Comparing the Security Performance of Network-Layer and Application-Layer Anycast

Martin Suchara
Princeton University
Department of Computer Science
Princeton, NJ 08544
msuchara@princeton.edu

Ioannis Avramopoulos
Deutsche Telekom Laboratories
Ernst-Reuter-Platz 7
10587 Berlin, Germany
ioannis.avramopoulos@telekom.de

Abstract

We provide a theoretical analysis of the security performance of two anycast techniques that could be used as a countermeasure against DNS attacks exploiting vulnerabilities in the interdomain routing system. We argue that that the performance of the two techniques – network and ideal application layer anycast – does not differ in practice. This is achieved by showing that the performance can only differ if a family of special subgraphs that we characterize appears in the interdomain network topology. Our result supports our earlier experimental findings. While experimentation will remain a crucial method to accurately evaluate the behavior of complex routing systems in the future, we hope that analysis such as this one can help to understand and design routing protocols with better security, reliability and performance properties.

1. Introduction

In this paper, we analytically compare the security performance against routing attacks of two alternative anycast implementations, one based on the network-layer and the other based on the application layer. In network-layer anycast [1], two or more servers are assigned the same IP address and the routing system is responsible for forwarding traffic from clients destined to this address to any of the servers. Network-layer anycast is typically implemented through the simultaneous origination of an IP prefix from multiple autonomous systems. In application-layer anycast [2], an endhost serving as a client must engage in an application-layer interaction to obtain one or more IP addresses of corresponding servers. Different from network-layer anycast in which all servers share the same address, in application-layer anycast each server has a separate address.

We consider two alternative application-layer anycast implementations, an *ideal* one and a *naïve* one. Using the comparison with the ideal implementation we demonstrate that the security performance of network-layer anycast is an upper bound on the security performance of any application-layer anycast implementation. We use the evaluation of the naïve implementation to demonstrate that application-layer anycast can do much worse in practice.

The motivation for this analytic comparison is an experimental finding that network-layer and ideal application-layer anycast have indistinguishable security performance [3]. A deeper understanding of this experimental finding can be used to guide practical design decisions for the domain name system (DNS) since the *backbone* of DNS is implemented as a combination of

network-layer and application-layer anycast. In addition to explaining the previous results, a contribution of this paper is an interesting analytic method for comparing the security performance of different route propagation techniques.

In section 2 we introduce terminology, formalize the operation of the interdomain routing system, and describe the operation of three anycast variants in our routing model: network-layer anycast (NLA), ideal application-layer anycast (ALA), and naïve application-layer anycast (Naïve-ALA). In section 3 we demonstrate that the security performance of NLA and ALA may differ. We show an example of a topology where a particular autonomous system uses a secure route in NLA and a malicious route in ALA. We show another example where a malicious route is used in NLA and a secure one in ALA. In section 4, we explain why the security performance of NLA and ALA is nearly identical in practice. Finally, in Section 5, we show that the fraction of autonomous systems in the network that possess a secure route in NLA (and hence also in ALA) is x/(x+y) where x is the number of distinct autonomous systems hosting anycast servers and y is the number of autonomous systems controlled by the adversary. We also explain why the fraction of the network secured by Naïve-ALA is only 1/(1+y). The results in section 5 match the experimental observations in [3].

2. Terminology and the Route Propagation Mechanisms

In this section we provide an overview of the key properties of the interdomain routing system, as well as a description of the three anycast variants whose performance we evaluate. This section also introduces the terminology used throughout the paper.

2.1 Model of the Routing System

The goal of the routing system is to establish IP prefix reachability in the network. The network consists of autonomous systems (AS) which export routes based on their business relationships [4]. The Internet can be modeled at the AS level as a graph that consists of nodes representing the ASes, and edges that are annotated by the business relationships of the neighboring ASes. These business relationships are *customer-to-provider, provider-to-customer,* and *peer-to-peer*.

A particular route r is a *customer route* for AS1 if the node that exports the route r to AS1 is its customer. It is a *peer or provider route* if the node is a peer or provider respectively. For convenience, our figures use horizontal edges to denote peer-to-peer edges and vertical or diagonal edges to denote customer-to-provider edges where the provider always appears higher than its customer.

