

No Direction Home: The True Cost of Routing Around Decoys

Amir Houmansadr
The University of Texas at Austin

Edmund L. Wong*
Meraki, Inc.

Vitaly Shmatikov
The University of Texas at Austin

Abstract—Decoy routing is a recently proposed approach for censorship circumvention. It relies on cooperating ISPs in the middle of the Internet to deploy the so called “decoy routers” that proxy network traffic from users in the censorship region. A recent study, published in an award-winning CCS 2012 paper [24], suggested that censors in highly connected countries like China can easily defeat decoy routing by selecting Internet routes that do not pass through the decoys. This attack is known as “routing around decoys” (RAD).

In this paper, we perform an in-depth analysis of the true costs of the RAD attack, based on actual Internet data. Our analysis takes into account not just the Internet topology, but also business relationships between ISPs, monetary and performance costs of different routes, etc. We demonstrate that even for the most vulnerable decoy placement assumed in the RAD study, the attack is likely to impose tremendous costs on the censoring ISPs. They will be forced to switch to much more costly routes and suffer from degradation in the quality of service.

We then demonstrate that a more strategic placement of decoys will further increase the censors’ costs and render the RAD attack ineffective. We also show that the attack is even less feasible for censors in countries that are not as connected as China since they have many fewer routes to choose from.

The first lesson of our study is that defeating decoy routing by simply selecting alternative Internet routes is likely to be prohibitively expensive for the censors. The second, even more important lesson is that a fine-grained, data-driven approach is necessary for understanding the true costs of various route selection mechanisms. Analyses based solely on the graph topology of the Internet may lead to mistaken conclusions about the feasibility of decoy routing and other censorship circumvention techniques based on interdomain routing.

I. INTRODUCTION

With recent advances in censorship technologies, evading censorship is becoming more challenging. New circumvention systems aim to make their traffic *unobservable* in order to (1) protect their users, and (2) prevent their services from being

blocked by censors. *Decoy routing* is a new approach to unobservable censorship circumvention, proposed independently in systems called DR [17], Telex [26], and Cirripede [15]. In contrast to traditional circumvention tools in which circumvention proxies run on end-host servers, decoy routing places these proxies—called *decoy routers*—at the routers of volunteer ISPs (in the rest of this paper, we will use the terms ISP and “autonomous system” interchangeably). To use a decoy routing system, a client connects to a non-blocked destination via a route containing a decoy router; the decoy router acts as a man-in-the-middle for the connection and proxies the traffic to the blocked destinations requested by the client.

Schuchard et al. [24] proposed the “*routing around decoys*” attack against decoy routing. In the rest of this paper, we will use the terms “RAD attack” and “RAD paper” to refer, respectively, to this attack and the paper in which it was published. The basis of the RAD attack is the observation that ISPs in the censorship region are likely to have multiple paths to any given destination. Therefore, censors can instruct the ISPs under their influence to exclusively select routes that do not pass through the ISPs known to deploy decoy routers.

The RAD attack is considered successful only if it manages to avoid the decoys while (mostly) maintaining the connectivity of the censoring ISPs to the rest of the Internet. Schuchard et al. analyze the Internet topology and show that—assuming that the decoy routers are placed in a small number of *randomly selected* autonomous systems—the RAD attack will maintain the censors’ connectivity.

Our contributions. In this paper, we take a closer look at the true costs of the RAD attack. We start by estimating the *quality* of the alternative routes selected by the RAD adversary, as opposed to their mere existence. In this analysis, we make the same random placement assumption as the RAD paper, even though it is heavily biased in favor of the RAD adversary (a random autonomous system is unlikely to transit others’ traffic, thus placing decoy routers in it serves little purpose).

The short summary of our findings is that the RAD attack is likely to impose huge monetary and performance costs on the censoring ISPs. The RAD paper observes that if decoy routers are placed at 2% of all autonomous systems, China—by far the easiest case for the RAD attack due to its high connectivity—would get disconnected only from 4% of the Internet [24, Fig. 2a]. While true, this is not the whole story. Our simulations show that:

- On average, the estimated latency of China’s Internet routes will increase by a *factor of 8*.

* Research described in this paper was performed at The University of Texas at Austin

- 44 of China’s customer autonomous systems will have to become “transit” autonomous systems, requiring vast re-organization and investment in their network infrastructure. By comparison, China today has only 30 transit autonomous systems.
- There will be dramatic changes in loads on China’s transit autonomous systems. For example, transit loads will increase by a *factor of 2800* for one autonomous system, while decreasing by 32% for another.
- 39% of China’s Internet routes will become longer; 12% will become more expensive.

A more strategic placement of decoy routers further amplifies the censors’ costs, even in terms of basic Internet connectivity. If decoy routers are placed in 2% of all autonomous systems, but the systems are chosen strategically rather than randomly, China will be disconnected from 30% of all Internet destinations, not 4% as calculated in the RAD paper.

We also analyze the feasibility of the RAD attack for other state-level censors. As intuitively expected, the costs of the RAD attack depend on the censoring country’s network infrastructure. Countries with less connectivity in the global Internet graph incur higher costs. For instance, a RAD attack against decoy routers strategically placed in 1% of all autonomous systems will disconnect China from 18% of all Internet destinations, whereas Venezuela and Syria will be disconnected from 54% and 87% of all destinations, respectively.

In addition to showing that “routing around decoys” is likely to be very costly, our study provides several lessons and recommendations. An important methodological lesson is that, when analyzing the feasibility and costs of attacks and defenses based on Internet routing, it is not enough to simply look at the topology of the Internet graph. The edges in this graph are not all equal, they have vastly different costs and performance characteristics. Relationships between autonomous systems, such as customer-provider, peer, etc., matter a lot. Therefore, any analysis of decoy routing and alternatives must be based on all available fine-grained data about individual nodes and edges in the Internet graph.

Organization. In Section II, we provide background information on the Internet ASes, decoy routing, and the RAD attack. In Section III, we describe how the RAD attack works. In Section IV, we explain the costs that must be incurred by censors to carry out a RAD attack. In Section V, we suggest strategic decoy placements. In Section VI, we describe our data sources and the simulation setup. In Section VII, we estimate the costs of the RAD attack. We conclude with lessons and recommendations in Section VIII.

II. BACKGROUND

A. Internet topology

The Internet is a globally distributed network composed of more than 44,000 [3] autonomous systems. An autonomous system (AS) is “a connected group of one or more IP prefixes run by one or more network operators which has a *single* and *clearly defined* routing policy” [14].

While the details of business agreements between ASes can be complex, the widely accepted Gao model [11] ab-

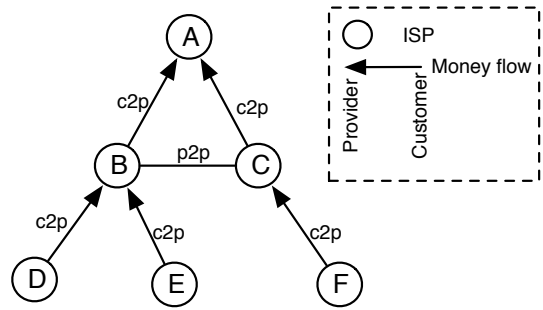


Fig. 1: A sub-tree of the Internet topology graph.




Fig. 2: The CDF of customer cone size (the maximum customer cone size, which is 22,664, is not shown).

stracts business relationships into the following three main types [1]:

- **Customer-to-provider (c2p):** An AS A is a customer of a connected neighbor AS B (the provider) if A pays B to transit A ’s traffic to Internet destinations that A cannot reach otherwise. Similarly, B has a **provider-to-customer (p2c)** relationship with A .
- **Peer-to-peer (p2p):** Two ASes are peers if they exchange Internet traffic between each other and each other’s customers free of charge, due to a mutual business agreement.
- **Sibling-to-sibling (s2s):** Two ASes are siblings if they belong to the same organization. Sibling ASes do not charge each other for the transit traffic.

Figure 1 illustrates these relationships.

