How Small Groups Can Secure Interdomain Routing

Ioannis Avramopoulos, Martin Suchara, and Jennifer Rexford
Princeton University

November 16, 2007

ABSTRACT

Although the Internet's routing system has serious security vulnerabilities, none of the existing proposals for a secure variant of BGP has been successfully deployed in practice. This is not surprising since deploying protocols that require the cooperation of tens of thousands of independently-operated networks is problematic. Instead, we argue that small groups should be the basis for securing BGP and we offer an alternative design in which interdomain routing is secured by a few (e.g., 5–10) participating ASes. We conduct extensive simulations on a realistic Internet topology to identify conditions for small groups to be effective. Even though the non-participants outnumber the group members by several orders of magnitude, the participants can achieve remarkable security gains by filtering compromised interdomain routes, cooperating to expose additional path diversity, inducing non-participants to select valid routes, and enlisting a few large ISPs to participate. We also propose two novel mechanisms that the group members can employ to achieve these goals, namely secure overlay routing and the cooperative announcement of each other's address space. Our experiments show that combining these two techniques allows small groups to secure interdomain routing.

1. INTRODUCTION

Today's Internet routing system is extremely vulnerable to attacks where adversarial networks announce routes for address blocks they do not own. In fact, "hijacking" another network's IP prefix is so easy that it often happens by accident [1, 2, 3, 4]. The consequences of prefix hijacking, and other forms of bogus routes announcements, are serious because the packets destined to the victim prefix are instead delivered to the adversary, who may drop the traffic, impersonate the destination, modify the payload, or snoop on the communication. For example, during the infamous "AS 7007" incident, a significant fraction of all Internet traffic was mistakenly directed to a small ISP for several hours [2].

The best way to defend against prefix hijacking is the subject of much debate. The role of secure routing protocols, in particular, has received considerable attention. The debate has been dominated by a "purist" philosophy that advocates the ubiquitous deployment of a secure version of the Border Gateway Protocol (BGP), the Internet's *de facto* interdomain routing protocol. The purist approach seems natural, if not mandatory, since BGP is the glue that holds the disparate parts of the Internet together. Purist solutions are advocated in public forums, such as the RPSEC working group of the IETF [5] and the North American Network Operators Group [6]. In fact, the debate focuses primarily on *which* secure routing protocol should be adopted (e.g., S-BGP or soBGP) [7], rather than *whether* a single solution should prevail. The Internet policy community has also discussed the possibility that the U.S. government might mandate S-BGP deployment [8].

Although ensuring that routing-protocol messages are authorized is clearly useful, the purist approach is problematic for both technical and economic reasons:

- Ubiquitous deployment of a secure routing protocol requires the cooperation of more than 25,000 Autonomous Systems (ASes). The large number of ASes prevents market forces from driving deployment, and government intervention may be both hard to realize (due to the global nature of the Internet) and undesirable (since it may stifle innovation).
- Smaller groups of like-minded ASes are much more likely to have aligned incentives that enable a partial deployment of a security solution. In a small group, one large company may have sufficient incentive to finance the participation of other members, or all of the ASes in the group (say, of large backbone providers) may decide to share the cost for their mutual gain.
- Groups benefit from deploying customized security solutions. No one interdomain security solution satisfies all of the security objectives, and, therefore, different groups may want to strike different trade-offs, based on their customer requirements and deployment costs.

In this paper, we argue that small groups of cooperating ASes should be the starting point for securing interdomain communication.

Interdomain communication needs to be protected against attacks on availability, confidentiality, and integrity. Ultimately, ensuring confidentiality and integrity requires end-to-end mechanisms, such as end-host encryption and authentication. As such, in designing and evaluating secure routing techniques, we focus primarily on improving end-to-end *availability*, though some of our solutions also improve the confidentiality and integrity of communication. Rather than *guaranteeing* availability—something that is inherently difficult to do, even for full deployments of S-BGP—we focus on significantly raising the bar for the adversaries to disrupt the delivery of traffic to the group members. For example, we would like to limit the number of places where an adversary can launch a successful attack, or require several colluding adversaries before an attack can succeed.

The problem is challenging because the non-participating ASes—who make no effort to detect or avoid the routes announced by the adversaries—outnumber the group members by several orders of magnitude. This imposes several serious

constraints on the space of solutions. First, the group members must use the conventional (insecure) version of BGP to exchange routing information with the non-participants, meaning the participants cannot completely upgrade to a new, secure protocol. Second, a non-participant may unwittingly propagate an adversary's route announcement to the group members, reducing the likelihood the participants learn *any* valid routes. Third, the traffic exchanged between group members often traverses non-participating ASes, who may unintentionally direct traffic toward the adversaries.

In this paper, we argue that the proposed secure variants of BGP are not equipped to overcome these obstacles. On the surface, secure origin BGP (soBGP) [9] is the most promising starting point for a partial deployment, since it is backwards compatible with BGP. Rather than authenticating the BGP messages themselves, soBGP has the routers verify the contents of BGP announcements against a *registry*, populated with information about prefix ownership and the AS topology. However, our evaluation under realistic AS-level topologies shows that small-scale deployment of the registry does not offer significant security gains. The group members can improve availability by applying more accurate techniques for detecting invalid routes, such as control-plane anomaly detectors or data-plane probing techniques. However, our experiments show that even a *perfect* detector would not make small groups effective at circumventing the adversary, even if several large ISPs participate in the group.

Our experiments suggest that to be effective, the members of the group must take two additional actions. First, they must cooperate to expose additional path diversity, to ensure that they have valid routes to the destination. Second, they must be proactive in inducing non-participants to select valid routes. In this paper, we present the design and evaluation of two novel mechanisms that the group can use to achieve these goals:

- **Secure overlay routing**: To circumvent the adversaries, the group members form a secure overlay network we call an SBone. In contrast to conventional overlays, an SBone connects *networks* rather than end hosts, collects path-quality measurements that are robust to adversaries, and avoids mapping virtual links on to compromised paths through the Internet.

- **Hijacking the hijacker**: To prevent non-participants from directing traffic toward the adversaries, all participating ASes originate BGP announcements for the prefixes the group wants to protect, and then forward the traffic over the secure overlay to the legitimate destination. "Shouting" the group's prefixes substantially improves availability, in exchange for a small increase in routing-table size and path lengths.

Our experiments show that these two techniques, combined with accurate detectors and the support of a few large ISPs, allow a small group (e.g., of 5 to 10 ASes) to achieve remarkable security gains at reasonable cost.

To quantify the effectiveness of small groups, we perform extensive simulations on a snapshot of the Internet's AS-level topology, annotated with the inferred business relationships between neighboring ASes. Simulation is necessary for an accurate evaluation because the composition of the group, the location of the adversaries, the connectivity between ASes, and the routing policies all have a profound influence on whether ASes learn and select legitimate routes. To date, synthetic models that accurately capture both the Internet's structure and BGP routing policies remain elusive, leading us to simulate security solutions on the existing Internet topology, rather than an abstract model. To understand the influence of group composition, we consider several models of group formation, including random group memberships and participation by ASes based on their node degree. In practice, we envision that groups of ASes will form based on shared incentives or the desire of large ISPs to offer enhanced security as a value-added service.

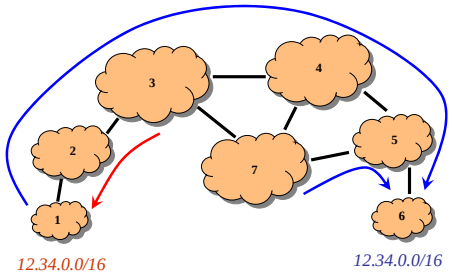

Figure 1: Announcement of prefix 12.34.0.0/16 from two origins partitions the network into two subsets.

The remainder of the paper is organized as follows. The next section presents a brief overview of prefix-hijacking attacks. Section 3 shows that small-scale deployments of soBGP are not very effective, and Section 4 shows that even perfect techniques for detecting invalid routes are not sufficient. In Section 5, we show how to improve availability for communication *between the participating ASes* through secure overlay routing. Then, in Section 6, we show how to coax the *non-participating ASes* into directing traffic towards the group members. Section 7 elaborates on how our solutions defend against sub-prefix hijacking attacks. Section 8 discusses how a small group of like-minded ASes should deploy our solutions in practice. Section 9 presents related work, and Section 10 concludes the paper with a discussion of future research directions. An Appendix expands on our economic argument that small groups should form the basis of secure interdomain communication [10].

