Stony Brook University

Measuring Routing Policies on the Internet

A Thesis presented

by

Mohammad Haseeb Niaz

to

The Graduate School

in Partial Fulfillment of the

Requirements

for the Degree of

Master of Science

in

Computer Science

Stony Brook University

May 2015

Stony Brook University

The Graduate School

Mohammad Haseeb Niaz

We, the thesis committee for the above candidate for the

Master of Science degree, hereby recommend

acceptance of this thesis

Phillipa Gill - Thesis Advisor
Assistant Professor, Computer Science Department

Aruna Balasubramanian - Thesis Committee Member
Assistant Professor, Computer Science Department

Samir R. Das - Thesis Committee Member
Professor, Computer Science Department

This thesis is accepted by the Graduate School

Charles Taber
Dean of the Graduate School

Measuring Routing Policies on the Internet

by

Mohammad Haseeb Niaz

Submitted to the Department of Computer Science in partial fulfillment of the requirements for the degree of
Master of Science in Computer Science

Stony Brook University

May 2015

Abstract

Models of Internet routing are critical for studies of Internet security, reliability and evolution, which often rely on simulations of the Internet's routing system. Accurate models are difficult to build and suffer from a dearth of ground truth data, as ISPs often treat their connectivity and routing policies as trade secrets. In this environment, researchers rely on a number of simplifying assumptions and models proposed over a decade ago, which are widely criticized for their inability to capture routing policies employed in practice. This thesis makes the following two contributions:

- **Investigating Interdomain Routing Policies.** First we put Internet topologies and models under the microscope to understand where they fail to capture real routing behavior. We measure data plane paths from thousands of vantage points, located in eyeball networks around the globe, and find that between 14-35% of routing decisions are not explained by existing models. We then investigate these cases, and identify root causes such as selective prefix announcement, misclassification of undersea cables, and geographic constraints. Our work highlights the need for models that address such cases, and motivates the need for further investigation of evolving Internet connectivity.

- **Study of attacks against decoy router deployments.** Second, we use our understanding of Internet Routing and tools developed, to study the effect of decoy router deployments over the Internet and come up with a way to figure out an optimal decoy router deployment in terms of the number of deployments required, in order to target a censoring AS. We introduce the notion of Helper ASes, which are autonomous systems that do not host a decoy router in their network but help out by making specially crafted routing announcements to force traffic from censoring countries through an AS hosting a decoy router. By introducing the Helper ASes into the decoy routing ecosystem, we are able to greatly reduce the number of deployments required for the system to be effective.

Dedication

I dedicate this thesis to my parents, Shahid and Farah. I hope that this achievement will complete the dream that you had for me all those years ago when you chose to give me the best education you could.

Contents

List of Figures

List of Tables

Acknowledgements

I would never have been able to finish my thesis without the guidance of my advisor, help from friends, and support from my family and wife.

I would like to thank my advisor Phillipa Gill for her continued support and guidance. This thesis benefited greatly from the ample guidance I received from Italo Cunha, David Choffnes and Ethan Katz-Bassett.

My completion of this thesis could not have been accomplished without the support and contributions of my friends, Ruwaifa, Laraib, Ankit, Rajesh and Udit. Thank you guys for providing the much needed time away from work.

Finally, to my caring, loving, and supportive wife, Maham: my deepest gratitude. Your encouragement when the times got rough are much appreciated and duly noted. I would not have been able to complete this thesis without your continuous love and encouragement. My heartfelt thanks.

1 Introduction

Research on existing and new protocols on the Internet is challenging because key aspects of the network topology is hidden from public view by interdomain routing protocols, and deploying new protocols at Internet scale requires convincing large numbers of autonomous networks to participate. As a result, networking researchers rely on assumptions, models, and simulations to evaluate new protocols [17, 34], network reliability [27, 50], and security [7, 20, 32].

Our existing models of interdomain routing [15], however, have important limitations. They are built and validated on the same incomplete topology datasets, typically routes observed via route monitors such as RouteViews and RIS. These vantage points expose a large fraction of paths from global research & education networks (GREN) and core networks, but they are incomplete in two keys ways. First, they expose few paths to and from eyeball and content networks. Second, they do not expose less preferred paths that would be used if the most preferred next-hop AS were not available. This is because active paths are established by the Border Gateway Protocol, an information-hiding protocol that exposes only the most preferred paths toward a prefix that an AS is willing to make available to its neighbor. As a result, they do not capture partial peering, more complex routing policies based on traffic engineering, or load balancing and the rich peering mesh which exists near the edge of the network [46]. The ability to observe the relative preferences of these alternative paths is crucial for understanding the

complex routing policies employed in practice.

While limitations of our existing models are well known [35, 38, 46]–and are even being addressed in recent work [19]–we lack a solid understanding of how much these limitations impact our ability to accurately model the interdomain routing system. Recent work has attempted to address this issue by observing destination-based routing violations in control plane data [36] and by surveying a population of network operators about their policies [16], however, these approaches are limited in terms of scale and their ability to observe behavior at the network edge. In this thesis, we develop a better understanding of the inter domain routing policies and then leverage this knowledge to evaluate the potential of interdomain routing to aid censorship circumvention via decoy routing.

1.1 Revisiting generally held assumptions and models of Internet routing

Prior work has largely focused on verifying the inferred relationship between pairs of ASes [19, 33]. Few studies have considered path-level validation and those that do are often hindered by a lack of ground truth to compare against. Mazloum *et al.* consider the converse problem and identify contradictions to routing models based on violations of destination-based routing in the control-plane [36]. Madhyastha *et al.* use a set of traceroutes from PlanetLab nodes as ground truth to val-

idate path prediction [35], however, these paths are heavily biased towards routes in academic and research networks.

In our work, we take a systematic approach to understanding how our models of routing policies hold in practice. To accomplish this, we leverage a combination of data-plane measurements covering the network edge ("eyeball networks") and control-plane experiments which allow us to directly measure relative preference of routing options. We create a methodology that takes into account numerous potential causes of violations to our assumptions including sibling ASes, complex AS relationships, prefix-specific routing policies, and the impact of geography. We use this methodology to investigate the prevalence of each of these sources of error in AS-level paths observed via measurements of the data and control planes.

With these measurements, we revisit generally held assumptions and models of Internet routing. Our goal is *not* to measure a complete Internet topology; rather, we seek to improve our understanding of routing decisions made by ASes when routing their traffic. Towards this goal we make the following observations for our measured paths:

- Hybrid and partial transit relationships (*e.g.*, those explored in [19]) contribute a surprisingly small amount to unexpected routing decisions. Details can be found in Section 3.3.1.
- Various ASes do not have the same export policies for all their prefixes and

end up selectively announcing prefixes to their neighbors. We find that per-prefix routing policies explain 10-20% of unexpected routing decisions, where an AS chooses a longer or more expensive path than our model predicts. Details can be found in Section 3.3.4.

- We find that some large content providers like Akamai and Netflix are destinations for a large fraction (21% and 17%, respectively) of unexpected routing decisions.
- Routing decisions vary based on geography. We find paths that traverse multiple continents deviate from our models more, owing to undersea cable ASes which are not accounted for in our models of AS relationships, and a tendency for ASes to prefer non-international paths when endpoints are in the same country. This is outlined in Section 3.5

Our results highlight areas where more investigation would yield the largest payoff in terms of improving our accuracy when modeling AS relationships and routing policies. We also identify key areas, specifically investigating prefix-specific routing policies, where additional vantage points and looking glass servers could improve the fidelity of our AS topology data.

1.2 Interdomain routing and censorship

Recent years have seen the increased us of online information controls, such as censorship, by governments around the globe to limit access to information [8,

12, 39, 41, 42]. Indeed, the situation is only getting worse, with Freedom house reporting 36 of the 65 countries they survey experiencing decreasing levels of Internet freedom between 2013 and 2014 [1].

The rise of these technologies has not gone unnoticed with the research community actively working to circumvent censorship and enable free access to information. Circumvention technologies have tended to focus on disguising the content being transferred, either taking a "look-like-something" approach to mimic popular existing protocols [23, 37] or "look-like-nothing" [13] approach making the traffic appear basically random. In this thesis, we focus on a circumvention technique that leverages properties of the interdomain routing system to circumvent censorship. Decoy routing systems which, instead of disguising content, aim to disguise the destination of the censored connection using in-path [22, 28, 52], and more recently on-path [51] proxies referred to as *decoy routers*. The idea is that a client using decoy routing will connect to any destination IP on the Internet that is reached via a path containing a decoy router. The decoy router then tunnels their traffic towards the desired (*i.e.*, censored) destination.

While decoy routing is effective at hiding the destination of a connection, the identity of networks running decoy routing systems are not assumed to be secret (or may be discovered via probing [49]). This has led researchers to propose "routing around decoys" (RAD) attacks [49] where networks leverage path diversity on the Internet to select routes not containing decoy routers. However, more recently

Houmansadr *et al.* [24] highlight that not all network paths are equal, with some incurring higher monetary costs or performance policies over those chosen when implementing RAD attacks.

