

Let the Market Drive Deployment: A Strategy for Transitioning to BGP Security

Phillipa Gill
University of Toronto

Michael Schapira
Princeton University

Sharon Goldberg
Boston University

Abstract

With a cryptographic root-of-trust for Internet routing (RPKI [17]) on the horizon, we can finally start planning the deployment of one of the secure interdomain routing protocols proposed over a decade ago (Secure BGP [22], secure origin BGP [37]). However, if experience with IPv6 is any indicator, this will be no easy task. Security concerns alone seem unlikely to provide sufficient local incentive to drive the deployment process forward. Worse yet, the security benefits provided by the S*BGP protocols do not even kick in until a large number of ASes have deployed them.

Instead, we appeal to ISPs' interest in increasing revenue-generating traffic. We propose a strategy that governments and industry groups can use to harness ISPs' local business objectives and drive global S*BGP deployment. We evaluate our deployment strategy using theoretical analysis and large-scale simulations on empirical data. Our results give evidence that the market dynamics created by our proposal can transition the majority of the Internet to S*BGP.

Categories and Subject Descriptors: C.2.2 [Computer-Communication Networks]: Network Protocols

General Terms: Economics, Security

1. INTRODUCTION

The Border Gateway Protocol (BGP), which sets up routes from autonomous systems (ASes) to destinations on the Internet, is amazingly vulnerable to attack [7]. Every few years, a new failure makes the news; ranging from misconfigurations that cause an AS to become unreachable [34, 29], to possible attempts at traffic interception [11]. To remedy this, a number of widely-used stop-gap measures have been developed to *detect* attacks [20, 25]. The next step is to harden the system to a point where attacks can be *prevented*. After many years of effort, we are finally seeing the initial deployment of the Resource Public Key Infrastructure (RPKI) [4, 27], a cryptographic root-of-trust for Internet routing that authoritatively maps ASes to their IP prefixes and public keys. With RPKI on the horizon, we

can now realistically consider deploying the S*BGP protocols, proposed a decade ago, to prevent routing failures by validating AS-level paths: Secure BGP (S-BGP) [22] and Secure Origin BGP (soBGP) [37].

1.1 Economic benefits for S*BGP adoption.

While governments and industry groups may have an interest in S*BGP deployment, ultimately, the Internet lacks a centralized authority that can mandate the deployment of a new secure routing protocol. Thus, a key hurdle for the transition to S*BGP stems from the fact that each AS will make deployment decisions according to its own local business objectives.

Lessons from IPv6? Indeed, we have seen this problem before. While IPv6 has been ready for deployment since around 1998, the lack of tangible local incentive for IPv6 deployment means that we are only now starting to see the seeds of large-scale adoption. Conventional wisdom suggests that S*BGP will suffer from a similar lack of local incentives for deployment. The problem is exacerbated by the fact that an AS cannot validate the correctness of an AS-level path unless all the ASes on the path deployed S*BGP. Thus, the security benefits of S*BGP only apply after a large fraction of ASes have already deployed the protocol.

Economic incentives for adoption. We observe that, unlike IPv6, S*BGP can impact routing of Internet traffic, and that this may be used to drive S*BGP deployment. These crucial observations enable us to avoid the above issues and show that global S*BGP deployment is possible even if local ASes' deployment decisions are *not* motivated by security concerns! To this end, we present a prescriptive strategy for S*BGP deployment that relies solely on Internet Service Providers' (ISPs) local economic incentives to drive global deployment; namely, ISP's interest in attracting revenue-generating traffic to their networks.

Our strategy is prescriptive (Section 2). We propose guidelines for how (a) ASes should deploy S*BGP in their networks, and (b) governments, industry groups, and other interested parties should invest their resources in order to drive S*BGP deployment forward.

1. Break ties in favor of secure paths. First, we require ASes that deploy S*BGP to actually use it to inform route selection. However, rather than requiring security be the first criterion ASes use to select routes, we only require secure ASes to *break ties* between equally-good routes in favor of secure routes. This way, we create incentives for ISPs to deploy S*BGP so they can transit more revenue-generating customer traffic than their insecure competitors.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SIGCOMM'11, August 15-19, 2011, Toronto, Ontario, Canada.
Copyright 2011 ACM 978-1-4503-0797-0/11/08 ...\$10.00.

2. Make it easy for stubs to adopt S*BGP. 85% of ASes in the Internet are *stubs* (i.e., ASes with no customers) [9]. Because stubs earn no revenue from providing Internet service, we argue for driving down their deployment costs by having ISPs sign BGP announcements on their behalf or deploy a simplex (unidirectional) S*BGP [26] on their stub customers. In practice, such a simplex S*BGP must either be extremely lightweight or heavily subsidized.

3. Create market pressure via early adopters. We propose that governments and industry groups concentrate their regulatory efforts, or financial incentives, on convincing a small set of *early adopters* to deploy S*BGP. We show that this set of early adopters can create sufficient market pressure to convince a large fraction of ASes to follow suit.

1.2 Evaluation: Model and simulations.

To evaluate our proposal, we needed a model of the S*BGP deployment process.

Inspiration from social networks? At first glance, it seems that the literature on technology adoption in social networks would be applicable here (e.g., [30, 21] and references therein). However, in social networks models, an entity’s decision to adopt a technology depends only on its immediate *neighbors* in the graph; in our setting, this depends on the number of secure *paths*. This complication means that many elegant results from this literature have no analogues in our setting (Section 9).

Our model. In contrast to earlier work that assumes that ASes deploy S*BGP because they are concerned about security [8, 5], our model assumes that ISPs’ local deployment decisions are based solely on their interest in increasing customer traffic (Section 3).

We carefully designed our model to capture a few crucial issues, including the fact that (a) traffic transited by an ISP can include flows from any pair of source and destination ASes, (b) a large fraction of Internet traffic originates in a few large content provider ASes [24], and (c) the cost of S*BGP deployment can depend on the size of the ISP’s network. The vast array of parameters and empirical data relevant to such a model (Section 8) mean that our analysis is *not* meant to *predict* exactly how the S*BGP deployment process will proceed in practice; instead, our goal was to evaluate the efficacy of our S*BGP deployment strategy.

Theorems, simulations and examples. We explore S*BGP deployment in our model using a combination of theoretical analysis and simulations on empirical AS-level graphs [9, 3] (Sections 5-7). Every example we present comes directly from these simulations. Instead of artificially reducing algorithmic complexity by subsampling [23], we ran our simulations over the full AS graph (Section 4). Thus, our simulations ran in time $O(N^3)$ with $N = 36K$, and we devoted significant effort to developing parallel algorithms that we ran on a 200-node DryadLINQ cluster [38].

1.3 Key insights and recommendations.

Our evaluation indicates that our strategy for S*BGP deployment can drive a transition to S*BGP (Section 5). While we cannot predict exactly how S*BGP deployment will progress, a number of important themes emerge:

1. Market pressure can drive deployment. We found that when S*BGP deployment costs are low, the vast majority of ISPs have incentives to deploy S*BGP in order to dif-

ferentiate themselves from, or keep up with, their competitors (Section 5). Moreover, our results show this holds even if 96% of routing decisions (across all source-destination AS pairs) are *not* influenced by security concerns (Section 6.6).

2. Simplex S*BGP is crucial. When deployment costs are high, deployment is primarily driven by simplex S*BGP (Section 6).

3. Choose a few well-connected early adopters. The set of early adopters cannot be random; it should include well-connected ASes like the Tier 1’s and content providers (Section 6). While we prove that it is NP-hard to even *approximate* the *optimal* set of early adopters (Section 6.1), our results show that even 5-10 early adopters suffice when deployment costs are low.

4. Prepare for incentives to disable S*BGP. We show that ISPs can have incentives to *disable* S*BGP (Section 7). Moreover, we prove that there could be deployment oscillations (where ASes endlessly turn S*BGP on and off), and that it is computationally hard to even *determine* whether such oscillations exist.

5. Minimize attacks during partial deployment. Even when S*BGP deployment progressed, there were always some ASes that did not deploy (Section 5, 6). As such, we expect that S*BGP and BGP will coexist in the long term, suggesting that careful engineering is required to ensure that this does not introduce new vulnerabilities into the interdomain routing system.