2. PREFIX HIJACKING

BGP, the interdomain routing protocol of the Internet, is responsible for establishing reachability to destinations that are specified as address blocks, also called *prefixes*. Each AS may legitimately announce or *originate* in BGP one or more prefixes. Then, BGP establishes paths so that these prefixes are reachable from the rest of the ASes. Each router maintains and further announces a single *best* path per prefix. A BGP-speaking router trusts that the information its neighbors propagate has been generated in a legitimate fashion. This assumption of trust can be exploited by adversaries, who may use BGP attacks to breach the availability, confidentiality, and integrity of interdomain communication.

Each prefix is usually announced by a single AS. However, this condition may be violated in practice either for legitimate reasons [11], such as origination of anycast prefixes, or illegitimate ones, such as malicious attacks. In the latter case, the adversary gains control of the address block, and

may either act as a sink and discard the received traffic, or act as a man-in-the-middle and forward packets to the legitimate destination. The term *prefix hijacking* usually refers to the origination of a victim prefix by an adversary instead of the legitimate origin AS. We will refer to this attack as a *simple origination attack*. An AS that selects a malicious route will propagate it to its neighbors, who may select it as well. For example, in Fig. 1 prefix 12.34.0.0/16 is initially announced by AS 6, and all source ASes point their routing tables toward AS 6. If AS 1 also announces the same prefix, the network is partitioned in two subsets of ASes according to the origin AS they have chosen for the prefix: ASes 2 and 3 point their routing tables toward AS 1, whereas ASes 4, 5, and 7 point their routing tables toward AS 6.

Another type of attack is the *path-spoofing attack* in which the adversary announces a forged AS path to the victim prefix so that the adversarial AS appears upstream of the legitimate origin AS. Path spoofing is an intelligent attack, motivated by the adverary's desire to evade detection. However, the attack increases the length of the AS path, which may make some ASes less likely to select the malicious route.

In a third variant of prefix hijacking, the adversary breaks the victim's prefix into multiple sub-prefixes and originates those instead. In this way, although the network will maintain routes to both the original prefix and the sub-prefixes, because of the *longest prefix matching* rule used in the data plane, traffic will be directed to the adverary-controlled sub-prefixes. We refer to this attack as *sub-prefix hijacking*.

There are other variants of prefix hijacking, such as *wormhole* [12] attacks. Wormhole attacks are a countermeasure the adversary can employ against secure routing protocols. Wormhole attacks are not discussed in this paper because in our evaluation scenarios the adversary is able to employ strictly more effective attacks.

3. DEPLOYING SOBGP IN SMALL GROUPS IS INEFFECTIVE

Secure routing protocols such as S-BGP and soBGP have been designed assuming ubiquitous deployment. In this section, we consider a partially deployable variant of soBGP, a protocol that has been widely discussed as an alternative to S-BGP [7]. Using simulations we demonstrate that small-scale deployments of soBGP provide only limited benefits to the adopters. These results motivate our exploration of the conditions that enable small groups to be effective.

3.1 Partial soBGP Deployment

The soBGP protocol is designed around a cryptographically-secured registry of routing information. The registry contains information about the prefixes each AS is authorized to advertise in BGP, as well as the pairs of ASes that are BGP neighbors. BGP advertisements are validated against the registry to ensure, first, that the origin AS in the advertisement has been authorized to advertise the corresponding prefix and, second, that all links in the AS-path of the advertisement match links included in the registry.

We consider a cryptographically secured registry of routing information similar to the one used by soBGP. However, our registry is *partial* in that it contains the routing information for only a subset of the ASes in the Internet. We call the set of ASes that publish information in the registry the *participants*, and all other ASes the *non-participants*. For each participant, the registry contains a list of prefixes the AS is allowed to originate as well as a list of the AS's neighbors. Besides having their information published, participants use the registry to validate BGP advertisements. A participating AS discards a BGP advertisement if it contains information that *contradicts* the registry. For example, the AS will discard a route to a registered prefix if the origin AS number is wrong. To evade detection, the adversary must spoof the origin of the route and possibly additional hops in the AS-path (such as the second-to-last-hop), to avoid contradicting the registry. Although the resulting advertisements do not contradict the registry, the longer AS-path makes them less likely to be selected by other ASes.

3.2 Evaluation Methodology

Our experiments evaluate the effectiveness of the aforementioned partial deployment of soBGP. The simulation techniques, and the data sets, that are introduced in this section are used throughout the paper. Our experiments simulate the propagation of BGP route announcements on the AS-level Internet topology, as well as how the announcements affect the routing tables of each AS. Route propagation is profoundly influenced by AS business relationships, such as customer-provider, peer-peer or sibling-sibling. Since the goal of ISPs is to generate profit, and since customers have to pay their providers, ASes prefer routes learned from customers over routes learned from peers or providers; if multiple routes of the same class are available, the AS prefers shorter AS paths over longer ones. The business relationships also determine whether an AS exports the chosen route to its neighbors. An AS exports a route learned from its customer to all of its neighbors, whereas a provider or a peer route is exported only to customers.

Our simulations use and extend BSIM [13], which provides a convenient environment to simulate policy-based route propagation on an arbitrary AS topology. BSIM accurately captures the influence of business relationships on how ASes select and export routes. As input, we used an AS topology (annotated with the inferred business relationships) from June 2007 available from CAIDA [14]. This is considered to be one of the most accurate and most complete AS topologies available. The topology is constructed from snapshots of the routing tables from RouteViews servers [15], and it contains $25,304$ ASes. Routing table inspection at the end of each experiment allows us to determine what fraction of ASes selected valid routes to the victim prefix.

To measure the impact of an attack, we compute the average number of ASes that accept a route to the legitimate origin AS over a sequence of 100 experiments. In each experiment, the set of participant ASes and the adversarial AS are selected at random, and the victim is a randomly chosen member of the group. The variance of the quantities we measure can be high. This is because the outcome of each experiment critically depends on the location of adversary and the composition of the group. For example, if the adversary is the victim's sole provider, *all* ASes select the attacker's route, but the opposite configuration offers perfect security. Using the mean as the performance metric enables us to estimate how difficult it is for the adversary to mount a successful attack against the defending group. As the mean increases or decreases the number of places where the adversary can launch a successful attack increases or decreases correspondingly. Fortunately, as we demonstrate

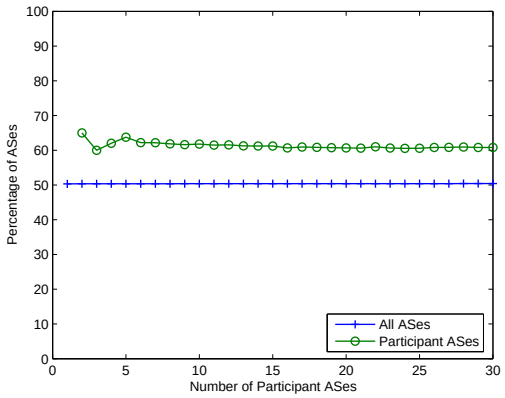

Figure 2: Up to 30 randomly selected member ASes participate in a routing registry. Figure depicts the percentage of ASes able to reach the prefix of the victim.

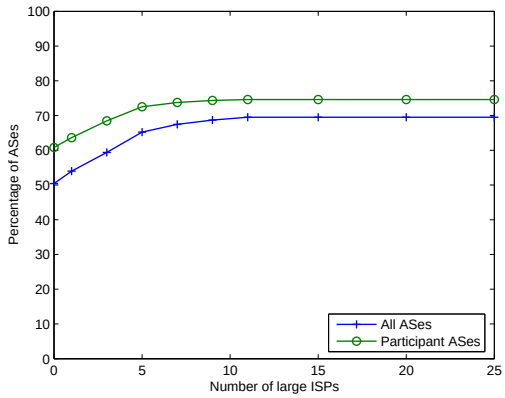

Figure 3: Up to 25 large ISPs plus 30 randomly selected ASes participate in a routing registry. The figure depicts the percentage of ASes able to reach the prefix of the victim.

later in the paper, our proposed solutions not only increase the mean but also decrease the variance.