In this work, we extend the analysis of Houmansadr *et al.* to incorporate the observation that not all destinations are created equal. Specifically, the popularity of content online is known to follow a highly skewed popularity distribution with the disconnection of the most popular content incurring significant costs in terms of user impact. After revisiting the potential for RAD, we explore the design space for active routing-based defenses against RAD attacks. These defenses leverage novel ideas presented in the networking community in the past few years [27, 30] to allow a destination to steer traffic destined for a specific IP prefix towards paths containing decoy routers.

2 Background on interdomain routing

In this thesis it is relevant to understand how routing is accomplished between networks on the Internet. When we discuss routing on the Internet, we consider routing between ISPs or autonomous systems (ASes). The de facto routing protocol for routing between ASes is the Border Gateway Protocol (BGP). BGP allows ASes to learn routes to reachable prefixes in a distributed manner and allows them to autonomously choose routes according to their own constraints. ASes then announce the chosen paths to a given destination to its neighboring ASes. Since making these announcement is a commitment by the AS to transit traffic, an AS may not announce all of its routes to every neighboring AS. The set of routes that are announced to a neighbor depend greatly on the type of business relation the AS has with its neighbor.

Since much of the information related to inter domain routing is considered trade secrets in the real world, studies of inter domain routing are often criticized for lacking key features of the system in practice [46]. Indeed, in the absence of ground truth, studies are left performing robustness tests to understand the impact of inaccuracies on their results. While prior work has attempted to define more accurate models of the system [14, 38], existing models remain entrenched [17, 20]. In this section, we overview the existing efforts in measuring and modeling the interdomain routing system and highlight why now, more than ever, it is critical to revisit this problem.

2.1 Modeling interdomain routing policies

The now standard model of routing policies was developed by Gao and Rexford [14, 15] based on seminal work by Griffin, Sheppard, and Wilfong [21] and Huston [25, 26]. In this model, ASes connect to each other based on business relationships:

1. Customer-Provider, where the customer pays the provider.
2. Peer-to-Peer, where the ASes engage in settlement-free peering and exchange traffic at no cost.

When choosing network paths, an AS a selects a path to destination AS s based on a combination of cost (based on business relationships) and length of the path. Also, a path that traverses a Provider-Customer edge or a Peer-to-Peer edge can not later traverse a Customer-Provider or Peer-to-Peer edge. This is know as the valley-free policy.

Path selection. AS a selects a path to d from the set of simple paths it learns from its neighbors as follows:

- **Local Preference.** Paths are ranked based on their next hop: customer is chosen over peer which is chosen over provider.
- **Shortest Paths.** Among the paths with the highest local preference, prefer the shortest ones.

- **Tie Break.** If there are multiple such paths, node a breaks ties: if b is the next hop on the path, choose the path where hash, $H(a, b)$ is the lowest.[1]

This standard model of local preference [15] captures the idea that an AS has incentives to prefer routing through a customer (that pays it) over a peer (no money is exchanged) over a provider (that it must pay).

Export Policies. This standard model of export policies captures the idea that an AS will only load its network with transit traffic if it is paid to do so [15]:

- AS b announces a path via AS c to AS a iff at least one of a and c are customers of b.

This model is sometimes augmented with the assumption that ASes only consider the next hop AS on the path when making their routing decisions. This simplifies analysis and makes debugging more tractable [27]. Simulation studies also often restrict path selection to the shortest among all paths satisfying Local Preference to induce unique routing decisions [17, 18].

[1] In practice, this is done using the distance between routers and router IDs. Since we do not incorporate this information in our model we use a randomized tie break which prevents certain ASes from "always winning".

2.2 Efforts to improve the models

Prior efforts to improve the models include efforts to better infer export policies by incorporating key factors like traffic engineering into the model [35]. Muhlbauer *et al.* propose a more realistic model that moves away from the typical AS-level model and propose a finer policy granularity where route selection is done on a per-prefix bases [38]. They argue that the routing policies are far more complex than the commonly held assumptions based on business relationships [15]. They introduce the notion of next hop atoms to capture the preferences an AS may use to choose its best path. A next hop atom is the set of neighboring ASs that an individual AS would use as the next hop for its best routes towards a given set of destinations. This way, when using the model to predict paths, policies are applied based on business relationships first and then the best candidate at each hop is chosen using the next hop atoms when there are multiple candidates.

While there have been other studies that evaluate our assumptions about routing policies [15]. Existing models of routing policies remain entrenched [17, 20]. These and variations of it have been used in many studies (*e.g.*, [7, 17, 20, 29, 50]), it is well known that this model fails to capture many aspects of the interdomain routing system [35, 38, 46]. These aspects include AS relationships that vary based on the geographic region [19] or destination prefix, and traffic engineering via hot-potato routing or load balancing.

Prior work has used traceroute measurements and BGP data to address some

of these issues (*e.g.*, [35, 38]); however, these measurements only offer a glimpse into ASes' routing preferences. Namely, they expose only the set of paths that are in use at the time of measurements. They do not expose less preferred paths that would be used if the most preferred next-hop AS were not available. As a result, they do not capture partial peering, more complex routing policies based on traffic engineering or load balancing and the rich peering mesh which exists near the edge of the network [46].

Despite the importance of studies of the inter domain routing system to networking research, this field stands on uncertain footing with existing data known to be incomplete (*e.g.*, challenges observing peering links [11]). These datasets have poor or no coverage of paths used by edge networks serving residential users [11].

On a smaller scale, network operators were surveyed about their routing policies to better understand how our models correspond to practice [16], but the scale and representativeness of a survey approach makes generalizing these observations infeasible.

3 Interdomain routing policies

We aim to understand the gap between interdomain routing models and empirically observed behavior on the Internet. Our methodology focuses on measurement techniques to gain better visibility into interdomain routing policies. We measure paths between edge networks and content providers to understand routing on paths that carry the bulk of the Internet's traffic [48].

3.1 Data-plane measurements

It is well known that a disproportionately large amount of Internet traffic originates from a few popular content providers [31, 48] towards large populations of end users. However, there is little empirical data about the paths this traffic takes [31]. We target our data plane measurements to cover these paths. Note that it is not our goal to explain routing decisions for the entire Internet. Rather, we focus on the more tractable task of measuring a subset of important Internet paths (those carrying most traffic) from a diverse set of vantage points, and putting those paths under the microscope to understand how and why they differ from predicted paths based on routing models.

Selecting content providers. We consider a list of the top applications from Sandvine [48] and top Web sites from Quantcast [43] and arrive at a list of 34 DNS names representing 14 large content providers.

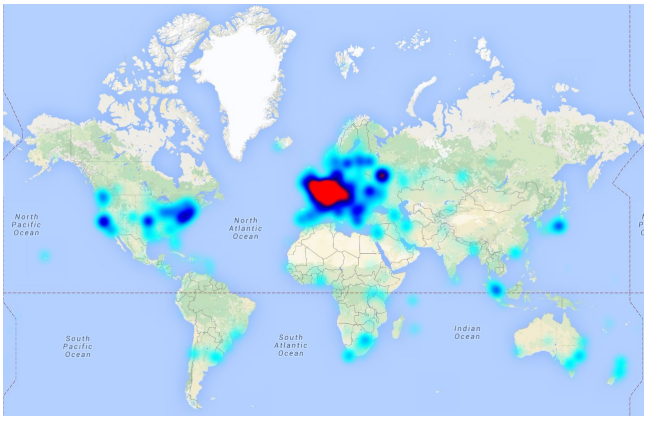

(a) Ripe Atlas probes

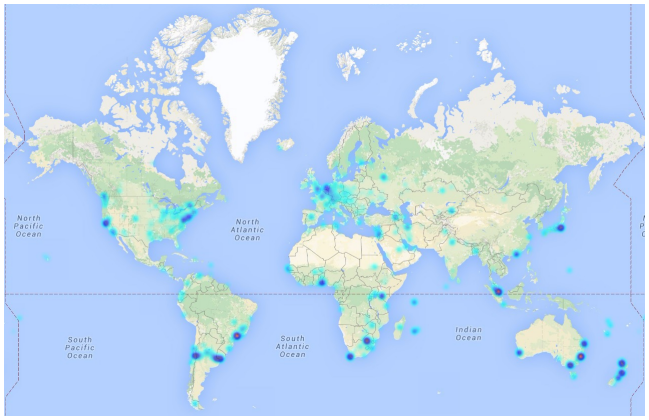

(b) Chosen probes

Figure 1: Distribution of probes around the world.

