset. If such a link exists, it is possible that traffic traverses that link instead of the detour. Table 4 shows the number of detours between ASes that also have peering relations compared to total number of detours without filtering peered paths. We find that 17.4% of the detours are avoided due to peering relations. We do not count these as detours in our analysis.

### Top Detour Origins and Prefixes:

To understand more about the nature of these detours, we focus on the origin and destination ASes. In Table 5 we show the common detour origins and country where the AS was approximated to origin the detour from. Next is the percentage of detours out of the total that

started from given origin. Following the percentage, is the most frequent destination that was visited from the origin, and lastly is the percentage of detours that went to most common destination from the said origin. We observe that most commonly these were access provider ASes. Similarly, in Table 6 we show top impacted pre-fxes.

### Country-wise analysis:

To provide an understanding on number of detours per peer in each country we normalize the data by dividing the number of detours by number of peers in the country. The reason to normalize data is simple, RouteViews and RIPE RIS peers are not evenly distributed among different countries. Therefore it is possible that more detours are seen in countries that have more peers due to more visibility. An average number of detours per peer per country provides better insight. Out of 30 countries, only 12 countries observed a detour. Figure 8 shows average number of detours per country. Russia showed most number of average detours. Understanding the total number of detours in different countries is important but it does not reflect if detours seen in different countries have different characteristics. In the next section we focus on characterizing these detours.

## 7.2 Characterizing Detours

To characterize detours we define two metrics:

### 1. Detour Dynamics


- Flap Rate:** Measure of *stability* of a detour; how many times a detour disappeared and reappeared.
- Duty Cycle:** Measure of *uptime* of a detour throughout the month measurement period.

**Table 5:** Top Detour Origin ASNs for all detoured paths

Top Detour Origin AS	Total %	Frequent Detour Destination AS	% to frequent destination
3356 (Level 3 Communications,BR)	8.39%	32787 (Prolexic-Technologies DDoS Mitigation Network)	30.99%
12956 (Telefonica International Wholesale Services,BR)	5.74%	262182 Media Networks Latin America	46.33%
6939 (Hurricane Electric,US)	4.99%	45932 (Net Sys International Limited)	15.9%

**Table 6:** Top Detoured prefixes and corresponding percentages

Prefix Affected	Total %	Frequent Detour Destination AS	% to frequent destination
199.253.181.0/24 (Internet Systems Consortium,US)	0.51%	766 (Entidad Publica Empresarial Red)	100%
167.220.28.0/23 (Microsoft,US)	0.51%	6584 (Microsoft Corp)	100%
199.6.5.0/24 (Internet Systems Consortium,US)	0.51%	766 (Entidad Publica Empresarial Red)	77.11%

**Figure 8:** Average number of detours per country

- Persistence:** Total number of continuous hours a prefix was seen detoured.

Before using the above metrics to characterize the detours, we perform data pruning to avoid skewing of data towards ASes that have more peers that provide BGP feeds to RouteViews and RIPE RIS. Also, ASes with multiple peers and similar views can contribute duplicate detours to our dataset. 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 36 (out of 79) peers. We now continue our characterization of detours by looking at **detour dynamics**. Specifically we focus on flap rate and duty cycle, defined as follows:

$$FlapRate = \frac{TotalTransitions}{TotalTime} \times 100$$


$$DutyCycle = \frac{TotalUptime}{TotalTime} \times 100$$

These metrics provide insights into the life cycle of detours by measuring route uptime and stability. BGP


route flapping is a known problem and has been studied in [23] by looking at BGP updates and RFC 2439 provides methods to dampen these. However, in context of this paper duty cycle and flap rate are calculated from the RIBs. We extract detours from the RIBs and evaluate when they disappear and reappear.

