# EVALUATING SECURITY-ENHANCED INTERDOMAIN ROUTING PROTOCOLS

A Thesis
Presented to
The Academic Faculty

by

Robert D. Lychev

In Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy in the
College of Computing

Georgia Institute of Technology
August 2014

# EVALUATING SECURITY-ENHANCED INTERDOMAIN ROUTING PROTOCOLS

Approved by:


Professor Nick Feamster, Advisor
College of Computing
*Georgia Institute of Technology*

Professor Alexandra Boldyreva, Advisor
College of Computing
*Georgia Institute of Technology*

Senior Research Scientist, Russ Clark
College of Computing
*Georgia Institute of Technology*

Professor Sharon Goldberg
Department of Computer Science
*Boston University*

Professor Michael Schapira
School of Computer Science and Engineering
*The Hebrew University of Jerusalem*



Date Approved: June 17, 2014

# ACKNOWLEDGEMENTS

Many people have kindly helped me during my journey towards gaining intellectual maturity and becoming a *good* member of the fascinating cult of academia. I do not know if I have really succeeded in reaching either goal, but I suppose that is an achievement in itself.

First, I would like to thank Nick Feamster and Alexandra Boldyreva, my mentors and advisors at Georgia Institute of Technology (GT). Their advice and award-winning patience were crucial throughout my graduate studies.

I was very lucky to get an opportunity to visit Sharon Goldberg at Boston University (BU) for approximately two and a half years, where I have also collaborated with Michael Schapira. They essentially became my advisors outside of Georgia Tech, and I thank them greatly for their efforts in helping me become what I am now.

Prior to GT, I have worked with Narayanan Menon, Kevin Fu, and Siman Wong during my bachelor's and master's studies at University of Massachusetts Amherst (UMass). I thank them for steering me in the right direction in those early days.

I would very much like to thank my mother for raising me and bringing me to the United States many years ago where I could pursue a career in research. Without her, I would definitely be a different person, having to deal with the uncertainties and dangers of life in Russia. I also thank my father and many of my other relatives for their continuous support.

Last but not least, I could not have made it this far without support from my friends and lab mates at GT, BU, and UMass. I will not list their names here, but they all know who they are. I thank them all greatly and wish them best of luck in all their future endeavors.

# TABLE OF CONTENTS

# LIST OF TABLES

# LIST OF FIGURES

# SUMMARY

The Internet consists of over 50 thousand smaller networks, called Autonomous Systems (ASes) (e.g., AT&T, Sprint, Google), that use the Border Gateway Protocol (BGP) to figure out how to reach each other. One way or another, we all rely on BGP because it is what glues the Internet together, but despite its crucial role, BGP remains vulnerable to propagation of bogus routing information due to malicious attacks or unintentional misconfigurations.

The United States Department of Homeland Security (DHS) views BGP security as part of its national strategy for securing the Internet, and there is a big push to standardize a secure variant of BGP (S*BGP) by the Internet Engineering Task Force (IETF). However, S*BGP properties and their impact on the Internet's routing infrastructure, especially in partial deployment, have not yet been fully understood.

To address this issue, in this thesis we use methodologies from applied cryptography, algorithms, and large scale simulations to study the following three key properties with respect to their deployment:

1. **provable security guarantees**
2. **stability in full and partial deployment with or without attackers**
3. **benefits and harm resulting from full and partial deployment**

With our analysis we have discovered possible security weaknesses in previously proposed secure BGP variants and suggest possible fixes to address them. Our analysis also reveals that security benefits from partially deployed S*BGP are likely to be meagre, unless a significant fraction of ASes deployed it. At the same time, complex interactions between S*BGP and the insecure, legacy BGP can introduce new

vulnerabilities and instabilities into the Internet's routing infrastructure. We suggest possible strategies for mitigating such pitfalls and facilitating S*BGP deployment in practice.

# CHAPTER I

# INTRODUCTION

## *1.1 Background and Motivation*

One way or another, as members of modern society we all depend on the Internet and its proper functionality. Be it for checking e-mail, checking the weather, shopping, finding a job, finding a life partner, or scheduling a doctor's appointment, we rely on the Internet every day, often oblivious to how and why it works. The Internet has become a magic black box, that assists us in our daily lives, and we often do not realize how much we depend on it and how vulnerable we become if it does not work as we expect it due to unintentional mistakes or malicious activities.

The Internet is a distributed system, *i.e.,* it is a connected network of many entities that require coordination in order for them to cooperate and work together as a single entity capable of carrying out many complex tasks, such as financial transactions, secure communication, search, etc. The root of many problems we hear about in the news stem from some form of miscoordination, often resulting from unintentional misconfigurations or deliberate attacks. Focusing mostly on the latter, one could argue that most attacks we hear about come from breaches of confidentiality (sensitive communication intended to be secret between entities that trust each other becomes available to entities which this information was not intended for) and authentication (innocent entities falling victims to bogus information that seems to come from trusted or reliable sources). The focus of this thesis is to study security vulnerabilities of the latter type with respect to one of the most vital parts of the Internet, its routing infrastructure. Arguably, communication protocols constitute the fundamental, underlying building block of many Internet applications, and Internet's

global routing infrastructure is what allows for this communication to happen across the globe by setting up routes that messages between various Internet entities could travel along.

Currently, the Internet is composed of over 50 thousand smaller networks, called Autonomous Systems (ASes) (*e.g.,* AT&T, Akamai, and Google). ASes are independent entities that make a profit either by forwarding other ASes' messages, or storing and providing content (*e.g.,* videos, music) or by other means that may require constant, reliable access to the Internet. To figure out how to reach each other, ASes use the Border Gateway Protocol (BGP). Evey machine connected to the Internet has an associated address, referred to as an IP address, that is typically 32 bits long (although it could also be 128 bits with respect to the IPv6 architecture instead of the current IPv4 architecture that is primarily used). BGP is a distributed protocol that allows ASes to exchange routing information about reachable IP prefixes—blocks of contiguous IP addresses—via route announcements with neighboring ASes. Each BGP route announcement contains a list of every AS en route to a destination AS (*i.e.,* the AS that owns a particular IP prefix), and every AS maintains a list of possible routes to all prefixes owned by distant ASes learned this way. Upon receipt of a new routing announcement, each AS applies its routing policy to select a single, most preferred, route to each destination, and then announces that route to its neighbors, who then select their most preferred routes and propagate routing information to their neighbors, and so on, so that reachability information is distributed globally. Neighboring ASes establish bilateral business relationships between each other which determine who provides connectivity to whom and affects route preferences and export policies (local rules that determine which routes a particular AS announces to which neighbors) of each AS.