Vantage points (VPs). We leverage the RIPE Atlas platform [45] which provides a large collection of probes located around the world for our traceroute measurements. RIPE Atlas has broad global coverage, but is known to have a disproportionate fraction of probes in Europe, as can be seen in Figure 1(a). To avoid a bias towards European ASes, we determine how many probes we would like to use and

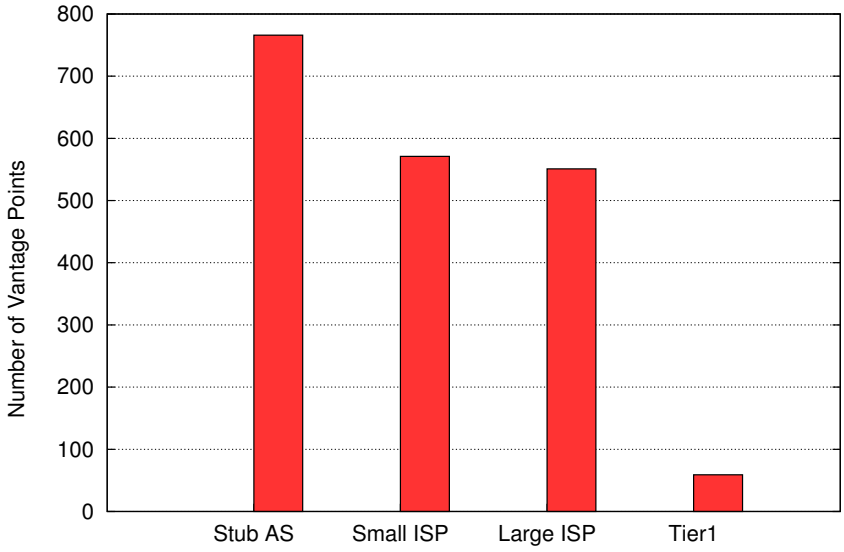

Figure 2: Distribution of the RIPE Atlas probes used in this study.

evenly divide the number of probes across all continents. We then start choosing probes from countries within each continent in a round robin manner focusing on distributing probes in different ASes within each country. We do this for each continent until we have allocated the target number of probes and this way we are able to get a more uniform distribution, as can be seen in Figure 1(b).

Furthermore we also investigate the location of the RIPE Atlas VPs in the AS level topology, shown in Figure 2, and observe that many of them are located near the network edge in stub and small ISP networks.. In order to to do this categorization , we use the approach described in [40].

To measure paths to content providers, each RIPE Atlas node performs a DNS

lookup for each of the 34 content DNS names, and then performs a traceroute to the resolved IP. In our experiments we use 1,998 RIPE Atlas probes,[2] located in 633 ASes, distributed according to our sampling methodology. Combined, these probes perform 28,051 traceroutes to 218 destination ASes. The number of destination ASes is large relative the number of content providers because large numbers of content servers are hosted outside the provider's network (*e.g.*, inside ISPs).

From traceroutes to routing decisions. We convert the traceroute-based IP-level paths into AS paths using the method described by Chen et al. [11]. Since interdomain routing is destination based, we can observe routing decisions for all ASes along the path to a given destination. We thus observe routing decisions for a total of 746 ASes.

3.2 Comparison with existing models

We compare paths observed in our our data- and control-plane measurements with CAIDA's topology of inferred inter-AS relationships. We aggregate 5 topologies (Oct 14 to Feb 15) inferred using the method presented by Luckie *et al.* [33]. We aggregate these snapshots of the AS level topology to mitigate the impact of transient link failures on the topology used in our analysis. When inferences conflicted, we took the majority poll of inferred relationships while assigning higher weight to more recent inferences. We use this topology to compute all paths that

[2]We targeted 2,000 probes but two did not return any data and had to be discarded.

satisfy the Gao-Rexford (GR) local preference model described in Section 2.

We compare the measured paths with all paths satisfying the GR model of local preference computed using CAIDA's inferred relationships. We consider two properties: (1) whether the measured path satisfies the GR model of local preference, and (2) whether the measured path has the same length as the shortest paths satisfying the GR model of local preference. Based on this we classify routing relationships as either obeying GR local preference; *i.e.*, using the neighbor with the Best Relationship type (**Best**), routing based on shortest path (**Short**), or a combination of the two. These sets should be treated as disjoint, with ASes that obey both Best and Short path policies appearing only in the **Best/Short** category. Observations which follow **neither** of these properties are considered inconsistent with existing models (*i.e.*, violations).

3.3 How often do models hold?

We now consider how empirically observed AS paths compare with those predicted by models using AS relationships inferred in [33]. We then investigate how often deviations can be explained by known sources of inaccuracies.

Encouragingly, we find that a majority of routing decisions (65%) are correctly inferred by the commonly used Best/Short model; however, a significant fraction (35%) are not. Figure 3 characterizes the observed routing decisions based on

whether the path chosen is Best or Short. We find only a small number of cases (8%) where decisions can neither be explained by Best or Shortest path selection. In the following sections, we explore the reasons behind these decisions that differ from model-based predictions.

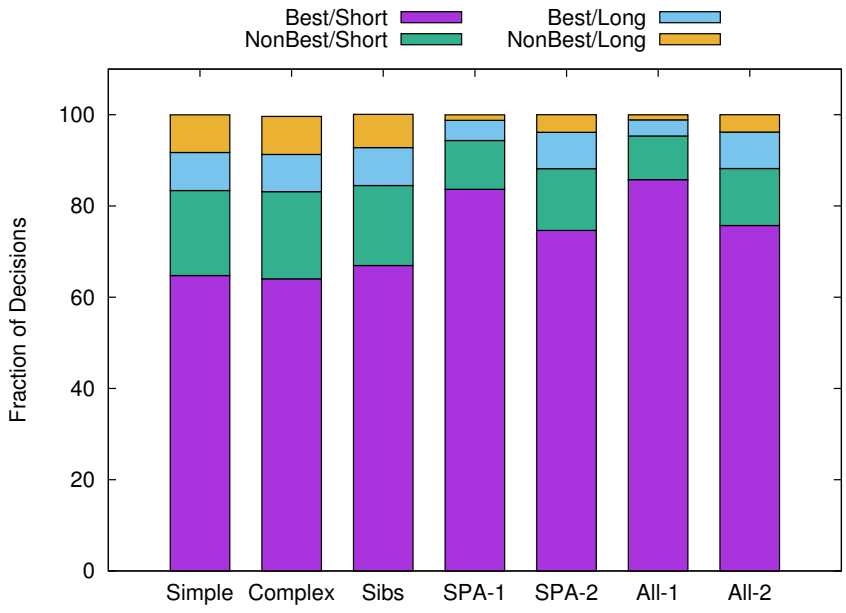

.

Figure 3: Breakdown of routing decisions observed taking into account the complex relationships.

3.3.1 Complex routing relationships

A well known limitation of existing routing policy models is the simplification of relationships into either customer-provider or settlement-free peering relationships. Recent work by Giotsas *et al.* aims to address this limitation by augmenting existing relationship inferences with cases of hybrid relationships (*i.e.*, ASes

whose arrangements vary based on location) and partial transit relationships (*i.e.*, ASes who will behave as providers, but only for a subset of prefixes) [19]. Figure 3 (Complex) shows the breakdown of routing decisions observed taking into account these complex relationships. Interestingly, we find that taking these relationships has nearly no impact on the classification in our dataset.

3.3.2 Internet eXchange Points (IXPs)

IXPs have been a phenomenon of considerable interest in the interdomain routing community for the past five years [5, 6]. While, the inferred topology we use only includes a set of 25 IXPs, the method we use for mapping traceroute paths to AS paths [11] incorporates a more extensive list [47]. We validate that the AS paths we use in our analysis are free of IXP ASes by correlating them with public data from a prior study of IXPs [6].

We find that 28% of Content traceroute paths contain IP addresses allocated to IXPs according to the public data. Of the IXP IP addresses, only 19% of them can be mapped to an ASN, consistent with the fact that many IXPs operate at layer 2. Further, none of the AS paths contain IXP ASes, indicating that existing methods provide good coverage of IXP ASes. Thus, none of the unexplained routing decisions we observe are attributed to IXP ASes.

3.3.3 Sibling ASes

The mapping between AS numbers and organizations is not one-to-one [9]. Many organizations manage multiple AS numbers, either for geographic regions (*e.g.*, Verizon with ASNs 701, 702, and 703) or due to mergers (*e.g.*, Level 3 (AS 3356) and Global Crossing (AS 3549)). As ISPs continue to grow and consolidate, it is critical that our models of routing are able to take these "sibling" relationships into account. While some studies have taken sibling relationships into account [9], the inferred topology we leverage, as well as large scale algorithmic frameworks [18] have not dealt with this special case.

We take a similar approach to Cai et al. [9] to identify AS siblings, but our approach differs in two key ways. First, we focus only on e-mail addresses in whois data, which previous work identified as the field with best precision and recall [9]. Second, we use DNS SOA records to identify different e-mail domains that belong to the same organization. For example, dish.com and dishaccess.tv share the dishnetwork.com authoritative domain. We also remove groups where the e-mail address is hosted by a popular e-mail provider (*e.g.*, hotmail.com), or regional Internet registry (*e.g.*, ripe.net). This results in a total of 94 sibling groups identified in our traceroute data set.