Paper organization. Section 2 presents our proposed strategy for S*BGP deployment. To evaluate the proposal, we present a model of the deployment process in Section 3. In Section 5-7 we explore this model using theoretical analysis and simulations, and present an in-depth discussion of our modeling assumptions in Section 8. Section 9 presents related work. The full version of this paper [2] contain implementation details for our simulations, proofs of all our theorems, and supplementary data analysis.

2. S*BGP DEPLOYMENT STRATEGY

2.1 S*BGP: Two possible solutions.

With RPKI providing an authoritative mapping from ASes to their cryptographic public keys, two main protocols have been proposed that prevent the propagation of bogus AS path information:

Secure BGP (S-BGP) [22]. S-BGP provides *path validation*, allowing an AS a_1 that receives a BGP announcement $a_1 a_2 \dots a_k d$ to validate that every AS a_j actually sent the announcement in the path. With S-BGP, a router must cryptographically sign each routing message it sends, and cryptographically verify each routing message it receives.

Secure Origin BGP (soBGP) [37]. soBGP provides a slightly weaker security guarantee called *topology validation*, that allows an AS to validate that a path it learns physically exists in the network. To do this, soBGP requires neighboring ASes to mutually authenticate a certificate for the existence of a link between them, and validate every path it learns from a BGP announcement against these cryptographic certificates.

Because our study is indifferent to attacks and adversaries, it applies equally to each of these protocols. We refer to

them collectively as S*BGP, and an AS that deploys them as *secure*.

2.2 How to standardize S*BGP deployment.

To create local economic incentives for ISPs to deploy S*BGP, we propose that Internet standards should require ASes to deploy S*BGP as follows:

2.2.1 Simplex S*BGP for stubs.

For stubs, Internet access is a cost, rather than a revenue source, and it seems unlikely that security concerns alone will suffice to motivate stubs to undertake a costly S*BGP deployment. However, because stubs propagate only *outgoing* BGP announcements for *their own IP prefixes* we suggest two possible solutions to this problem: (1) allow ISPs to sign on behalf of their stub customers or (2) allow stubs to deploy simplex (unidirectional) S*BGP. Indeed, the latter approach has been proposed by the Internet standards community [26].

Simplex S-BGP. For S-BGP, this means that stubs need only sign outgoing BGP announcements for their own IP prefixes, but not validate incoming BGP announcements for other IP prefixes¹. Thus, a stub need only store its own public key (rather than obtaining the public keys of each AS on the Internet from the RPKI) and cryptographically sign only a tiny fraction of the BGP announcements it sees. Simplex S-BGP can significantly decrease the computational load on the stub, and can potentially be deployed as a software, rather than hardware, upgrade to its routers.

Simplex soBGP. For soBGP, this means that a stub need only create certificates for its links, but need not need validate the routing announcements it sees. Simplex soBGP is done offline; once a stub certifies his information in the soBGP database, its task is complete and no router upgrade is required.

The objective of simplex S*BGP is to make it easy for stubs to become secure by lowering deployment costs and computational overhead. While we certainly allows for stubs (*e.g.*, banks, universities) with an interest in security to move from simplex S*BGP to full S*BGP, our proposal does not require them to do so.

Impact on security. With simplex S*BGP, a stub lacks the ability to validate paths for prefixes other than its own. Since stubs constitute about 85% of ASes [9], a first glance suggests that simplex S*BGP leads to significantly worse security in the global Internet.

We argue that this is not so. Observe that if a stub s has an immediate provider p that has deployed S*BGP and is correctly validating paths, then no false announcements of fully secure paths can reach s from that provider, unless p *himself* maliciously (or mistakenly) announces false secure paths to s . Thus, in the event that stubs upgrade to simplex S*BGP and all other ASes upgrade to full S*BGP, the only open attack vector is for ISPs to announce false paths to their *own* stub customers. However, we observe the impact of a single misbehaving ISP is small, since 80% of ISPs have less than 7 stub customers, and only about 1% of ISPs have more than 100 stub customers [9]. Compare this to the

¹A stub may even choose to delegate its cryptographic keys to its ISPs, and have them sign for him; while this might be a good first step on the path to deployment, ceding control of cryptographic keys comes at the cost of reduced security.

insecure status quo, where an arbitrary misbehaving AS can impact about half of the ASes in the Internet (around 15K ASes) on average [14].

2.2.2 Break ties in favor of fully secure paths.

In BGP, an AS chooses the path to a given destination AS d based on a *ranking* on the outgoing paths it learns from its neighbors (*e.g.*, Appendix A). Paths are first ranked according to interdomain considerations (local preference, AS path length) and then according to intradomain considerations (*e.g.*, MEDs, hot-potato routing)².

Secure paths. We say that a path is secure iff *every* AS on that path is secure. We do this because an AS cannot validate a path unless every AS on the path signed the routing announcement (S-BGP) or issued certificates for the links on the path (soBGP).

Security as part of route selection. The next part of our proposal suggests that once an AS has the ability to validate paths, it should actually use this information to inform its routing decisions. In principle, an AS might even modify its ranking on outgoing paths so that security is its highest priority. Fortunately, we need not go to such lengths. Instead, we only require secure ASes to *break ties* between equally good interdomain paths in favor of secure paths. This empowers secure ISPs to attract customer traffic away from their insecure competitors. To ensure that a newly-secure AS can *regain* lost customer traffic, we require that original tie-break criteria (*e.g.*, intradomain considerations) be employed in the case of equally good, *secure* interdomain paths. Thus, the size of the set of equally-good interdomain paths for a given source-destination pair (which we call the *tiebreak set*) gives a measure of competition in the AS graph.

Route selection at stubs. For stubs running simplex S*BGP, we consider both the case where they break ties in favor of secure paths (*i.e.*, because they trust their providers to verify paths for them) and the case where they ignore security altogether (*i.e.*, because they do not verify paths) (Section 6.7).

Partially secure paths. We do not allow ASes to prefer partially-secure paths over insecure paths, to avoid introducing *new* attack vectors that do exist even without S*BGP (*e.g.*, attack in Appendix B).

We shall show that S*BGP deployment progresses quite effectively even if stubs ignore security and tiebreak sets are very small (Section 6.7-6.6).

2.3 How third parties should drive deployment.

Early adopters. To kick off the process, we suggest that interested third parties (*e.g.*, governments, regulators, industry groups) focus regulation, subsidies, or external financial incentives on convincing a set of *early adopter* ASes to deploy S*BGP. One regulatory mechanism may be for the government to require their network providers to deploy S*BGP first. In the AS graph ([9, 3]), providers to the government include many Tier 1 ISPs who may be difficult or expensive to persuade via other means.

ISPs upgrade their stubs. Next, we suggest that a secure ISP should be responsible for upgrading all its insecure stub customers to simplex S*BGP. To achieve this,

²For simplicity, we do not model intradomain routing considerations. However, it should be explored in future work.

interested third parties should ensure that simplex S*BGP is engineered to be as lightweight as possible, and potentially provide additional subsidies for ISPs that secure their stubs. (ISPs also have a local incentives to secure stubs, *i.e.*, to transit more revenue-generating traffic for multi-homed stubs (Section 5.1).)

3. MODELING S*BGP DEPLOYMENT

We evaluate our proposal using a model of the S*BGP deployment process. For brevity, we now present only the details of our model. Justification for our modeling decisions and possible extensions are in Section 8.

3.1 The Internet and entities.

The AS graph. The interdomain-routing system is modeled with a labeled AS graph $G(V, E)$. Each node $n \in V$ represents an AS, and each edge represents a physical link between ASes. Per Figure 1, edges are annotated with the standard model for business relationships in the Internet [13]: *customer-provider* (where the customer pays the provider), and *peer-to-peer* (where two ASes agree to transit each other’s traffic at no cost). Each AS n is also assigned weight w_n , to model the volume of traffic that *originates* at each AS. For simplicity, we assume ASes divide their traffic evenly across all destination ASes. However, our results are robust even when this assumption is relaxed (Section 6.8).