To gain some basic intuition about how BGP works, let us draw a very informal analogy between BGP and how regular mail works. Suppose you want to send a

package to your friend who lives on a different continent. Generally, you would pay one company such as FedEx or UPS to deliver your package, and check its status on line until it is delivered to your friend's door steps. As you check its progress, you notice that your package makes many stops before reaching its final destination. At each stop, it is processed by a local office to decide where to ship it next. This process continues, until it reaches the closest office to where your friend lives, at which point a truck is summoned to deliver your package to your friend's address. Now, imagine what would happen if at each stop your package made along the way, the shipment company responsible for delivering that package changed. For example, you would bring your package and pay to FedEx, who would then hand it off to UPS, who would then hand it off to Maersk, who would then hand it off to some other company that you have never heard of that actually ends up bringing the package to your friend's place. How would each company know where to forward your package and how much to charge for their services? Each shipping company would have to exchange information about which destinations they can reach, how long it would take to reach those destination, and how much they would charge for shipment. This would be done behind the scenes in such a manner that you would only have to pay one company, and you would not know a priori which route your package would take, but you would have some guarantee that it would get to its destination somehow. On the Internet, your every message is like a package in the real world, and every AS is like a shipment company. Every message you send is likely to go through the hands of multiple ASes before reaching its destination, and you would only have to pay your Internet Service Provider, *e.g.,* AT&T, Comcast, Verizon, etc., for connectivity. BGP is a way for ASes to have conversations and exchange information about how they can reach different destinations.

Thus, BGP is what glues the Internet together; its current version, version 4, is the de facto standard for routing across the Internet [88]. Thanks to BGP, our

messages, regardless of where we are currently located, can reach any host on the Internet, so that, for example, we can view a website located on a different continent, retrieve YouTube videos, read e-mail, or purchase something from Amazon. This is why, in some way, we all depend on BGP working properly. Despite its crucial role, however, BGP remains vulnerable to propagation of bogus routing information due to unintentional misconfigurations as well as malicious attacks. For example, suppose a shipment company says that it can reach a certain destination, that it in fact cannot reach, obtains your package along the way, and then either throws it away or opens it to access its contents. Something like that could easily happen now to anyone's messages on the Internet. The reason is that BGP works more or less by word of mouth, so there is no way to check the integrity of information disseminated via BGP. This protocol was originally designed without security in mind. It was supposed to work like an honors system because it was intended to be used by parties that trust each other. However, this assumption is certainly no longer true, and, in fact, has not been true for awhile.

Such incidents do happen quite frequently due to unintentional misconfigurations [77, 86, 34] as well as deliberate manipulation of routing information intended to attract traffic [24, 83, 85, 19]. For a pictorial example of an attack on BGP, let us consider the network in Figure 1 running BGP, where we focus only on ASes 0, 1, 2, and 3. Assume that each AS's ranking of routes is as depicted beside it, and that each AS announces all of its selected, favorite routes to all of its neighbors. Figure 1(a) depicts the scenario when there is no attacker and ASes 1, 2, and 3 select the direct routes to destination $d$ in accordance with their route preferences. Figure 1(b) depicts the scenario when AS 0 launches a route authentication attack by announcing a direct, fake route to destination $d$, $(0, d)$. Due to their route preferences, ASes 1, 2 and 3 fall for this attack and select to route through the attacker. This allows the attacker to intercept and peak at their traffic, or even just drop it.

One crucial requirement of BGP is that it converges to a stable routing state. This means that after some point in its execution, the selected, most preferred route of each AS remains constant, provided that the underlying topology and ASes' routing policies do not change. BGP is not always guaranteed to converge, however, and we say that it diverges when routing policies of ASes interact in ways that lead to persistent routing oscillations in which some ASes endlessly change their route selection, even when network topology and ASes routing policies stay constant. Network instability can be very disruptive and harm network performance, because every time an AS switches routes, it may delay, mis-order, or even drop, some fraction of the traffic it is carrying [63, 65].

Let us consider an example of an attack on BGP that could lead to persistent routing oscillations in Figure 1(c). Here, due to route preferences of ASes 1, 2 and 3, the underlying network becomes an instance of the Bad Gadget network [51], which is known to be unstable. To see this, suppose that ASes 1 and 3 think they are using routes $(1, 0, d)$ and $(3, 1, 0, d)$ respectively, while AS 2 thinks it's using route $(2, 0, d)$. This is unstable because AS 1 would prefer to use route $(1, 2, 0, d)$ instead, and so it will change its route selection. This would result in 3 selecting route $(3, 0, d)$ due to its route preferences. By symmetry, this situation will repeat endlessly, where ASes 1, 2 and 3 will take turns selecting longer routes through 2, 3, and 1. Note, that in the network depicted in Figure 1(c) without AS 0 launching its attack, even though each of the ASes 1, 2, and 3 prefers the longer routes to $d$ via AS 0, these routes will not become available as the link $(0, d)$ does not exist. Thus, without the attack, this network would be stable and each of these ASes would select the direct route to $d$.

To appreciate the significance of this attack with respect to our mail analogy from above, imagine how frustrating it would be if a gift package that is supposed to arrive on a friend's birthday gets delayed in transit or, possibly even lost, because FedEx keeps on changing its mind on how it prefers to ship its packages to the town where

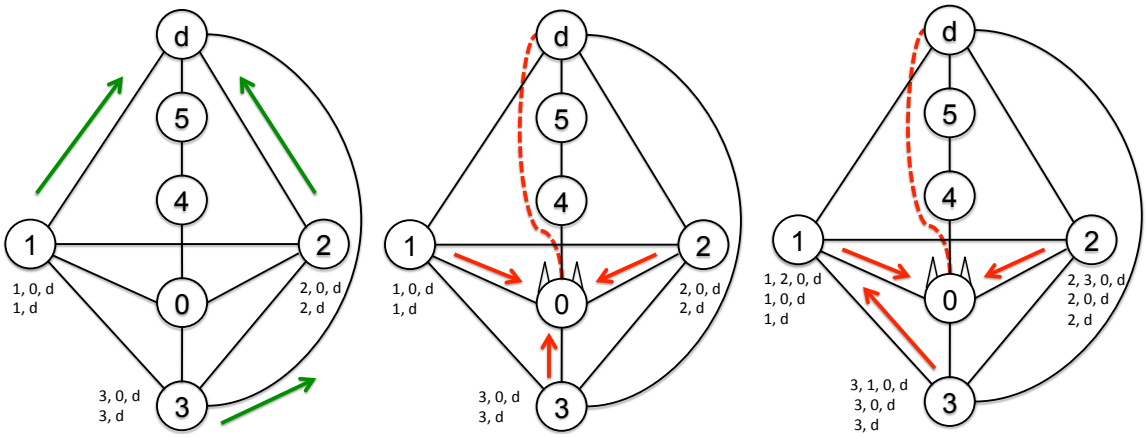


(a) Normal Operation (b) Authentication Attack (c) Destabilization Attack

**Figure 1:** In (a), ASes 1, 2, and 3 all select direct routes to destination $d$. In (b), AS 0 launches a route authentication attack by pretending to be directly connected to destination $d$, thereby attracting traffic from ASes 1, 2 and 3. In (c) AS 0 can destabilize this network with its attack due to route preferences of ASes 1, 2, and 3.

this friend lives.