For every non Best/Short decision that an AS makes, we check whether the AS chose a path via a sibling. If the path is a via a sibling, we mark this decision as satisfying the Best condition. Figure 3 (Sibs) shows the result of this change—3%

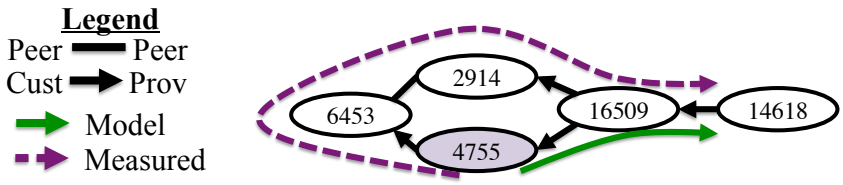

Figure 4: Example of routing decisions made towards Amazon's AS 14618.

more decisions are classified as Best/Short.

3.3.4 Prefix-specific policies

Interdomain routing is often abstracted to the level of a destination AS. However, in practice routing is based done on IP prefixes which may be subject to different export policies by their originating AS (*e.g.*, forwarding prefixes hosting enterprise-class services to a more expensive provider). While Giotsas *et al.* consider partial transit [19], which is a type of prefix-specific policy, they do not explicitly consider per-prefix policies as implemented by origin ASes.

When analyzing paths with unexplained routing policies, we find that selective prefix announcements can explain many counterintuitive routing decisions. One particular case concerned AS 14618, an AS maintained by Amazon, which has Amazon's AS 16509 as its sole upstream provider. Interestingly, we found that ASes routing to AS 14618 would prefer a longer and more costly path to Amazon's AS 16509, even though many of them had a direct customer-provider relationship with AS 16509. Figure 4 illustrates an example of these strange decisions, with AS 4755 preferring a 4-hop path via a provider over a 2-hop path via

a customer towards AS 14618. Closer inspection of Routeviews' data revealed that AS 16509 was not actually exporting prefixes from AS 14618 to all of its providers. Instead, it announced these prefixes to six providers, forcing its other providers to select longer and sometimes more costly paths, over direct customer-provider links.

Determining when networks do selective prefix announcements definitively is challenging without extensive looking glass server capabilities. However, in the absence of complete looking glass server coverage, we use two criteria to identify origin-based prefix specific policies based on correlation with BGP data obtained from Routeviews/RIPE [3,44]. Given an origin AS (O), neighbor N and prefix P:

- **Criteria 1** do not assume the edge $N - O$ exists for prefix P unless we observe O announcing P to N in the BGP data.
- **Criteria 2** is similar to Criteria 1, except that we require that we observe at least one prefix announced from O to N before applying Criteria 1.

The first criteria can be seen as being more aggressive whereas the second aims to ensure that our observation is actually caused by selective prefix announcement and not poor visibility.

Figure 3 (SPA-1, SPA-2) shows the breakdown of routing decisions using Criteria 1 and 2 above, respectively. We find that prefix-specific policies account for a significant fraction (10-19%) of unexpected routing decisions.

Validation. In order to validate cases of prefix-specific policies, we try to find a Looking Glass server hosted by the neighboring AS. The looking glass server gives us insight as to what prefixes were actually announced to that AS. We say that an instance of selective prefix announcement is valid when we see that only a subset of prefixes from the originating AS were announced to its next hop AS (the one hosting the looking glass server). There were a total of 630 cases of prefix-specific policies involving 149 unique neighboring ASes. We were able to find looking glass servers in 28 of the neighboring ASes. Using these looking glass servers we manually verify 100 cases of prefix-specific policies and confirm that applying Criteria 1 was correct 78% of the time.

3.4 Sources of Violations

In this section we investigate which source and destination ASes account for most of the routing decisions which deviate from our model. Figure 5 (a) and (b) shows the cumulative fraction of routing decisions which violate either the Best or Short condition (*i.e.*, the AS chooses a path that is longer or more expensive than we would expect). If violations were evenly distributed across ASes, the curves would fix $y = x$; otherwise, some ASes are responsible for a disproportionately larger (or smaller) fraction of violations. We find this effect is present in both plots, but

more prominently for destination ASes. We focus on the latter.

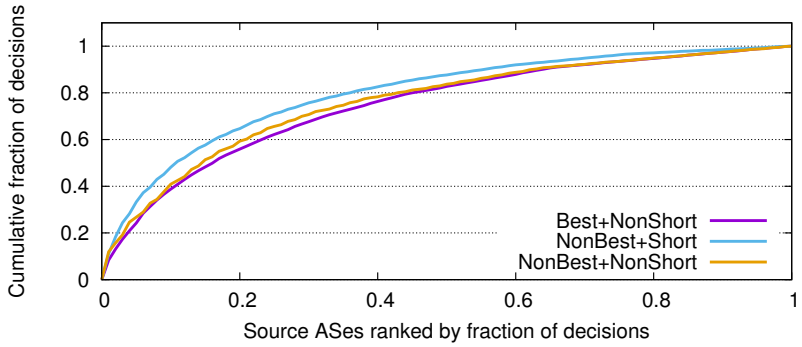

(a) Distribution of violations across source ASes.

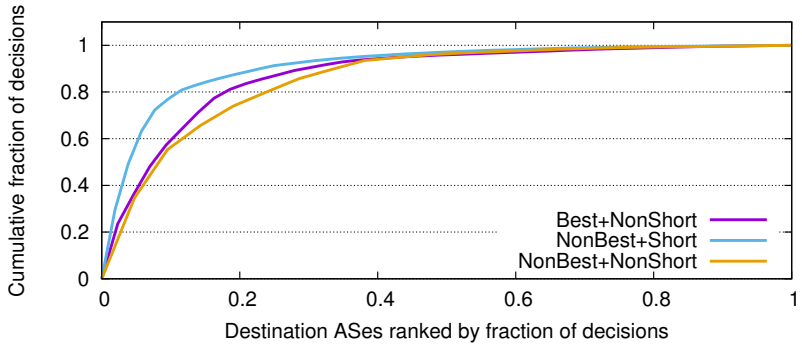

(b) Distributions of violations across destination ASes.

Figure 5: CDF plot of the fraction of violations (x-axis) explained by source and destinations ASes (y-axis). Violations observed in our dataset are skewed significantly toward Akamai and Netflix. The skew for source ASes is less prominent.

Destination ASes owned by Akamai account for 21% of violations. Of these, Cogent (AS174) is the most common source, responsible for 3.4% of violations. These Cogent-Akamai violations tend to occur when Cogent prefers a peer-to-peer path through a Tier-1 AS over a longer customer route. Netflix's AS is the

destination on 17% or paths with violations. Of these, nearly a quarter (24%) are due to a stale inter-AS link in CAIDA's topology, which included a direct link between AS3549 and Netflix that no longer exists.

For source ASes, the distribution is less skewed. Cogent and Time Warner are the top two sources, responsible for 4.1% and 2.2% of violations, respectively.

3.5 Impact of Geography

We next consider the role of geography on routing decisions. First, we isolated traceroutes that stay within a continent (Continental traceroutes), *i.e.*, all hops stay inside a given continent based on geolocating router IP addresses. We use the geolocation data from [10], which offers good coverage of infrastructure IPs such as routers. Figure 6 shows the breakdown of decisions in the continental traceroutes (45% of those in our dataset). The fraction of decisions explained by Best/Short for continental traceroutes is significantly greater than for transcontinental ones.

3.5.1 Domestic paths.

Next we focused on traceroutes where we infer that the all entities on the entire traceroute path stayed within a single country, but there is a better *multinational Best/Short path* (in the CAIDA data), which we define to be a path with at least one AS registered (via whois data) in a country outside the source and destination AS's country. It is important to note here that we are talking about where

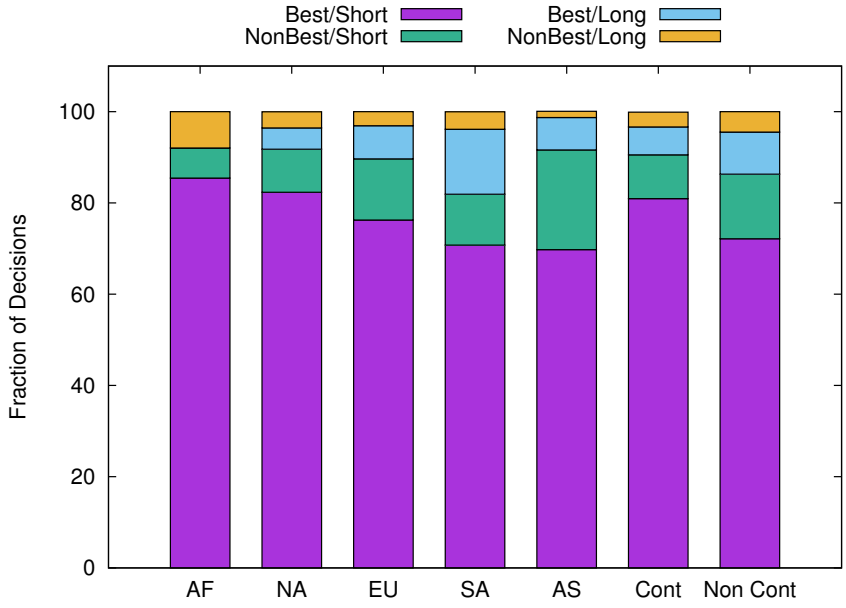

Figure 6: Breakdown of traceroutes that stay within a continent

ASes are registered and make no claims on where the path actually goes. We find that more than 40% of non-Best/Short decisions can be explained by avoiding alternative multinational paths. Table 1 details the non-Best/Short decisions explained by ASes preferring domestic routes. There can be multiple reasons behind the fact that a large number of non-Best/Short decisions can be explained by this table. One obvious conclusion is that many ASes prefer intra-country routes over cheaper path that takes routes outside of the country. In some cases we can attribute this behavior to wrong inference of BGP relationships. For example, 132267 is a small ISP in Bangladesh, but it is shown to be a provider of Level3 (Tier1 AS) in CAIDAs inference data.