An AS’s **customer cone** includes the AS itself plus all ASes that can be reached from that AS through provider-to-customer links.¹ In other words, A ’s customer cone includes A , A ’s customers, A ’s customers’ customers, and so on. Figure 2 shows the CDF of customer cone size for all 44,064 Internet ASes.

An **edge AS** is an AS whose customer cone has size 1, i.e., it has no customers. A **transit AS** is an AS whose customer cone is greater than 1, i.e., it transits other ASes’ traffic to the rest of the Internet.

Internet routes are based on paths between ASes (inter-

¹<http://as-rank.caida.org/?mode0=as-intro#customer-cone>

domain routes) which are established via BGP, the Border Gateway Protocol [22]. A **path** is a sequence of neighbor ASes that connect the source AS to the destination AS in the Internet topology graph. A path is **valid** if, for every transit AS on the path, there exists a customer [1] who is its immediate neighbor. A path is **invalid** if at least one transit AS is not paid by a neighbor in the path [1, 10]. Valid paths are also referred to as valley-free (**VF**). Correspondingly, we refer to invalid paths as non-valley-free (**NVF**). Figure 3 shows examples of valid and invalid paths.

Valley-freeness is not a requirement of the BGP protocol, i.e., BGP routers are technically able to advertise NVF paths. However, as described above, a NVF path will impose undesired monetary costs on some transit ISP because it will not earn money (or may even have to pay money) for transiting the traffic of another ISP. Therefore, ISPs widely refrain from advertising NVF paths.

B. Decoy routing

Decoy routing is a new architecture for censorship circumvention which was proposed in three independent works: DR [17], Telex [26], and Cirripede [15]. In contrast to traditional circumvention techniques [2, 5, 7, 8, 9, 16, 20, 25] that operate on computer servers located outside censorship regions, decoy routing systems are deployed on a number of routers in the middle of the Internet, called *decoy routers*, by ASes that we refer to as *decoy ASes*. Instead of making direct connections to the circumvention endpoints, e.g., proxies, a decoy routing client makes a TLS [6] connection to arbitrary, non-blocked Internet destinations, known as *overt destinations*. The client selects overt destinations so that the routes to these destinations pass through decoy routers and steganographically signals the decoy router to treat these connections as circumvention connections. The decoy router intercepts the client’s traffic and proxies the connection to the *covert destination* requested by the client. To a censor observing the client’s traffic, the client appears to be communicating with a non-blocked, overt destination, while the client is actually communicating with a forbidden, covert destination.

In DR [17] and Telex [26], the decoy router itself proxies covert connections, whereas in Cirripede [15] decoy routers deflect the traffic to external proxies. Also, while Telex and Cirripede require clients to probe for overt destinations that happen to have decoy routers on routes leading to them, DR assumes that clients obtain the secret locations of decoy routers through out-of-band channels. The proposed decoy routing designs also use different signaling techniques: Cirripede uses the initial sequence number of the TLS connection, whereas Telex uses the TLS nonce. Further details on the design of decoy routing systems can be found in the original papers [15, 17, 26].

How to select ASes for decoy placement has been studied in three papers. Houmansadr et al. [15] and Cesareo et al. [4] analyzed the placement of decoy routers in a non-adversarial setting, while Schuchard et al. [24] analyzed the placement of decoy routers in the presence of a censor capable of changing routing decisions—see Section II-C.

C. Routing around decoys (RAD)

Schuchard et al. [24] introduced the “routing around decoys” (RAD) attack against decoy routing systems. The RAD attack is conducted by a *routing-capable adversary*, i.e., a censoring regime who can modify the standard routing decisions of the ISPs under its influence in order to ensure that their Internet traffic does not pass through any decoy ASes. The ASes controlled by a RAD adversary discard all BGP paths that contain even one decoy AS and choose alternative, decoy-free paths. In order to launch the RAD attack, the RAD adversary needs to know which ASes deploy decoy routers. This can be done, for example, via probing schemes proposed in the RAD paper.

The main intuition behind the RAD attack is as follows. For any given source and destination, the Internet topology is likely to provide multiple interdomain paths. Consequently, a RAD adversary can compel its ASes to avoid paths that contain decoy ASes without sacrificing much of its Internet connectivity. If censorship results in a significant loss or degradation of Internet connectivity in the censorship region, it causes significant collateral damage and is less likely to be in the censors’ interest. Therefore, *the RAD attack is considered successful only if the RAD adversary can avoid all decoy ASes while maintaining its connectivity with most of the Internet.*

To improve the RAD adversary’s connectivity, the RAD paper assumes that the ASes under the adversary’s control share interdomain paths with each other *regardless of their business relations*. In other words, an AS controlled by a RAD adversary can use the paths known to any other AS controlled by the same RAD adversary.

The RAD paper considers several censoring regimes as possible RAD adversaries, including China, Iran, and Syria. As the RAD paper suggests, China is the most powerful RAD adversary due to its significant connectivity.

III. INTERDOMAIN ROUTING IN RAD

The BGP [22] protocol is the de facto standard used by ASes to construct interdomain paths. The RAD attack forces ASes under the RAD adversary’s control to change how they make BGP routing decisions. We refer to the resulting protocol as RBGP.²

A. BGP routing

A BGP router maintains a database with the paths to different Internet destinations and advertises some of these paths to the routers of the neighbor ASes, as determined by the ASes’ business relationships (see Section II-A). For instance, a BGP router of a transit AS advertises all known paths to its customers’ routers in order to earn money by transiting their traffic. On the other hand, a BGP router should not advertise its paths to the provider ASes, otherwise the AS that owns the router would end up paying its providers for transiting their traffic (such paths are NVF, as explained in Section II-A).

A BGP router is likely to know multiple paths to a given Internet destination (identified by its IP address prefix). BGP

²The name should not be confused with the R-BGP protocol of Kushman et al. [18].




Fig. 3: Sample AS paths.

routers use a list of *decision factors*, shown in Table I, to identify the best path. These factors are applied in order, with each factor filtering out the set of paths left by the previously applied factor. For example, the B_2 factor is applied only to the paths that are considered best according to the B_1 factor. The router applies the factors until only one path remains, i.e., the best path. For instance, suppose that for a certain destination a BGP router knows four paths, two of which pass through its provider neighbors and the other two pass through its peer neighbors (we explain the difference between providers and peers in Section II-A). In this case, the B_3 factor filters out the two paths that route through providers, and the B_4 factor is applied only to the two paths that route through peers.

We only focus on two of the decision factors from Table I since they are highly influenced by the RAD attack. The description of the other factors can be found in the BGP specification [22].

B_3 Business preference (highest Local-Pref) This factor selects routes with the best benefit for the router’s AS. This benefit is usually monetary. Typically, B_3 prefers paths that route through a customer, then those that route through a peer, and finally those that route through a provider. This is due to the AS business relationships described in Section II-A, e.g., routing through a peer is free while routing through a provider costs money.

B_4 Shortest AS path The fourth decision factor is the path length, i.e., the number of ASes in the path from the source AS to the destination prefix. Path length affects the quality of service of the connection, hence it comes immediately after the business preference factor. A path composed of more ASes is susceptible to higher network latencies, lower throughputs, and more frequent network failures.

B. RBGP routing

The RAD attack changes how BGP routers choose AS paths. BGP routers controlled by a RAD adversary use a modified list of decision factors to select the best path to a given destination; we call such routers *RBGP* routers. An RBGP router has two objectives that distinguish it from a standard BGP router.

Avoiding decoy routers: Because the main intention of a RAD adversary is to avoid paths that contain decoy routers, an RBGP router simply discards all paths that pass through at least one decoy AS.

Traffic re-routing: If a RAD AS does not have a decoy-free path to a given destination, the RAD paper suggests that it can

use decoy-free paths known to other RAD ASes, regardless of the business relation between these ASes. In other words, a RAD AS who knows a decoy-free path to a given destination transits the traffic of other RAD ASes to that destination even if this contradicts the standard BGP decision factors.