When an AS learns multiple routes for the same prefix from its neighbors, a decision process is used to select the best route. The routes are ordered as follows. First, customer routes are preferred to peer routes which are preferred to provider routes. If the resulting ordering is not total, shorter routes are preferred to longer ones. Finally, if the ordering is still not total, a route originated by an AS with the smallest AS number is preferred. An AS *selects* the most preferred route. If the selected route is a customer route it is *exported* to all neighbors. An AS exports a selected provider or peer route only to its customers. As a result of these export policies, routes are valley free [4].

2.2 Secure and Insecure Nodes

Some of the nodes in the graph host legitimate servers, and other nodes host clients. The goal of the clients is to contact the servers. We define *white nodes* as the nodes that host the legitimate servers. Another subset of nodes represents the autonomous systems that are controlled by an adversary, and try to prevent the clients from contacting the honest servers. These nodes are referred to as *black nodes*. The black nodes can succeed by making false routing announcements of the address prefixes owned by the white nodes in an attempt to attract the client requests addressed to the servers of the white nodes. If a black node announces an address prefix and another node selects such a route, its clients will be unable to reach any legitimate server with address in that prefix range. This is because each node selects a single route per prefix.

A *white* and a *black route* are routes announced by a white and a black node respectively. We define a *secure node* as a node whose clients are able to access a legitimate server by using a white route, and an *insecure node* as a node whose clients are unable to contact a legitimate server because the node uses a black route. The *security performance* of an anycast variant is defined as the fraction of secure nodes in the network. Next, we will formalize the operation of network and application-layer anycast.

2.3 Network and Application-layer Anycast

In *network-layer anycast (NLA)* the same IP prefix is announced by each white and each black node. The route propagation and decision process described above is used to establish prefix reachability in the network. If a particular node selects a route originated by a black node, it is insecure. If the node selects a route originated by a white node, it is secure. Therefore, the factors that influence the security performance are the number and location of the black and white nodes. By considering a random selection of the location of the black and white nodes, we are able to estimate the average security performance over all possible selections.

In *ideal application-layer anycast (ALA)*, each white node announces a unique IP prefix. We assume that each black node announces all the prefixes announced by all the white nodes. The route propagation and decision process described earlier is used to establish prefix reachability in the network. It follows that if the network is connected, every node in the topology selects one route for each unique IP prefix. We say that the node is secure if it selects at least one white route, and it is insecure if all the selected routes are black. This is justified by an assumption that the nodes in application-layer anycast are equipped with an oracle[1] that distinguishes black and white routes with perfect accuracy [3].

In *naïve application-layer anycast (Naïve-ALA)*, the prefix announcements and route propagations are done in the same way as in the ideal application-layer anycast. However, nodes are not equipped with an oracle, and instead select routes among the available ones at random. Therefore, a node is secure if the route it selects among the available ones is white, and insecure if it is black.

[1] We note that the nodes do not use the knowledge gained from the oracle to modify the route propagation mechanism, i.e., a route can be announced even if the application-layer oracle allows the node to detect that it is malicious. This is motivated by the fact that the application-layer oracle may be able to use more information than what is available during the route decision process at the network layer.

To aid with the theoretical analysis of ideal application-layer anycast, a technique that requires announcement and propagation of a number of different IP prefixes, we introduce the following terminology. The route propagation where a single white node W1 announces some IP prefix p1 and all black nodes announce the same prefix p1 is denoted as *ALA-W1*. A node X is secure in ALA if there exists a white node W belonging to the set of white nodes such that X selects a white route in ALA-W, and insecure if such node W doesn't exist. Similarly, node X is secure in Naïve-ALA if for a node W randomly selected among the set of white nodes, X selects a white route in ALA-W, and insecure otherwise.

3. Security of Network and Application Layer Anycast can Differ

First we show that the security performance of ideal application-layer anycast (ALA) is not always better than the security performance of network-layer anycast (NLA). To show this, it suffices to find an example of a topology where some node X selects a black route in ALA and a white route in NLA. An example of such a topology is depicted in Figure 1. Nodes B1 and B2 are black nodes, and nodes W1 and W2 are white nodes. In this example we assume that the AS numbers of the nodes are ordered as follows: B1 < W1 and B2 < W2.