We distinguish three types of ASes:

Content providers. Content providers (CPs) are ASes whose revenue (*e.g.*, advertising) depends on reliably delivering their content to as many users as possible, rather than on providing Internet transit. While a disproportionately large volume of Internet traffic is known to originate at a few CPs, empirical data about Internet traffic volumes remains notoriously elusive. Thus, based on recent research [24, 35] we picked five content providers: Google (AS 15169), Facebook (AS 32934), Microsoft (AS 8075), Akamai (AS 20940), and Limelight (AS 22822). Then, we assigned each CP weight w_{CP} , so that the five CPs originate a x fraction of Internet traffic (equally split between them), with the remaining $1 - x$ split between the remaining ASes.

Stubs. Stubs are ASes that have no customers of their own and are not CPs. Every stub s has unit weight $w_s = 1$. In Figure 1, ASes 34376 and 31420 are stubs.

ISPs. The remaining ASes in the graph (that are not stubs or CPs) are ISPs. ISPs earn revenue by providing Internet service; because ISPs typically provide transit service, rather than originating traffic (content), we assume they have unit weight $w_n = 1$. In Figure 1, ASes 25076, 8866 and 8928 are ISPs.

3.2 The deployment process.

We model S*BGP deployment as an infinite round process. Each round is represented with a state S , capturing the set of ASes that have deployed S*BGP.

Initial state. Initially, the only ASes that are secure are (1) the ASes in the set of early adopters and (2) the direct customers of the early adopter ISPs that are stubs. (The stubs run simplex S*BGP.) All other ASes are insecure. For example, in Figure 1, early adopters ISP 8866 and CP 22822 are secure, and stub 31420 runs simplex S*BGP because its provider is secure.




Figure 1: Destinations (left) 31420, (right) 22822.

Each round. In each round, *every* ISP chooses an action (deploy S*BGP or not) that improves its utility relative to the current state. We discuss the myopic best-response strategy that ISPs use to choose their actions in Section 3.3. Once an ISP becomes secure, it deploys simplex S*BGP at *all* its stub customers (Section 2.3). Because CPs do not earn revenues by providing Internet service, some external incentive (*e.g.*, concern for security, subsidies) must motivate them to deploy S*BGP. Thus, in our model, a CP may only deploy S*BGP if it is in the set of early adopters.

Once ASes choose their actions, paths are established from every source AS i to every destination AS d , based on the local BGP *routing policies* of each AS and the state S of the AS graph. We use a standard model of BGP routing policies, based on business relationships and path length (see Appendix A). Per Section 2.3, we also assume that routing policies of secure ASes require them to break ties by preferring fully secure paths over insecure ones, so that the path to a given destination d depends on the state S . Paths to a destination d form a tree rooted at d , and we use the notation $T_n(d, S)$ to represent the subtree of ASes routing through AS n to a destination d when the deployment process is in state S . Figure 1 (right) shows part of the routing tree for destination 22822; notice that $T_{8866}(22822, S)$ contains ASes 31420, 25076, 34376.

Termination. We proceed until we reach a stable state, where no ISP wants to deploy (or disable) S*BGP.

3.3 ISP utility and best response.

We model an ISP’s utility as related to the volume of traffic it transits for its customers; this captures the fact that many ISPs either bill their customers directly by volume, or indirectly through flat rates for fixed traffic capacities. Utility is a function of the paths chosen by each AS. Because path selection is a function of routing policies (Appendix A) and the state S , it follows that *the utility of each ISP is completely determined by the AS weights, AS graph topology, and the state S .*

We have two models of ISP utility that capture the ways in which an ISP can transit customer traffic:

Outgoing utility. ISP n can increase its utility by forwarding traffic to its customers. Thus, we define outgoing utility as the amount of traffic that ISP n routes to each destination d via a customer edge. Letting $\hat{D}(n)$ be the set of such destinations, we have:

$$u_n(S) = \sum_{d \in \hat{D}(n)} \sum_{i \in T_n(d, S)} w_i \quad (1)$$

Let’s use Figure 1 to find the outgoing utility of ISP $n = 8866$ due to destinations 31420 and 22822. Destination 31420

is in $\hat{D}(n)$ but destination 22822 is not. Thus, two CPs (Google AS 15169 and Limelight 22822), and 3 other ASes (*i.e.*, AS 8928, 25076, 34376) transit traffic through $n = 8866$ to destination $d = 31420$, contributing a $2w_{CP} + 3$ outgoing utility to $n = 8866$.

Incoming utility. An ISP n can increase its utility by forwarding traffic *from* its customers. Thus, we define incoming utility as the amount of traffic that ISP n receives via customer edges for each destination d . We restrict the subtree $T_n(d, S)$ to branches that are incident on n via customer edges to obtain the *customer subtree* $\hat{T}_n(d, S) \subset T_n(d, S)$, we have:

$$u_n(S) = \sum_d \text{Destns} \sum_{i \in \hat{T}_n(d, S)} \text{Sources} w_i \quad (2)$$

Let’s compute outgoing utility of $n = 8866$ due to destinations 31420 and 22822 in Figure 1. For destination 31420, ASes 25076 and 34376 are part of the customer subtree $\hat{T}_n(d, S)$, but 15169, 8928 and 22822 are not. For destination $d = 22822$, ASes 31420, 25076, 34376 are part of the customer subtree. Thus, these ASes contribute $2 + 3$ incoming utility to ISP $n = 8866$.

Realistically, ISP utility is some function of both of these models; to avoid introducing extra parameters into our model, we consider each separately.

Myopic best response. We use a standard game-theoretic update rule known as *myopic best response*, that produces the most favorable outcome for a node in the next round, taking other nodes’ strategies as given [16]. Let $(\neg S_n, S_{-n})$ denote the state when n ‘flips’ to the opposite action (either deploying or undeploying S*BGP) that it used in state S , while other ASes maintain the same action they use in state S . ISP n changes its action in state S iff its *projected utility* $u_n(\neg S_n, S_{-n})$ is sufficiently high, *i.e.*,

$$u_n(\neg S_n, S_{-n}) > (1 + \theta) \cdot u_n(S) \quad (3)$$

where θ is a threshold denoting the increase in utility an ISP needs to see before it is willing to change its actions. Threshold θ captures the cost of deploying BGP security; *e.g.*, an ISP might deploy S*BGP in a given round if S*BGP deployment costs do not exceed $\theta = 5\%$ of the profit it earns from transiting customer traffic. Since θ is multiplicative, it captures the idea that deployment costs are likely to be higher at ISPs that transit more traffic. The update rule is myopic, because it focuses on increasing ISP n ’s utility in the next round only. It is best-response because it does *not* require ISP n to speculate on other ASes’ actions in future rounds; instead, n takes these actions as given by the current state S .

Discussion. Our update rule requires ASes to predict their future utility. In our model, ASes have full information of S and G , a common approach in game theory, which enables them to project their utility accurately. We discuss the consequences of our update rule, and the impact of partial information in Sections 8.1-8.2.

4. SIMULATION FRAMEWORK

Computing utility $u_n(S)$ and projected utility $u_n(\neg S_n, S_{-n})$ requires us to determine the path from *every* source AS to *every* destination AS, for *every* ISP n ’s unique projected

state $(\neg S_n, S_{-n})$. Thus, our simulations had complexity $O(|V|^3)$ on an AS graph $G(V, E)$. To accurately simulate our model, we chose *not* to ‘sample down’ the complexity of our simulations:

Projecting utility for each ISP. If we had computed the utility for only a few sampled ISPs, this would reduce the number of available secure paths and artificially prevent S*BGP deployment from progressing.

Simulations over the entire AS graph. Our proposal is specifically designed to leverage the extreme skew in AS connectivity (*i.e.*, many stubs with no customers, few Tier 1s with many customers), to drive S*BGP deployment. To faithfully capture the impact of this skew, we computed utility over traffic from *all* sources to *all* destination ASes. Furthermore, we ran our simulations on the full empirical AS graph [9], rather than a subsampled version [23], or a smaller synthetic topology [28, 39], as in prior work [8, 5]. We used the Cyclops AS graph (with its inferred AS relationships) from Dec 9, 2010 [9], with an additional 16K peering edges discovered at Internet exchange points (IXPs) [3], as well as an additional peering-heavy AS graph described in Section 6.8.