There are two possible strategies the adversary can use to maximize the impact of an attack against a victim prefix belonging to a participant. In the first, the adversary launches a simple origination attack, listing its own AS number as the origin AS. This attack is intended to affect the *non-participants*, as the participants can easily detect that the offending route contains contradictory information. In the second, the adversary spoofs the shortest possible path that does not contradict the registry (i.e., the shortest AS path from any non-participant to the true origin AS). This attack is intended to evade detection by the *participants*. We assume that an AS has accepted a route to the adversary if it is affected by at least one of these attacks. This approach is justified by the fact that the adversary can combine the two attacks. For example, the adversary can use one of the attacks and if a particular AS of interest does not accept the adversarial route, the adversary can withdraw the first announcement and then use the other attack.

3.3 Simulation Results

In the first experiment, we consider a registry that has been formed by a random group of participants. Random participation is an apt deployment model when the locations of the participants are chosen based on criteria other than securing interdomain communication, for example, when the multiple sites of a single organization decide to act jointly. The adversary attacks a randomly chosen victim participating in the registry from a single randomly chosen AS. Fig. 2 shows the percentage of participant and non-participant ASes that have accepted a route leading to the legitimate origin AS for the victim prefix.

The percentage of non-participant ASes that are able to reach the victim remains flat at approximately 50% as participation grows up to 30 members. In the absence of any protection mechanisms the victim would also be reachable on average by 50% of the ASes.[1] Therefore, if the participation is random, the partial registry is unable to help the victim, which remains as vulnerable as if there were no participants to prevent the propagation of the adversarial routes. This result can be explained by the following observation. Since most ASes in the Internet are stubs, the random selection of participants implies that stubs are most likely to be selected. However, stubs are *terminal* points that do not propagate any routes. Therefore, randomly selected participants can at most help themselves avoid the adversarial routes. Furthermore, although the percentage of participant ASes that are able to reach the victim prefix is higher, it does not exceed 60%. This result can be explained by the limited upstream connectivity of the stub participants that have only a few alternate routes to choose from. In fact, inspection of the AS topology used in the simulation reveals that approximately 35% of the stubs are connected to a single upstream provider, having a single route per destination prefix, and that approximately 40% of the stubs are connected to two upstream providers, therefore, having at most two routes per destination prefix.

It is also worth noting that, in addition to the poor average security gain that the partial registry attains, the security gain has high variability. For example, in a group of 10 randomly selected members, the legitimate origin of the victim prefix is reachable by *every* other member in only 22% of the simulation instances, whereas in 5% of the simulation instances *no* member can reach the legitimate origin.

In the next experiment, we consider a different participation model. The random set of participants is helped by enlisting in the group a set of *deputies*, which are large ASes having high node degree. Because of their rich connectivity, deputy ASes have the highest potential to prevent the propagation of adversarial routes by filtering these routes and propagating legitimate routes instead. Fig. 3 shows the percentage of participant and non-participant ASes that have accepted a route to the victim prefix leading to the legitimate origin AS. The origin has been selected at random among a set of 30 participants. We decided to fix the number of participants to 30 as this corresponds to a large group of participants providing an upper bound on the effectiveness of any smaller group. The adversary attacks from a single randomly chosen AS. We consider a group of deputy

[1]This observation is easy to verify using the symmetry between the victim and adversarial ASes.

ASes ranging in size from 0 to 25. Although the benefits are significant in comparison to the results obtained by random deployment, performance is poor even if the number of deputy ASes is large. For example, more than 25% of the ASes still cannot reach the legitimate origin, even if the group includes the the 25 highest-degree nodes. The corresponding percentage for non-participants is even larger. It is also worth noting that the formation of groups of 25 or more large ISPs is not realistic because it requires significant coordination by networks that are otherwise business competitors and have been notoriously reluctant to form sizable coalitions.

4. PERFECT FILTERING OF INVALID ROUTES IS NOT SUFFICIENT

In Section 3, the group members do not enjoy significant security benefits, in part because they cannot accurately detect invalid routes. In this section, we explore whether more accurate detection techniques are enough to make small groups effective. We first explain how participating ASes can detect and filter invalid routes using existing proposals for detecting control-plane and data-plane anomalies. Then we present simulation experiments that show that, while ideal filters offer clear security benefits, the adversary is still able to inflict substantial damage.

4.1 Accurate Detection of Invalid Routes

Although protocols like S-BGP and soBGP rely on cooperation in creating and maintaining registries, several new solutions have been proposed so that individual ASes can detect (and potentially filter) invalid routes. Some solutions detect invalid routes in the control plane (i.e., by analyzing BGP announcements), whereas others operate in the data plane (i.e., by monitoring the forwarding path).

Control-plane techniques: Anomaly-detection techniques running in the control plane can identify suspicious routes based on a history of past BGP announcements [16, 17, 18, 19]. These techniques essentially allow individual ASes to use historical data to construct a more complete *de facto* registry of prefix ownership and AS-level connectivity. Although early anomaly-detection techniques only detected simple origination attacks [16, 17], recent solutions are able to detect more sophisticated path-spoofing attacks launched by intelligent adversaries [18, 19]. Although these techniques provably detect spoofed paths, they are vulnerable to false alarms that mistakenly flag valid routes as suspicious.

Data-plane techniques: An alternate approach to detecting suspicious routes is to detect anomalies in the forwarding path to the (alleged) destination. For example, a significant change in the number of hops in the path, or differences in the end-host properties, would suggest that a hijack has occurred [20, 21]. However, these techniques are vulnerable to intelligent adversaries who actively try to evade detection. More sophisticated data-plane techniques are possible when the source and destination ASes cooperate to detect availability problems along the path. For example, the communicating ASes could employ passive-monitoring techniques, like coordinated sampling [22] and stealth probing [23], to monitor loss and delay of the data packets traversing the forwarding path.

When data-plane measurements or control-plane anomalies suggest that the current path is invalid, a participating AS can simply filter the offending route, and select an alternate path (if one is available). The proposed detection techniques vary in how well they identify invalid routes, withstand intelligent adversaries, and avoid false positives. In the rest of this section, we evaluate the effectiveness of small groups that are armed with a *perfect* detector of invalid routes. As such, our results provide an *upper bound* on how well the proposed detection techniques might work in practice, when deployed in a small group of ASes. Since our results show that even a perfect detector is not sufficient to make small groups effective, we do not evaluate the individual detection techniques.

In conducting the evaluation, we must consider how an intelligent adversary would adapt the attack strategy in the face of these detection techniques. In particular, the adversary no longer has any incentive to spoof the AS path, since the participating ASes can easily detect and discard the invalid route. Instead, the adversary's best strategy is to launch a simple origination attack, in the hope of enticing as many non-participants as possible to select an invalid route. The impact of the attack is determined by the number of ASes that accept a route leading to the adversary, or are left with no route at all.

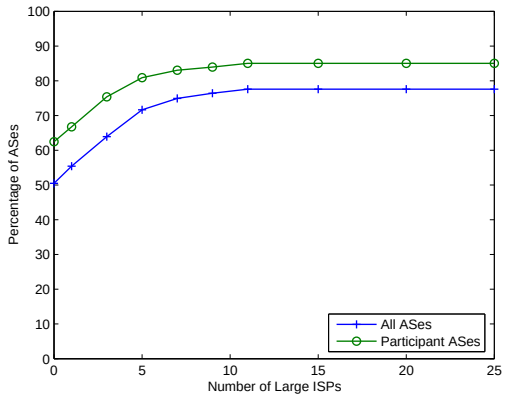

Figure 4: Ideal filters are employed at 30 member ASes and up to 25 additional large ISPs. Figure depicts the percentage of ASes able to reach the prefix of the victim.

4.2 Simulation Results

In this simulation, we evaluate the effectiveness of ideal filtering when the victim AS is selected at random among a set of 30 randomly-chosen participants. As in the previous section, we decided to fix the number of participating ASes at 30, as this corresponds to a relatively large group of participants providing an upper bound on the effectiveness of any smaller group. The adversary attacks from a single randomly chosen AS, and we consider the benefits of enlisting 0 to 25 large ISPs as deputies.

Fig. 4 shows the percentage of participant and non-participant ASes that select a legitimate route to the victim prefix. When 25 deputy ASes augment the group of 30 participants, around 15% of the group members cannot reach the victim. This represents a significant improvement over the partial soBGP deployment evaluated in the previous section, where more than 25% of the participants could not

reach the victim. Non-participants also benefit when the group members and deputy ASes can accurately filter invalid routes. Compared to the previous section, the percentage of non-participants able to reach the victim prefix increased from 65% to 75% assuming 10 deputy ASes participate. The non-participants benefit because the deputy ASes select valid routes, essentially blocking the propagation of invalid routes. This decreases the likelihood that a non-participant inadvertently selects an invalid route.