In NLA a single prefix is announced by W1, W2, B1 and B2. Node Z picks the customer route originated by W2 which is exported to node X. Finally, node X picks the shortest available provider route, which is the route originated by W1, i.e., a white route. In ALA-W1, only nodes W1, B1 and B2 originate a certain prefix. Z picks the route originated by B1, and X chooses between the two shortest routes of B1 and W1. The black route B1 is chosen because B1 < W1. Similarly in ALA-W2 only nodes W2, B1 and B2 originate some other prefix. Z picks the route of W2, and X picks the route of B2 because B2 < W2. Therefore, X does not have a white route in ALA.

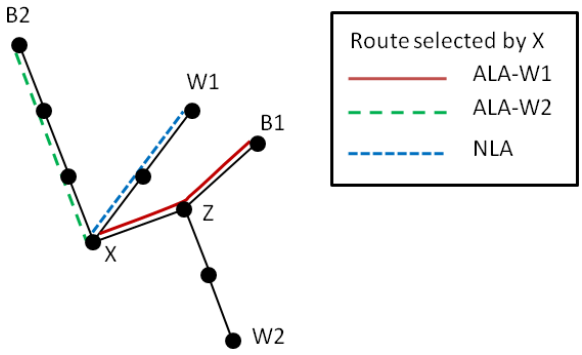

Fig 1: Node X selects a black route in ALA-W1 and ALA-W2, but it selects a white route in NLA. The AS numbers are ordered as follows: B1 < W1 and B2 < W2.

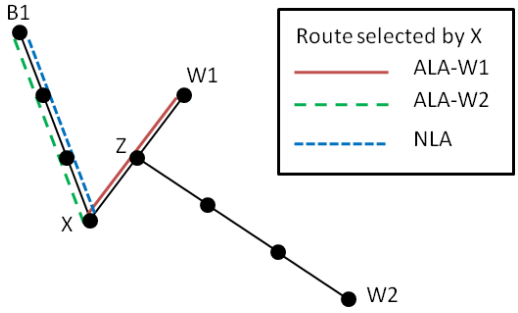

Fig 2: Node X selects a black route in NLA but it selects a white route in ALA-W1.

On the other hand, the security performance of network-layer anycast (NLA) is not always better than the security performance of ideal application-layer anycast (ALA) for a particular AS in some network. In this case, it suffices to find an example of a topology where some node X picks a white route in ALA and a black route in NLA. An example of such a topology is in Figure 2. Node B1 is a black node, and nodes W1 and W2 are white nodes.

In NLA a single prefix is announced by W1, W2, and B1. Node Z picks the customer route of W2 which is exported to node X. Node X then picks the shorter provider route of B1, i.e., a black route. In ALA-W1, only nodes W1, and B1 originate a certain prefix. Node Z picks the route of W1, and X also picks the shorter route of W1, i.e., a white route. Therefore, X has a white route in ALA.

The differences in security performance on specific topologies are not surprising given the qualitative differences between NLA and ALA. However, this contrasts with the experimental findings in [3] which demonstrate that the security performance of the two schemes is nearly identical on a realistic model of the Internet topology. We reconcile this seeming disparity in the next section by observing that one of a family of specific subgraphs has to appear in the topology for the security performance to differ. After characterizing these esoteric subgraphs it becomes clear that the number of nodes whose security performance differs will be small in any realistic topology.

4. Characterizing Security Performance Differences in NLA and ALA

In this section we prove that if the security performance of a node differs in NLA and ALA, then one of a family of special subgraphs must appear in the topology. We further argue that the likelihood that such a subgraph appears in a realistic topology is small, and therefore the security performance of NLA and ALA is in practice nearly identical.