To address BGP's security vulnerabilities, there have been many proposals [49, 103, 9, 55, 106, 102, 26, 27, 81] that, in particular, focus on authenticity of BGP route announcements. For example, S*BGP protocols, such as S-BGP [59], SoBGP [103], BGPSEC [69], rely on a public key infrastructure (PKI) (*e.g.,* [6, 70]), with each AS holding a certified public-secret key pair, and use digital signatures to ensure proper integrity/authentication of route announcements. There is now a big push to standardize a secure variant of BGP by the Internet Engineering Task Force (IETF) [69], but it is still not clear which security-enhanced routing protocol(s) should be deployed on the Internet and, most importantly, why.

## *1.2 Summary of Contributions*

What has been missing so far in the research community is a general approach for evaluating and comparing the advantages and disadvantages of any security-enhanced routing protocol. When evaluating any proposal that addresses BGP security vulnerabilities, we posit that it is necessary to investigate the following three crucial

questions that we study in this thesis:

1. **What provable security guarantees does it provide?**
2. **Can its deployment destabilize BGP routing and result in persistent routing oscillations with or without attackers?**
3. **What are the overall security benefits gained by deploying it and are these benefits worth the extra efforts to deploy it?**

The purpose of each question is to evaluate various attacker's goals and how well a security-enhanced routing protocol could prevent any of these goals at different levels of granularity.

The aim of the first question is to find out if a security-enhanced routing protocol can guarantee that an attacker can succeed in having ASes accept faulty route announcements with only negligible probability. Such guarantees are essential in providing assurance that ASes will not select bogus routes when a security-enhanced routing protocol is in full deployment (*i.e.,* every AS executes the same protocol). This is because without such guarantees, it is not clear why a particular security-enhanced routing protocol should be deployed or not.

The aim of the second question is to determine the conditions under which it could be guaranteed that an attacker cannot succeed in destabilizing a network during or as a result of a full or partial deployment (*i.e.,* some ASes executing the old, insecure protocol, and the other ASes executing the new security-enhanced protocol) of a security-enhanced routing protocol. Because routing instabilities can be very disruptive, without such guarantees, deploying a protocol may either have no purpose or even bring harm.

Finally, the aim of the third question is to figure out how to quantify the success of the attacker in attracting traffic by having ASes select bogus routes through the attacker instead of selecting legitimate routes that avoid the attacker, when a security-enhanced routing protocol is in full, partial or no deployment. This would allow one to

figure out how much overall benefit a particular deployment scenario on the Internet topology provides, and make judgments about its significance by comparing it to the no deployment scenario.

Due to the scale and complexity of the Internet, it is expected that any S*BGP protocol would have to go through stages of partial deployment for possibly a very long time. This is why it is especially crucial to be able to quantify the tradeoffs between security benefits (guarantees) and harm (*e.g.,* new security vulnerabilities, instabilities) when security-enhanced routing protocols are introduced only partially.

In this thesis we address these three questions with respect to S*BGP protocols.

### 1.2.1 Provable Security Analysis

Provable security analysis is standard methodology used in the cryptographic community to evaluate protocols' security guarantees with respect to well-defined classes of attacks, but it is not commonly used to analyze protocols in the networking community and has not been applied to routing protocols before. We provide provable security analysis of well known security-enhanced routing protocols S-BGP [59] and SoBGP [103].

We have designed a general security model for path-vector routing protocols, that captures many BGP security vulnerabilities, and then used it to show that S-BGP provides protection from many, albeit not all, threats, even against adaptive, colluding attackers capable of controlling all communication between ASes. We have also described necessary and sufficient modification to S-BGP that could provide protection against all attacks captured by our security model. In addition, we have shown that SoBGP does not provide the same level of security guarantees as S-BGP. Finally we have considered various relaxations to our security model and protocol modifications that could result in security and efficiency improvements of S-BGP and SoBGP, and we provided sufficient and necessary conditions for S-BGP to have analogous security

guarantees when PKI is only partially deployed (not all ASes have certified keys).

### 1.2.2 Network Stability Analysis

As was demonstrated in Figure 1(c), a stable network can be destabilized by a single fixed-route attacker (attacker that announces the same bogus information to its neighbors throughout the attack). We have also shown that it is possible to destabilize a stable network by introducing only partial deployment of S-BGP or BGPSEC. BGPSEC is a security-enhanced BGP variant which is currently being standardized by the IETF [69], and its essential operations are similar to S-BGP.

To address these phenomena, we have studied conditions under which convergence to a unique stable state can be guaranteed even in presence of fixed route attackers and partial deployments of S-BGP and BGPSEC. In our studies we have considered various routing models to account for different routing policies of ASes. Our convergence results hold regardless of the number and locations of the fixed-route attackers, of the specific fixed-route attacks launched, and specifics of S-BGP/BGPSEC deployment. These results constitute an important building block for performing empirical evaluation of the security benefits of various BGPSEC deployment scenarios described next.

### 1.2.3 Quantifying Security Benefits and Complications

We have studied the conditions under which security benefits from full or partial deployment of BGPSEC would be significant enough to justify its deployment on the Internet. BGPSEC is designed to be deployed on top of the Resource Public Key Infrastructure (RPKI). which provides protection from attacks believed to be the cause of most of the Internet routing outages, called prefix hijacks, and is currently being deployed [70]. BGPSEC can prevent more sophisticated attacks than RPKI.

We have studied BGPSEC security benefits over RPKI with theoretical analysis and large-scale simulations of many partial BGPSEC deployments with respect to

multiple routing models and underlying Internet AS-level topologies. We have shown that security prioritization in route selection plays a crucial role in the security benefits, routing complexities, and vulnerabilities that could arise when BGPSEC is only partially deployed. We have found that if network operators do not prioritize security above all other considerations in their routing policies (which is the likely scenario [42]), partial BGPSEC deployments result in only marginal security improvements vis-a-vis the benefits provided by the RPKI. We have also demonstrated that, other than routing instabilities mentioned above, partial deployment of BGPSEC can result in counterintuitive situations where having more ASes deploy BGPSEC can cause its security benefits to decrease. Finally, we have put forth guidelines on how to deal with these difficulties in practice.

## *1.3 Road Map*

We discuss related work in Chapter 2. Chapters 3, 4 and 5 are dedicated to presenting the main contributions of this thesis, namely, provable security analysis, network stability analysis, and quantifying security benefits and complications of various security-enhanced BGP variants respectively. We conclude with a discussion of our contributions and their practical implications as well as propose directions for future work in Chapter 6.

## *1.4 Bibliographic Notes*

Our investigation of the first question has appeared in CCS 2012 [22], while our investigation of the second and third questions have appeared in PODC 2013 and SIGCOMM 2013 [71, 73].

# CHAPTER II

# RELATED WORK

In this chapter we discuss previous work related to the main questions we study in this thesis.