Continent	Non-Best/Short Decisions explained
Asia	40.1%
Africa	62.5%
N. America	10.9%
Oceana	62.9%
S. America	66.6%

Table 1: Summary of Non-Best/Short decisions explained by ASes preferring intra-country routes.

3.5.2 Undersea cables.

Undersea cable ASes are a critical component of Internet topologies that previous work overlooks. While some cables are jointly owned by large ISPs, e.g., Pan-American Crossing, Americas-II (owned by AT&T, Sprint, and many others), we observed that others, e.g., EAC- C2C (PACNET), are operated by independent organizations using their own allocated ASNs and IP prefixes. Because these cable operators only provide point-to-point transit along the cables (i.e., they do not originate traffic and peer in locations proportional to cable landings), they resemble high-latency, high-cost IXPs and thus confuse existing AS relationship models. As such, we need techniques to identify cable ASes and correct their relationships in inferred topologies.

We use a list of undersea cables from the TeleGeography Submarine Cable Map [2] to identify ASes for undersea cable operators. Overall, cable-ASes appear on less than 2% of paths but most of the decisions (51%) involving cable-ASes caused deviations from Best/Short paths. Table 2 shows fraction of each type of decision explained by undersea cable ASes.

Violation type	Pct. of decisions explained
Non-Best & Short	3%
Best & Long	6.5%
Non-Best & Long	4.5%

Table 2: Fraction of decisions of each type that can be attributed to undersea cables.

4 Defending Against Routing Around Decoy Attacks with Path Steering

We aim to use our understanding of inter domain routing and recent work outlining the use of BGP poisonings to influence AS paths over the Internet [27] to try to come up with a deployment scheme fro decoy routers. The goal is to either (1) make it infeasible for the censor to route around decoy routers or (2) increase the collateral cost of routing around them to be unacceptably high.

4.1 Decoy routing

Decoy routing operates using a proxy, referred to as a decoy router, located on the path between the client and a benign destination (sometimes referred to as the *overt* destination). The benign destination can be any IP address reachable by the client over a network path containing at least one decoy router. To leverage decoy routing, the client establishes a TLS connection to the benign destination and uses steganography (*e.g.*, on TLS header values) to signal to the decoy-router that this connection should be treated differently and directed towards a *covert* destination. This redirection of traffic is highlighted in Figure 7

One of the key features of decoy routing, over traditional proxying systems, is that it drastically increases the set of IP addresses a censor would need to block to limit access to the system [28]. If a censor is focused on blocking IP addresses, they would need to block all IP addresses accessible over paths containing de-

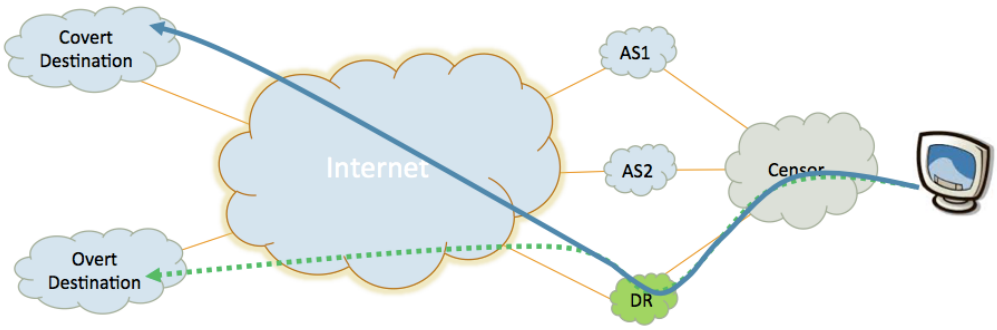

Figure 7: The figure shows how a decoy router redirects traffic destined towards an overt destination to a covert destination.

coy routers. If a decoy router is located in a large transit AS this may boil down to blocking nearly all destination IPs on the Internet. Further, since the clients are accessing benign IP addresses which are not affiliated with decoy routing in any explicit way, this system provides a certain amount of deniability that is not present in traditional proxy systems.

A key challenge faced by decoy routing is deployment. Specifically, it requires buy in from ISPs (ideally large ones) to deploy decoy routers. While deployment still remains a challenge, new proposals to implement decoy routing in an *on-path* manner [51] rather than *in-path* [52] dramatically lowering the barriers to deployment. This is because on-path systems, which simply require tapping a link and receiving a copy of traffic, are considered much safer than in-path devices which operate on live network traffic and may impact connectivity if they malfunction.

4.2 Routing around decoy (RAD) attacks

IP address based blocking is a standard approach for blocking proxy systems. However, since decoy routers are deployed in specific networks and act on all traffic traversing the network path, the censoring ISP needs to modify its BGP decision process to exclude paths containing networks running decoy routers, this is referred to as routing around decoys (RAD). The censoring AS can then make sure that all BGP announcements for prefixes that go over an AS known to host a decoy router are dropped and alternate paths are used.

Earlier papers argue about the possibility for censors to route around decoy deployments. Schuchard *et al.* [49] present the RAD attack and show how a country as well connected as China can easily perform this attack. However, they do not take into account the business relationships of ASes within the country when they perform their analysis. Rather they describe that if any AS within the country receives a path for a prefix that does not go over a decoy router, all other ASes within the country can use that AS to reach the affected prefix. This would essentially mean that ASes could end up transiting traffic for other ASes that are not actually paying, leading to a violation of the valley-free policy. Houmansadr *et al.* point this out and perform the analysis again by taking into account the business relations between ASes [24]. In that paper, the authors point out that although China can launch a RAD attack, but the real cost of it would be prohibitively expensive for them in terms of the number of destination ASes that become unreachable.

There are essentially three main factors that impact whether the RAD attack would work or not. These are:

1. The cost of the alternate path used when avoiding one that has a decoy router on it. The adversary might have to choose a more expensive route *e.g.*, routing via a provider when they have an available customer or peer path.

2. The risk of disconnecting from some networks since it is not always guaranteed that alternate paths would exist to be used.

3. Inaccessibility of content. Apart from being disconnected from networks, important traffic from popular content providers could get effected.

We work on extending the analysis of Houmansadr *et al.* to better understand the cost to route around decoys and present a candidate defence leveraging "helper ASes" to prevent RAD attacks.

4.3 Modeling the Effects of RAD Attacks and Defenses

Decoy routing relies on the fact that an adversary is able to filter and block traffic based on the source and destination IP addresses of packets but is not able to do

so based on which routers the traffic flows through while it is in transit. Since the adversary cannot use IP based filtering to block decoy routers deployed within other autonomous systems, it has to look into ways of blocking traffic from going through the AS hosting a decoy router. This is where BGP based filtering comes into play.

Figure 8 shows a simplistic scenario where a censoring AS performs a RAD attack. **Case 1** shows how advertisements passing through ASes known to deploy decoy routers are simply dropped and an alternate path is used to reach the destination.

Case 2 highlights the case where a destination becomes completely unreachable. This happens when all paths to a given destination contain a decoy router, and are thus dropped, causing the destination to become unreachable.