Theorem 1: If the security performance of some node is worse in ALA than in NLA, then the topology must contain the following subgraph. There exist nodes X, Z, white nodes W1, W2, and black node B1, and valley-free provider routes r_W1, r_W2, r_B1 from node X to nodes W1, W2, B1 respectively such that:

(1) r_W1, r_W2 and r_B1 are routes of length at least two, r_W2 and r_B1 are imported to X from the same provider and r_W1 is imported from another distinct provider
(2) either r_W1 is shorter than r_W2 or the length of the two routes is the same and W1 < Y
(3) either r_B1 is shorter than r_W1 or the length of the two routes is the same and B1 < W1
(4) node Z is the last common node on the routes r_B1 and r_W2 from X to B1 and W2, and node Z prefers the part of the route r_W2 from Z to W2 to the part of r_B1 from Z to B1
(5) there exists valley-free route r_B from X to some black node (that may or may not be node B1), and the route is imported to X from some other provider than route r_B1

Remark 1: The existence of a single instance of this subgraph implies that the security performance only differs for node X, and the peers and customers of X that accept its routes. In addition to satisfying the conditions above, the security performance can only differ for node X, and the peers and customers of X that accept its routes.

Remark 2: It follows from the statement of Theorem 1 that the likelihood that a subgraph satisfying all conditions (1) through (5) simultaneously appears in a topology with a random assignment of black and white nodes is small. Moreover, appearance of these subgraphs only affects the security performance of node X and the customers and peers of X that accept its routes. Therefore, the performance of ALA is rarely worse than that of NLA in practice. Some

examples of subgraphs satisfying conditions (1) through (5) that result in a worse security performance of node X in ALA than in NLA are shown in Figure 3.

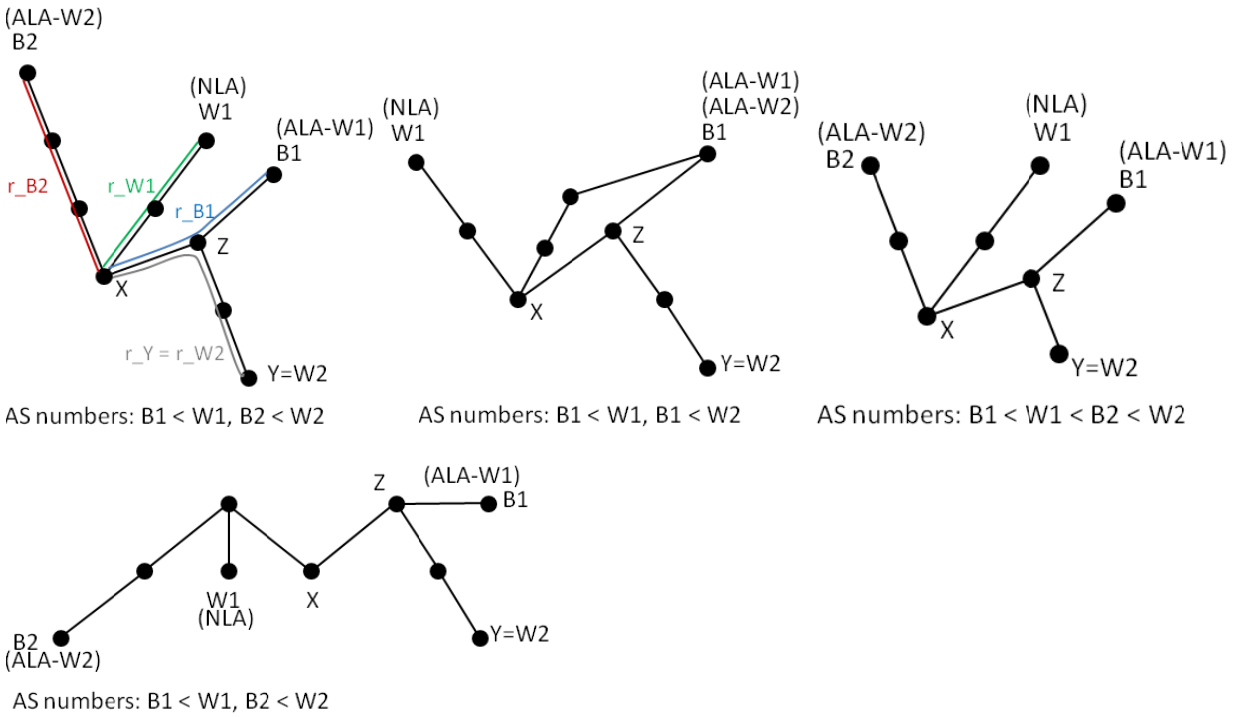

Fig 3: Examples of subgraphs with worse performance in ALA than in NLA. Anycast name in parenthesis indicates origin of the route chosen by X in that anycast variant.