Despite the security gains attainable over partial soBGP deployment, we believe that too many ASes must participate before significant benefits are achieved. Accurate detection of invalid routes and the support of large ISPs are helpful, but not sufficient. Since the non-participants cannot detect invalid routes, they propagate these routes to the participants in lieu of legitimate routes. As such, some participants do not learn *any* valid route. Many of these ASes are stub networks with just one or two upstream providers, which significantly limits the number of BGP routes they learn. Although these ASes can detect that the routes they learn are invalid, they do not have enough options to select a valid alternative. In the next section, we present a technique that overcomes this limitation by providing the participating ASes with additional routes. This new technique offers significantly better security gains, even for groups as small as 5–10 ASes.

5. SECURE OVERLAY ROUTING

This section describes how a small group of ASes can effectively secure interdomain communication between its members. The group members form a secure overlay network that offers alternate overlay paths a participating AS can use when no valid BGP route is available. First, we introduce the Security Backbone (SBone), which protects intra-group traffic despite *deployment gaps* separating the participants. Then, we present simulations that quantify the substantial security gains for a realistic Internet topology. Although primarily designed to protect *availability*, the SBone can also enhance the confidentiality and integrity of communication, as discussed at the end of the section.

5.1 Security Backbone (SBone)

The participating ASes must effectively handle deployment gaps, i.e., non-member networks that are uncooperative or even hostile. To enable participants to circumvent availability problems, we connect the group members by a mesh of virtual links forming an overlay network. We call this overlay network a Security Backbone (SBone). The SBone differs from traditional overlays (e.g., RON [24]) in that it is an overlay of *networks* rather than individual end hosts or servers. In traditional overlays, the participating hosts have little or no control over the routes their upstream providers pick. In contrast, because the SBone is created by the administrators of the participating ASes, it has visibility into (and control over) BGP routing. In fact, the SBone may run directly on the routers in the participating ASes.

The virtual links are created by connecting members networks with IP tunnels that encapsulate and decapsulate the data packets. For each pair of group members X and Y, we create two tunnels: one from X to Y, and one from Y to X. The SBone improves availability in two main ways. First, multiple interdomain paths may exist between the endpoints of each virtual link, and traffic can be directed over any of these paths. That is, an SBone node could switch a virtual link from one underlying interdomain path to another, after detecting an availability problem on the original path. Second, if all of the underlying paths have availability problems, an SBone node can direct traffic through an intermediate SBone node via the overlay network.

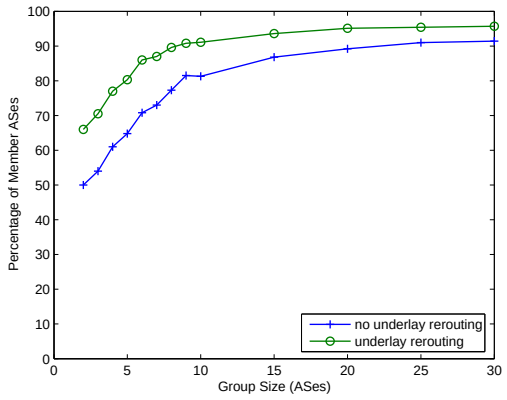

Figure 5: Percentage of up to 30 randomly selected ASes that are able to reach the legitimate origin through an overlay. Case with underlay rerouting allows a participant to control the selection of its BGP routes.

The SBone relies on a monitoring system to detect availability problems and to disseminate the measurement results to other nodes. The SBone could apply any of the monitoring techniques outlined earlier in Section 4.1, but techniques tailored to detecting availability attacks are especially appealing. For example, the SBone nodes could easily apply highly accurate availability-monitoring techniques that require cooperation between the communicating end points [22, 23], or active probes sent by separate probe machines. These data-plane monitoring techniques allow the SBone to react to a broader range of attacks, including adversaries that propagate valid BGP advertisements while maliciously dropping or delaying the data packets—a problem traditional secure routing protocols like S-BGP cannot handle.[2]

5.2 Random Deployment

We evaluate how effective the SBone is in preventing routing attacks using simulation. As before, our simulator is based on BSIM, and the AS topology is the same as in the previous experiments. In the simulation, we assume that the participating nodes are equipped with ideal filters that are able to distinguish the routes leading to the adversary. Fig. 5 shows the percentage of group members that are able to reach the legitimate origin AS through an overlay path. The origin AS, the defending group, and the adversarial AS have been selected at random among the set of all ASes. The size of the defending group reaches up to 30 members. We consider a case in which the group members are not able to select the BGP path of a virtual link, a scenario akin to a setting in which the group members are overlay hosts that do

[2]In fact, *secure* availability monitoring techniques could protect against especially insidious adversaries that try to bias the data-plane measurements to evade detection [23].

not have visibility into BGP, and a case in which the group members have control over BGP. The latter case is shown to outperform the former. Note, however, that participation of large groups is required to attain significant security gains in either case. A refinement is required so that deployment of the SBone in small groups attains significant security gains.

5.3 Reinforcement through Tier-1 ASes

In the previous experiment, the participant ASes were selected at random from the set of all ASes. Since the majority of ASes are stubs of limited upstream connectivity, this resulted in limited path diversity, and hence large groups were required to attain significant security benefits. Therefore, we consider a participation model similar to the one used in Sections 3 and 4 in which the random group of participants is reinforced by a few tier-1 ASes acting as deputies to assist the group. Because of their rich connectivity, deputy ASes are able to expose rich path diversity that is able to overcome the limitations the random group of participants faces. In the next experiment, we demonstrate that enlisting one or more such deputy ASes in a randomly selected group can thwart the adversary's power even if the overall size of the defending group that includes the deputies is small. Note that the participation of large ASes in the defending group could be arranged in exchange for pay.

Fig. 6(a) shows the percentage of group members that are able to reach the legitimate origin of a randomly selected victim prefix. The deputy ASes are excluded from the set of possible origins and the adversary has also been selected at random. We assume that the adversary attacks both the participant's prefix and the tunnel endpoints of the overlay formed by the defending group. Both the participants and deputy ASes are equipped with filters that are able to distinguish the routes leading to the adversary. We consider four cases such that the group of deputy ASes consists of 0, 1, 3, and 5 members. We find that the percentage of participants able to reach the victim prefix exceeds 95% provided that the group of deputy ASes consists of 3 or more members. This result holds irrespective of the size of the group of participants. We also find that if the size of the group of participants is small, the security gains improve substantially as the size of the group of deputy ASes increases from 0 to 5 members; larger groups of participants are able to achieve significant security gains without the help of the deputies.

5.3.1 Five Adversaries

So far, we have assumed that the adversary attacks from a single AS, and we have shown that the security benefits attainable by partially deployed secure routing protocols are poor even against the smallest adversarial group. In this section, we show that the SBone is resilient not only against a single adversary but also against larger adversarial groups. In particular, we consider an adversarial group that consists of 5 members.

Fig. 6(b) shows the percentage of group members able to reach a victim prefix belonging to a randomly selected participant under the same assumptions as in the experiment shown in Fig. 6(a). We observe that large adversarial groups have a significant impact on those defending groups that do not enlist deputy ASes. Furthermore, the effect of large adversarial groups diminishes as the number of deputy ASes that are enlisted increases. For example, in a defending group of 10 members that enlist 5 deputy ASes, 90% of the members are able to reach the legitimate origin of the victim prefix.