### 2.0.1 BGP Security Enhancements and Analyses

It is well known that BGP attacks and miconfigurations can result in serious routing outages on the Internet, and its security has undergone much scrutiny [25, 74, 16, 79]. Over the past few decades many security enhancement to BGP have been proposed, including but not limited to [49, 104, 9, 55, 106, 102, 26, 27]. They incorporate additional measures to handle authenticity/ integrity and authorization issues in BGP, such as, in particular, integrity of the route announcements. Secure BGP (S-BGP) protocol [60, 62] stands out as the most comprehensive attempt to secure the Internet's routing infrastructure to date, and its current variant, *i.e.,* BGPSEC, is being standardized by the IETF [69] to run on top of the RPKI, which is being currently deployed [70]. RPKI provides the functionality of a PKI and protection form prefix hijacks, and it allows for its cryptographic operations to be done mostly off line. BGPSEC can prevent more sophisticated attacks than RPKI, but it requires its cryptographic operations to be done on line.

Most existing proposals and analyses, however, do not go further than pointing out specific attacks and suggesting possible fixes. For example, although a survey of BGP security [25] throughly discusses such threats as message tampering, session termination, prefix hijacking, prefix deaggregation, subversion of rout information, route flapping, etc., it is not immediately clear what precisely an adversary is capable of doing when attacking BGP and what its goals are. For example, can an attacker

peek on communication and collude, and when is the attack considered successful? Even though the proposed solutions may seem plausible, there is no provable way of quantifying their security guarantees. For example, the proposal for secure path-vector routing described in [55] without provable security analysis was later shown to suffer from attacks that could be mounted by 60% of AS's on the Internet in [78]. Although this vulnerability was mentioned in [55], there was no way of formally quantifying its seriousness.

Provable security treatment is not uncommon in practical communication protocols, such as SSH [17, 84] and Kerberos [13, 21]. However, to the best of our knowledge, the only attempt to use similar methodology in the context of securing BGP has been done in [26, 29, 107, 96]. Compared to our study in Chapter 3, the security models considered in [26, 29] are weaker, in the sense that they do not capture route validity attacks (collusions are not captured in [26]). In [96], the authors propose a cryptographic extension to BGP to address a certain type of route withdrawal attacks aimed at preventing certain routes being available to certain ASes due to the use of such mechanisms as Route Flap Damping and Minimum Route Advertisement Interval. These mechanisms are used to limit the negative impact of route flapping, and the attack considered in [96] is not exactly against BGP, bur rather against these mechanisms. In [107], the authors focus on a slightly different type of attacks that stem from ASes violating contractual promises in the context of BGP routing. This work proposes a mechanism for verifying if such promises are fulfilled in a matter that does not make ASes' private policies and contractual obligations public.

### 2.0.2 BGP Convergence

Much work has been done on the problem of BGP convergence, including but not limited to [53, 101, 52, 39, 38, 51, 95, 50], with respect to many different routing models. However, none of this work has considered the problem of network stability

in presence of attackers and/or partially deployed security enhancements to BGP. To the best of our knowledge, such scenarios have been considered only in [94, 32]. In [94] the authors explored various network destabilizing attacks in practical settings, but did not focus on provable stability guarantees. Investigation in [32] considers BGP convergence in presence of attackers with respect to a more restricted routing model than the one we study in this thesis in Chapter 4, because it does not capture convergence in presence of partially deployed secure BGP variants.

### 2.0.3 Quantifying Impact of Secure BGP Variants

Previous proposals of new security enhancements to BGP have focused on scenarios where ASes will reject insecure routes [12, 28]. However, such analyses is appropriate for studying only the full deployment scenario, where every AS has already deployed S*BGP [46, 25, 14], and does not apply to partial deployment scenarios which are more likely to occur in practice. In addition, previous work on incentivizing S*BGP adoption [40, 28] suggests that S*BGP and BGP may have to coexist potentially for a very long time.

The partial deployment scenarios we study in Chapter 5 have been suggested before in practice [87] and in the research communities [12, 40, 28]. Like in our work, Investigation in [46] also quantifies security benefits of S*BGP deployments as the fraction of source ASes that avoid having their traffic intercepted by the attacking AS. However, this work also considers only the full deployment scenario, and thus does not analyze the possible complications that can arise when S*BGP is deployed only partially.

# CHAPTER III

# EVALUATING PROVABLE SECURITY GUARANTEES

## *3.1 Introduction*

In this chapter we address the question of whether security-enhanced routing protocols such as S-BGP [59], BGPSEC [69] and SoBGP [104], can provably guarantee that an attacker can succeed in having ASes accept bogus route announcements with only negligible probability. Recall that without such guarantees, it is not clear at all why a particular security-enhanced routing protocol should or should not be deployed. Provable security [75, 90] is the method we use to address this question.

Unlike the cyclic trial-and-error approach to security, provable security allows us to have protocols, whose security is provably guaranteed, as long as the assumption about the underlying hard problem remains true for computationally bounded adversaries. In general, this approach consists of the following components:

1. a formal definition of a protocol's syntax
2. a formal definition of the security task in question that includes a precise description of adversarial capabilities and when the adversary is considered successful
3. a reduction proof showing that the only way to break the protocol according to the definition is by breaking the underlying problem, believed to be hard

Such treatment requires precise notation and definitions at each of the above steps, and in Section 3.2 we introduce some notation and definitions that are not common in the networking literature, but are rather standard in the cryptographic literature.

One of the main challenges in our analysis is to provide a security model that would capture many BGP vulnerabilities, while being concise and easy to use for doing security analysis. To this end, we designed a very general model of interdomain

networks, path-vector protocols and a path-vector protocol security definition Sections 3.3 and 3.5. We demonstrate how BGP and S-BGP fit our path-vector model in Section 3.4, and then use our security model to show that S-BGP provides protection from many, albeit not all, threats, even against adaptive, colluding attackers capable of controlling all communication between ASes in Section 3.6. We next describe necessary and sufficient modification to S-BGP that could provide protection against all attacks captured by our security model in Section 3.7.

Reliance on full PKI deployment of any kind as well as the use of public-key cryptography are expensive measures, and in Sections 3.8 and we carry out similar analysis of S-BGP but in partial PKI-deployment scenarios, *i.e.,* where not all ASes have keys. This setting also captures full PKI deployment scenarios where, for performance reasons, not all ASes want to execute parts of the protocol that require the use of public-key cryptographic operations.

Although the results of sections 3.6-3.8 focus on S-BGP, where applicable, we also comment on how to extend them to BGPSEC, whose essential operations are very similar to S-BGP.

Beside S-BGP and BGPSEC, SoBGP [104] is another effort to secure BGP that received much attention in the community and we use our security model to analyze SoBGP in Section 3.9. Our analysis of SoBGP shows that it does not provide the same level of guarantees as S-BGP, so we do not analyze it with respect to partial PKI-deployment scenarios in this thesis.

Finally, in Section 3.10 we propose and analyze SoBGP and a more light-weight variant of S-BGP with respect to weaker but more operationally realistic threat models than the one considered in Section 3.5. We then we conclude with a high-level discussion of our results and their practical implications in Section 3.11.

## *3.2 Preliminaries*

In this section we introduce some basic notation and definitions that we will be using throughout this chapter.