Lemma 1.1: If a node imports a route from some neighbor in any of the experiments ALA-W1, ALA-W2, etc. then it imports some route from the same neighbor in NLA.
Proof: This is shown for each of the cases ALA-W1, ALA-W2, etc. separately by induction on the length of the routes chosen in that case. □

Lemma 1.2: If node X imports route r_V originated by node V in NLA, and if node V originates a prefix in ALA-W, then X imports route r_V or a more preferred route for the same prefix from the same neighbor in ALA-W.
Proof: Since node V originates a prefix in ALA-W, we observe that each successive node on the original route r_V receives some route for the prefix from the same neighbor from which it received route r_V in NLA. We also note that the received route is not necessarily the most preferred route and the node may select and export a different, more preferred, route. However, a route cannot be received from a neighbor with a better business relationship than in NLA (from Lemma 1.1, suppose such a route was received, then r_V would not be available to X in NLA). Therefore, the route that propagates along r_V must be either r_V, or a route shorter than the original r_V, or the origin's AS number must be lower. This route must be eventually imported by X in ALA-W. □

Proof of Theorem 1:
To prove the theorem it suffices to prove conditions (i) through (vii) below. Our approach is to prove that if the performance for node X is worse in ALA than in NLA, then conditions (i) through (vii) must be satisfied.

(i) node X in the subgraph selects a white route r_W1 that is originated by white node W1 in NLA and a black route r_B1 originated by black node B1 in ALA-W1,
(ii) both routes r_W1 and r_B1 are routes of length at least two imported to X from two distinct neighbors with the same business relationship with X,
(iii) node X receives a white route r_W2 to node W2 ≠ W1 in NLA from the same neighbor from which it receives r_B1 in ALA-W1, i.e., in the subgraph there must be some node Z on r_B1 that in NLA prefers the valley free route r_W2 to the route received on r_B1, and yet node X in the subgraph prefers r_W1 to r_W2 in NLA,
(iv) routes r_W1, r_B1 and r_W2 are provider routes for node X,
(v) either r_W1 is shorter than r_W2 or the length of the two routes is the same and W1 < Y,
(vi) either r_B1 is shorter than r_W1 or the length of the two routes is the same and B1 < W1,
(vii) node X receives a black provider route in ALA-W2 that is imported from a different neighbor than r_B1, and must be preferred by X to r_W2.

Condition (i):
We claim that if the performance in ALA is worse for some node X, there WLOG exist two white nodes W1 and W2, node X selects a white route in NLA that is originated by W1 (let's denote the route r_W1) and selects a black route originated by B1 in ALA-W1 (let's denote the route r_B1). If there is just one white node or any other part of the claim is not satisfied, the performance of ALA and NLA cannot differ. In the remainder of the proof, let's assume that the performance of the direct provider of X does not differ. This can be achieved by renaming the provider as X.

Condition (ii):
First we show that r_W1 and r_B1 are exported to X from two nodes with the same business relationship with X, and not from a single node. We show that the other cases cannot arise. The relationship with the node from which r_W1 is exported in NLA cannot be less preferred than r_B1 because some route is exported to X from the first hop on r_B1 in NLA (Lemma 1.1) and then r_W1 would not be selected in NLA. The case where the relationship with the node from which r_W1 is exported is more preferred also cannot arise because some route is exported to X from the first hop on r_W1 in ALA-W1 (Lemma 1.2), but then X doesn't choose r_B1 in ALA-W1. Finally, r_W1 and r_B1 cannot be exported from the same node because then that node has to select route r_W1 in NLA and r_B1 in ALA-W1, contradicting our assumption that the performance of the provider of X does not differ.