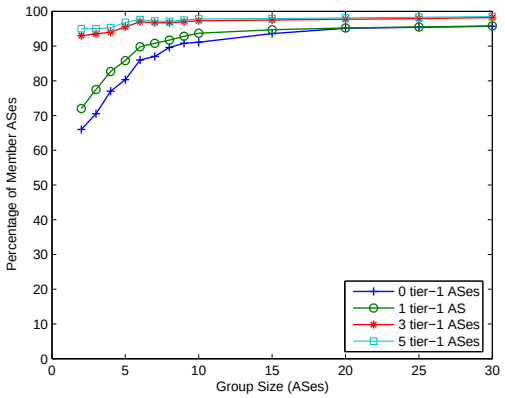

(a) 1 Adversarial AS

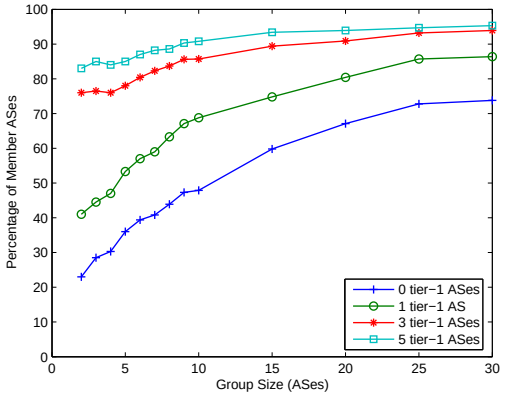

(b) 5 Adversarial ASes

Figure 6: Up to 5 tier-1 ASes assist the group of participants. Figure depicts the percentage of group members that have an overlay path to the prefix of the victim.

5.3.2 Comparison with Secure Routing Protocols

We use the security benefit attainable by an ideal secure routing protocol as a baseline for comparison of the benefit attainable by the SBone under large adversarial groups. Fig. 7 shows the percentage of member ASes able to reach a victim prefix belonging to a randomly selected member AS that is attacked by 5 adversarial ASes. We consider two cases. In the first case, an ideal secure routing protocol has been used to protect the victim and, in the second case, the victim is protected by the SBone. We assume that the group of participants consists of 10 members. We consider four cases such that the group of deputy ASes consists of 0, 1, 3, and 5 tier-1 ASes. We find that the SBone is able to protect $20-30\%$ more participants than the ideal secure routing protocol.

It is also worth noting that the SBone has significantly lower variability in its performance than the ideal secure routing protocol. For example, in the case of 5 deputy ASes, the standard deviation corresponding to the SBone is 8%,

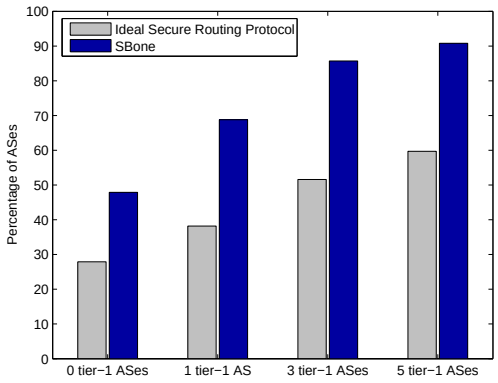

Figure 7: Comparison of the security performance of an ideal secure routing protocol and the SBone against 5 adversarial ASes. Up to 5 tier-1 ASes assist the group of 10 participants. Figure depicts the percentage of the participants able to reach the victim prefix.

whereas the standard deviation of the ideal secure routing protocol is 15%. Therefore, the security benefit of the SBone is more predictable than the benefit an ideal secure routing protocol would attain.

5.4 Enhancing Confidentiality and Integrity

The SBone also enables routing strategies that protect confidentiality and integrity by preventing adversaries from receiving sensitive traffic. For example, the SBone allows the group to select overlay paths that proactively avoid untrusted ASes owned by business competitors or known not to implement best common practices. Furthermore, in a manner complementary to filtering unwanted routes, group members can employ routing strategies that mitigate the risks of attacks. For example, the group members can proactively spread traffic over multiple paths, reducing the overall amount of traffic carried over any single path, and, therefore, substantially increasing the amount of resources an adversary would have to invest to observe all the traffic.

6. SHOUT

In this section, we present Shout, a technique that the defending group can employ to secure traffic originating from non-participating ASes that is destined to the participating ASes. Shout extends the benefits of the SBone to non-participating ASes that are entirely agnostic of the protection mechanisms that the participants apply. The ability to secure traffic originating from a potentially large number of non-participating ASes is in the vested interest of the participants as in this way the value they receive from their participation in the system increases. Noting that to receive the full benefits of the SBone an AS must eventually become a participant, Shout is intended for deployment during the transient period in which the SBone is incrementally building up to include a target group of members.

6.1 Hijacking the Hijacker

Shout is a destination-driven technique that attracts traffic from sources that may not be participating in the system and that may even be ignorant of the existence of the system. The ability to offer the service despite non-participating traffic sources decreases the degree of participation required for the system to be effective and facilitates the formation of small groups. Shout coaxes non-participants into picking routes leading to nearby participants instead of routes leading to adversarial ASes. Shout competes with the adversary to attract traffic from the non-participants using the adversary's own armory. Shout *hijacks the hijacker* by having the defending group of ASes simultaneously originate in BGP a participant's prefix. In this way, even if the adversary attacks the prefix receiving the protection of the group, the non-participants will prefer the routes leading to the group members over the adversary's routes.

The simultaneous origination of the prefix from multiple ASes may cause the network to direct traffic from different sources attempting to communicate with the same IP address to different destination machines. Although there are cases where a behavior like this is a limitation, sometimes this behavior is desirable. An example is *anycast*, a service in which multiple hosts share the same IP address. Traffic destined to this address may be delivered to any of the corresponding hosts. Anycast is well suited for applications that do not maintain state across multiple packets, such as single request-response applications like DNS. In fact, anycast has been deployed to improve the resilience of root DNS servers [25]. Anycast becomes problematic in connection-oriented multi-packet transfers requiring state across different packets, like in TCP.[3]

For those applications that anycast is unsuited for, the traffic that enters the group must be delivered from the AS that first receives the traffic to the participant AS hosting the prefix. This delivery is carried out by the SBone. The effectiveness of the combined arrangement of, first, using Shout to coax non-participants into picking routes leading to participants and, then, using the SBone to deliver the traffic to its destination was evaluated using simulation. As before, our simulator is based on BSIM, and we use the same AS topology as in the previous experiments. Fig. 8(a) shows the percentage of all ASes in the Internet that are able to reach a participant's prefix when the corresponding participant AS is selected at random and the adversarial AS is also selected at random. We assume that the adversary attacks both the participant's prefix and the tunnel endpoints of the overlay formed by the defending group. We consider four cases in which 0, 1, 3, and 5 tier-1 ASes are enlisted in the defending group. If the group has 10 or more members and 3 or more tier-1 deputies, more than 95% of the ASes in the Internet are able to reach the victim despite the attack.

The case of an adversarial group that consists of 5 members is shown in Fig. 8(b). As in the case of the SBone, large adversarial groups have a significant impact in those defending groups that do not enlist deputy ASes. Furthermore, the effect of large adversarial groups diminishes as the number of deputy ASes that are enlisted increases. For example, in a defending group of 10 members that enlists 5 deputy ASes, 85% of the group members are able to reach the legitimate origin.

Finally, the security benefits of the SBone are compared with the security benefits of ideal secure routing protocols in

[3]This is because routing changes in the middle of an ongoing connection can lead the packets to a different host, which may not have the appropriate state to continue the transfer.

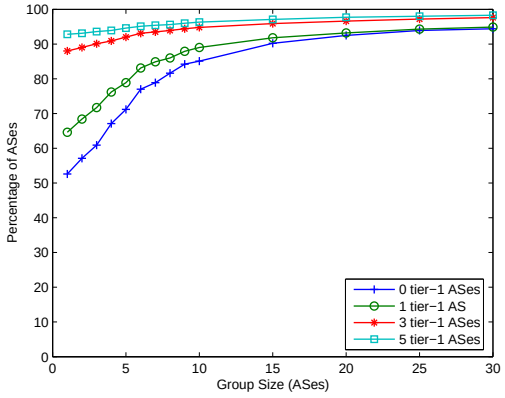

(a) 1 Adversarial AS

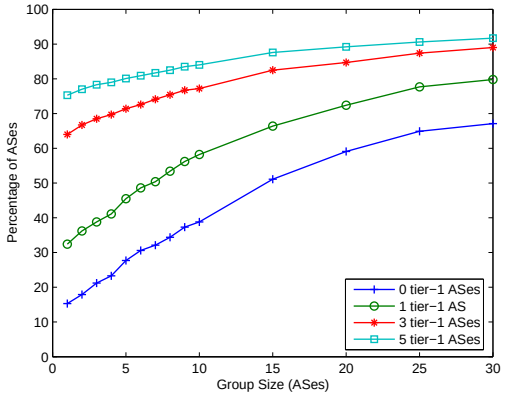

(b) 5 Adversarial ASes

Figure 8: Up to 5 tier-1 ASes and up to 30 participants use Shout and SBone. Figure depicts the percentage of ASes in the entire Internet that have a path to the prefix of the victim.