### 3.2.1 Notation and Conventions.

We denote by $\{0,1\}^*$ the set of all binary strings of finite length. If $x, y$ are strings then $(x, y)$ denotes the concatenation of $x$ and $y$ from which $x$ and $y$ are uniquely decodable. $\mathbb{N}$ is the set of non-negative integers. If $\kappa \in \mathbb{N}$ then $1^\kappa$ denotes the string consisting of $\kappa$ consecutive "1" bits. If $\mathsf{S}$ is a finite set, then $s \xleftarrow{\$} \mathsf{S}$ denotes that $s$ is selected uniformly at random from $\mathsf{S}$. If $\mathcal{A}$ is a randomized algorithm and $n \in \mathbb{N}$, then $a \xleftarrow{\$} \mathcal{A}(i_1, i_2, \ldots, i_n)$ denotes that $a$ is assigned the outcome of the experiment of running $\mathcal{A}$ on inputs $i_1, i_2, \ldots, i_n$. The empty string is denoted by $\varepsilon$. An adversary is an algorithm, and by convention, the running-time of an adversary includes that of its overlying experiment. All algorithms are assumed to be randomized and efficient, *i.e.,* polynomial in the size of the input.

### 3.2.2 PKI and Signature Schemes

Whenever the use of public keys is required, in this chapter we implicitly assume that a *public key infrastructure (PKI)* is supported, in the sense that the public keys are valid, bound to users' identities, and are publicly known.

A *digital signature scheme* $\mathcal{SS} = (\mathsf{Kg}, \mathsf{Sign}, \mathsf{Ver})$ with associated *message space* $\mathsf{MsgSp}$ is defined by three algorithms. The randomized *key generation* algorithm $\mathsf{Kg}$ takes the security parameter $1^k$ and outputs a public–secret key pair: $(pk, K) \xleftarrow{\$} \mathsf{Kg}(1^k)$. The *signing* algorithm $\mathsf{Sign}$, that could be randomized, takes the secret key and message $M \in \mathsf{MsgSp}$ and outputs a signature: $\sigma \xleftarrow{\$} \mathsf{Sign}(K, M)$. The deterministic *verification* algorithm $\mathsf{Ver}$ takes the public key, a message and a signature and outputs a bit $b \in \{0,1\}$ indicating whether the signature is deemed valid or not: $b \leftarrow \mathsf{Ver}(pk, M, \sigma)$. For correctness, it is required that for every $(pk, K)$ output by

$\mathsf{Kg}(1^k)$ and every $M \in \mathsf{MsgSp}$ we have that $\mathsf{Ver}(pk, M, \mathsf{Sign}(K, M)) = 1$.

The traditional security notion for a signature scheme $\mathcal{SS} = (\mathsf{Kg}, \mathsf{Sign}, \mathsf{Ver})$ considers an experiment $\mathbf{Exp}^{\text{uf-cma}}_{\mathcal{SS}}(A)$ associated with an adversary $A$. First, a pair of keys is generated: $(pk, K) \xleftarrow{\$} \mathsf{Kg}(1^k)$. Then, $A$ is given $pk$ and the signing oracle, and it has to output a message and a forgery: $(M, \sigma) \xleftarrow{\$} A^{\mathsf{Sign}(K,\cdot)}(pk)$. The adversary wins and the experiment returns 1 if and only if $\mathsf{Ver}(pk, M, \sigma) = 1$, $M \in \mathsf{MsgSp}$ and $A$ never queried $M$ to $\mathsf{Sign}(K, \cdot)$. We say that $\mathcal{SS}$ is *uf-cma-secure* if $\Pr\left[\mathbf{Exp}^{\text{uf-cma}}_{\mathcal{SS}}(A) = 1\right]$ is negligible in $k$ for all efficient algorithms $A$.

### 3.2.3 Certification Schemes

To the best of our knowledge, the certification scheme primitive has not been explicitly defined, but it has been considered as parts of other protocols, *e.g.,* certified encryption and digital signature schemes in [20].

A two-party *certification protocol* $\mathcal{CP} = (\mathsf{Kg}_{\mathsf{CA}}, (\mathsf{CA}, \mathsf{U}), \mathsf{Vercert})$ is defined by a key generation algorithm, a pair of possibly interactive and randomized algorithms executed between the certification authority and a user (in this thesis, an AS), and a verification algorithm. The protocol is associated with an ID space $\mathsf{IDSp}$ and data space $\mathsf{DSp}$.

- $\mathsf{Kg}_{\mathsf{CA}}$ takes the security parameter $1^k$ and outputs a public-secret key pair ($pk_{\mathsf{CA}}$, $K_{\mathsf{CA}}$) for the CA.

- $\mathsf{CA}$ takes as input a secret key $K_{\mathsf{CA}}$, the identity of user $ID \in \mathsf{IDSp}$ and data $D \in \mathsf{DSp}$. For the purpose of this thesis, a node's ID is the unique AS number given to the AS associated with that node by the Internet Assigned Numbers Authority (IANA) [2], as is done for every AS on the Internet.

- $\mathsf{U}$ takes as input the public key $pk_{\mathsf{CA}}$, the identity $ID \in \mathsf{IDSp}$ and data $D \in \mathsf{DSp}$. As result of the interaction, the outputs of both parties are $\perp$, if something

went wrong, or $(ID, D, \mathsf{cert})$, where $\mathsf{cert}$ is an issued certificate. We write $((ID, D, \mathsf{cert}), (ID, D, \mathsf{cert})) \xleftarrow{\$} (\mathsf{CA}(K_{\mathsf{CA}}, ID, D),\ \mathsf{U}(pk_{\mathsf{CA}},\ ID,\ D))$ for the result of an honest interaction.

- $\mathsf{Vercert}$ takes as input $(pk_{\mathsf{CA}}, ID, D, \mathsf{cert})$ and outputs a bit.

The correctness requirement states that for any pair $(pk_{\mathsf{CA}},\ K_{\mathsf{CA}})$ output by $\mathsf{Kg}_{\mathsf{CA}}(1^k)$, any $ID \in \mathsf{IDSp}$ and $D \in \mathsf{DSp}$, the result of certification $((ID,\ D,\ \mathsf{cert}), (ID, D, \mathsf{cert})) \xleftarrow{\$} (\mathsf{CA}(K_{\mathsf{CA}}, ID, D),\ \mathsf{U}(pk_{\mathsf{CA}}, ID, D))$ passes verification, *i.e.,* $\mathsf{Vercert}\ (pk_{\mathsf{CA}},\ ID,\ D, \mathsf{cert}) = 1$.

We now define the security of the certification protocol $\mathcal{CP} = (\mathsf{Kg}_{\mathsf{CA}}, (\mathsf{CA}, \mathsf{U}), \mathsf{Vercert})$ with $\mathsf{IDSp}, \mathsf{DSp}$, and we call this notion *unforgeability under chosen-data attack*. Consider the following experiment $\mathbf{Exp}^{\text{uf-cda}}_{\mathcal{CP}}(A)$ associated with an adversary $A$.