Next we show that the distance of B1 and X must be at least 2 hops, and the distance of W1 and X must be at least 2 hops. Assume that B1 and X are neighbors. Then B1 exports route r_B1 directly to X both in NLA and ALA-W1. Since r_B1 and r_W1 are routes with the same business relationship with X, route r_W1 can only be preferred in NLA if W1 and X are also neighbors and W1 < B1. But then r_B1 cannot be preferred in ALA-W1. We can use a similar argument to lowerbound the distance of W1 and X.

Condition (iii):
First we show that node X receives from its two neighbors route r_W1 and route r_Y ≠ r_B1 to some node Y in NLA. Therefore, route r_Y must be selected by some node Z on the original route r_B1, and X prefers r_W1 to r_Y, and Z prefers r_Y to the route received on r_B1. This is shown by contradiction. If the case above does not arise, then node X receives routes r_W1 and r_B1 and prefers r_W1 in NLA. Therefore, the performance in NLA and ALA-W1 can only differ because node X receives routes r_B1 and r_Y ≠ r_W1 in ALA-W1. Moreover, X must prefer r_W1 to r_Y. However, this contradicts Lemma 1.2.

Next, we show that node Y is a white node W2 distinct from W1. Suppose Y is a black node. By property (iii) node Z receives route r_Y in NLA, and then by Lemma 1.2, Z also receives r_Y or a more preferred route in ALA-W1. But then Z cannot choose r_B1 in ALA-W1. Repeating the same argument for node Z that receives route from node Y = W1 we conclude Y ≠ W1.

Condition (iv):
We show that r_W1, r_B1 and r_W2 are provider routes for X. Routes r_B1 and r_W2 are exported to X from the same neighbor, and by property (ii) r_W1 and r_B1 are exported to X from two nodes with the same business relationship. Therefore, it suffices to show that r_B1 is a provider route. Suppose not. Then route export policies imply both r_B1 and r_W2 are customer routes for Z. We know that node Z prefers r_W2 to the route received on r_B1 in NLA, i.e., r_W2 is shorter than the route received on r_B1, or the AS number of W2 is lower. We also know that the route received on r_B1 must be r_B1, or be shorter than r_B1, or have the same length but a lower AS number (it is a route that consists of only customer edges, and if the route received in NLA differs from the route in ALA-W1, some node on the original r_W1 must prefer the new route, i.e., it must be shorter or must have a lower origin AS number). Then X must prefer r_W2 to r_B1. We also know that X prefers r_W1 to r_W2, so it must prefer r_W1 to r_B1. But according to Lemma 1.2, X receives route r_W1 or a more preferred route in ALA-W1, contradicting the fact that X selects route r_B1 in ALA-W1.

Condition (v):
We claim that either r_W1 is shorter than r_W2, or the length of the two routes is the same and W1 < W2. The statement is true because X prefers r_W1 to r_W2 and by property (ii) both routes are imported to X from a neighbor with the same business relationship.

Condition (vi):
We claim that either r_B1 is shorter than r_W1, or the length of the two routes is the same and B1 < W1. X prefers r_B1 to the route received on r_W1 in ALA-W1. By Lemma 1.2 the route received on r_W1 is either r_W1 or a more preferred route. The claim follows because both routes are imported to X from a neighbor with the same business relationship by property (ii).

Condition (vii):
Finally, we show that node X must receive a black provider route in ALA-W2. Moreover, the route must be imported to X from a different neighbor than the original route r_B1, and must be preferred by X to r_W2. By Lemma 1.2 node X receives r_W2 or a more preferred route in ALA-W2, and thus a more preferred black route needs to be received by X from another neighbor. □

The case when the performance of NLA is worse than that of ALA is addressed by Theorem 2.

Theorem 2: If the security performance of some node is worse in NLA than in ALA, then the topology must contain the following subgraph. There exist nodes X, Z, white nodes W1, W2, and black node B1, and valley-free provider routes r_W1, r_W2, r_B1 from node X to nodes W1, W2, B1 respectively such that:

(1) r_W1, r_W2 and r_B1 are routes of length at least two, r_W1 and r_W2 are imported to X from the same provider and r_B1 is imported from another distinct provider,
(2) either r_B1 is shorter than r_W2 or the length of the two routes is the same and B1 < Y,
(3) either r_W1 is shorter than r_B1 or the length of the two routes is the same and W1 < B,
(4) node Z is the last common node on the routes r_W1 and r_W2 from X to W1 and W2, and node Z prefers the part of the route r_W2 from Z to W2 to the part of r_W1 from Z to W1.