The AS graph $G(V, E)$ had $|V| = 36K$; to run $O(|V|^3)$ -simulations at such a scale, we parallelized our algorithms on a 200-node DryadLINQ cluster [38] that could run through a single simulation in 1-12 hours. (Details of our implementation are in the full version.)

5. CASE STUDY: S*BGP DEPLOYMENT

We start by showing that even a small set of early adopters can create enough market pressure to transition the vast majority of ASes to S*BGP.

Case study overview. We focus on a single simulation where the early adopters are the five CPs (Google, Facebook, Microsoft, Limelight, Akamai, see Section 3.1), and the top five Tier 1 ASes in terms of degree (Sprint (1239), Verizon (701), AT&T (7018), Level 3 (3356), Cogent (174)). Every ISP uses an update rule with a relatively low threshold $\theta = 5\%$, that the five CPs originate $x = 10\%$ of the traffic in the Internet, and that stubs *do* break ties in favor of secure routes. We now show how even a small set of ten early adopters (accounting for less than 0.03% of the AS graph) can convince 85% of ASes to deploy S*BGP, and secure 65% of all paths in the AS graph.

5.1 Competition drives deployment.

We start by zooming in on S*BGP deployment at two competing ISPs, in a scenario we call a DIAMOND.

Figure 5: Two ISPs, AS 8359 and AS 13789, compete for traffic from Sprint (AS 1239) to their stub customer, AS 18608. Sprint is an early adopter of S*BGP, and initially the three other ASes are insecure. Both ISPs offer Sprint equally good two-hop customer paths to the stub, and AS 8359 is chosen to carry traffic by winning the tie break. In the first round, AS 13789 computes its projected utility, and realizes it can gain Sprint’s traffic by adopting S*BGP and upgrading its stub to simplex S*BGP. (See Section 8.2 for more discussion on how ISPs compute projected utility.) By the fourth round, AS 8359 has lost so much utility (due to traffic lost to ASes like 13789) that he decides to deploy S*BGP.




Figure 2: The number of ASes that deploy S*BGP each round.




Figure 3: Normalized utility of ISPs in Fig. 5 and 6.




Figure 4: Projected and actual utility before deploying S*BGP normalized by starting utility.




Figure 5: A Diamond: ISPs 13789 and 8359 compete for traffic from Sprint (AS 1239).

Of course, Figure 5 is only a very small snapshot of the competition for traffic destined to a single stub AS 18608; utility for each ISPs is based on customer traffic transited to *all* destinations in the AS graph. Indeed, this DIAMOND scenario is quite common. We counted more than 6.5K instances of the DIAMOND, each involving two ISPs, a stub, and one of our early adopters.

5.2 Global deployment dynamics.

Figure 2: We show the number of ASes (*i.e.*, stubs, ISPs and CPs) and the number of ISPs that deploy S*BGP at each round. In the first round, 548 ISPs become secure; because each of these ISPs deploy simplex S*BGP in their stubs, we see that over 5K ASes become secure by the end of the first round. In subsequent rounds, hundreds of ISPs deploy S*BGP in each round; however, the number of newly secure stubs drops dramatically, suggesting that many ISPs deploy S*BGP to regain traffic lost when their stubs were secured by competitors. After the 17th iteration, the process tapers off, with fewer than 50 ASes becoming secure in each round. The final surge in deployment occurs in round 25, when a large AS, 6939, suddenly became secure, causing a total of 436 ASes to deploy S*BGP in the remaining six rounds. When the process terminates, 85% of ASes are secure, including 80% of the 6K ISPs in the AS graph.

5.3 Longer secure paths sustain deployment.

In Figure 2 we observed rapid, sustained deployment of S*BGP in the first 17 iterations. This happens because longer secure paths are created as more ASes deploy S*BGP, thus creating incentives for S*BGP at ASes that are far away from the early adopters:

Figure 6: We once again encounter AS 8359 from Figure 5. We show how AS 8359’s decision to deploy S*BGP in round 4 allows a new ISP (AS 6371) to compete for traffic. In round 5 AS 6731 sees a large increase in utility by becoming secure.




Figure 6: A newly created four-hop secure path.

This occurs, in part, because AS 6371 can now entice six of the early adopters to route through him on a total of 69 newly-secure paths. Indeed, when AS 6371 becomes secure, he continues the chain reaction set in motion by AS 8359; for instance, in round 7 (not shown), AS 6371’s neighbor AS 41209 becomes secure in order to offer Sprint a new, secure four-hop path to one of 41209’s own stubs.

5.4 Keeping up with the competition.

Two behaviors drive S*BGP deployment in a DIAMOND. First, an ISP becomes secure to steal traffic from a competitor, and then the competitor becomes secure in order to regain the lost traffic. We can watch this happening for the ISPs from Figure 5 and 6:

Figure 3: We show the utilities of ISPs 8359, 6731, and 8342 in each round, normalized by *starting utility* *i.e.*, the utility before the deployment process began (when all ASes, including the early adopters, were still insecure). As we saw in Figure 5, AS 8359 deploys S*BGP in round 4 in order to regain traffic he lost to his secure competitors; here we see that in round 4, AS 8359 has lost 3% of his starting utility. Once AS 8359 deploys S*BGP, his utility jumps up to more than 125% of his starting utility, but these gains in utility are only temporary, disappearing around round 15. The same is true in round 6 for AS 6371 from Figure 6. By round 15, 60% ISPs in the AS graph are already secure (Figure 2), and our ISPs can no longer use security to differentiate themselves, causing their utility to return to within 3% of their starting utility.

This is also true more generally:

Figure 4: For each round i , we show the median utility and median projected utility for ISPs that become secure in round $i+1$, each normalized by starting utility. (Recall from (3) that these ISPs have projected utility at least $1+\theta$ times their utility in round i .) In the first 9 rounds, ISPs mainly

deploy S*BGP to steal traffic from competitors; that is, their projected utility in the round before they deploy S*BGP is at least $1 + \theta = 105\%$ times their starting utility. However, as deployment progresses, ASes increasingly deploy S*BGP in order to recover lost traffic and return to their starting utility; that is, in rounds 10-20 ISP utility drops to at least $\theta = 5\%$ less than starting utility, while projected utility approaches starting utility ($y=1$).

5.5 Is S*BGP deployment a zero-sum game?

Our model of S*BGP deployment is indeed a zero-sum game; we assume that ISPs compete over a fixed set of customer traffic. Thus, when the vast majority of ASes have deployed S*BGP, ISPs can no longer use security to distinguish themselves from competitors (Figure 3). At the termination of this case study, only 8% of ISPs have an increase in utility of more than $\theta = 5\%$ over their starting utility. On the other hand, 85% of ASes now benefit from a (mostly) secure Internet. Furthermore, like ASes 8359 and 6731 in Figure 3, many of these secure ASes enjoyed a prolonged period of increased utility that could potentially help defray the costs of deploying S*BGP.

It is better to deploy S*BGP. One might argue that a cynical ISP might preempt the process by *never* deploying S*BGP. However, a closer look shows that its almost always in the ISPs interest to deploy S*BGP. ISPs that deploy S*BGP usually return to their starting utility or slightly above, whereas ISPs that do *not* deploy S*BGP lose traffic in the long term. For instance, AS 8342 in Figure 6 never deploys S*BGP. As shown in Figure 3, when the deployment process terminates, AS 8342 has lost 4% of its starting utility. Indeed, another look at the data (not shown) shows that the ISPs that remain insecure when the process terminates lose on average 13% of their starting utility!

6. CHOOSING EARLY ADOPTERS

Next, we consider choosing the set of ASes that should be targeted to become early adopters of S*BGP.

6.1 It's hard to choose early adopters.

Ideally, we would like to choose the *optimal* set of early adopters that could cause the maximum number of other ASes to deploy S*BGP. We show that this is NP-hard by presenting a reduction to the 'set cover' problem (proof in the full version):

THEOREM 6.1. *For an AS graph $G(V, E)$ and a parameter $1 \leq k \leq |V|$, finding a set of early adopter ASes of size k that maximizes the number of ASes that are secure when the deployment process terminates is NP-hard. Approximating the solution within a constant factor is also NP-hard.*

As such, we use simulations³ of the deployment process to investigate heuristic approaches for choosing early adopters, including AS degree (*e.g.*, Tier 1s) and volume of traffic originated by an AS (*e.g.*, content providers).