First, the CA's keys are generated: $(pk_{\mathsf{CA}}, K_{\mathsf{CA}}) \xleftarrow{\$} \mathsf{Kg}_{\mathsf{CA}}(1^k)$. $A$ gets $pk_{\mathsf{CA}}$ and after that can repeatedly output $(ID, D)$ so that $ID \in \mathsf{IDSp}, D \in \mathsf{DSp}$ and for each such pair participate in $(\mathsf{CA}(K_{\mathsf{CA}}, ID, D),\ A(pk_{\mathsf{CA}},\ ID,\ D))$ on behalf of the user interacting with the CA.

The experiment outputs 1 if and only if $A$ at some point returns $(ID', D',\ \mathsf{cert}')$ so that $ID' \in \mathsf{IDSp}, D' \in \mathsf{DSp}$, $\mathsf{Vercert}(pk_{\mathsf{CA}},\ ID', D',\ \mathsf{cert}') = 1$ and $\mathsf{CA}$ never output $(ID', D', \mathsf{cert}'')$, for any $\mathsf{cert}''$.

We define $A$'s advantage $\mathbf{Adv}^{\text{uf-cda}}_{\mathcal{CP}}(A)$ in this experiment to be $\Pr\left[\mathbf{Exp}^{\text{uf-cda}}_{\mathcal{CP}}(A) = 1\right]$. We say that $\mathcal{CP}$ is *uf-cda-secure* if $\mathbf{Adv}^{\text{uf-cda}}_{\mathcal{CP}}(A)$ is negligible in $k$ for all efficient algorithms $A$. Note that one could define a stronger security notion, but that would be an overkill for the purposes of our application.

**Construction 3.2.1.** *Let $\mathcal{SS} = (\mathsf{Kg}, \mathsf{Sign}, \mathsf{Ver})$ be a signature scheme with $\mathsf{MsgSp}$. We define the corresponding $\mathcal{CP}_s = (\mathsf{Kg}, (\mathsf{CA}, \mathsf{U}), \mathsf{Vercert})$ with $\mathsf{IDSp}, \mathsf{DSp}$ so that for every $ID \in \mathsf{IDSp}$ and $D \in \mathsf{DSp}$, $(ID, D) \in \mathsf{MsgSp}$. $(\mathsf{CA}, \mathsf{U})$ is then as follows. The CA sends $\mathsf{cert} = \mathsf{Sign}(K_{\mathsf{CA}}, (ID, D))$ to the user. The user verifies*

$\mathsf{Ver}(pk_{\mathsf{CA}}, (ID, D), \mathsf{cert})$ *and, if correct, both output* $\mathsf{cert}$*:* $(ID, D, \mathsf{cert})$*, otherwise they both output* $\perp$*.* $\mathsf{Vercert}(pk_{\mathsf{CA}}, ID, D, \mathsf{cert})$ *returns* $\mathsf{Ver}(pk_{\mathsf{CA}}, (ID, D), \mathsf{cert})$*.*

**Theorem 3.2.2.** *Let* $\mathcal{SS} = (\mathsf{Kg}, \mathsf{Sign}, \mathsf{Ver})$ *be a signature scheme with message space* $\mathsf{MsgSp}$ *and let* $\mathcal{CP}_s = (\mathsf{Kg}, (\mathsf{CA}, \mathsf{U}), \mathsf{Vercert})$ *be its corresponding certification scheme with identity and data spaces* $\mathsf{IDSp}, \mathsf{DSp}$ *as per Construction 3.2.1. Then,* $\mathcal{CP}_s$ *is uf-cda-secure if* $\mathcal{SS}$ *uf-cma-secure.*

*Proof.* In this proof we show that for every adversary $A$ attacking unforgeability of $\mathcal{CP}$, there exists adversary $B$ attacking unforgeability of $\mathcal{SS}$ such that $\mathbf{Adv}^{\text{uf-cma}}_{\mathcal{SS}}(B) = \mathbf{Adv}^{\text{uf-cda}}_{\mathcal{CP}}(A)$ and the resources of $B$ are that of $A$.

Let $A$ be an adversary attacking the uf-cda security of $\mathcal{CP}_s$. We construct an adversary $B$ attacking the uf-cma security of $\mathcal{SS}$ as follows. $B$ is given $pk_{\mathsf{CA}}$ and the signing oracle $\mathsf{Sign}(K_{\mathsf{CA}}, \cdot)$. For every $(ID, D)$ output by $A$, $B$ runs $(\mathsf{CA}, A)$ with $A$ on behalf of the CA. To compute $\mathsf{cert}$, $B$ queries $(ID, D)$ to its signing oracle and returns the result to $A$. When $A$ halts and outputs a forgery $(ID', D', \mathsf{cert}')$, $B$ also halts and outputs $((ID', D'), \mathsf{cert}')$.

It is easy to see that the view of $A$ in the simulated experiment has the same distribution as that in $\mathbf{Exp}^{\text{uf-cda}}_{\mathcal{CP}}(A)$ and that $B$ wins, whenever $A$ wins, *i.e.,* $B$'s forgery is valid whenever the same is true for $A$. Finally, we observe that $A$ and $B$ make the same number of equal-length signing queries and have the same running time. □

## *3.3 Interdomain Network and Path-Vector Protocol Models*

In this section we define syntaxes for interdomain networks and path-vector protocols that we will be using in this chapter. The models we use in our analysis here are slightly different from the models we will use in Chapters 4-5 because the focus of this chapter is not on global network effects such as path-vector protocol convergence to bogus or non-bogus routes (if at all), but rather on local guarantees concerning

authenticity and integrity of routing announcements exchanged between ASes. We will emphasize these distinctions where necessary.

### 3.3.1 A Model of Interdomain Networks

We model an *interdomain network* as a tuple $\mathcal{I} = (\mathbf{G} = (\mathsf{ASes}, \mathsf{link}), \mathsf{Prefixes}, \mathsf{OrforPr}, \mathsf{relation}, \mathsf{preferto}, \mathsf{policy})$.

- $\mathbf{G}$ is a finite, connected graph consisting of a set of nodes, $\mathsf{ASes}$, that represent autonomous systems (ASes) and a set of edges, defined by a function $\mathsf{link}\colon \mathsf{ASes} \times \mathsf{ASes} \to \{0,1\}$ that returns 1 if an only if those ASes are *neighbors*.

- $\mathsf{Prefixes}$ is a set of strings in $\{0,1\}^*$ representing *prefixes*, which specify sets of IP addresses.

- The origin-for-prefix function $\mathsf{OrforPr}\colon \mathsf{Prefixes} \to \mathsf{ASes}$ takes a prefix and returns an AS designated to own that prefix (called *origin*).