To understand if country where detours occur plays a role in detour dynamics, next we drill into country specific detours. Figure 9 shows 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 dataset; they account for 93% of detours. We see a triangular pattern with some outliers. Large number of detours show high duty cycle and low flap rate. We divide each figure into 4 quadrants based on average flap rate and average duty cycle of all detours. We name quadrants anti-clockwise starting from top right. US detoured paths appear more stable (lower flap rate and higher duty cycle) in  $II^{nd}$  quadrant. On the other hand, Russian and Brazilian detoured paths fall mostly in the  $I^{st}$ ,  $III^{rd}$  and  $IV^{th}$  quadrant. Russian detours in general showed lower duty cycle than US and Brazil. We also present a similar scatter plot for all the non US, BR and RU detours in Figure 10. In this case we observed detours mostly in extreme ends on  $II^{nd}$  and  $III^{rd}$  quadrant indicating two categories of detours, either long lasting or very rare events. A network operator can use information like this and decide which quadrant detours are more interesting to focus on. While all of detours may need attention, we believe detours with low duty cycle and low flap rate may need immediate attention. We talk more about this in Section 7.3.

Next, we examine the **persistence** of detours. Figure 11 shows the number of consecutive days a detour was visible by any peer. Note that persistence is measured in number of consecutive hours hence captures different characteristics than duty cycle which measures




**Figure 9:** Flap Rate vs DC for US, RU and BR prefixes




**Figure 10:** Flap Rate vs DC for Non US, RU and BR prefixes

uptime throughout the dataset. We see a U-shaped pattern in Figure 11, 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 12 by plotting CDF of duration in hours. We see that most detours are short-lived, with about 92% lasting less than 72 hours, defined as *transient* detours. Finally, we 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.



**Figure 11:** Persistence of definite detoured paths as seen by all peers



**Figure 12:** Distribution of detour duration



**Figure 13:** Average number of transient detours per country

### 7.3 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 13. We detected transient detours in only 8 countries where Russia topped the list. In comparison to Figure 8 Italy and India showed more average number of transient detours than US. Figures 14 and 15 show a distribution of ASes that initiate detours and prefixes affected by detours. We observe that 4 ASes originate 50% of the transient detours and only 30 prefixes account for 50% of the transient detours. Similar to Table 5, shown in Table 7 are the most common transient detour origins and Table 8 shows top impacted prefixes by transient detours. AS9002, RETN-AS, started the most number of transient detours in our dataset. We note that in *ASWatch* [16] authors gathered ground-truth data from security blogs which enlisted AS9002 as a malicious AS. Another previously known malicious AS that appeared in our findings was AS49934 as a detour destination for 7 Russian prefixes. AS49934 is currently unassigned. It was assigned in Ukraine between 2009-10-14 and 2016-01-03 and was known to announce bogus prefixes and host bots.



Finally, we look at *flash* detours. These are detours that appeared only once and were observed in only one

**Table 7: Top Transient Detour Origin ASNs**

Transient Detour Origin AS	Total %	Frequent Detour Destination AS	% to frequent destination
9002 (RETN-AS RETN Limited, RU)	22.64%	2914 (NTT America)	99.07%
6939 (Hurricane Electric, IT)	10.94%	8551 (Bezeq International)	100%
1299 (TELIANET, IT)	10.87%	8708 (RCS-RDS)	100%

**Table 8: Prefixes affected the most by transient detoured BGP paths**

Prefix Affected	Total %	Frequent Detour Destination AS	% to frequent destination
178.79.218.0/23 (Limelight Networks, Inc, IT)	5.5%	8551 (Bezeq International, IL)	100%
185.19.164.0/22 (Digi Italy S.R.L, IT)	5.5%	8708 (RCS-RDS, RO)	100%
46.21.30.0/24 (Tekka Digital, IT)	5.5%	8758 (Iway, CH)	67.08%

**Figure 14:** Distribution of ASes that originated a transient detour. The top 4 Detour Origin ASes account for 50% of all transient detours**Figure 15:** Distribution of prefixes that experienced a transient detour. About 30 prefixes account for 50% of all transient detours

RIB of a peer. Flash detours account for 26% of the transient detours, 328 prefixes (6% of all prefixes that suffered detour) experienced at least one flash detour.

Owners of the prefix which suffered flash detours might be interested to know such findings. While 328 prefixes suffered flash detours in our dataset, due to space limitation we point out a few interesting ones in Table 9.

The list in Table 9 raises serious concerns. Data from government agencies, banks, insurance companies can easily be subject to wiretapping once it leaves national boundaries. Based on our control-plane only data, it is not possible to verify if these institutions were attacked or not. Nevertheless, we believe our findings will mo-

tivate network operators to look more closely into why their prefix detoured and if they intended it to happen.