Fig. 9. If the group of the participants consists of 10 members and the group of the adversary consists of 5 members, we find that the SBone is able to protect $20-40\%$ more participants than the ideal secure routing protocol.

6.2 Performance and Scalability

We use simulation to measure the impact of Shout on the performance and scalability of interdomain routing. The impact on performance is measured by the extent to which end-to-end paths are prolonged, whereas the impact on scalability is measured by the increase in routing table sizes. We show that the increase in both quantities is small.

Shout forces the data traffic to take detours prolonging the end-to-end path, even when no adversary is launching an attack. Within each randomly-created group, we select a random destination AS, and enlist either 0, 1, 3, and 5 tier-1 ASes as deputies. We compute the length of the path from each (source) AS to the destination without Shout enabled. Then, we simulate the operation of Shout and determine the path the traffic would take to the destination AS, where part of the path takes the traffic from the source to the group and the remainder traverses the virtual link to the destination.

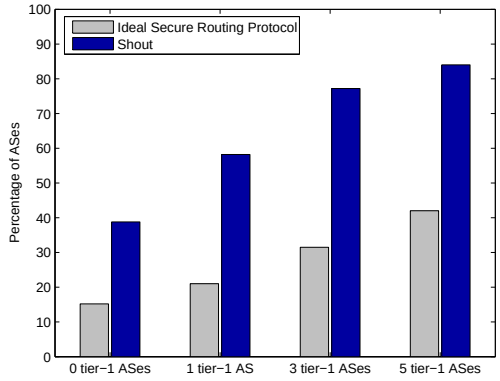

Figure 9: Comparison of ideal secure routing protocol and the SBone reinforced by Shout. Up to 5 tier-1 ASes assist the group of 10 participants. Figure depicts the percentage of all ASes able to reach the victim.

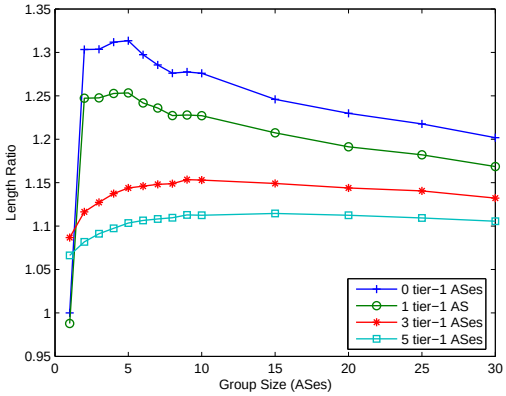

Figure 10: Impact of Shout on the lengths of paths used to reach the overlay origin. Ratios of the original and resulting hop counts depicted.

Fig. 10 shows the ratio of the path length after activating Shout over the path length before activating Shout, averaged over all source ASes. As expected, the ratio decreases as the number of deputy ASes increases. In a group where no tier-1 ASes participate, the ratio is at most 1.35. In a group in which 5 tier-1 ASes participate, the ratio drops below 1.15.

Considering now the impact of Shout on the routing-table size, we note that by originating a prefix from multiple ASes, Shout increases the number of alternate BGP routes for that prefix. Although in BGP each router selects exactly one route for each prefix, the alternate routes are stored in the RIB, increasing the size of the RIB at those routers. To evaluate the impact of Shout on RIB size, we counted the number of routes each AS stores for the "shouted" prefix. The average RIB size never increased by more than 5%, across all simulation instances for up to 30 group members and 0, 1, 3 or 5 deputy ASes.

Our experiments show that the increase in path length and routing-table size are modest. Still, in some cases, the group members may not want to incur the small perfor-

mance penalty or carry the extra traffic. Fortunately, the Shout mechanism can be employed *reactively* upon detecting a prefix-hijacking attack, using a control-plane anomaly detector like the ones described earlier in Section 4.1. When one of the group members, or a separate anomaly-detection system, detects a hijack, the group members can be instructed to activate Shout for the affected prefix. As long as the detector has a low false-positive rate, and few if any false negatives, invoking Shout reactively is both effective and efficient. Fortunately, several such control-plane anomaly detectors already exist [18, 19]. Ultimately, the decision of whether or not to run Shout reactively depends on what trade-off the group wants to strike between efficiency/performance and delay in reacting to attacks.

6.3 Discussion

The goal of Shout is to increase the resources an adversary must expend to disrupt the communication of a participant with the non-participating ASes. Shout protects the direction of communication from the non-participating traffic sources to the participating destinations. Our system does not protect the reverse direction of communication. Despite this limitation, the resources the adversary must expend to attack a participant are substantial. The reason is since traffic destined to the participant is secured, the adversary must be in a position to attack the traffic originating from the participant. However, to attack this traffic the adversary must be able to attack a potentially large number of traffic destinations. Therefore, Shout substantially raises the bar for the adversary to perform a successful attack against a participant. It is worth noting that our system does not attempt to level the degree of protection offered to the communication between participants with the degree of protection offered to the communication between a participant and a non-participating AS. We believe that this compromise is justifiable by the fact that once a non-participant decides to participate, he can receive the full benefits of the system.

7. DEFENDING AGAINST SUB-PREFIX HIJACKING

Thus far our adversary attacked a prefix by announcing it in BGP. Here we examine a scenario in which the adversary breaks or *deaggregates* the victim's prefix into multiple subprefixes and originates those instead. Although the network will maintain routes to both the original prefix and the subprefixes, the *longest prefix matching* rule will ensure that traffic is directed to the adversary's subprefixes. We refer to this attack as *sub-prefix hijacking*. There is a limit to the granularity that the adversary can deaggregate a prefix. This limit is imposed by filtering rules employed by ISP networks that discard routes to prefixes more specific than a /24.

7.1 Defending the Participants

The tunnel endpoints of the SBone overlay can be easily protected against sub-prefix hijacking. To protect a tunnel endpoint, the corresponding participant can announce the address of the endpoint with a /24 prefix, preventing the adversary from announcing a more specific prefix covering the endpoint. If a network participates in more than one overlay simultaneously, it is possible to use the same tunnel endpoint address in each of the overlays. Therefore, each network needs to only announce one /24 prefix irrespective of the number of the overlays it participates in.

Since the tunnel endpoints of the SBone overlay are resilient to sub-prefix hijacking, the delivery of traffic among participants is also resilient to this attack. This is true because traffic sent by participants is delivered via the secure overlay, with the packets encapsulated with the IP address of the recipient's tunnel end-point. Therefore, at full participation, to attack the intra-group traffic, the adversary must resort to the prefix hijacking attacks we considered in the previous sections. Next, we describe how one can protect the traffic sent by non-participants against sub-prefix hijacking attacks.

7.2 Defending the Non-Participants

One way to protect non-participant traffic destined to the participants is to proactively deaggregate the prefixes of the participants into sub-prefixes and use Shout to advertise the sub-prefixes instead of the aggregate prefixes. However, such a proactive countermeasure would create extra routes leading to a potentially significant increase of the BGP routing-table size. In the following, we present a reactive technique that is able to relieve the routing system from the extra routes. This technique is useful because it can obviate the need for the current practice in which ASes deaggregate their prefixes proactively [26]. This is well aligned with the recent efforts to limit the growth of the BGP routing tables.

7.2.1 Monitoring BGP Advertisements

To protect traffic sent by non-participants and destined to the participants against sub-prefix hijacking, we use the participants to monitor BGP advertisements and detect the offending advertisements. Assuming each participant knows the prefixes every other participant should advertise, if a BGP advertisement announcing a subprefix of the known prefixes is detected by any participant, a sub-prefix hijacking must have occurred. If any participant detects this event, it notifies the rest of the members, who respond to the attack by advertising the same sub-prefix. Based on the results of the previous section regarding the impact of Shout on the routing-table size, this countermeasure does not increase the routing table entries by more than 5% of the number of subprefixes the adversary has already advertised.

7.2.2 Evading the Monitors is Ineffective

To prevent the countermeasures from taking effect, the adversary will attempt to conceal the attack from the participants. However, the adversary has limited control over how the sub-prefix advertisement propagates in the system. Control over the propagation of the advertisement can be exercised using the following trick: BGP has a mechanism for avoiding loops in which a BGP advertisement is discarded by an AS if that AS already appears in the AS-path. Relying on this mechanism, the adversary can prevent the participants from receiving the offending announcement by adding the *union* of the neighboring ASes of each participant to the AS-path of the announcement. In this way, loop detection will be triggered at the neighboring ASes preventing the announcement from reaching the participants.