For instance, if a Chinese AS does not have a decoy-free path to a certain destination, it can re-route traffic to that destination through one of the other 198 ASes in China, e.g., a customer AS or an AS with which it has no business relationship. This is a key factor in the success of the RAD attack, because it increases the number of alternative paths available to the RAD ASes. The resulting routes may be invalid (NVF) routes, as defined in Section II-A. While the RAD paper does not describe in detail how re-routing is performed, it suggests the use of network engineering tools such as MPLS VPN tunnels [23, Section 3.1] across all ASes controlled by the RAD adversary. In the rest of this paper, we will argue that, regardless of the networking technique used to implement re-routing, it will be extremely costly to the ASes involved.

To achieve the two objectives described above, an RBGP router uses a different list of decision factors (compared to BGP) for finding the best path to a given Internet destination. This list is shown in Table II. It adds two new decision factors: R_1 (**Ignore if the route includes decoy ASes**) and R_2 (**Prefer VF routes over NVF routes**). The latter factor is necessary because NVF routes are much more costly than VF routes.

IV. THE COSTS OF RAD ROUTING

The non-standard decision factors used by RBGP impose additional costs on the ASes controlled by the RAD adversary. These costs fall into several categories: (1) collateral damage (e.g., social unrest) caused by the fact that significant parts of the Internet become unreachable; (2) collateral damage due to the significantly lowered quality of service for the customers of the RAD-controlled ASes; (3) monetary costs for buying and deploying new networking equipment; and (4) monetary costs due to switching to more expensive Internet routes.

Intuitively, all of these costs stem from one main reason. The standard list of decision factors used by conventional BGP routers aims to minimize ASes’ routing costs and to maximize the quality of service for their network traffic. Therefore, any change to these decision factors is likely to increase their costs, decrease quality of service, or both.

In the following, we describe the negative impacts of RAD routing, arranged by type.

1. Degraded Internet reachability (*Reachability*) Avoiding paths that contain decoy routers may disconnect

TABLE I: BGP’s decision factors for choosing the best path (in order).

B_1	Ignore if next hop unreachable
B_2	Prefer locally originated networks
B_3	Business preference (highest Local-Pref)
B_4	Shortest AS path
B_5	Prefer lowest Origin
B_6	Prefer lowest MED
B_7	Prefer eBGP over iBGP
B_8	Prefer nearest next hop
B_9	Prefer lowest Router-ID or Originator-ID
B_{10}	Prefer shortest Cluster-ID-List
B_{11}	Prefer lowest neighbor address

RAD-controlled ASes from an Internet destination unless the RAD adversary can find a decoy-free path to that destination. By definition, a large number of disconnected destinations means that the attack has failed (see Section II-C).

2. Less-preferred paths (Business) As explained in Section III-A, one of the first decision factors that standard BGP routers consider is the business relationship between the router’s AS and the first AS of a candidate interdomain path (the decision factor B_3). In RBGP, however, two other decision factors, R_1 and R_2 , have higher priority. As a result, it is likely that for some destination the RBGP router selects a path with a lower business preference compared to what a standard BGP router would have selected.

For example, suppose that a router chooses between two paths to some destination: path A goes through a provider and contains no decoy ASes, while path B goes through a peer and contains a decoy AS. A standard BGP router would have selected path B because it is cheaper, but an RBGP router will select the more expensive path, A .

3. Longer paths (Length) As explained in Section III-A, one of the top standard decision factors of BGP is the length of the available paths (factor B_4). Based on this factor, a standard BGP router prefers the path that contains the fewest transit ASes. This helps maximize quality of service for routed traffic because longer paths may have higher latency and are more susceptible to network failures. For RBGP routers, B_4 is lower in the preference order, which may cause them to select longer paths than BGP routers.

4. Higher path latencies (Latency) Longer routes are not the only cause of higher latencies. The alternative paths selected by RBGP are likely to pass through less popular transit ASes that offer lower capacity, causing packets to experience higher latencies. This is confirmed by our simulations in Section VII, which show that, even when an RBGP path has the same length as the corresponding BGP path, it usually has higher latency.

5. Non-valley-free routes (Valley) As explained above,

TABLE II: RBGP’s decision factors for choosing the best path (in order).

R_1	Ignore if the route includes decoy ASes
R_2	Prefer VF routes over NVF routes
R_3	Ignore if next hop unreachable
R_4	Prefer locally originated networks
R_5	Business preference (highest Local-Pref)
R_6	Shortest AS path
R_7	Prefer lowest Origin
R_8	Prefer lowest MED
R_9	Prefer eBGP over iBGP
R_{10}	Prefer nearest next hop
R_{11}	Prefer lowest Router-ID or Originator-ID
R_{12}	Prefer shortest Cluster-ID-List
R_{13}	Prefer lowest neighbor address

RBGP routers may be forced to selected non-valley-free (NVF) paths in order to avoid decoy ASes. Such paths are extremely expensive, which is why they are shunned by normal BGP routers.

Suppose that for a given Internet destination, a RAD AS A has no decoy-free BGP path and must use the path known to another RAD AS B . In this example, either A has to pay B for transiting A ’s traffic (A would not have had to pay B if A had used standard BGP), or else B has to pay the expenses for transiting A ’s traffic (e.g., to B ’s provider). Additionally, the source AS A may have to pay its own provider in order to transit traffic to B . The monetary costs of Valley are likely to be much worse than Business costs.

6. New transit ASes (NewTransit) The RAD attack relies on the fact that the ASes under the adversary’s control transit traffic for each other (see Section II-C). However, only a small fraction of ASes under the control of a typical RAD adversary are transit ASes and thus have the requisite network equipment and resources.

For instance, China has 199 ASes, but only 30 of them are transit ASes. For the RAD attack to be successful, the RAD adversary needs to transform many of the edge ASes into transit ASes. Changing a typical edge AS to a transit AS is extremely costly since it requires the purchase and installation of sophisticated networking equipment.

7. Massive changes in transit load (TransitLoad) Transit ASes earn money by transiting other ASes’ traffic. On the other hand, transiting this traffic imposes significant fixed and variable costs, including equipment, network management, etc.

Our simulations in Section VII shows that the RAD attack significantly changes the transit load of the transit ASes under the RAD adversary’s control. Due to the routing changes caused by the RAD attack, some transit ASes lose a large fraction of their transit traffic (and thus lose money), while other transit ASes must handle tremendous increases in their transit load.

V. PLACING DECOY ROUTERS

For a decoy routing system to become operational, it must be deployed by several autonomous systems (*decoy ASes*) who are economically or politically motivated to assist in censorship circumvention. The number of the decoy ASes as well as their location in the Internet are important factors determining whether a decoy routing system can withstand the RAD attack.

The original RAD paper simulated the RAD attack for two specific placements of decoy ASes: *top-tier* and *random*. The former placement assumes that the decoys are deployed in top-tier Internet ASes, while the latter assumes that the ASes for decoy deployment are chosen randomly from the set of all 44,000 ASes. Analysis in the RAD paper suggests that the RAD attack fails against the top-tier placement because it results in disconnecting the RAD adversary from large parts of the Internet. The RAD paper observes, however, that top-tier placement is expensive and may not be practically feasible.

For the random placement, the RAD paper shows that if decoys are deployed in a small, random fraction of all ASes, the RAD adversary is disconnected only from a small part of the Internet—mainly from the decoy ASes themselves—thus the RAD attack is considered successful.

We believe that the random decoy placement analyzed in [24] is biased in favor of the RAD adversary and does not reflect how the RAD attack would fare against a realistic decoy deployment strategy. Based on the AS ranking statistics, available from CAIDA,³ we observe that 86.2% of all ASes are edge ASes, i.e., the size of their customer cone is 1 (see Section II-A). Therefore, the random decoy placement considered in [24] is likely to place decoys primarily into edge ASes. Obviously, evading an edge AS disconnects the RAD adversary only from that AS because it is not on the path to any other AS.