4.4 Ability of censoring countries to RAD

We start our analysis by considering the top 14 countries for censorship, based on Freedom House [1], but focus on China because they are the only one that can actually route around decoy router deployments. None of the other 13 countries in our list were nearly as highly connected as China and even a simple random deployment of decoy routers over the Internet causes these countries to lose connectivity to a significantly high fraction of the ASes on the Internet if they try to

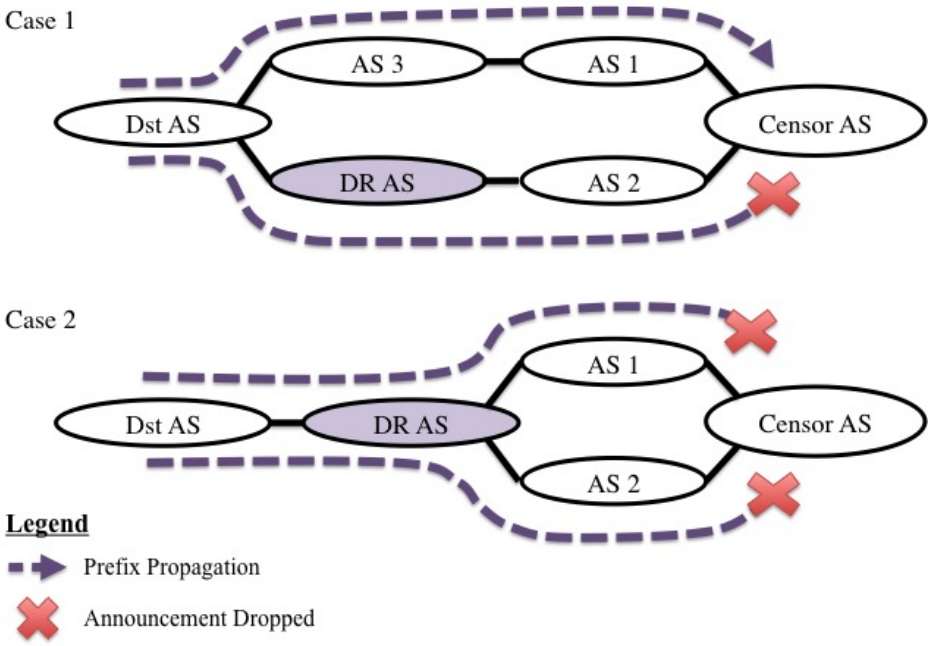

Figure 8: The figure shows how a censoring AS would route traffic when performing a 'Routing Around Decoys' attack.

perform a RAD attack [24, 49]. For this reason we focus our analysis on China.

4.4.1 Initial data collection and tools used

Extending the AS topology around China. In order to be able to study how different decoy router deployments affect connectivity for China, we needed a better picture of the ASes connected to Chinese ASes. We used Routeviews [3] data to identify Chinese edge ASes (connected to international AS's) and ring ASes (international ASes connected to edge ASes).

Modifying our simulator for RAD. For the analysis of paths taken by an AS

performing a RAD attack we needed a way to compute the alternate paths the AS would take. We extend the simulator used in our policy work to compute paths to destinations via algorithmic simulations [18] for our policies work and modify it to take in as input a set of ASes that host decoy routers and then compute paths to all routable ASes from a given source by avoiding paths that traverse an AS hosting a decoy router.

4.4.2 Modeling popularity of destinations.

We replicate Houmansadr's *et al.* simulations [49] and try to figure out what fraction of the traffic gets affected by the random deployment of decoy routers. It has been shown in earlier studies that a Zipf distribution can be used to model traffic over the Internet [4]. We start by retrieving a list of the top 1,000 websites for China from Alexa and then use Zipf parameters to figure out relative weight of a site at each rank. It turns out that only 114 out of the top 1,000 websites were hosted outside China, and hence could get affected by RAD attack. Since the rest were hosted within China, a RAD attack had no effect on their connectivity.

Figure 10 shows what fraction of the traffic gets affected for various deployment percentages in ASes with customer cone size of at least five. Even with a 100% deployment in ASes with customer cone size of at least five (which means 2,965 ASes decide to host a decoy router in their network), only 88 sites become unreachable and they account for just 7.5% of the total number of hits from China

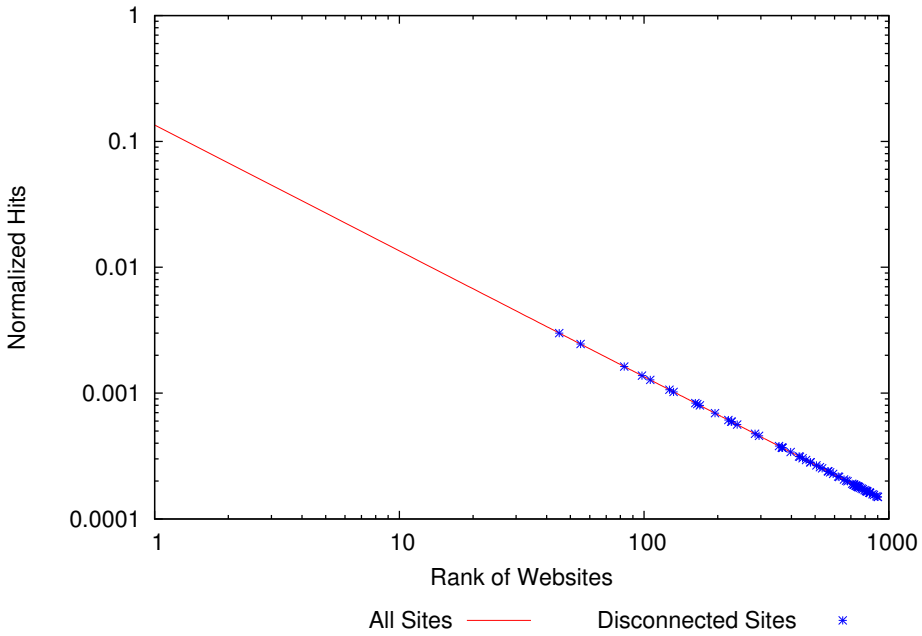

Figure 9: The figure shows the fraction of hits that every disconnected site has from within China. The hit fractions are computed using a Zipf distribution.

(computed using a Zipf distribution). Figure 9 shows that these 88 sites lie at the lower end of the top 1,000 list.

Houmansadr *et al.* have shown in their work [49] how deploying decoy routers in less than 10% of ASes in the Internet can cause China to lose connectivity to about 45% of the total ASes. We replicated the simulations that the authors had done and try to figure out what fraction of the traffic does get affected by these random deployments. We suspected that many of the disconnected ASes were of little importance (in terms of the amount of traffic) to China and hence did not incur as big of a cost.

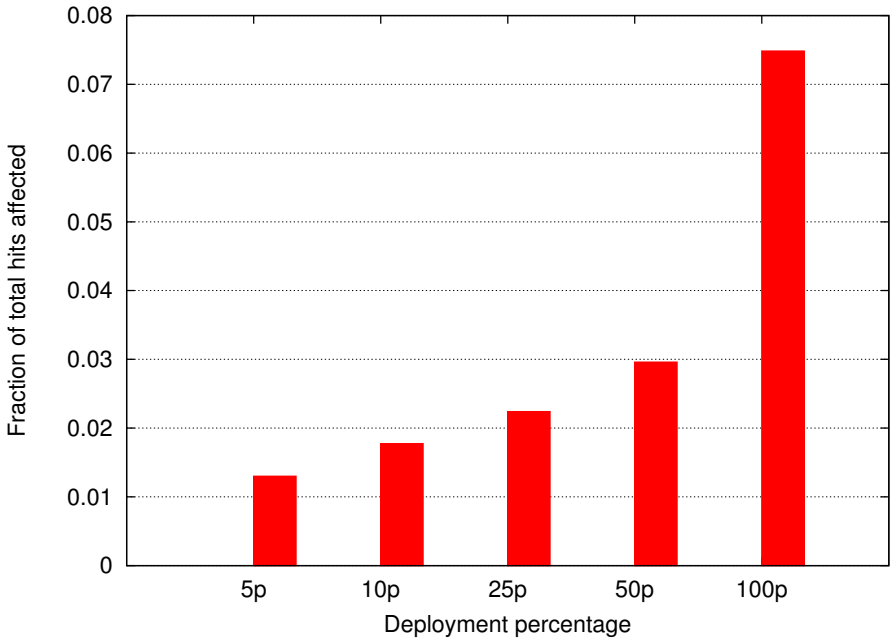

Figure 10: Fraction of traffic hits affected by various deployments of DRs in ASes with a customer cone size of at least five.

We conclude that even though China does lose connectivity to a sizable number of ASes when it tries to route around decoys, the actual impact on the country's traffic is very small, so small in fact that it is unlikely to deter China from perform a RAD attack.

4.5 Helping get Decoy Routers on Network Paths

Our goal is to design a deployment scheme which minimizes the number of ASes that need to deploy decoy routers while making RAD infeasible/costly for censors. Towards this goal, we consider the possibility of enlisting willing ASes to serve

as helpers to the decoy routing system via specially crafted BGP announcements.

4.5.1 Defining the entities

Censoring ASes are ASes that belong to countries that are known to censor content over the Internet. We only take into consideration the larger edge ASes which are able to perform the RAD attack due to their increased connectivity with other ASes. The smaller ASes connected behind these larger ASes not taken into consideration.

Helper ASes are ASes belonging to organizations that want to help circumvent censorship but do not host a decoy router in their network. Any network with the capability of making BGP announcements could theoretically become a helper AS. Ideally helper ASes could be networks that announce high value prefixes, *e.g.*, Google, so that dropping their announcements or blocking them would be considered expensive for the censor.