In order to conceal the attack from the participants, the adversary must limit the impact of the attack on the non-participants. To measure the damage that the adversary can induce, we used simulation on the AS topology one more

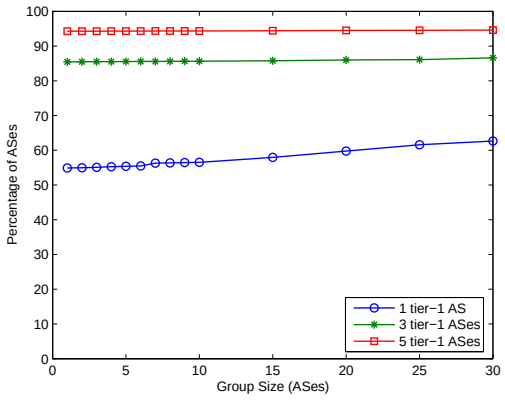

Figure 11: Percentage of ASes able to reach the legitimate origin of a victim prefix against which the adversary mounts a sub-prefix hijacking attack.

time. Fig. 11 shows the percentage of ASes that are able to reach a randomly selected victim prefix against which the adversary mounts a sub-prefix hijacking attack. We assume that the adversary forges the offending BGP advertisement so that the attack is concealed from the set of ASes that are neighbors of the participants. We consider three cases in which 1, 3, and 5 tier-1 ASes comprise the set. It is shown that to conceal the attack from 5 tier-1 ASes, the adversary affects less than 5% of the ASes in the Internet.

Fig. 11 shows that the impact of the adversary's attack decreases rapidly as the number of participating tier-1 ASes increases. This can be explained by the fact that tier-1 ASes have a large number of neighbors, rapidly increasing the number of ASes that must discard the offending advertisement to prevent detection. In fact, inspection of the topological map of the Internet reveals that the number of ASes that discard the malicious advertisement through loop detection is on order of several thousand. This observation gives rise to a simple countermeasure to thwart the adversary's ability to conceal a sub-prefix hijacking attack using the aforementioned trick. In this countermeasure, filters prevent the propagation of BGP advertisements containing more than a few hundred ASes in the AS-path. In practice, AS-paths of this length only appear due to configuration errors (or attacks), and anecdotal evidence suggests that certain ISPs already filter these long BGP advertisements. Moreover, such advertisements have been known to cause some routers to crash.

8. DEPLOYMENT CONSIDERATIONS

In this section we explore practical issues that concern deployment of our system. So far we have been mostly considering the average level of security attainable when the members of a group are selected at random. However, the security attainable by any given group depends on the relative location of the group members. Although the desired security level may not be achievable using the resources of the group itself, enlisting one or more large ISPs, similarly as we did in our experiments, may be sufficient to achieve the target objective. We argue that it is important to consider both the location of the members and the goals of the group when enlisting deputy ASes.

SBone: Let's consider intra-group communication. We observe that the security of a path between two participants improves as their relative distance decreases. This is intuitive since as the number of AS hops in the path increases, the probability that an intermediate AS selects a compromised route increases as well. Therefore, if the group members are located close to one another, it may be possible to effectively secure intra-group communication without the help of additional deputy ASes. In contrast, groups that contain remotely located members can ensure secure intra-group communication by enlisting deputy ASes, effectively decreasing the relative distance of the remote group members.

Shout: Let's consider communication of non-participants and participants. We observe that, in contrast to intra-group communication, diversity in the location of the members improves the effectiveness of Shout. Therefore, enlisting deputy ASes can be helpful for groups of limited geographic presence. The choice of which deputy ASes to enlist critically depends on the footprint of their topology. For example, a group may prefer deputies with a strong presence in areas that contain vulnerable non-participants that are of particular interest to the group members.

We envision that the group can enlist deputy ASes in exchange for a fee. Since our solution allows creation of multiple co-existing groups, each group can enlist the nodes that are the most beneficial for the particular configuration and desired level of protection. To that end, the groups can employ the simulation techniques we introduced in this paper to find the best deputies to enlist.

9. RELATED WORK

Our research relates to earlier attempts to secure the interdomain routing system. Previous solutions have focused primarily on *cryptographic techniques* that ensure the validity of the route announcements [27, 9, 28, 12], or on *anomaly detectors* that detect (and in some cases filter) suspicious routes [16, 20, 21, 17, 19]. The cryptographic techniques are expensive and offer limited (if any) gains in small-scale deployments. Filtering invalid routes based on anomaly-detection techniques is effective, but only for larger group sizes (e.g., 50 or more ASes, including large ISPs). Our solution leverages these anomaly detectors as one way to detect, and avoid, compromised routes. Yet we show that, to be effective, small groups also need to increase path diversity and proactively coax non-participants to pick valid routes.

Our research is part of a recent body of work exploring interdomain security solutions that are successful in smaller-scale deployments. For example, the work in [29] argues that large ISPs can offer increased path diversity as a service, to offer their customers higher availability. In our work, we provide additional path diversity through secure overlay routing, and also propose Shout to help the *non-participants* reach the group members. As another example, the work in [30] models the conditions (in terms of costs and security benefits) that would lead all ISPs to adopt one of the various secure routing protocols. In contrast, we quantify the effectiveness of the protocols in small-scale deployments, and identify new mechanisms that enable small groups to be effective. Finally, the paper builds on our own earlier work [10] that proposed the secure overlay routing technique in Section 5.1 and outlined the economic argument presented in the Appendix. However, our earlier paper did not present

any simulation experiments or propose techniques for coaxing the non-participants to select valid routes.

10. CONCLUSION

In this paper, we argued that small groups of cooperating ASes should form the basis of solutions for secure interdomain routing, with an emphasis on improving end-to-end availability. We evaluated and compared several techniques that small groups can employ, to identify four conditions that enable them to be effective. In particular, the participating ASes should: (i) apply accurate techniques for detecting and filtering compromised routes, (ii) cooperate to expose additional path diversity, (iii) actively induce the non-participating ASes to select valid routes, and (iv) enlist a few large ISPs to join the group. All four conditions are important—omitting any of them results in a much less effective solution. We also proposed and evaluated a novel approach, based on secure overlays and cooperative BGP announcements, that achieves these four goals, allowing groups as small as 5–10 ASes to enjoy substantial security benefits.

As small groups become effective in securing interdomain routing, other ASes may want to join the group. As more ASes join, both parts of our solution become more effective—the overlay exposes even greater path diversity and the cooperative BGP announcements reach an even larger fraction of the remaining non-participants. Interestingly, as the group grows even larger, the "shout" mechanism becomes increasingly less necessary because most important communication stays within the group. However, as the group grows in size, the assumption that all group members trust one another becomes less reasonable. At that point, it becomes important to protect the honest members of the group from malicious *participants*. For example, the members of the group could agree to deploy a secure routing protocol *within the overlay network*, in addition to the mechanisms we have already discussed for protecting against adversaries outside of the group. In fact, the deployment of a secure routing protocol within the group, coupled with the group's growing size, could present a viable incremental deployment path for traditional cryptographic solutions, like S-BGP.

11. REFERENCES

[1] R. Mahajan, D. Wetherall, and T. Anderson, "Understanding BGP misconfiguration," in *Proc. ACM SIGCOMM*, Aug. 2002.

[2] V. J. Bono, "7007 explanation and apology," Apr. 1997. `http://www.merit.edu/mail.archives/nanog/1997-04/msg00444.html`.

[3] P. Boothe, J. Hiebert, and R. Bush, "How prevalent is prefix hijacking on the Internet?," Feb. 2006. `http://www.nanog.org/mtg-0602/boothe.html`.

[4] R. Blog, "Con-Ed steals the 'net." `http://www.renesys.com/blog/2006/01/conedstealsthenet.shtml`.

[5] `http://www.ietf.org/html.charters/rpsec-charter.html`.

[6] `http://www.nanog.org/`.

[7] *S-BGP/soBGP Panel: What Do We Really Need and How Do We Architect a Compromise to Get It?,* *http://www.nanog.org/mtg-0306/sbgp.html*, June 2003.

[8] *ARIN IX Public Policy Meeting,* *http://www.arin.net/meetings/minutes/ARIN_IX/ppm_minutes.html*, Apr. 2002.