We argue that, in any realistic deployment, *decoy routers should be placed in transit ASes, not edge ASes, even in the absence of a RAD adversary*. The larger the customer cone of an AS, the better it serves as a decoy AS, for two reasons: (1) an AS with a larger customer cone is on the path to more ASes, thus the RAD attack is likely to disconnect the adversary from these “downstream” ASes, too, and (2) even in the absence of a RAD adversary, placing decoys on ASes with larger customer cones provides better unobservability for decoy routing clients and gives them more options for choosing their overt destinations.

For example, suppose that a decoy routing system is installed only in a single edge AS. In this case, its clients’ options for overt destinations are limited to the destinations belonging to that single AS. Therefore, a user who frequently visits destinations within the decoy AS may raise the censor’s suspicion that the user is engaging in decoy routing. On the other hand, if decoys are installed in a transit AS with a customer cone of 5, then a decoy routing client can choose overt destinations from 5 ASes, resulting in better connectivity and better unobservability.

Based on these observations, we propose the following strategic decoy placement strategies, which are much more

likely to defeat the RAD attack than the random placement considered in [24].

Sorted placement (sorted): In this approach, decoy ASes are chosen from among the ASes that transit more traffic for the RAD adversary. Specifically, we sort ASes based on the number of times they appear on the BGP routes of the RAD adversary’s ASes. We then choose decoy ASes from the top of this sorted list. We exclude all ASes controlled by the RAD adversary, i.e., Chinese ASes if China is the adversary.

We propose two types of sorted placements. In the *sorted-with-ring* placement, decoy ASes are chosen from the set of all ASes not directly controlled by the RAD adversary (i.e., non-Chinese ASes in the case of China). In the *sorted-no-ring* placement, we additionally exclude all ASes that have a direct business relationship with the RAD adversary, since they are less likely to deploy decoy routers. We use the term *ring ASes* for the ASes that are not controlled by the adversary, but have a business relationship. From our data sources (see Section VI), we identified 551, 69, and 5 ring ASes for China, Venezuela, and Syria, respectively.

Strategic random placement (random): Instead of selecting random ASes from the set of all ASes, as suggested in [24], our random placement strategy selects ASes from the set of all ASes with a given customer cone size. In a *random-C* placement strategy, decoy ASes are chosen randomly from the set of all ASes with a customer cone size larger than or equal to C . Our *random-1* strategy is thus the exact random strategy suggested in [24] (since 1 is the minimum value for the customer cone size). Similar to the *sorted* placement, we further subdivide *random-C* placement into two types: *random-with-ring-C* and *random-no-ring-C*. Both exclude adversary-controlled ASes, and the latter additionally excludes all ring ASes that have a direct business relationship with an adversary-controlled AS.

VI. SIMULATION SETUP AND DATA SOURCES

We use simulation to estimate the various costs imposed by RBGP routing on the RAD adversary, described in Section IV. Our simulator uses CBGP [21], a popular BGP simulator, as its engine, and a Python interface to interact with CBGP and query for BGP routes between ASes. The rest of the simulations are performed in Python.

We use several sources of Internet measurements in our simulations:

- Geo location: We use the “GeoLite Country” dataset from GeoLite’s geolocation database⁴ to map IP addresses to countries.
- AS relations: We use CAIDA’s inferred AS relationship dataset,⁵ which is based on [11], to model the relationships between ASes.
- AS ranking: We use CAIDA’s AS rank dataset⁶ to infer the customer cones of individual ASes.

⁴<http://dev.maxmind.com/geoip/legacy/geolite>

⁵<http://www.caida.org/data/active/as-relationships/>

⁶<http://as-rank.caida.org/>

³<http://as-rank.caida.org/>

TABLE III: Comparing the Internet connectivity of state-level censors.

Country	Number of ASes controlled	Number of ring ASes
China	199	551
Venezuela	44	69
Syria	3	5

- Latency: We use iPlane’s⁷ [19] “Inter-PoP links” dataset to estimate BGP and RBGP path latencies. This dataset contains daily latency measurements between different points-of-presence (PoP) of ASes.
- Network origin: We use iPlane’s “Origin AS mapping” dataset to map IP address prefixes to the corresponding ASes.

VII. SIMULATION RESULTS

The success of the RAD attack depends on the placement of decoys in ASes. Therefore, we evaluate the costs of the attack for different placement strategies described in Section V. In all cases, we assume that the RAD adversary knows the identities of all ASes that deploy the decoys. Obviously, this assumption favors the adversary.

A RAD adversary is a censorship authority who controls a large number of ASes and forces them to modify their BGP decisions as described in Section III-B. Intuitively, a RAD adversary’s Internet connectivity is proportional to the number of ASes it controls and the number of its ring ASes (see Section V). The larger these numbers, the more alternative routes are likely to be available to the RAD adversary for any given Internet destination. As mentioned before, the RAD attack is successful only if it does not disconnect the adversary’s ASes from many ASes in the rest of the Internet.

This suggests that China is the most powerful RAD adversary because it controls a large number of ASes (199) and is connected to more ring ASes than other state-level censors (see Table III). We demonstrate this by comparing China’s success as a RAD adversary with other censoring countries, such as Venezuela (44 ASes) and Syria (3 ASes).

Figure 4 shows the percentage of ASes that become unreachable as a consequence of the RAD attack, assuming sorted-no-ring decoy placement. This shows that China significantly outperforms Syria and Venezuela in maintaining its connectivity with the rest of the Internet.

For the rest of the simulations, we only report the results for China. The simulations were performed for two different scenarios:

- China-World: China is the RAD adversary; decoy ASes are chosen, using different placement strategies from Section V, from all 44,000 ASes excluding the 199 ASes located in China (we additionally exclude the 551 ring ASes of China in the case of no-ring placements, as described in Section V). The costs of the RAD attack are then estimated for connections from China to all Internet destinations across the world, excluding the Chinese destinations.

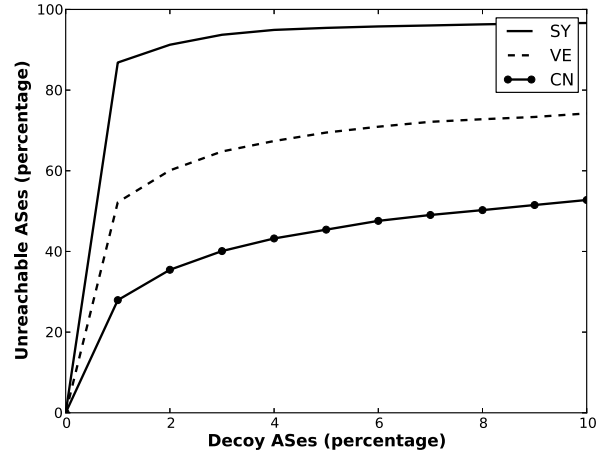


Fig. 4: Loss of connectivity for different RAD adversaries assuming the sorted-no-ring decoy placement strategy.

- China-US: China is the RAD adversary; decoy ASes are selected only from the 13,299 ASes located in the United States. This scenario represents a geographically limited deployment of decoy routers. In this case, the costs of the RAD attack are only estimated for the Internet destinations inside the US. As above, China’s ring ASes are excluded in the no-ring deployments.

A. Loss of connectivity

Figure 5 shows the percentage of Internet ASes that become unreachable from China under different placement strategies and for different numbers of decoy ASes. As described above, for the China-US scenario both decoy ASes and destination ASes are only selected from the US-based ASes, while for the China-World scenario they are selected from all non-Chinese ASes.

The random-no-ring-1 placement is exactly the placement studied in the RAD paper [24], where it was called “random” placement. Following the RAD paper [24], our simulations confirm that random-no-ring-1 mainly disconnects China from the decoy ASes only. This happens because the majority of the Internet ASes have small customer cones (see Figure 2) and random placement is likely to choose many of these ASes.

When decoy ASes are selected from among the non-edge ASes, China’s connectivity drops significantly. For instance, for the random-no-ring-5 placement (i.e., choosing transit ASes with a minimum customer cone of 5), placing decoys in only 5% of global ASes disconnects China from around 43% of all Internet ASes, versus 7% for the random-no-ring-1 placement.