Decoy Router ASes These are ASes that host a decoy router in their network. Theoretically any AS wanting to help could host a decoy router but they must provision to transit the circumvented traffic. Ideally large transit networks with a large customer cone should host a decoy router. This way dropping paths traversing the large AS would cut the censor off from its customers.

4.5.2 Strawman solution: Using simple poisoning to combat RAD

Ring ASes are international ASes that are adjacent to the edge ASes of the censoring country. This strawman approach relies on the assumption that an AS in the ring AS set would deploy a decoy router. The idea here is to have the helper AS poison ring ASes such that the only remaining path for the censor is via the ring AS which deploys decoy routing. This approach is illustrated in Figure 11. This could essentially create incentives for other ASes in the ring to deploy decoy routers in order to gain revenue from traffic (similar incentives have been explored for BGP security in [17]). On the other hand, ASes might not choose to deploy decoy routers because doing so might cause the censor to end their contract and look for a replacement provider.

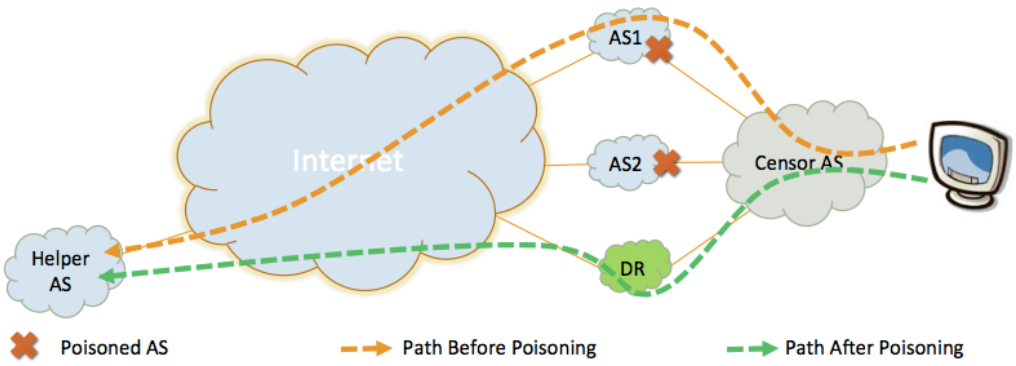

Figure 11: The figure shows how a helper AS could poison ring ASes to force traffic through a decoy router.

The key problem with a simplistic poisoning of ring ASes is that customers under the poisoned ASes might lose connectivity as a result of the poisoning. For

that we propose sub prefix and super prefix announcements approach. The idea is to make 2 announcements, a longer more specific prefix and a shorter prefix encompassing the longer one poisoning the ring ASes on the longer prefix (From the helper AS). This forces the censoring AS to route using the more specific prefix that it is receiving from one of its ring AS (the one hosting the decoy router). Since a larger prefix encompassing this prefix was also announced, customers under the poisoned ring ASes can still route to the helper AS with no loss of connectivity.

4.5.3 Path steering

In this section, we consider the idea of poisoning BGP announcements to influence how ASes choose to route towards the helper AS. To provide intuition about how poisoning can impact network paths, we look at LIFEGUARD [30], which uses poisoning to allow a server to influence paths chosen towards its network. They do so by poisoning the AS they want to avoid thereby shifting transiting traffic away from it.

Path steering is a method for poisoning, that instead of avoiding a given AS, shifts the destination‘s traffic onto a given target AS (in our case the target ASes would be ASes that deploy decoy routers). Our aim is to figure out the set of all ASes that a given helper can steer its path over by performing poisoning. To this end we modify our simulator to perform iterative poisonings to get all ASes that the path towards the helper AS could be steered over. Using these paths we

figure out which ASes are best candidates for decoy router deployments. We go a step further by analyzing multiple Helper ASes and figure out the least number of decoy router deployments that would serve multiple helper ASes. The steps used to iteratively prune all reachable ASes for a given helper AS and target AS are shown below.

1. A path is computed from Censor AS T to Helper AS A and poisoning is done iteratively to explore all ASes that can be reached.

2. First the initial path is computed and added into a queue (This is depth 0, where depth x means x hops away from the destination).

3. Then we poison at depth 1 and keep doing so until we get to the poison limit and add each path back into the queue with the number of poisoning performed to get that path.

4. The depth counter is incremented and now all paths with poison counter $<$ the limit AND depth count $<$ current depth counter are extracted and the whole cycle is repeated until all paths in queue max out the poison counter.

The whole idea is to move away from the previously proposed random deployments that required 100s to 1000s of ASes deploying decoy routers in order to allow users within a censoring AS to benefit from them. We propose this targeted approach that introduces the notion of helper ASes and drastically lowers

the number of ASes that would need to deploy decoy routers.

Deciding which ASes get to be helpers. To analyze the effect of our approach, we choose to target China since it is the only censoring country capable of performing a RAD attack. China has a large number of autonomous systems but ISPs serving end users were generally behind AS4134 (China Telecom Backbone) and hence AS4134 was chosen as the target. To get a list of possible helpers for China, the top 1,000 list from Alexa (for China) was used and all websites were mapped to the ASes hosting them. These ASes are good candidates to be helper ASes as blocking them would mean higher collateral in terms of lost connectivity. Since it is unlikely that China would allow ASes within the country to deploy censorship circumvention tools, therefore we filter all sites hosted within China and we are left with 64 ASes that could potentially be helper ASes. After running our simulations on these 64 ASes, we find that 94% of the helper ASes are able to steer their path over more than one AS. The trend on the number of ASes these 64 helper ASes can steer their path over is shown in Figure 12

Figuring out potential decoy router candidates. An ideal decoy router candidate is an AS that is located in the network in such a place that multiple helper ASes can leverage its decoy router deployment by steering their paths through it. All ASes identified to be steerable over by the 64 helper ASes are candidates for decoy router deployments. This shows us how a single helper AS can drastically narrow down the number decoy deployments required. We also find 12 ASes that

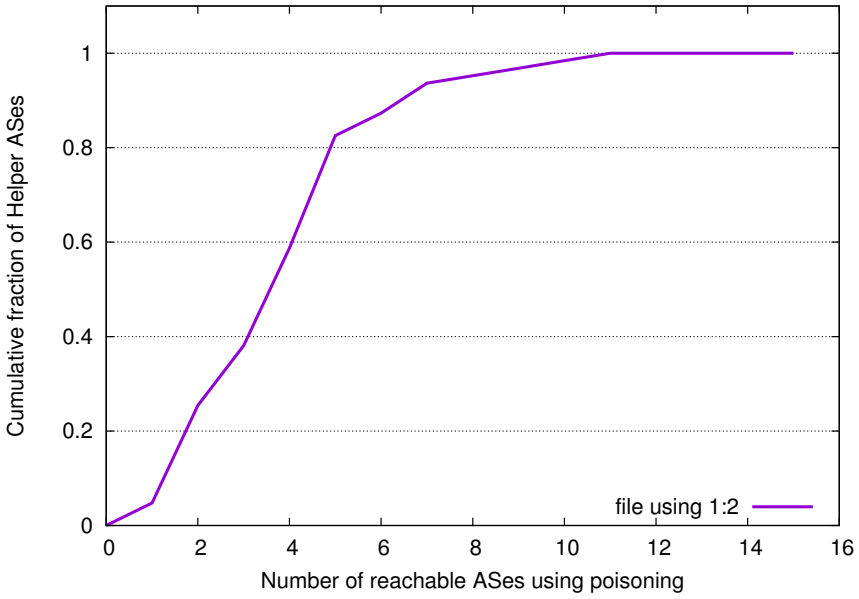

Figure 12: The figure shows the CDF of the number of ASes reachable by poisonings made by the 64 identified helper ASes.

are steerable towards from multiple helper AS. Therefore Deploying DR in these ASes would result in a much higher collateral damage for China as compared to other random deployments.

5 Conclusion

We set out to better understand how the age old interdomain routing models fare in today‘s Internet. To that end, we investigated how interdomain paths predicted by the routing models differ from empirically observed routes. It is important to note here that we cannot measure all Internet paths, so our results are biased toward those we can. In order to mitigate the impact of this limitation, we selected routes that likely carry most Internet traffic (paths to content). We found that while a majority of paths in our dataset agree with models, about a third do not. We explained a significant fraction of these differences due to factors such as complex relationships, sibling ASes, selective prefix announcements, undersea cables and incorporated them into our model. Limited path visibility constrains our ability to identify the root cause behind the remaining inconsistent routing decisions. Improving this visibility should help explain significantly more cases of modeling failures.

We noticed recent works related to routing around decoy were not taking into account some important aspects of the routing models while performing the evaluation of their ideas. We go on to use the tools and understanding we developed to extend their analysis of defenses against the RAD attack and study the effects of various deployments of decoy routers. We introduced the notion of helper ASes and showed how introducing them into the decoy routing ecosystem can actually drastically reduce the number of decoy routers required to target a censoring AS.

6 References

[1] Freedom in the World. https://freedomhouse.org/.