[9] R. White, "Securing BGP through secure origin BGP," *The Internet Protocol Journal*, vol. 6, no. 3, 2003.

[10] Short workshop paper, citation removed to protect author anonymity.

[11] X. Zhao, D. Pei, L. Wang, D. Massey, A. Mankin, S. F. Wu, and L. Zhang, "An analysis of BGP multiple origin AS (MOAS) conflicts," in *Proc. ACM SIGCOMM Internet Measurement Workshop*, Nov. 2001.

[12] Y.-C. Hu, A. Perrig, and M. Sirbu, "SPV: A secure path vector routing scheme for securing BGP," in *Proc. ACM SIGCOMM*, Sept. 2004.

[13] J. Karlin, S. Forrest, and J. Rexford, *PGBGP simulator.* `http://www.cs.unm.edu/~karlinjf/pgbgp/`.

[14] `http://as-rank.caida.org/data/`.

[15] `http://www.routeviews.org/`.

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

[17] J. Karlin, S. Forrest, and J. Rexford, "Pretty Good BGP: Improving BGP by cautiously adopting routes," in *Proc. IEEE ICNP*, Nov. 2006.

[18] J. Karlin, S. Forrest, and J. Rexford, "Protecting BGP from invalid paths," Tech. Rep. TR-CS-2007-12, University of New Mexico, Computer Science, Aug. 2007.

[19] J. Qiu, L. Gao, S. Ranjan, and A. Nucci, "BGP prefix hijacking and path spoofing detection," in *Proc. SecureComm*, Sept. 2007.

[20] C. Zheng, L. Ji, D. Pei, J. Wang, and P. Francis, "A light-weight distributed scheme for detecting IP prefix hijacks in real-time," in *Proc. ACM SIGCOMM*, Aug. 2007.

[21] X. Hu and Z. M. Mao, "Accurate real-time identification of IP prefix hijacking," in *Proc. IEEE Symposium on Security and Privacy*, May 2007.

[22] N. Duffield and M. Grossglauser, "Trajectory sampling for direct traffic observation," *IEEE/ACM Trans. Networking*, vol. 9, pp. 280–292, Jun. 2001.

[23] I. Avramopoulos and J. Rexford, "Stealth probing: Efficient data-plane security for IP routing," in *Proc. USENIX Annual Technical Conference*, May/Jun. 2006.

[24] D. Andersen, H. Balakrishnan, F. Kaashoek, and R. Morris, "Resilent overlay networks," in *Proc. ACM Symposium on Operating System Principles*, Oct. 2001.

[25] T. Hardie, "Distributing authoritative name servers via shared unicast addresses," RFC 3258, IETF, Apr. 2002.

[26] P. Smith, R. Evans, and M. Hughes, "RIPE routing working group recommendations on route aggregation," Document ripe-399, RIPE, Dec. 2006.

[27] S. Kent, C. Lynn, and K. Seo, "Secure Border Gateway Protocol (Secure-BGP)," *IEEE Journal on Selected Areas in Communications*, vol. 18, pp. 582–592, Apr. 2000.

[28] T. Wan, E. Kranakis, and P. van Oorschot, "Pretty secure BGP (psBGP)," in *Proc. Network and Distributed System Security Symposium*, Feb. 2005.

[29] D. Wendlandt, I. Avramopoulos, D. Andersen, and J. Rexford, "Don't secure routing protocols, secure data delivery," in *Proc. ACM SIGCOMM HotNets Workshop*, Nov. 2006.

[30] H. Chan, D. Dash, A. Perrig, and H. Zhang, "Modeling adoptability of secure BGP protocols," in *Proc. ACM SIGCOMM*, Sept. 2006.

[31] R. Cornes and T. Sandler, *The Theory of Externalities, Public Goods and Club Goods.* Cambridge University Press, second ed., 1996.

[32] M. Olson, *The Logic of Collective Action.* Harvard University Press, 1971.

APPENDIX

A. ECONOMIC CASE FOR SMALL GROUPS

In this appendix, we motivate our design decision to support the formation of multiple coexisting small groups as a basis for providing secure interdomain communication. First, we argue that a *purist* solution that requires the ubiquitous deployment of a secure routing protocol prevents market forces from driving adoption. Then, we argue for a *pluralist* approach that supports customized security solutions for groups of various sizes and is consistent with market forces.

A.1 Economics of Groups and Goods

Secure interdomain communication requires action in *groups*. Although there are techniques that an AS acting alone can use to reduce the likelihood of attacks (such as applying protective filters to routing protocol messages and data packets), these techniques cannot ensure confidentiality, integrity, and availability for interdomain communication. Symmetric encryption instead, the simplest technique to ensure confidentiality, requires bilateral cooperation to establish a security association and encrypt/decrypt the data. The formation of *groups* of ASes is, therefore, essential for interdomain communication security.

Because the ASes in the commercial Internet are independent, payoff-maximizing entities, it is important to consider the economic incentives of individual ASes to join groups that provide interdomain communication security services. From this perspective, interdomain communication security is an economic *good* that the group provides to its members by deploying common security mechanisms. Depending on the interaction among the group members goods are, in general, classified as (1) purely public, (2) purely private, and (3) impurely public, with different economic implications. Pure public goods are *non-rival*, i.e., consumption of the good by one member does not diminish the availability of the good to other members, and *non-excludable*, i.e., the privilege of consumption of the good is unrestricted. An example of a pure public good is public television broadcasting. In contrast, pure private goods are rival and excludable, for example, recorded music sold in music stores. Impure public goods are partially rival or partially excludable, such as cable television broadcasting.

The appropriate classification of a good depends on the group's incentive structure for production and consumption. In fact, technological innovations can transform a good from one class to another. For example, encryption changed television broadcasting from a pure public good to an impure public good. As another example, peer-to-peer file-sharing applications are rapidly transforming recorded music from a pure private good to a pure public good. The rest of this section argues, first, that the purist view treats secure interdomain communication as a pure public good, which prevents market forces from driving adoption, and, second, that a pluralist approach in which smaller groups coexist better matches the economic incentives of those groups.

A.2 Purism is not Economically Viable

The ubiquitous deployment of a secure routing protocol that purism requires is unappealing because it implies *non-excludability*. Consider, for example, an exclusion mechanism based on fees. The option of charging a fee to prospective customer networks for connecting them to your secure routing protocol, implies the possibility of networks that decline to pay the fee, therefore, leading to partial, non-ubiquitous deployment of the protocol. I.e., the option of receiving an economic gain through the deployment of a secure routing protocol is inconsistent with the ubiquity requirement of purism. In the absence of other sources of revenue (e.g., advertising), non-excludability leads to *market failure*, i.e., no supply of the good, or a level of provision that is grossly inefficient. This is the situation today, where no secure interdomain routing protocol is deployed, despite a pressing need for better security.

Avoiding market failure under non-excludability typically requires government intervention, such as regulation [31]. However, resorting to regulatory action to mandate the ubiquitous deployment of a secure routing protocol is unnecessary, and in fact may stifle the creation and deployment of superior alternatives. Instead, we believe it is possible for market forces to drive the deployment of security mechanisms, including the existing and novel secure routing protocols, just not based on the purist view. In the rest of this section, we advocate *pluralism*, i.e., the coexistence of multiple groups of various sizes, and discuss market-based incentive structures for secure communication.

A.3 Pluralism is Incentive Compatible

As mentioned earlier, whether a group will form to counteract a threat will ultimately depend on the economic incentives of individual ASes to join. The theory of collective action [32] argues that small and medium-sized groups are more effective in providing public goods than large ones. In a small group, one large member may have sufficient incentive to provide the good by himself, essentially financing the participation of the other members. For example, a large corporation may finance the deployment of encryption devices at smaller business partners for business-to-business transactions. In a medium-sized group, a good can be provided by strategic interaction and bargaining. For example, large backbone providers may form a coalition to deploy a secure routing protocol to protect their customers.

Accommodating independent variable-sized groups, instead of mandating the formation of a single large group that purism advocates, would enable market forces alone to drive the provision of communication-security goods, in two main ways. First, as noted above, a small or medium-sized group may provide a pure public good based solely on alignment of incentives or bargaining. Second, exclusion mechanisms can be leveraged to provide goods as impurely public or purely private. For example, a coalition of networks that had deployed a secure routing protocol may charge non-member networks to use its routes.