Figure 5 further shows that deploying decoys in the ring ASes of China amplifies the costs of the Chinese RAD attack. Another observation based on Figure 5 is that, while global decoy deployment is more effective, even regional deployment causes China to lose much of its connectivity.

⁷<http://iplane.cs.washington.edu/data/data.html>

Figure 5 also estimates popularity-weighted reachability after the RAD attack (Figures 5c and 5f). Each AS is weighted by the number of IP addresses that belong to it, and routes are weighted according to the weights of the ASes on the route.

In the rest of the simulations, we only consider the “no-ring” placements (i.e., we do not select decoy ASes from among the ring ASes).

B. Non-valley-free paths

The key technique suggested by the RAD paper is to re-route traffic between different adversary-controlled ASes in order to take advantage of more alternative routes (see Section III-B). As discussed in Section IV, routing through NVF paths is extremely costly. Figure 6 shows the percentage of paths that become NVF (the denominator includes only reachable destinations). In all cases, a large fraction of destinations are only reachable via NVF paths. Deploying decoys in ASes with larger customer cones amplifies this effect.

Table 7 shows the average number of Chinese transit ASes that must transit NVF traffic. This estimates how many links of the NVF paths are “inside the valley.”

C. Costly valley-free paths

We now demonstrate that even valley-free (VF) paths selected by the Chinese ASes as part of the RAD attack are more costly than the paths that would have been selected in the absence of the attack.

Using less-preferred paths (Business): Figure 8 shows the percentage of VF paths that become more expensive as a consequence of using RBGP (this is the Business cost described in Section IV). This ratio varies between 6% and 21% depending on the placement strategy.

Note that, in the case of `random-no-ring-1` placement, this ratio declines as the number of decoy ASes increases. The reason is that as the number of decoy ASes increases, more destinations are reachable only via (even costlier) NVF paths, as shown in Figure 6.

Longer paths (Length): In Section IV, we discussed the effects of longer paths on the quality of service. Figure 9 shows the percentage of VF paths that become longer when RBGP is used instead of BGP. This percentage varies between 20% and 43% depending on the placement strategy. The average increase in path length varies from 1.12 to 1.40.

Higher latencies (Latency): We now show that even when RBGP selects paths of the same length as the corresponding BGP paths, the RBGP paths are likely to have significantly higher latency. The reason for this increase is that RBGP paths are forced to use less popular transit ASes which have less network capacity (see Section IV).

To estimate latency, we use the following metric. For two neighbor ASes A and B , we define $eLat$ as:

$$eLat(A, B) = \frac{1}{n_A * n_B} \sum_{i=1}^{n_A} \sum_{j=1}^{n_B} Lat(A_i, B_j)$$

where A_i represents the i th point-of-presence (PoP) of the AS A and n_A is the number of A ’s PoPs. $Lat(X, Y)$ returns the

measured latency between two PoPs X and Y from iPlane’s “Inter-PoP links” dataset (see Section VI). For a BGP/RBGP path composed of k ASes $\{T_1, \dots, T_k\}$, we define $eLat$ to be the sum of $eLat$ for all neighbor ASes in the path:

$$eLat(\{T_1, \dots, T_k\}) = \sum_{i=1}^{k-1} eLat(T_i, T_{i+1})$$

The raw $eLat$ metric is a coarse estimate that may not represent the actual latency of a given path. That said, we can use the *relative increase* in $eLat$ due to the RAD attack, i.e., the ratio between $eLat$ for an RBGP path and $eLat$ for the corresponding BGP path, to estimate the increase in actual latency, without knowing the exact value of the former.

The iPlane dataset does not contain the latencies for every PoP pair and every AS. Therefore, we only estimate latencies for the paths where the latency of each individual link is available in the dataset.

Figure 10 show that the RAD attack causes a significant increase in the $eLat$ metric. For instance, for the `random-no-ring-1` placement (the random placement strategy considered in the RAD paper, with decoys placed in only 1% of ASes), launching the RAD attack makes the routes from China to Internet destinations over 4 times slower. The impact is even worse when decoys are placed more strategically and/or in more ASes.

The fluctuations in the graphs are caused by the limitations of the iPlane dataset, which prevent us from estimating latency for some of the paths (i.e., some of the paths chosen by Chinese ASes to avoid a particular decoy placement “disappear” from the measurements).

D. The need for infrastructural changes

Launching the RAD attack requires China to make dramatic changes to its network infrastructure.

Edge ASes acting as transit ASes (NewTransit): The RAD attack fundamentally assumes that all Chinese ASes are capable and willing to transit traffic for each other (see Section III-B). However, as discussed earlier, the majority of the Internet ASes are edge ASes and do not have the requisite network equipment and resources to transit other ASes’ traffic.

Our simulations show that the RAD attack requires many edge ASes to be converted into transit ASes, requiring huge re-organization and investment in their network infrastructure. China currently has 199 ASes, of which only 30 are transit ASes. Figure 11 shows the number of Chinese edge ASes that must become transit ASes in order to launch the RAD attack. For example, a `random-no-ring-1` placement in the `China-World` scenario with decoys in 2% of all ASes requires 59 edge ASes to be converted into transit ASes, almost doubling the number of transit ASes in China.

Converting a typical edge AS into a transit AS is highly non-trivial. Besides the monetary costs of purchasing and deploying new networking equipment, the organizational policies of edge ASes present significant obstacles. For example, would a university-owned ISP built for educational purposes or an ISP owned by a private, international company be willing—or even capable, if forced by the government—to act a transit AS?

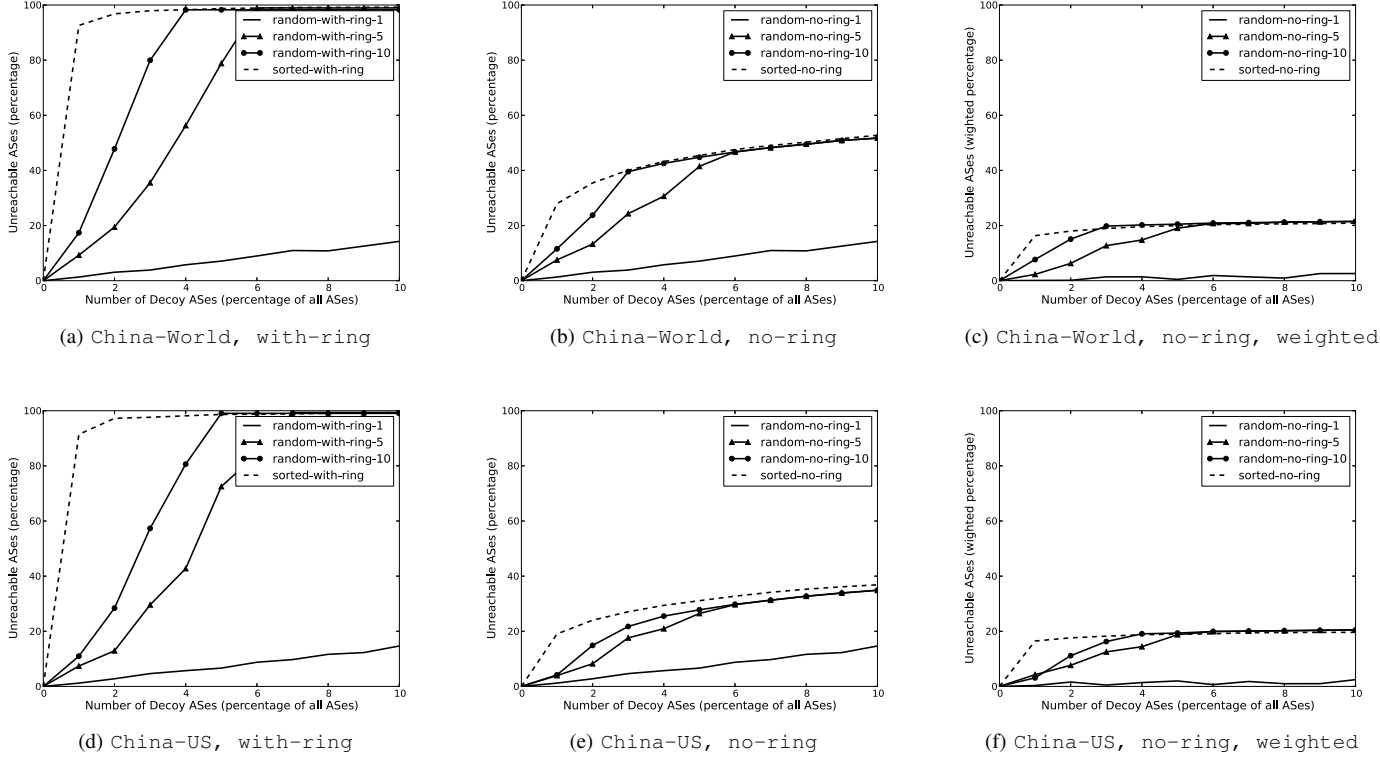


Fig. 5: The percentage of unreachable destination ASes.