[2] TeleGeography Submarine Cable Map. http://www.submarinecablemap.com/.

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

[4] L. A. Adamic and B. A. Huberman. Zipfs law and the internet. *Glottometrics*, 3(1):143–150.

[5] B. Ager, N. Chatzis, A. Feldmann, N. Sarrar, S. Uhlig, and W. Willinger. Anatomy of a large european IXP. In *SIGCOMM'12*, 2012.

[6] B. Augustin, B. Krishnamurthy, and W. Willinger. IXPs: Mapped? In *IMC'09*, 2009.

[7] H. Ballani, P. Francis, and X. Zhang. A study of prefix hijacking and interception in the Internet. In *SIGCOMM*, 2007.

[8] M. A. Brown. Rensys Blog: Pakistan hijacks YouTube. http://www.renesys.com/blog/2008/02/pakistan_hijacks_youtube_1.shtml.

[9] X. Cai, J. Heidemann, B. Krishnamurthy, and W. Willinger. Towards an AS-to-organization map. In *USENIX ATC*, 2010.

[10] B. Chandrasekaran, M. Bai, M. Schoenfield, A. Berger, N. Caruso, G. Economou, S. Gilliss, B. Maggs, K. Moses, D. Duff, K. Ng, E. G. Sirer, R. Weber, and B. Wong. Alidade: Ip geolocation without active probing. *Department of Computer Science, Duke University, Technical Report*, 2015.

[11] K. Chen, D. Choffnes, R. Potharaju, Y. Chen, F. Bustamante, D. Pei, and Y. Zhao. Where the sidewalk ends: Extending the internet AS graph using traceroutes from P2P users. In *CoNEXT '09*, 2009.

[12] A. Dainotti, C. Squarcella, E. Aben, K. Claffy, M. Chiesa, M. Russo, and A. Pescapé. Analysis of country-wide internet outages caused by censorship. In *Proceedings of the 2011 ACM SIGCOMM conference on Internet measurement conference*, pages 1–18. ACM, 2011.

[13] R. Dingledine, N. Mathewson, and P. Syverson. Tor: The second-generation onion router. Technical report, DTIC Document, 2004.

[14] L. Gao, T. Griffin, and J. Rexford. Inherently safe backup routing with BGP. *IEEE INFOCOM*, 2001.

[15] L. Gao and J. Rexford. Stable Internet routing without global coordination. *Trans. Netw.*, 2001.

[16] P. Gill, S. Goldberg, and M. Schapira. A survey of interdomain routing policies. *ACM CCR*, 2014.

[17] P. Gill, M. Schapira, and S. Goldberg. Let the market drive deployment: A strategy for transistioning to BGP security. *SIGCOMM'11*, 2011.

[18] P. Gill, M. Schapira, and S. Goldberg. Modeling on quicksand: dealing with the scarcity of ground truth in interdomain routing data. *SIGCOMM Comput. Commun. Rev.*, 42(1):40–46, Jan. 2012.

[19] V. Giotsas, M. Luckie, B. Huffier, and K. Claffy. Inferring Complex AS Relationships. In *ACM IMC*, November 2014.

[20] S. Goldberg, M. Schapira, P. Hummon, and J. Rexford. How secure are secure interdomain routing protocols? In *SIGCOMM'10*, 2010.

[21] T. Griffin, F. B. Shepherd, and G. Wilfong. The stable paths problem and interdomain routing. *Trans. Netw.*, 2002.

[22] A. Houmansadr, G. Nguyen, M. Caesar, and N. Borisov. Cirripede: circumvention infrastructure using router redirection with plausible deniability. In *ACM CCS 2011*, 2011.

[23] A. Houmansadr, T. J. Riedl, N. Borisov, and A. C. Singer. I want my voice to be heard: Ip over voice-over-ip for unobservable censorship circumvention.

[24] A. Houmansadr, E. L. Wong, and V. Shmatikov. No direction home: The true cost of routing around decoys. In *Proceedings of the 2014 Network and Distributed System Security (NDSS) Symposium*, 2014.

[25] G. Huston. Peering and settlements - Part I. *The Internet Protocol Journal (Cisco)*, 2(1), March 1999.

[26] G. Huston. Peering and settlements - Part II. *The Internet Protocol Journal (Cisco)*, 2(2), June 1999.

[27] U. Javed, I. Cunha, D. R. Choffnes, E. Katz-Bassett, T. Anderson, and A. Krishnamurthy. Poiroot: Investigating the root cause of interdomain path changes. In *SIGCOMM*, 2013.

[28] J. Karlin, D. Ellard, A. W. Jackson, C. E. Jones, G. Lauer, D. P. Mankins, and W. T. Strayer. Decoy routing: Toward unblockable internet communication. In *USENIX workshop on free and open communications on the Internet*, 2011.

[29] J. Karlin, S. Forrest, and J. Rexford. Nation-state routing: Censorship, wiretapping, and BGP. *CoRR*, 2009.

[30] E. Katz-Bassett, C. Scott, D. R. Choffnes, I. Cunha, V. Valancius, N. Feamster, H. V. Madhyastha, T. Anderson, and A. Krishnamurthy. LIFEGUARD: Practical repair of persistent route failures. In *SIGCOMM*, 2012.

[31] C. Labovitz, S. Iekel-Johnson, D. McPherson, J. Oberheide, and F. Jahanian. Internet inter-domain traffic. In *SIGCOMM'10*, 2010.

[32] M. Lad, D. Massey, D. Pei, Y. Wu, B. Zhang, and L. Zhang. PHAS: A prefix hijack alert system. In *Proc. USENIX Security Symposium*, 2006.

[33] M. Luckie, B. Huffaker, A. Dhamdhere, and V. Giotsas. AS relationships, customer cones, and validation. In *ACM Internet Measurement Conference*, 2013.

[34] R. Lychev, S. Goldberg, and M. Schapira. Is the juice worth the squeeze? BGP security in partial deployment. In *SIGCOMM'13*, 2013.

[35] H. Madhyastha, E. Katz-Bassett, T. Anderson, A. Krishnamurthy, and A. Venkataramani. iPlane Nano: Path prediction for peer-to-peer applications. In *Usenix NSDI*, 2009.

[36] R. Mazloum, M. Buob, J. Auge, B. Baynat, D. Rossi, and T. Friedman. Violation of Interdomain Routing Assumptions. In *Passive and Active Measurement Conference*, March 2014.

[37] H. Mohajeri Moghaddam, B. Li, M. Derakhshani, and I. Goldberg. Skypemorph: Protocol obfuscation for tor bridges. In *Proceedings of the 2012 ACM conference on Computer and communications security*, pages 97–108. ACM, 2012.

[38] W. Mühlbauer, A. Feldmann, O. Maennel, M. Roughan, and S. Uhlig. Building an AS-topology model that captures route diversity. In *SIGCOMM*, 2006.

[39] H. Noman. Dubai free zone no longer has filter-free Internet access. ONI Blog: http://opennet.net/blog/2008/04/dubai-free-zone-no-longer-has-filter-free

[40] R. Oliveira, D. Pei, W. Willinger, B. Zhang, and L. Zhang. Quantifying the completeness of the observed internet AS-level structure. *UCLA Computer Science Department - Techical Report TR-080026-2008*, Sept 2008.

[41] ONI research profile: Burma. http://opennet.net/research/profiles/burma, 2012.

[42] ONI research profile: Indonesia. http://opennet.net/research/profiles/indonesia, 2012.

[43] Quantcast. http://www.quantcast.com.

[44] RIPE Network Coordination Center. RIPE Routing Information Service. http://www.ripe.net/data-tools/stats/ris/routing-information-service.

[45] RIPE Atlas, 2013. https://atlas.ripe.net/.

[46] M. Roughan, W. Willinger, O. Maennel, D. Perouli, and R. Bush. 10 lessons from 10 years of measuring and modeling the Internet's autonomous systems. *JSAC*, 2011.

[47] M. A. Sánchez, J. S. Otto, Z. S. Bischof, D. R. Choffnes, F. E. Bustamante, B. Krishnamurthy, and W. Willinger. Dasu: pushing experiments to the internet's edge. In *NSDI*, 2013.

[48] Sandvine. Fall 2012 global internet phenomena, 2012.

[49] M. Schuchard, J. Geddes, C. Thompson, and N. Hopper. Routing around decoys. In *Proceedings of the 2012 ACM conference on Computer and communications security*, pages 85–96. ACM, 2012.

[50] J. Wu, Y. Zhang, Z. M. Mao, and K. Shin. Internet routing resilience to failures: Analysis and implications. In *CoNEXT*, 2007.

[51] E. Wustrow, C. M. Swanson, and J. A. Halderman. Tapdance: end-to-middle anticensorship without flow blocking. In *USENIX Security Symposium. USENIX*, page 118, 2014.

[52] E. Wustrow, S. Wolchok, I. Goldberg, and J. A. Halderman. Telex: Anticensorship in the network infrastructure. In *USENIX Security Symposium*, 2011.