6.2 The parameter space.

We consider how the choice of early adopters is impacted by assumptions on (1) whether or not stubs running simplex

³Since there is no sampling involved, there is no variability between simulations run with the same set of parameters.

S*BGP break ties based on security, (2) the AS graph, and (3) traffic volumes sourced by CPs.

Outgoing utility. Also, recall that we have two models of ISP utility (Section 3.3). In this section, we dive into the details of the *outgoing utility* model because it has the following very nice property:

THEOREM 6.2. *In the outgoing utility model, a secure node will never have an incentive to turn off S*BGP.*

As a consequence of this theorem (proof in the full version), it immediately follows that (a) every simulation must terminate, and (b) we can significantly reduce compute time by *not* computing projected utility for ISPs that are already secure. (We discuss complications that arise from the incoming utility model in Section 7.)

Deployment threshold θ . Our update rule (3) is such that ISPs change their actions if they can increase utility by at least θ . Thus, to gain insight into how 'difficult' it is to convince ISPs to deploy S*BGP, we assume that each ISP uses the same threshold θ , and sweep through different values of θ (but see also Section 8.2).

6.3 Comparing sets of early adopters.

We next explore the influence of different early adopters:

Figure 7 (top): We show the fraction of ASes that adopt S*BGP for different values of θ . We consider no early adopters, the top 5-200 ISPs in terms of degree, the five CPs, five CPs in combination with the top five ISPs, and 200 random ISPs.

There are incentives to deploy S*BGP. For low values of $\theta < 5\%$, we observe that there is sufficient competition over customer traffic to transition 85% of ASes to S*BGP. Moreover, this holds for almost every set of early adopters we considered. (Note that in the unrealistic case where $\theta = 0$, we see widespread S*BGP deployment even with *no* early adopters, because we assume the stubs break ties in favor of secure paths. But see also Section 6.7.) Furthermore, we find that the five CPs have approximately the same amount of influence as the case where there are no early adopters; we investigate this in more detail in Section 6.8.

Some ISPs always remain insecure. We find 20% of the 6K ISPs in the AS graph [9, 3] never deploy S*BGP, because they are never subject to competition for customer traffic. This highlights two important issues: (1) some ISPs may never become secure (*e.g.*, ASes whose customers are exclusively single-homed) (2) S*BGP and BGP will coexist in the long term.

Choice of early adopters is critical. For higher values of $\theta \geq 10\%$, it becomes important to choose ISPs with high customer degree as early adopters. In fact, Figure 7 shows a set of 200 random ASes has significantly lower influence than a set containing only the five top ASes in terms of degree. For large values of $\theta \geq 30\%$, a larger set of high-degree early adopters is required, with the top 200 ASes in terms of degree causing 53% of the ASes to deploy S*BGP for $\theta = 50\%$. However, to put this observation in some perspective, recall that $\theta = 30\%$ suggests that the cost of S*BGP deployment exceeds 30% of an ISP's profit margin from transiting customer traffic.

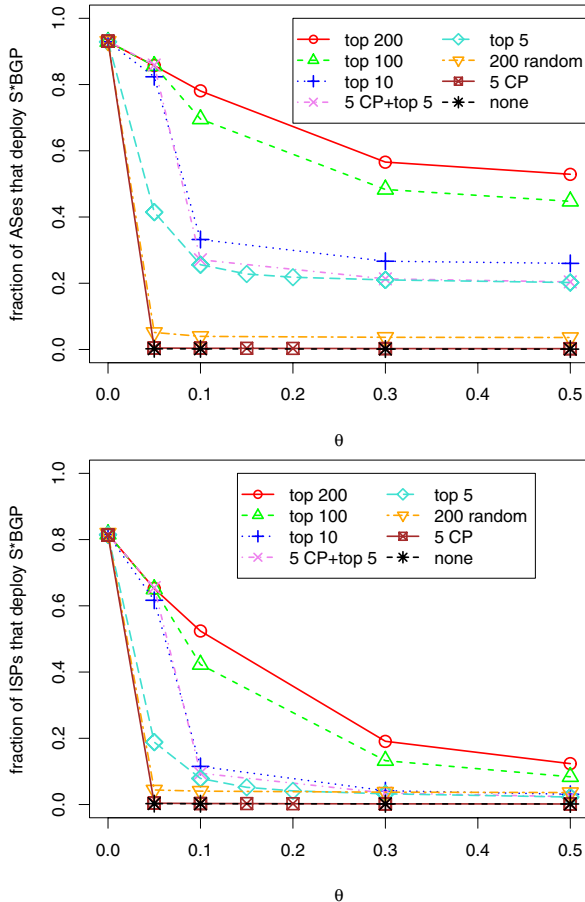


Figure 7: Fraction of ASes (top) and ISPs (bottom) that deploy S*BGP for varying θ and early adopters.

6.4 How much security do we get?

We count the number of secure paths at the end of the deployment process, as a measure of the efficacy of S*BGP deployment. (Of course, this is *not* a perfect measure of the AS graph’s resiliency to attack; quantifying this requires approaches similar to [14, 8], an important direction for future work.) We find that the fraction of secure path is only slightly lower than f^2 , where f is the fraction of ASes that have deployed S*BGP (figure in the full version). (The f^2 follows from the fact that for a path to be secure, both its source AS and its destination AS must be secure.)

6.5 Market pressure vs. simplex S*BGP

The cause of for global S*BGP deployment differs for low and high values of the deployment threshold θ :

Figure 7 (bottom): We show the fraction of ISPs (not ASes) that deploy S*BGP for the early adopter sets and varying values of θ . For low values of θ , market pressure drives a large fraction of ISPs to deploy S*BGP. In contrast, for higher values of θ very few ISPs deploy S*BGP, even for large sets of well-connected early adopters. In these cases, most of the deployment is driven by ISPs upgrading their stub customers to simplex S*BGP. For example, for the top 200 ISPs, when $\theta = 50\%$, only a small fraction of secure ASes (4%) deploy S*BGP because of market pressure, the vast majority (96%) are stubs running simplex S*BGP.

6.6 The source of competition: tie break sets.

Recall that the *tiebreak set* is the set of paths on which an AS employs the security criterion to select paths to a destination AS (Section 2.2.2). A tiebreak set with multiple paths presents opportunities for ISPs to compete over traffic from the source AS.

We observe that tiebreak sets are typically very small in the AS graph under the routing policies of Appendix A (figure in the full version). Moreover, only 20% tiebreak sets contain more than a single path.

This striking observation suggests that even a very limited amount of competition suffices to drive S*BGP deployment for low θ .

6.7 Stubs don’t need to break ties on security.

So far, we have focused on the case where secure stubs break ties in favor of secure paths. Indeed, given that stubs typically make up the majority of secure ASes, one might expect that their routing decisions can have a major impact of the success of the S*BGP deployment process. Surprisingly, we find that this is not the case. Indeed, our results are insensitive to this assumption, for $\theta > 0$ and regardless of the choice of early adopter (Figure shown in full version). We explain this by observing that stubs both (a) have small tiebreak sets, and (b) transit no traffic.

Security need only effect a fraction of routing decisions! Thus, only 15% of ASes (*i.e.*, the ISPs) need to break ties in favor of secure routes, and only 23% of ISP tiebreak sets contain more than one path. Combining these observations, we find that S*BGP deployment can progress even if only $0.15 \times 0.23 = 3.5\%$ of routing decisions are effected by security considerations!

6.8 Robustness to traffic and connectivity

6.8.1 Varying parameters.

To understand the sensitivity of our results we varied the following parameters:

1. Originated traffic volumes. We swept through different values $x = \{10\%, 20\%, 33\%, 50\%\}$ for the fraction of traffic originated by the five CPs (Section 3.1); recent work suggests a reasonable range is $x = 10\text{-}20\%$ [24].