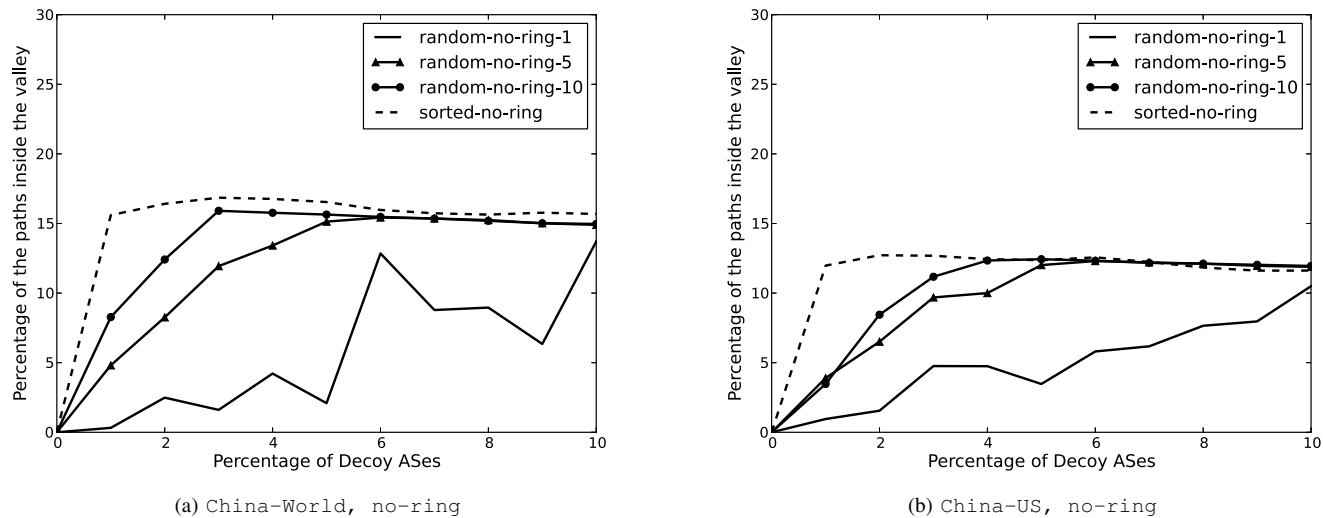


Fig. 6: The percentage of paths that become NVF due to the RAD attack.

Increased load on existing transit ASes (TransitLoad): Transit ASes are significantly affected by changes in their transit loads. Our simulations show that the RAD attack dramatically changes transit loads on many Chinese transit ASes. Since we only consider the traffic that leaves China, our estimates are conservative.

The information on traffic volumes between Internet ASes is not public. To simulate changes in transit loads, we assume that traffic volume between two ASes AS_1 and AS_2 is proportional to the number of IP addresses they respectively possess:

$$L(AS_1, AS_2) = IP_s(AS_1) \times IP_s(AS_2)$$


Fig. 7: The average path length inside the valley.

(a) China-World, no-ring


Placement/Percent	1	2	3	4	5	6	7	8	9	10
random-no-ring-1	1.84	1.99	2.01	1.81	1.88	1.89	1.88	1.81	1.96	2.00
random-no-ring-5	1.88	1.85	1.97	1.96	1.99	2.00	2.00	2.00	2.00	2.00
random-no-ring-10	1.98	1.95	1.99	1.99	1.99	2.00	2.00	2.00	2.00	2.00
sorted-no-ring	1.98	1.99	1.99	2.00	2.00	2.00	2.00	2.00	2.00	2.00

(b) China-US, no-ring

Placement/Percent	1	2	3	4	5	6	7	8	9	10
random-no-ring-1	1.92	1.93	1.92	1.89	1.88	1.87	1.92	1.84	1.96	1.92
random-no-ring-5	2.17	1.94	1.98	1.90	1.97	1.97	1.98	1.97	1.97	1.97
random-no-ring-10	1.84	2.01	1.91	1.97	1.97	1.97	1.98	1.97	1.97	1.98
sorted-no-ring	1.99	1.98	1.99	1.97	1.97	1.97	1.97	1.97	1.97	1.97



(a) China-World, no-ring



(b) China-US, no-ring

Fig. 8: The percentage of less-preferred paths due to the RAD attack.

where $IPs(A)$ is the number of IP addresses owned by the AS A .

We add $L(AS_1, AS_2)$ to the load of every transit AS on the path from AS_1 to AS_2 . In other words, we model the transit load of a transit AS as the sum of traffic volumes for all paths that cross this AS.

This model may not be accurate for some ASes since the higher number of IP addresses does not necessarily imply higher traffic volumes. However, it provides us with a simple estimate of transit loads in the absence of public data on actual traffic volumes. Furthermore, the inaccuracy is averaged across all paths, thus overestimates and underestimates cancel out to some extent.


Using this model for each Chinese transit AS A , we compute the *transit load increase* factor, which is the ratio of A 's transit load after the RAD attack over A 's transit load before the attack (we exclude traffic that does not leave China). Table IV shows the maximum value of the transit

load increase factor over all 30 transit ASes in China, for the China-World and China-US scenarios.


The RAD attack significantly increases loads on some transit ASes because they are forced to transit additional traffic, e.g., that of NVF paths. Some of the increases are so drastic that we believe it is extremely unlikely that existing transit ASes will be able to handle them. For example, assuming a random-no-ring-1 placement with decoys deployed on 2% of ASes in the China-World scenario, there is a Chinese transit AS that must transit roughly *122 times* more traffic due to the RAD attack.

Tables V and VI show the median transit load increase factor for the most affected 10% and 20% of transit ASes, respectively. The increase factor grows rapidly with the number of decoy ASes and with better decoy placements since both force Chinese ASes to route more traffic over NVF paths.

The RAD attack also causes some transit ASes to lose transit traffic, which is the source of their revenue. Table VII




(a) China-World, no-ring




(b) China-US, no-ring

Fig. 9: The percentage of VF paths with increased length.



(a) China-World, no-ring



(b) China-US, no-ring

Fig. 10: The average increase in estimated latency due to the RAD attack.


shows the minimum values of the transit load increase factor. For the `random-no-ring-1` placement, there is a transit AS that loses 30% of its transit load. Tables VIII and IX show the median and average changes in transit load, respectively.

Transit load does not increase monotonically with the number of decoy ASes. On the one hand, increasing the number of decoy ASes increases load imbalance and forces more traffic to shift to better-connected transit ASes. On the other hand, increasing the number of decoy ASes makes more destination ASes unreachable (see Figure 5) and thus reduces overall transit traffic. Furthermore, the results for the `random` simulations are reported for different, randomly selected decoy placements, which may have slightly different effects on the distribution of transit loads.


VIII. LESSONS AND RECOMMENDATIONS

1. The RAD attack proposed by Schuchard et al. [24] is **extremely costly** to the censors, even for the simple decoy placement considered in the RAD paper. The costs include collateral damage due to the loss of connectivity to many Internet destinations and much lower quality of service for the remaining destinations, monetary costs of buying and deploying new networking equipment needed to re-route massive amounts of traffic and convert edge ASes into transit ASes, and monetary costs caused by switching to less-preferred and, in particular, non-valley-free paths.