Remark: Similarly as with Theorem 1, the the likelihood that a subgraph will satisfy conditions (i) through (vi) simultaneously is small, and only the performance of node X and its customers and peers may differ. Therefore, the performance of NLA is rarely worse than that of ALA in practice. However, because the analog of property (vii) is missing, the requirements of Theorem 2 are less stringent than the requirements of Theorem 1. Examples of subgraphs satisfying conditions (i) through (vi) of Theorem 2 that result in a worse security performance of node X in NLA than in ALA are shown in Figure 4.

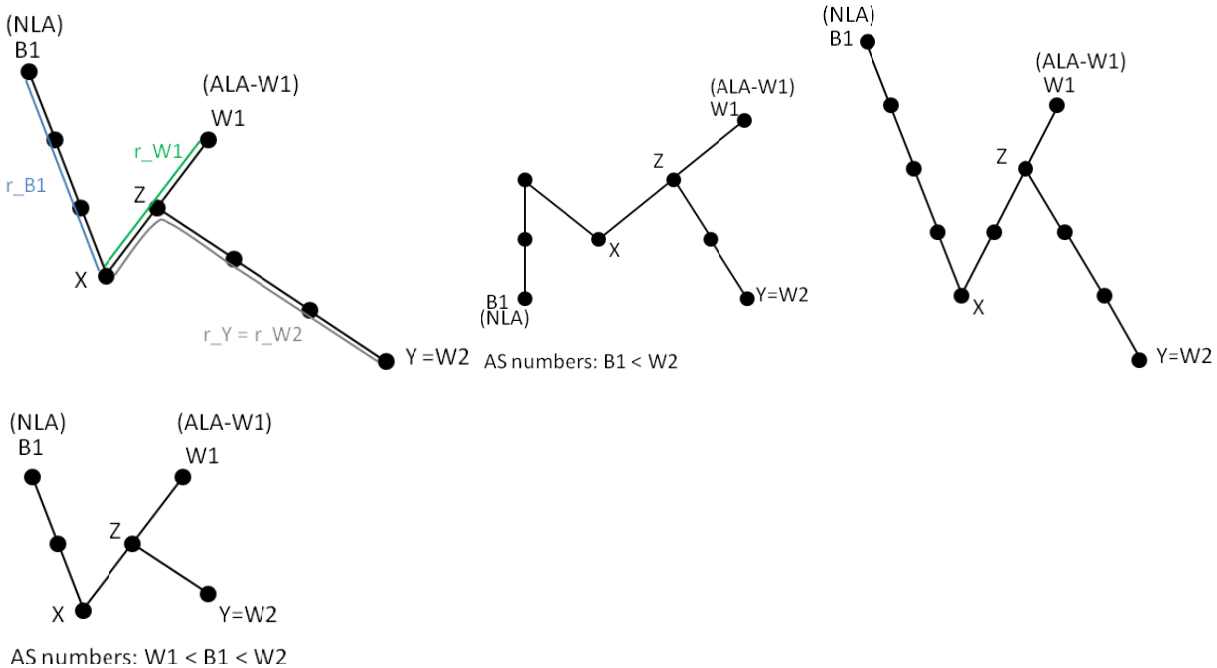

Fig 4: Examples of subgraphs with worse performance in NLA than in ALA. Anycast name in parenthesis indicates origin of the route chosen by X in that anycast variant.

Proof of Theorem 2:

It suffices to prove properties (i) to (vi) below. The proofs are analogous to the proofs of properties (i) to (vi) respectively in Theorem 1.