- $\mathsf{relation}\colon \mathsf{ASes} \times \mathsf{ASes} \to \mathrm{BR}$ is a function that takes two ASes and returns their business relationship if they are neighbors and $\perp$ otherwise Note that $\mathsf{link}$ may be redundant given $\mathsf{relation}$, but we keep the former to maintain a general graph definition. Here BR defines the set of all possible pair-wise business relationships in $\mathcal{I}$ between neighbors. For example, as we will consider in Section 3.10 and Chapters 4-5, the neighbors could have $(\mathsf{peer}, \mathsf{peer})$ or $(\mathsf{cust}, \mathsf{prov})$ relationships [38]. However, we do not assume any relationships in this chapter because such details are not essential to our analysis here.

  Before defining the last two components of $\mathcal{I}$, we provide some comments and auxiliary definitions.

  Note that $\mathcal{I}$ implicitly defines the set of origins $\mathsf{Origins} \subseteq \mathsf{ASes}$ as the image set of function $\mathsf{OrforPr}$. We denote the set of neighbors of an AS $N$ as $\mathsf{Neighbors}(N)$.

- A *route* in $\mathcal{I}$ is a sequence of ASes $(N_n, N_{n-1}, \ldots, N_2, N_1)$, for some $n \in \mathbb{N}$ and $N_i \in \mathsf{ASes}$ for all $1 \leq i \leq n$, such that $N_1 \in \mathsf{Origins}$. Here $N_1$ is the destination of traffic and $N_i$ is a possible source of traffic for every $2 \leq i \leq n$. Unless otherwise specified, for convenience, ASes on routes will be indexed in increasing order right-to-left, starting with the origin. We say that $N_i$ is up- or down-stream from AS $N_j$ on a particular route, if $i < j$ or $i > j$ respectively.

- A *subroute* of some route $R = (N_n, \ldots N_2, N_1)$ is a sequence of ASes $(N_i, \ldots N_1)$, for any $1 \leq i \leq n$, that is defined as the $i$ right-most entries of $R$. A route is said to be *feasible* if for every pair of consecutive ASes $(N_{i+1}, N_i)$ in that route, $\mathsf{link}(N_{i+1}, N_i) = 1$ for $n < i \leq 1$, *i.e.,* the ASes are neighbors. A route $(N_n, \ldots N_2, N_1)$ is said to be *to* some prefix $P \in \mathsf{Prefixes}$ if $\mathsf{OrforPr}(P) = N_1$.

- The function $\mathsf{preferto}$ specifies total and transitive binary relations $\mathsf{preferto}_N$ on routes to the same prefix in $\mathsf{Prefixes}$ for each AS $N \in \mathsf{ASes}$.

- $\mathsf{policy}$ specifies functions $\mathsf{policy}_N$ that define export policy rules for each AS $N \in \mathsf{ASes}$. $\mathsf{policy}_N$ takes a route to some prefix $P$ together with the output of $\mathsf{relation}$ on $N$ and the first AS on that route (the second parameter is ignored if $N$ owns $P$) and outputs a set of ASes to which $N$ is allowed to export (*i.e.,* advertise) that route. With this syntax we consider only next-hop export policy functions whose outputs depend on the routes and business relationships of neighbors on those routes of the AS exporting the route, since they are believed to quite reasonably approximate the export policy rules that ASes on the Internet of today use to advertise their routes to different neighbors [38]. We will comment on how our results could be extended for more complicated export policy functions in Section 3.7.

We say that $N_i \in \mathsf{ASes}$ *prefers* some route $R$ to some other route $R'$, both to the same prefix $P$, if $R$ $\mathsf{preferto}_{N_i}$ $R'$, and we say that a route $R = (N_{n-1}, \ldots, N_2, N_1)$ to

prefix $P \in \mathsf{Prefixes}$ is AS $N_n$'s $j^{\text{th}}$ *most preferred* route to $P$, for some $j \geq 1$, if there are exactly $j-1$ distinct routes $R' = (M_\ell, \ldots, M_1, N_1)$ to $P$ such that $R'$ $\mathsf{preferto}_{N_n}$ $R$. We say that $R$ is $N_n$'s most preferred route to $P$ if $j = 1$. For any AS $N_n$, for any route $R = (N_{n-1}, \ldots, N_2, N_1)$ to some prefix $P$, $R$ $\mathsf{preferto}_{N_n}$ $\varepsilon$ if and only if $\mathsf{OrforPr}(P) = N_1$ unless $\mathsf{OrforPr}(P) = N_n$, in which case $\varepsilon$ is $N_n$'s most preferred route to $P$.

A route $R = (N_n, \ldots N_2, N_1)$ is *valid* if it is feasible and consistent with $\mathsf{policy}$ of every AS on that route, *i.e.,* $N_i \in \mathsf{policy}_{N_i-1}((N_{i-1}, \ldots N_2, N_1), \mathsf{relation}(N_{i-1}, N_{i-2}))$, for all $2 \leq i \leq n$.

Our model of an interdomain network is certainly a simplification of the Internet of today. For example two neighboring ASes could have multiple distinct business relationships at different locations, and any AS's route preference and export policy rule could also be a function of the prefix corresponding to the route. However, such details are not necessary to study the essential attacks on the current Internet's routing infrastructure. Furthermore, our network model can be easily extended to incorporate extra features, possibly at the expense of making the analysis more complicated. For instance, one could consider a graph where each AS represents a border gateway (*i.e.,* a router at the border of neighboring ASes), and one could require for the preference relation and the policy function to be defined for each prefix.

### 3.3.2 A Model of Path-Vector Protocols

Let $\mathcal{I} = (\mathbf{G} = (\mathsf{ASes}, \mathsf{link}), \mathsf{Prefixes}, \mathsf{OrforPr}, \mathsf{relation}, \mathsf{preferto}, \mathsf{policy})$ be an interdomain network. An interactive and stateful *path-vector protocol* $\mathcal{PV} = (\mathsf{Init}, \mathsf{An})$ is defined by two algorithms.

- $\mathsf{Init}$ is an optional randomized algorithm run by an AS (or a CA) that takes the security parameter $1^k$ and generates the corresponding public and secret keys for that AS (or the CA).

- $\mathsf{An}$ is a stateful and possibly randomized, interactive multiparty algorithm run between the AS and possibly the CA. Each AS $N \in \mathsf{ASes}$ is given inputs ($N$, $\mathsf{Neighbors}(N)$, $\mathsf{relation}_N$, $\mathsf{preferto}_N$, $\mathsf{policy}$, $\mathbf{P}_N, pk_{\mathsf{CA}}, \mathbf{pk}$), where $\mathsf{relation}_N$ outputs $\mathsf{relation}(N, N')$ for all $N' \in \mathsf{Neighbors}(N)$ and $\perp$ otherwise. $\mathbf{P}_N \subseteq \mathsf{Prefixes}$ is the set of prefixes $N$ owns, $pk_{\mathsf{CA}}$ is the optional public key of the CA and $\mathbf{pk}$ denotes the optional set of public keys of all ASes in $\mathsf{ASes}$. The optional CA takes as inputs $(\mathcal{I}, pk_{\mathsf{CA}})$. During the execution, $N_i$ sends messages known as *route announcements* to $N_j \in \mathsf{Neighbors}(N_i)$, in accordance with $\mathsf{policy}_{N_i}$, of the form $(N_i, N_j, R, P, W, Aux)$, where $R$ is a route to $P \in \mathsf{Prefixes}$ known as the *path attribute*, $W \in \{0,1\}$ is the withdrawal flag, and $Aux \in \{0,1\}^*$ holds any additional information. Upon receipt of a route announcement, $N_j$ can *reject* it by outputting $\perp$. We say that $N_j$ *accepts* a message if $N_j$ does not reject it.