## 8. DISCUSSION

In this paper we present a first attempt to characterize detours in the Internet. We sampled BGP routing tables from 416 peers around the world over the entire month of January 2016 to investigate international detours. We see about 18.9K distinct entries in RIBs that show a detour. More than 90% 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. Our work is the first to present different types of detours, namely, persistent and transient. We also present novel insights on their characteristics such as detour dynamics in different countries, top impacted prefixes and detour origins.

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 Washington state government were de-

**Table 9:** Some prefixes affected by flash detours

Prefix Affected	Owner	Detour Destination
170.61.199.0/24	Mellon Bank, US	28513 (Uninet, MX)
192.230.0.0/20	Washington State Department of Information Services, US	7660 (Asia Pacific Advanced Network, JP)
212.11.152.0/21	Moscow Mayor Office, RU	2603 (NORÐNet, NO)
208.79.7.0/24	Security Equipment Inc, US	53185 (William Roberto Zago, BR)
161.151.72.0/21	The Prudential Insurance Company of America, US	2510 (Infoweb Fujitsu, JP)

toured through Japan; and for some detours from Russia malicious AS49934 appeared as detour destination.

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

#### Dataset Contributions:

We make the geolocation and detours detection data available to the community via a public RESTful API interface. The motivation to do so is as follows. 1) Network operators can easily query our database and check if their prefix suffered a detour. 2) Internet measurement researchers can use this information to study various BGP anomalies such as route leaks, detecting malicious ASes, etc. Our results on AS and prefix geolocation are available at <http://geoinfo.bgpmn.io> and detours results can be accessed at <http://detours.bgpmn.io>.

## 9. CONCLUSIONS AND FUTURE WORK

There is an increasing need, fueled by new national regulations in Europe and Australia, for ISPs to ensure that personal information belonging to their users does not leave the country. It is unclear whether such regulations cover data in transit as well as storage, but data can certainly be sniffed while in transit, violating the original intent. Such regulations may place a substantial burden on ISPs to prove that such data remains within a country for its entire lifetime, even when it moves. It is still far from clear what the implications are on ISP operations. Currently we do not have the tools to monitor data in transit and state with confidence that data has not left a country, even briefly.

Our work does not solve this problem. Rather, it lays the ground for an important conversation about the challenges new regulatory frameworks will pose to researchers, industry and network operators. Our work investigates only a small part of the problem, namely finding the subset of paths where we can detect international detours with some confidence. Our work provides some answers, but also brings attention to the problem and will hopefully stimulate more work in this new direction. The gauntlet was thrown and we expect a lot more research in this area.

Within its scope, we believe our work was executed carefully by taking into account measurements from

both control and data planes. We show that for the cases we were able to study there is agreement between the two planes. This is a significant result. Equally significant, our work has also illuminated the difficulties in expanding the scope within the existing measurement infrastructures. One of the main difficulties we encountered for example, is finding measurement points with both control (BGP peers) and data (RIPE probes) monitors to correlate results. This problem cannot be easily solved, it would take substantial effort to scale the existing infrastructures by an order of magnitude or more. Another important obstacle is lack of knowledge about peering relationships between ASes. This is also a hard problem to solve, since such relationships are not readily disclosed. It is interesting, however, to contemplate the issue if regulatory requirements require such disclosures.

Based on our results, we believe that it will be hard to solve this problem without substantial data plane monitor deployment to corroborate control plane measurements. ISPs and IXPs may be required to install sophisticated data plane probe infrastructures and geolocation databases may have to become far more accurate for infrastructure IP addresses in order to detect international detours with some certainty. Control plane monitoring is still very important as it provides efficient global monitoring and can immediately flag potential anomalies where data plane monitoring should be directed. Our work shows that it is effective and should be expanded.

In the future we plan to continue to build a system that detects international detours in real time. It is very apparent that we need to include both control and data plane measurements and study algorithms that take input from both. Our first goal is to provide ISPs with a tool to alert when a detour has taken place, followed by information about it (origin and destination AS, duration, source and amount of data in the ISP that followed the detour). We also plan to study emerging regulatory requirements and provide feedback about the challenges they pose.