(i) node X in the subgraph selects a black route r_B1 that is originated by black node B1 in NLA and a white route r_W1 originated by white node W1 in ALA-W1
(ii) both routes r_W1 and r_B1 are routes of length at least two imported to X from two distinct neighbors with the same business relationship with X
(iii) node X receives a white route r_W2 to node W2 ≠ W1 in NLA from the same neighbor from which it receives r_W1 in ALA-W1, i.e., in the subgraph there must be some node Z on r_W1 that in NLA prefers the valley free route r_W2 to the route received on r_W1, and yet node X in the subgraph prefers r_B1 to r_W2 in NLA
(iv) routes r_W1, r_B1 and r_W2 are provider routes for node X
(v) either r_B1 is shorter than r_W2 or the length of the two routes is the same and B1 < Y
(vi) either r_W1 is shorter than r_B1 or the length of the two routes is the same and W1 < B

□

5. Quantifying the Security Performance of NLA, ALA and Naïve-ALA

So far, we have confirmed our earlier experimental results and concluded that in practice, the security performance of NLA and ALA is nearly identical. Our last task is to quantify the performance of the three anycast variants – NLA, ALA and Naïve-ALA.

Theorem 3: The fraction of the secure nodes in network-layer anycast (NLA) is x/(x+y) where x is the number of the white nodes and y is the number of black nodes.
Proof: We use a symmetry argument. It is easy to see that if z nodes chosen at random originate the same prefix, on average 1/z of the nodes in the network will choose a route belonging to each origin. In NLA we have x white nodes that originate a prefix and y black nodes that maliciously originate the same prefix. Therefore, the fraction of the autonomous systems that choose a white route must be x/(x+y). □

Theorem 4: The fraction of the secure nodes in naïve application-layer anycast (Naïve-ALA) is 1/(1+y) where y is the number of black nodes.
Proof: In naïve application-layer anycast each of the white nodes originates a different prefix, but all black nodes originate each of the prefixes. Therefore, if each node in the Internet chooses one prefix at random, the overall security performance must be the same as if only one white node originated a single prefix and there were y black nodes originating the same prefix. Therefore, using the same symmetry argument as in the proof of Theorem 3 we conclude the fraction of the nodes that select a white route must be 1/(1+y). □

6. Conclusion and Future Work

In this paper we provided a theoretical analysis of the security performance of network-layer and application-layer anycast, techniques which can be used as countermeasures against DNS attacks. This is achieved by showing that the performance can only differ if one of a family of subgraphs that we characterize appears in the interdomain network topology. Our result is important because it reconfirms the surprising experimental result which we report in [3].

In our ongoing work, we are extending our results in several directions. First, we are investigating the frequency with which the subgraphs characterized by Theorems 1 and 2 appear

in a realistic interdomain topology. This can be achieved by studying the AS-level Internet topology provided by CAIDA. Second, we are extending the theoretical results by characterizing the properties of nodes whose security performance differs. For example, a trivial corollary of Theorems 1 and 2 is that the security performance of a Tier-1 autonomous system cannot differ. A more challenging question concerns the security performance of autonomous systems classified into groups based on factors such as their node degree [5] or the number and distribution of their peering links [6].

While we believe that experimentation will remain a crucial method to accurately evaluate the behavior of complex routing systems in the future, analysis such as this one can help to understand and design routing protocols with better security, reliability and performance properties.

7. References

1. C. Partridge, T. Mendez, and W. Milliken, "Host anycasting service," RFC 1546, IETF, November 1993.

2. E. Zegura,M. Ammar, Z. Fei, and S. Bhattacharjee, "Application-layer anycasting: A server selection architecture and use in a replicated web service," IEEE/ACM Transactions on Networking, 8(4), August 2000.

3. I. Avramopoulos, and M. Suchara, "Protecting DNS from routing attacks: A comparison of two alternative anycast implementations," IEEE Security & Privacy, Issue on Securing the Domain Name System, September/October 2009.

4. L. Gao, and J. Rexford, "On inferring Autonomous System relationships in the Internet", IEEE/ACM Transactions on Networking, 9(6), December 2001.

5. R. Govindan, and A. Reddy, "An analysis of inter-domain topology and route stability," Proceedings of IEEE INFOCOM, April 1997.

6. L. Subramanian, S. Agarwal, J. Rexford, and R. Katz, "Characterizing the Internet hierarchy from multiple vantage points," Proceedings of IEEE INFOCOM, June 2002.