Even if the censors are willing to pay the monetary costs, evidence indicates that social costs may prevent them from



(a) China-World, no-ring



(b) China-US, no-ring

Fig. 11: The number of edge ASes that must become transit ASes.

TABLE IV: Maximum transit load increase factor for Chinese transit ASes due to the RAD attack.

(a) China-World, no-ring

Placement/Percent	1	2	3	4	5	6	7	8	9	10
random-no-ring-1	122.06x	2807.90x	807.97x	3388.97x	773.61x	14149.49x	3180.45x	3617.08x	3584.44x	9677.14x
random-no-ring-5	1718.21x	4588.29x	3402.40x	6418.70x	6338.64x	4688.07x	3972.97x	4173.69x	3128.00x	3030.92x
random-no-ring-10	1272.79x	4097.07x	5857.81x	3737.32x	4211.12x	4441.51x	4694.09x	3906.02x	3128.00x	2015.18x
sorted-no-ring	7744.57x	6507.31x	7895.25x	5814.86x	5850.94x	5864.12x	5125.12x	5117.52x	5075.41x	4920.45x

(b) China-US, no-ring

Placement/Percent	1	2	3	4	5	6	7	8	9	10
random-no-ring-1	294.73x	500.66x	1665.49x	1735.54x	1230.66x	1964.71x	2067.50x	2594.94x	2583.04x	3279.70x
random-no-ring-5	108.58x	3174.01x	3144.05x	409.45x	521.34x	3217.32x	422.18x	401.43x	388.16x	357.01x
random-no-ring-10	540.93x	472.35x	586.65x	596.57x	539.82x	3217.21x	432.20x	401.03x	379.72x	369.57x
sorted-no-ring	2474.72x	2499.81x	2502.29x	5269.66x	5269.66x	5270.44x	2978.76x	2965.68x	405.79x	398.96x

TABLE V: Median transit load increase factor for the most affected 10% of Chinese transit ASes due to the RAD attack.

(a) China-World, no-ring

Placement/Percent	1	2	3	4	5	6	7	8	9	10
random-no-ring-1	1.31x	2.26x	35.05x	394.80x	6.56x	106.29x	169.12x	105.93x	122.47x	47.60x
random-no-ring-5	215.27x	432.47x	1353.81x	1056.09x	887.89x	922.83x	922.83x	768.59x	728.39x	699.50x
random-no-ring-10	567.20x	1733.25x	1181.85x	1058.98x	957.31x	917.58x	882.66x	866.08x	728.81x	703.36x
sorted-no-ring	1933.21x	1748.12x	1697.72x	1616.68x	1540.24x	1499.73x	1457.66x	1440.96x	1428.41x	1723.51x

(b) China-US, no-ring

Placement/Percent	1	2	3	4	5	6	7	8	9	10
random-no-ring-1	2.31x	1.74x	2.51x	4.08x	14.25x	28.58x	241.57x	103.49x	27.23x	11.79x
random-no-ring-5	294.66x	159.13x	164.61x	483.25x	488.71x	446.56x	225.57x	108.33x	94.48x	96.31x
random-no-ring-10	261.42x	194.69x	281.52x	276.90x	542.18x	442.66x	430.17x	108.33x	105.63x	102.81x
sorted-no-ring	1426.64x	1353.49x	1334.47x	1356.43x	1345.60x	1329.11x	461.33x	426.44x	82.77x	82.77x

deploying disruptive censorship technologies. For example, the Great Firewall of China does not block many popular Internet services even though they are encrypted,⁸ due to their popularity among Chinese Internet users.

2. A strategic placement of decoy routers significantly raises

the costs for the RAD adversary. We propose several strategic decoy placement strategies.

3. The costs of the RAD attack vary significantly for **different state-level censors**. Countries with less Internet connectivity (i.e., those that have fewer internal ASes and are connected to fewer ring ASes) incur higher costs if they launch the RAD

⁸<http://www.google.com/transparencyreport/traffic/>

TABLE VI: Median transit load increase factor for the most affected 20% of Chinese transit ASes due to the RAD attack.

(a) China-World, no-ring

Placement/Percent	1	2	3	4	5	6	7	8	9	10
random-no-ring-1	1.00x	1.00x	1.08x	51.82x	1.05x	1.21x	1.10x	1.84x	5.73x	1.54x
random-no-ring-5	2.35x	1.74x	230.34x	3.87x	3.30x	3.31x	3.31x	3.46x	2.82x	2.82x
random-no-ring-10	1.41x	303.87x	3.51x	3.52x	3.39x	3.48x	2.88x	2.84x	2.82x	2.80x
sorted-no-ring	443.83x	397.87x	369.17x	348.17x	320.85x	312.49x	275.50x	270.01x	267.43x	350.41x

(b) China-US, no-ring

Placement/Percent	1	2	3	4	5	6	7	8	9	10
random-no-ring-1	1.01x	1.42x	1.01x	1.06x	6.96x	1.70x	102.50x	13.52x	1.18x	1.46x
random-no-ring-5	11.02x	20.30x	32.22x	35.48x	44.76x	32.12x	29.97x	22.55x	21.70x	21.07x
random-no-ring-10	33.23x	69.44x	51.74x	49.30x	42.17x	30.96x	23.12x	22.55x	21.21x	20.09x
sorted-no-ring	68.51x	67.28x	61.12x	66.04x	66.92x	64.93x	39.30x	31.68x	29.94x	29.11x

TABLE VII: Minimum transit load increase factor for Chinese transit ASes due to the RAD attack.

(a) China-World, no-ring

Placement/Percent	1	2	3	4	5	6	7	8	9	10
random-no-ring-1	0.98x	0.67x	0.94x	0.71x	0.62x	0.31x	0.46x	0.60x	0.49x	0.29x
random-no-ring-5	0.84x	0.93x	0.81x	0.70x	0.68x	0.65x	0.65x	0.65x	0.65x	0.65x
random-no-ring-10	0.88x	0.82x	0.67x	0.66x	0.66x	0.66x	0.65x	0.65x	0.65x	0.64x
sorted-no-ring	0.77x	0.76x	0.74x	0.73x	0.73x	0.72x	0.72x	0.72x	0.72x	0.71x

(b) China-US, no-ring

Placement/Percent	1	2	3	4	5	6	7	8	9	10
random-no-ring-1	0.89x	0.77x	0.71x	0.69x	0.69x	0.74x	0.63x	0.67x	0.58x	0.53x
random-no-ring-5	0.63x	0.70x	0.65x	0.78x	0.62x	0.59x	0.58x	0.57x	0.57x	0.57x
random-no-ring-10	0.91x	0.66x	0.60x	0.59x	0.60x	0.59x	0.58x	0.57x	0.57x	0.57x
sorted-no-ring	0.63x	0.62x	0.61x	0.61x	0.60x	0.60x	0.60x	0.58x	0.58x	0.58x

TABLE VIII: Median transit load increase factor for Chinese transit ASes due to the RAD attack.

(a) China-World, no-ring

Placement/Percent	1	2	3	4	5	6	7	8	9	10
random-no-ring-1	1.00x	1.00x	1.00x	1.00x	1.00x	1.00x	0.99x	1.00x	0.98x	0.99x
random-no-ring-5	1.00x	0.98x	0.98x	0.97x	0.95x	0.94x	0.94x	0.90x	0.90x	0.89x
random-no-ring-10	0.97x	0.98x	0.95x	0.95x	0.94x	0.94x	0.90x	0.90x	0.90x	0.89x
sorted-no-ring	0.98x	1.00x	1.00x	1.00x	1.00x	0.99x	0.99x	0.99x	0.95x	0.95x

(b) China-US, no-ring

Placement/Percent	1	2	3	4	5	6	7	8	9	10
random-no-ring-1	1.00x	0.99x	1.00x	0.99x	0.99x	1.00x	0.99x	1.00x	0.99x	0.99x
random-no-ring-5	0.99x	0.97x	0.95x	0.91x	0.90x	0.87x	0.86x	0.86x	0.85x	0.84x
random-no-ring-10	0.98x	0.95x	0.90x	0.88x	0.88x	0.87x	0.86x	0.86x	0.85x	0.84x
sorted-no-ring	0.99x	0.97x	0.97x	0.95x	0.95x	0.95x	0.88x	0.84x	0.84x	0.84x

TABLE IX: Average transit load increase factor for Chinese transit ASes due to the RAD attack.