2. Traffic destinations. Initially, we assume ASes uniformly spread their traffic across all potential destinations. We test the robustness of our results to this assumption by modeling traffic locality. We model locality by assuming ASes send traffic proportional to $1/k$ to destination ASes that are k hops away.

3. Connectivity of content providers. Published AS-level topologies are known to have poor visibility into peering links at the edge of the AS-level topology [31]. This is particularly problematic for CPs, who in recent years, have shifted towards peering with many other ASes to cut down content delivery costs [12]. Indeed, while the CPs known to have short path lengths [32], their average path length in our AS graph (with routing policies as in Appendix A) was 2.7 hops or more. Thus, for sensitivity analysis, we created a peering-heavy AS graph with 19.7K artificial peering edges from the five CPs to 80% of ASes found to be present at IXPs [3]. In our augmented AS graph, the average path length of the CPs dropped to about 2, and their degree increased to be as high as the largest Tier 1 ISPs.




Figure 8: AS 4755 incentives turn off S*BGP.

6.8.2 Impact of traffic volumes and connectivity

We now present an overview of our model’s robustness (additional detail in the full version):

1. *Originated traffic volumes vs. degree.* Surprisingly, when the five CPs source $x = 10\%$ of traffic, they are much less effective as early adopters than the top five Tier 1 ASes. Even though in the augmented topology the Tier 1s and CPs have about equal degree, the dominant factor here is traffic; even though the CPs *originate* 10% of traffic, the Tier 1s still *transit* 2-9X times more traffic.

2. *Localized interdomain traffic.* We validate that our results are robust to localized interdomain traffic using the 5 CP and top 5 as early adopters. For both the original and augmented topology, our results are robust even when ASes direct most of their traffic to nearby destinations.

3. *Impact of peering-heavy structure on simplex S*BGP.* Even in the augmented topology, where the CPs peer with large number of ASes, the Tier 1s consistently out perform the CPs by immediately upgrading their *stub customers* to simplex S*BGP. This suggests that having CPs to upgrade their *stub peers* to simplex S*BGP could potentially drive S*BGP deployment further.

6.9 Summary and recommendations.

We make two key observations regarding selection of early adopters. First, only a small number of ISPs suffice as early adopters when deployment thresholds θ are small. Second, to withstand high θ , Tier 1 ASes should be targeted. This is due to the high volumes of traffic they transit and the many stubs they upgrade to simplex S*BGP. Finally, we note that our results hold even if more than 96% of routing decisions are insensitive to security considerations!

7. OTHER COMPLICATIONS

Intuition suggests that a secure ISP will observe increased utility because secure ASes transit traffic through it. While this is true in the *outgoing utility* model (Theorem 6.2), it turns out that this is *not* the case for the *incoming utility* model. We now discuss complications that might arise because we require S*BGP to play a role in route selection.

7.1 Buyer’s Remorse: Turning off S*BGP.

We present an example of a severe obstacle to S*BGP deployment: an secure ISP that has incentive to turn *off* S*BGP. The idea here is that when an ISP n becomes secure, some of n ’s incoming traffic might change its path, and enter n ’s network along peer/provider edges instead of customer edges, thus reducing n ’s utility. This causes the secure ISP’s utility to satisfy Equation 3, resulting in the ISP opting to undeploy S*BGP.

Figure 8: We show that AS 4755, a Telecom provider in India, has an incentive to turn off S*BGP in its network. We assume content providers have $w_{CP} = 821$ which corre-

sponds to 10% of Internet traffic originating at the big five CPs (including Akamai’s AS 20940).

In the state S on the left, Akamai, AS 4755, and NTT (AS 2914) are secure, the stub customers of these two secure ISPs run simplex S*BGP, and all other ASes are insecure. Here, AS 4755 transits traffic sourced by Akamai from his provider NTT AS 2914, to a collection of twenty-four of its stub customers (including AS 45210). Akamai’s traffic does *not* increase AS 4755’s utility because it arrives at AS 4755 along a provider edge.

In the state $(\neg S_{4755}, S_{-4755})$ on the right, AS 4755 turns S*BGP off. If we assume that stubs running simplex S*BGP do *not* break ties based on security, then the only ASes that could potentially change their routes are the secure ASes 20940 and 2914. Notice that when AS 4755 turns S*BGP off, Akamai’s AS 20940 has *no secure route* to AS 4755’s stub customers (including AS 45210). As such, Akamai will run his usual tie break algorithms, which in our simulation came up in favor of AS 9498, a customer of AS 4755. Because Akamai’s traffic is now enters AS 4755 on customer edges, AS 4755’s incoming utility *increases* by a factor of 205% per each of the 24 stub destinations.

Turning off the entire network. Our simulations confirmed that, apart from Akamai changing its chosen path these twenty-four stubs, all other ASes use the same routes in state S and state $(\neg S_{4755}, S_{-4755})$. This means that AS 4755 has an incentive to turn off S*BGP in his *entire network*; no routes other than those ones Akamai uses to reach the twenty-four stubs are impacted by his decision. Indeed, we found that the utility of AS 4755 increase by a total of 0.5% (over all destinations) when he turns off S*BGP!

Turning off a destination. AS 4775 could just as well turn off S*BGP on a *per destination* basis, *i.e.*, by refusing to propagate S*BGP announcements for the twenty-four stubs in Figure 8, and sending insecure BGP messages for these destinations instead.

7.2 Turning off S*BGP can cause oscillations.

To underscore the seriousness of an ISP turning off S*BGP in his entire network, we now argue that a group of ISPs could *oscillate*, alternating between turning S*BGP on and off, and never arriving at a stable state. In the full version, we exhibit an example AS graph and state S that proves that oscillations could exist. Worse yet, we show that it is hard to even *determine* whether or not the deployment process will oscillate!

THEOREM 7.1. *Given an AS graph and state S , it is PSPACE-complete to decide if the deployment process will terminate at a stable state in the incoming utility model.*

Our proof, in the full version is by reduction to the PSPACE-complete problem of determining whether a space-bounded Turing Machine will halt for a given input string. The complexity class PSPACE consists of all decisions problems that can be solved using only polynomial *space*, but in unbounded *time*. PSPACE-complete problems (intuitively, the hardest problems in PSPACE) are at least as hard as the NP-complete problems, and widely believed to be even harder.

7.3 How common are these examples?

At this point, the reader may be wondering how often an AS might have incentives to turn off S*BGP.

Turning off an entire network? Figure 8 proves that cases where an ISP has an incentive to turn off S*BGP in its *entire network* do exist in realistic AS-level topologies [9]. However, we speculate that such examples will occur infrequently in practice. While we cannot provide any concrete evidence of this, our speculation follows from the fact that an ISP n obtains utility from many destinations. Thus, even if n has increased its utility by turning OFF S*BGP for destinations that are part of subgraphs like Figure 8, he will usually obtain higher utility by turning ON S*BGP for the other destinations that are not part of such subgraphs. (In Figure 8, this does not happen because the state S is such that only a very small group of ASes are secure; thus, no routes other than the ones pictured are effected by AS 4755’s decision to turn off S*BGP.)

Turning off a destination is likely. On the other hand, it is quite easy to find examples of *specific destinations* for which an ISP might want to turn off S*BGP. Indeed, a search through the AS graph found that at least 10% of the 5,992 ISPs could find themselves in a state where they have incentives to turn off S*BGP for at least one destination!

8. DISCUSSION OF OUR MODEL

The wide range of parameters involved in modeling S*BGP deployment means that our model (Section 3) cannot be *predictive* of S*BGP deployment in practice. Instead, our model was designed to (a) capture a few of the most crucial issues that might drive S*BGP deployment, while (b) taking the approach that simplicity is preferable to complexity.

8.1 Myopic best response.

For simplicity, we used a *myopic best-response* update rule that is standard in the game-theory literature [16]. In Section 5.5, we discussed the consequences of the fact that ISPs only act to improve their utility in the next round, rather than in long run. Another potential issue is that our update rule ignores the possibility that *multiple* ASes could deploy S*BGP in the transition from a round i to round $i + 1$, resulting in the gap between the projected utility, and the actual utility in the subsequent round. Fortunately, our simulations show projected utility $u_n(\neg S_n, S_{-n})$ is usually an excellent estimate of actual utility in the subsequent round. For example, in the case study of Section 5, 80% of ISPs overestimate their utility by less than 2%, 90% of ISPs overestimate by less than 6.7%. In the full version, we present additional results that show that this observation also holds more generally across simulations.