#### Acknowledgement

We would like to thank Randy Bush (IIJ) for helpful discussion on AS geolocation; Emile Aben (RIPE NCC) for helpful insights on the paper and using OpenIPMap;



Alessandro Improta and Luca Sani (Isolario Project, Italy) for providing code to crawl IXP websites; Roya Ensafi (Princeton) for help in geolocating infrastructure IP addresses.

## 10. REFERENCES

- [1] Caida ark dataset.  
<http://www.caida.org/projects/ark/>.
- [2] Caida as ranks. <http://as-rank.caida.org>.
- [3] Caida as relationships. <http://www.caida.org/data/as-relationships/>.
- [4] The caida internet topology data kit.  
<http://www.caida.org/data/internet-topology-data-kit/>.
- [5] iplane datasets, university of washington. <http://iplane.cs.washington.edu/data/data.html>.
- [6] Packet clearing house ixp datasets.  
[https://prefix.pch.net/applications/ixpdir/menu\\_download.php](https://prefix.pch.net/applications/ixpdir/menu_download.php).
- [7] Peering db 2.0 api. [https://prefix.pch.net/applications/ixpdir/menu\\_download.php](https://prefix.pch.net/applications/ixpdir/menu_download.php).
- [8] Ripe ncc openipmap. <https://github.com/RIPE-Atlas-Community/openipmap>.
- [9] Ruwaifa Anwar, Haseeb Niaz, David Choffnes, Ítalo Cunha, Phillipa Gill, and Ethan Katz-Bassett. Investigating interdomain routing policies in the wild. In *Proceedings of the 2015 ACM Conference on Internet Measurement Conference*, IMC '15, pages 71–77, New York, NY, USA, 2015. ACM.
- [10] Jim Cowie. The new threat: Targeted internet traffic misdirection, Nov 2013.  
<http://www.renesys.com/2013/11/mitm-internet-hijacking/>.
- [11] 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.
- [12] 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.
- [13] Bradley Huffaker, Amogh Dhamdhere, Marina Fomenkov, and Kc Claffy. Toward topology dualism: Improving the accuracy of as annotations for routers. In *Proceedings of the 11th International Conference on Passive and Active Measurement*, PAM'10, pages 101–110, Berlin, Heidelberg, 2010. Springer-Verlag.
- [14] Josh Karlin, Stephanie Forrest, and Jennifer Rexford. Nation-state routing: Censorship, wiretapping, and BGP. *CoRR*, abs/0903.3218, 2009.
- [15] Ethan Katz-Bassett, John P. John, Arvind Krishnamurthy, David Wetherall, Thomas Anderson, and Yatin Chawathe. Towards ip geolocation using delay and topology measurements. In *Proceedings of the 6th ACM SIGCOMM Conference on Internet Measurement*, IMC '06, pages 71–84, New York, NY, USA, 2006. ACM.
- [16] 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.
- [17] MaxMind LLC. Maxmind geoip country database. <http://dev.maxmind.com/geoip/legacy/geolite/>.
- [18] Matthew Luckie, Young Hyun, and Bradley Huffaker. Traceroute probe method and forward ip path inference. In *Proceedings of the 8th ACM SIGCOMM Conference on Internet Measurement*, IMC '08, pages 311–324, New York, NY, USA, 2008. ACM.
- [19] Jonathan A. Obar and Andrew Clement. Internet Surveillance and Boomerang Routing: A Call for Canadian Network Sovereignty. *SSRN Electronic Journal*, 2013.
- [20] Yuval Shavitt and Noa Zilberman. A study of geolocation databases. *CoRR*, abs/1005.5674, 2010.
- [21] 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.
- [22] kc claffy Young Hyun, Andre Broido. Traceroute and bgp as path incongruities.
- [23] Beichuan Zhang, Daniel Massey, and Lixia Zhang. Bgp dynamics during route flap damping. Technical report.
- [24] Yu Zhang, Ricardo Oliveira, Hongli Zhang, and Lixia Zhang. Quantifying the pitfalls of traceroute in as connectivity inference. In *Proceedings of the 11th International Conference on Passive and Active Measurement*, PAM'10, pages 91–100, Berlin, Heidelberg, 2010. Springer-Verlag.