(a) China-World, no-ring

Placement/Percent	1	2	3	4	5	6	7	8	9	10
random-no-ring-1	1.08x	1.54x	6.41x	61.24x	2.50x	25.49x	23.52x	52.67x	45.09x	19.66x
random-no-ring-5	33.40x	54.69x	199.41x	150.03x	254.56x	197.00x	197.00x	179.49x	144.25x	139.41x
random-no-ring-10	136.41x	248.79x	257.97x	187.01x	191.15x	194.39x	162.98x	173.49x	144.28x	96.10x
sorted-no-ring	378.03x	326.33x	365.64x	294.90x	290.39x	288.00x	261.12x	259.66x	257.47x	273.67x

(b) China-US, no-ring

Placement/Percent	1	2	3	4	5	6	7	8	9	10
random-no-ring-1	1.54x	2.74x	5.02x	13.55x	28.01x	18.33x	18.89x	30.23x	17.72x	25.68x
random-no-ring-5	15.13x	68.83x	110.74x	142.47x	133.46x	125.19x	72.50x	19.27x	18.29x	17.53x
random-no-ring-10	16.06x	57.49x	41.76x	33.86x	55.61x	125.26x	73.06x	19.74x	18.50x	17.48x
sorted-no-ring	135.88x	134.16x	133.20x	226.48x	226.16x	225.45x	118.10x	115.96x	19.02x	18.74x

attack. This implies that even a very limited deployment of decoy routers may be enough to deter relatively small state-level censors such as Syria from launching the attack.

4. While a global deployment of decoy routing is ideal (i.e., the China-World scenario), even a **regional deployment** (e.g., only in the U.S., as in the China-US scenario) is effective in defeating the RAD attack. This is an important finding because regional deployment is more practical than global deployment. For example, the U.S. government may mandate or incentivize U.S.-based ASes to deploy decoy routers to support the freedom of Internet in Syria.

5. Any **real-world deployment** of decoy routing systems requires decoys to be installed in multiple ASes. The networking community has faced similar challenges with the adoption of new networking protocols and technologies. Their solutions [12, 13] can be adapted to the problem of decoy routing.

In particular, techniques proposed for deploying secure BGP protocols may provide an inspiration. Gill et al. [12] suggest an initial deployment by “early-adopter” ASes who are incentivized by third parties. This initial deployment will eventually lead to a competition among ASes to install the new technology, as they aim to increase their revenue-generating traffic. Similarly, an initial deployment of decoy routers on a small number of transit ASes, perhaps incentivized by pro-freedom NGOs or governments, can “diffuse” decoy routing to other transit ASes who want to capture a share of the decoy routing traffic.

6. A **fine-grained, data-driven approach** is necessary for understanding the true costs of various route selection mechanisms. Analysis based solely on the graph topology of the Internet may lead to mistaken conclusions about the feasibility of decoy routing, as well as other censorship circumvention techniques based on Internet routing. Any analysis of decoy routing and alternatives must be based on all available data about individual nodes and links in the Internet connectivity graph.

ACKNOWLEDGMENTS

This research was supported by the Defense Advanced Research Projects Agency (DARPA) and SPAWAR Systems Center Pacific, Contract No. N66001-11-C-4018, and NSF grant CNS-0746888.

REFERENCES

- [1] “AS relationships,” <http://www.caida.org/data/active/as-relationships/>.
- [2] S. Burnett, N. Feamster, and S. Vempala, “Chipping away at censorship firewalls with user-generated content,” in *USENIX Security*, 2010.
- [3] “AS rank: AS ranking,” <http://as-rank.caida.org/>.
- [4] J. Cesareo, J. Karlin, J. Rexford, and M. Schapira, “Optimizing the placement of implicit proxies,” <http://www.cs.princeton.edu/~jrex/papers/decoy-routing.pdf>, 2012.
- [5] I. Clarke, T. W. Hong, S. G. Miller, O. Sandberg, and B. Wiley, “Protecting free expression online with Freenet,” *IEEE Internet Computing*, vol. 6, no. 1, pp. 40–49, 2002.

- [6] T. Dierks and E. Rescorla, “The Transport Layer Security (TLS) protocol — version 1.2,” Internet RFC 5246, 2008.
- [7] R. Dingledine and N. Mathewson, “Design of a Blocking-Resistant Anonymity System,” <https://svn.torproject.org/svn/projects/design-paper/blocking.html>.
- [8] R. Dingledine, N. Mathewson, and P. Syverson, “Tor: The second-generation onion router,” in *USENIX Security*, 2004.
- [9] N. Feamster, M. Balazinska, G. Harfst, H. Balakrishnan, and D. Karger, “Infranet: Circumventing Web censorship and surveillance,” in *USENIX Security*, 2002.
- [10] L. Gao and J. Rexford, “Stable Internet routing without global coordination,” *IEEE/ACM TON*, vol. 9, no. 6, pp. 681–692, 2001.
- [11] L. Gao, “On inferring autonomous system relationships in the Internet,” *IEEE/ACM ToN*, vol. 9, no. 6, pp. 733–745, 2001.
- [12] P. Gill, M. Schapira, and S. Goldberg, “Let the market drive deployment: A strategy for transitioning to BGP security,” in *SIGCOMM*, 2011.
- [13] S. Goldberg and Z. Liu, “The diffusion of networking technologies,” in *SODA*, 2013.
- [14] J. Hawkinson and T. Bates, “Guidelines for creation, selection, and registration of an autonomous system (AS),” 1996.
- [15] A. Houmansadr, G. Nguyen, M. Caesar, and N. Borisov, “Cirripede: Circumvention infrastructure using router redirection with plausible deniability,” in *CCS*, 2011.
- [16] A. Houmansadr, T. Riedl, N. Borisov, and A. Singer, “I Want My Voice to Be Heard: IP over Voice-over-IP for Unobservable Censorship Circumvention,” in *NDSS*, 2013.
- [17] J. Karlin, D. Ellard, A. Jackson, C. Jones, G. Lauer, D. Mankins, and W. Strayer, “Decoy routing: Toward unblockable Internet communication,” in *FOCI*, 2011.
- [18] N. Kushman, S. Kandula, D. Katabi, and B. Maggs, “R-BGP: Staying connected in a connected world,” in *NSDI*, 2007.
- [19] H. V. Madhyastha, T. Isdal, M. Piatek, C. Dixon, T. Anderson, A. Krishnamurthy, and A. Venkataramani, “iPlane: An information plane for distributed services,” in *OSDI*, 2006.
- [20] “Psiphon,” <http://psiphon.ca/>.
- [21] B. Quoitin and S. Uhlig, “Modeling the routing of an autonomous system with C-BGP,” *IEEE Network*, vol. 19, no. 6, pp. 12–19, 2005.
- [22] Y. Rekhter, T. Li, and S. Hares, “A Border Gateway Protocol 4 (BGP-4),” RFC 4271, 2006.
- [23] E. Rosen and Y. Rekhter, “BGP/MPLS IP Virtual Private Networks (VPNs),” RFC 4364 (Proposed Standard), 2006.
- [24] M. Schuchard, J. Geddes, C. Thompson, and N. Hopper, “Routing around decoys,” in *CCS*, 2012.
- [25] Q. Wang, X. Gong, G. Nguyen, A. Houmansadr, and N. Borisov, “CensorSpoof: Asymmetric communication using IP spoofing for censorship-resistant Web browsing,” in *CCS*, 2012.
- [26] E. Wustrow, S. Wolchok, I. Goldberg, and J. Halderman, “Telex: Anticensorship in the network infrastructure,” in *USENIX Security*, 2011.