8.2 Computing utility locally.

Because we lack information about interdomain traffic flows in the Internet, our model uses weighted counts of the subtrees of ASes routing through ISP n as a stand-in for traffic volumes, and thus ISP utility. While computing these subtrees in our model requires *global* information that would be unavailable to the average ISP (*e.g.*, the state S , the AS graph topology, routing policies), in practice, an ISP can just compute its utility by locally observing traffic flows through its network.

Computing projected utility. Computing projected utility $u_n(\neg S_n, S_{-n})$ in practice is significantly more complex. While projected utility gives an accurate estimate of actual utility when it is computed using global informa-

tion, ISPs may inaccurately estimate their projected utility when using only local information. Our model can accommodate these inaccuracies by rolling them into the deployment threshold θ . (That is, if projected utility is off by a factor of $\pm\epsilon$, model this with threshold $\theta \pm \epsilon$.) Thus, while our approach was to sweep through a common value of θ for every ISP (Section 6.2), extensions might capture inaccurate estimates of projected utility by randomizing θ , or even by systematically modeling an ISP’s estimation process to obtain a measure for how it impacts θ .

Practical mechanisms for projecting future traffic patterns. Because S*BGP deployment can impact route selection, it is crucial to develop mechanisms that allow ISPs predict how security will impact traffic patterns through its network. Moreover, if ISPs could use such mechanisms to estimate projected utility, they would also be an important driver for S*BGP deployment. For example, an ISP might set up a router that listens to S*BGP messages from neighboring ASes, and then use these message to predict how becoming secure might impact its neighbors’ route selections. A more sophisticated mechanism could use extended “shadow configurations” with neighboring ASes [1] to gain visibility into how traffic flows might change.

8.3 Alternate routing policies and actions.

Routing policies. Because our model of ISP utility depends on traffic volumes (Section 3.3), we need to a model for how traffic flows in the Internet. In practice, traffic flow is determined by the local routing policies used by each AS, which are arbitrary and not publicly known. Thus, we use a standard model of routing policies (Appendix A) based on business relationship and path length [14, 6].

Routing policies are likely to impact our results by determining (a) AS path lengths (longer AS paths mean it is harder to secure routes), and (b) tiebreak set size (Section 6.6). For example, we speculate that considering shortest path routing policy would lead to overly optimistic results; shortest-path routing certainly leads to shorter AS paths, and possibly also to larger tiebreak sets. On the other hand, if a large fraction of multihomed ASes always use one provider as primary and the other as backup (irrespective of the AS path lengths *etc.*) then our current analysis is likely to be overly optimistic. (Of course, modeling this is difficult given a dearth of empirical data on backup paths).

Choosing routing policies. An AS might cleverly choose its routing policies to maximize utility. However, the following suggests that this is intractable:

THEOREM 8.1. *When all other ASes’ routing policies are as in Appendix A, it is NP hard for any AS n to find the routing policy that maximizes its utility (in both the incoming and outgoing utility models). Moreover, approximating the optimal routing policy within any constant factor is also NP hard.*

The proof (in the full version) shows that this is NP-hard even if n has a single route to the destination, and must only choose the set of neighbors to which it announces the route. (Thus, the problem is tractable when the node’s neighbors set is of constant size.)

Lying and cheating. While it is well known that an AS can increase the amount of traffic it transits by manipulating its BGP messages [7], we avoided this issue because our focus

is on technology adoption by economically-motivated ASes, not BGP manipulations by malicious or misconfigured ASes.

9. RELATED WORK

Social networks. The diffusion of new technologies in social networks has been well studied in economics and game theory (*e.g.*, [30, 21] and references therein). The idea that players will myopically best-respond if their utility exceeds a threshold is standard in this literature (*cf.*, our update rule (3)). However, in a social network, a player’s utility depends only on its immediate *neighbors*, while in our setting it depends on the set of secure *paths*. Thus, while [21] finds approximation algorithms for choosing an optimal set of early adopters, this is NP-hard in our setting (Theorem 6.1).

Protocol adoption in the Internet. The idea that competition over customer traffic can drive technology adoption in the Internet has appeared in many places in the literature [10, 33]. Ratnasamy *et al.* [33] suggest using competition for customer traffic to drive protocol deployment (*e.g.*, IPv6) at ISPs by creating new mechanisms for directing traffic to ASes with IPv6. Leveraging competition is much simpler with S*BGP, since it directly influences routing decisions without requiring adoption of new mechanisms.

Multiple studies [19, 18, 36] consider the role of converters (*e.g.*, IPv4-IPv6 gateways) on protocol deployment. While S*BGP must certainly be backwards compatible with BGP, the fact that security guarantees only hold for fully-secure paths (Section 2.2.2) means that there is no reason to convert BGP messages to S*BGP messages. Thus, we do not expect converters to drive S*BGP deployment.

S*BGP adoption. Perhaps most relevant is Chang *et al.*’s comparative study on the adoptability of secure inter-domain routing protocols [8]. Like [8], we also consider how early adopters create local incentives for other ASes to deploy S*BGP. However, our study focuses on how S*BGP deployment can be driven by (a) simplex S*BGP deployment at stubs, and (b) the requirement that security plays a role in routing decisions. Furthermore, in [8] ISP utility depends on the security benefits offered by the partially-deployed protocol. Thus, the utility function in [8] depends on possible attacker strategies (*i.e.*, path shortening attacks) and attacker location (*i.e.*, random, or biased towards small ISPs). In contrast, our model of utility is based solely on economics (*i.e.*, customer traffic transited). Thus, we show that global S*BGP deployment is possible even if ISPs’ local deployment decisions are *not* driven by security concerns. Also, complementary to our work is [5]’s forward-looking proposal that argues that extra mechanisms (*e.g.*, secure data-plane monitoring) can be added to S*BGP to get around the problem of partially-secure paths (Appendix B). Finally, we note both our work and [5, 8] find that ensuring that Tier 1 ASes deploy S*BGP is crucial, a fact that is not surprising in light of the highly-skewed degree distribution of the AS graph.

10. CONCLUSION

Our results indicate that there is hope for S*BGP deployment. We have argued for (1) simplex S*BGP to secure stubs, (2) convincing but a small, but influential, set of ASes to be early adopters of S*BGP, and (3) ensuring that S*BGP influences traffic by requiring ASes to (at minimum) break ties between equally-good paths based on security.

We have shown that, if deployment cost θ is low, our proposal can successfully transition a majority of ASes to S*BGP. The transition is driven by market pressure created when ISPs deploy S*BGP in order to draw revenue-generating traffic into their networks. We also pointed out unexplored challenges that result from S*BGP’s influence of route selection (*e.g.*, ISPs may have incentives to disable S*BGP).

We hope that this work motivates the standardization and research communities to devote their efforts along three key lines. First, effort should be spent to engineer a lightweight simplex S*BGP. Second, with security impacting route selection, ISPs will need tools to forecast how S*BGP deployment will impact traffic patterns (*e.g.*, using “shadow configurations”, inspired by [1], with cooperative neighboring ASes) so they can provision their networks appropriately. Finally, our results suggest that S*BGP and BGP will coexist in the long term. Thus, effort should be devoted to ensure that S*BGP and BGP can coexist without introducing new vulnerabilities into the interdomain routing system.

Acknowledgments

This project was motivated by discussions with the members of the DHS S&T CSD Secure Routing project. We especially thank the group for the ideas about simplex S*BGP, and Steve Bellovin for the example in Appendix B.

We are extremely grateful to Mihai Budiu, Frank McSherry and the rest of the group at Microsoft Research SVC for helping us get our code running on DryadLINQ. We also thank Edwin Guarin and Bill Wilder for helping us get our code running on Azure, and the Microsoft Research New England lab for supporting us on this project. We thank Azer Bestavros, John Byers, Mark Crovella, Jef Guarente, Vatche Ishakian, Isaac Keslassy, Eric Keller, Leo Reyzin, Jennifer Rexford, Rick Skowrya, Renata Texiera, Walter Willinger and Minlan Yu for comments on drafts of this work. This project was supported by NSF Grant S-1017907 and a gift from Cisco.

11. REFERENCES

- [1] R. Alimi, Y. Wang, and Y. R. Yang. Shadow configuration as a network management primitive. In *Sigcomm*, 2008.
- [2] Anonymized. Let the market drive deployment: A strategy for transitioning to bgp security. Full version. Technical report, 2011.
- [3] B. Augustin, B. Krishnamurthy, and W. Willinger. IXPs: Mapped? In *IMC*, 2009.
- [4] R. Austein, G. Huston, S. Kent, and M. Lepinski. Secure inter-domain routing: Manifests for the resource public key infrastructure. draft-ietf-sidr-rpki-manifests-09.txt, 2010.
- [5] I. Avramopoulos, M. Suchara, and J. Rexford. How small groups can secure interdomain routing. Technical report, Princeton University Comp. Sci., 2007.
- [6] H. Ballani, P. Francis, and X. Zhang. A study of prefix hijacking and interception in the Internet. In *SIGCOMM*, 2007.
- [7] K. Butler, T. Farley, P. McDaniel, and J. Rexford. A survey of BGP security issues and solutions. *Proceedings of the IEEE*, 2010.
- [8] H. Chang, D. Dash, A. Perrig, and H. Zhang. Modeling adoptability of secure BGP protocol. In *SIGCOMM*, 2006.
- [9] Y.-J. Chi, R. Oliveira, and L. Zhang. Cyclops: The Internet AS-level observatory. *ACM SIGCOMM CCR*, 2008.
- [10] D. D. Clark, J. Wroclawski, K. R. Sollins, and R. Braden. Tussle in cyberspace: defining tomorrow’s Internet. *Trans. on Networking*, 2005.

- [11] J. Cowie. Rensys blog: China's 18-minute mystery. <http://www.renysys.com/blog/2010/11/chinas-18-minute-mystery.shtml>.
- [12] A. Dhamdhere and C. Dovrolis. The internet is flat: Modeling the transition from a transit hierarchy to a peering mesh. In *CoNEXT*, 2010.
- [13] L. Gao and J. Rexford. Stable Internet routing without global coordination. *Trans. on Networking*, 2001.
- [14] S. Goldberg, M. Schapira, P. Hummon, and J. Rexford. How secure are secure interdomain routing protocols. In *Sigcomm*, 2010.
- [15] T. Griffin, F. B. Shepherd, and G. Wilfong. The stable paths problem and interdomain routing. *Trans. on Networking*, 2002.
- [16] S. Hart. Adaptive heuristics. *Econometrica*, 2005.
- [17] IETF. Secure interdomain routing (SIDR) working group. <http://datatracker.ietf.org/wg/sidr/charter/>.
- [18] Y. Jin, S. Sen, R. Guerin, K. Hosanagar, and Z. Zhang. Dynamics of competition between incumbent and emerging network technologies. In *NetEcon*, 2008.
- [19] D. Joseph, N. Shetty, J. Chuang, and I. Stoica. Modeling the adoption of new network architectures. In *CoNEXT*, 2007.
- [20] J. Karlin, S. Forrest, and J. Rexford. Autonomous security for autonomous systems. *Computer Networks*, oct 2008.
- [21] D. Kempe, J. Kleinberg, and E. Tardos. Maximizing the spread of influence through a social network. In *ACM SIGKDD*, 2003.
- [22] S. Kent, C. Lynn, and K. Seo. Secure border gateway protocol (S-BGP). *JSAC*, 2000.
- [23] V. Krishnamurthy, M. Faloutsos, M. Chrobak, L. Lao, J.-H. Cui, and A. G. Percus. Sampling large internet topologies for simulation purposes. *Computer Networks (Elsevier)*, 51(15):4284–4302, 2007.
- [24] C. Labovitz, S. Iekel-Johnson, D. McPherson, J. Oberheide, and F. Jahanian. Internet inter-domain traffic. In *SIGCOMM*, 2010.
- [25] M. Lad, D. Massey, D. Pei, Y. Wu, B. Zhang, and L. Shang. Phas: Prefix hijack alert system. In *Usenix Security*, 2006.
- [26] M. Lepinski and S. Turner. Bgpsec protocol specification, 2011. <http://tools.ietf.org/html/draft-lepinski-bgpsec-overview-00>.
- [27] C. D. Marsan. U.S. plots major upgrade to Internet router security. *Network World*, 2009.
- [28] A. Medina, A. Lakhina, I. Matta, and J. Byers. BRITE: an approach to universal topology generation. In *MASCOTS*, 2001.
- [29] S. Misel. "Wow, AS7007!". Merit NANOG Archive, apr 1997. <http://www.merit.edu/mail.archives/nanog/1997-04/msg00340.html>.
- [30] S. Morris. Contagion. *Review of Economics Studies*, 2003.
- [31] R. Oliveira, D. Pei, W. Willinger, B. Zhang, and L. Zhang. Quantifying the completeness of the observed internet AS-level structure. *UCLA Computer Science Department - Technical Report TR-080026-2008*, Sept 2008.
- [32] F. Orbit. <http://www.fixedorbit.com/metrics.htm>.
- [33] S. Ratnasamy, S. Shenker, and S. McCanne. Towards an evolvable Internet architecture. In *SIGCOMM*, 2005.
- [34] Rensys Blog. Pakistan hijacks YouTube. http://www.renysys.com/blog/2008/02/pakistan_hijacks_youtube_1.shtml.
- [35] Sandvine. Fall 2010 global internet phenomena, 2010.
- [36] S. Sen, Y. Jin, R. Guerin, and K. Hosanagar. Modeling the dynamics of network technology adoption and the role of converters. *Trans. on Networking*, 2010.
- [37] R. White. Deployment considerations for secure origin BGP (soBGP). draft-white-sobgp-bgp-deployment-01.txt, June 2003, expired.
- [38] Y. Yu, M. Isard, D. Fetterly, M. Budiu, U. Erlingsson, P. K. Gunda, and J. Currey. Dryadlinq: a system for




Figure 9: A new attack vector.

- general-purpose distributed data-parallel computing using a high-level language. In *Usenix OSDI*, 2008.
- [39] E. Zegura, K. Calvert, and S. Bhattacharjee. How to model an internetwork. In *Infocom*, 1996.

APPENDIX

A. A MODEL OF ROUTING WITH BGP.

We follow [15] by assuming that each AS a computes paths to a given destination AS d based a *ranking* on outgoing paths, and an *export policy* specifying the set of neighbors to which a given path should be announced.

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

LP Local Preference. Paths are ranked based on their next hop: customer is chosen over peer which is chosen over provider.

SP Shortest Paths. Among the paths with the highest local preference, prefer the shortest ones.

SecP Secure Paths. If there are multiple such paths, and node a is secure, then prefer the secure paths.

TB Tie Break. If there are multiple such paths, node a breaks ties: if b is the next hop on the path, choose the path where hash, $H(a, b)$ is the lowest.⁴

This standard model of local preference [13] captures the idea that an AS has incentives to prefer routing through a customer (that pays it) over a peer (no money is exchanged) over a provider (that it must pay).

Export Policies. This standard model of export policies captures the idea that an AS will only load its network with transit traffic if its customer pays it to do so [13]:

GR2 AS b announces a path via AS c to AS a iff at least one of a and c are customers of b .

B. ATTACKS ON PARTIALLY SECURE PATHS

We show how preferring partially secure paths over insecure paths can introduce *new* attack vectors that do not exist even without S*BGP:

Figure 9: Suppose that only ASes p and q are secure, and that malicious AS m falsely announces the path (m, v) , and suppose that p 's tiebreak algorithm prefers paths through r over paths through q . Then, p has a choice between two paths; a partially-secure false path (p, q, m, v) , and an insecure true path (p, r, s, v) . If no AS used S*BGP, p would have chosen the true path (per his tiebreak algorithm); if p prefers partially secure paths, he will be fooled into routing to AS m .

⁴In practice, this is done using the distance between routers and router IDs. Since we do not incorporate this information in our model we use a randomized tie break which prevents certain ASes from "always winning".