Note that although export policy function of each AS is given as input to each AS, ASes cannot find out other ASes' decisions with respect to exporting arbitrary routes, because they are not provided with information in regards to the business relationships of remote ASes and what the feasible routes of remote ASes may be. We comment on how our results could be extended for scenarios when other ASes' export policy rules are not publicly known in Section 3.7.

We say that $\mathcal{PV}$ is *correct* for a class of networks $\mathcal{C}$ if when every AS in $\mathsf{ASes}$ follows $\mathcal{PV}$, every announcement during its execution is accepted for every network $\mathcal{I} \in \mathcal{C}$.

One could consider a stricter notion of correctness that would require path-vector protocols to be useful and allow ASes to learn routes to various destinations, *e.g.,* in practice path-vector protocols such as BGP are considered useful for the Internet only if they converge. As we will discuss in more detail in Chapter 4, we say that $\mathcal{PV}$ *converges* over $\mathcal{I}$, if after a finite number of sent route announcements every AS selects that AS's most preferred route, out of all routes it receives as announcements from

neighbors, to every prefix that the AS has a valid route to in $\mathcal{I}$, such that *subroute consistency* is satisfied. Subroute consistency requires that if $R_i = (N_{i-1}, \ldots, N_1)$ is the most preferred route selected by $N_i$ to $P$, then for every $1 < j < i$, subroute $R_j = (N_j, \ldots, N_1)$ of $R_i$ is the most preferred route selected by $N_j$ to $P$, for all $P \in \mathsf{Prefixes}$ and all $N_i \in \mathsf{ASes}$. We say that $\mathcal{PV}$ *diverges* over $\mathcal{I}$ if during its execution there is at least one AS in $\mathcal{I}$ that keeps on switching between different routes ad infinitum. If $\mathcal{PV}$ does not diverge, but subroute consistency is not satisfied, we say that $\mathcal{PV}$ neither converges nor diverges but comes to an inconsistent, stable state. This is relevant to our discussion of a particular class of attacks in partial PKI deployment scenarios in Section 3.8.

The convergence requirement may be unnecessarily complicated and we do not consider it in the correctness definition of $\mathcal{PV}$ protocols in this chapter. Thus, according our correctness requirement, some $\mathcal{PV}$ protocols may be technically correct while being useless in practice. However, as we will explain further in Section 3.5, in this chapter we focus only on the vulnerabilities of path-vector protocols preventable with cryptographic tools that deal strictly with honest ASes accepting bogus announcements, but do not consider network destabilizing attacks. We address such attacks in Chapter 4.

## *3.4 How BGP and S-BGP Work*

In this section we fist describe BGP using the language we developed in the previous section, and then show how S-BGP extends it to incorporate security features. Although in our model we do not require communication to be either concurrent or asynchronous, for the rest of the paper we assume only asynchronous communication as it captures delays and re-ordering ubiquitous in practice.

### 3.4.1 The Border Gateway Protocol

We present the essential aspects of the Border Gateway Protocol (BGP) that is used to establish routes on the Internet of today. Let $\mathcal{I} = (\mathbf{G} = (\mathsf{ASes}, \mathsf{link}), \mathsf{Prefixes}, \mathsf{OrforPr}, \mathsf{relation}, \mathsf{preferto}, \mathsf{policy})$ be an interdomain network. BGP uses no PKI and no CA, so the optional algorithm $\mathsf{Init}$ is never invoked. We now the describe the $\mathsf{An}$ algorithm.

Every AS $N \in \mathsf{ASes}$ maintains state in the form of a table $T_N$, called the *routing table*, which is initially empty. Each field $T_N[P]$ indexed by a prefix $P \in \mathsf{Prefixes}$, for which $\mathsf{OrforPr}(P) \neq N$, is a list consisting of routes to $P$ that $N$ has received as announcements from neighbors. Each route in $T_N[P]$ is ranked such that $T_N[P][i]$ contains $N$'s $i^{\text{th}}$ most preferred route to $P$.

If an AS's input $\mathbf{P}_N$ is nonempty (*i.e.,* $N \in \mathsf{Origins}$), then for every prefix $P \in \mathbf{P}_N$, $N$ sends an announcement $(N, N', (N), P, 0, \varepsilon)$, advertising access to $P$, to every neighbor $N' \in \mathsf{policy}_N((N), \varepsilon)$.

During BGP's execution, when an AS receives an announcement advertising a new route to some prefix, that announcement is ignored if the advertised route is already contained in that AS's routing table to that prefix, or if that AS is contained in the announced route. The latter condition is required to prevent routing loops. Otherwise, that AS determines the new route's rank in its routing table to the same prefix, records that route and its rank, and, if necessary, updates the ranks of the other routes to that prefix. If the announced route becomes the most preferred route to that prefix, that AS propagates that route to its neighbors in accordance with its export policy rules. If an AS receives an announcement that is a notification of a withdrawal of a route (*i.e.,* that route should not to be used by the receiving AS) that it has stored in its routing table, that AS deletes that entry from its table and propagates that route's withdrawal to its neighbors in accordance with its export policy rules. Let us now describe BGP more concretely.

For every route announcement $(N', N, R, P, W, \varepsilon)$ that $N$ receives from neighbor $N'$, if $R$ and $T_N[P]$ do not contain $N$ and $R$ respectively, $N$ sends a route announcement to every neighbor as per $\mathsf{policy}_N$ and updates $T_N[P]$ according to rules (1)-(3) below.

(1) If the announcement presents the most preferred route to $P$, *i.e.,* $W = 0$ and $R$ $\mathsf{preferto}_N$ $T_N[P][1]$, then $N$:

   (a) sends a route withdrawal announcement ($N$, $N'$, ($N$, $T_N[P][1]$), $P, 1, \varepsilon$) to every neighbor as per $\mathsf{policy}_N$ (although in practice withdrawals in this specific scenario may be implicit, we make them explicit here for clarity),

   (b) sends a route advertisement $(N, N',\ (N, R),\ P, 0, \varepsilon)$ to every neighbor as per $\mathsf{policy}_N$,

   (c) increments by one the rank of every route in $T_N[P]$ and makes an update $T_N[P][1] \leftarrow R$.

(2) If the announcement presents a route to $P$ that is not the most preferred, *i.e.,* $W = 0$ and $T_N[P][1]$ $\mathsf{preferto}_N$ $R$, then $N$ determines rank $i$ such that $R$ is the $i^{\text{th}}$ most preferred route out of all routes in $T_N[P]$, increments by one the rank of every route in $T_N[P]$ that is less preferred than $R$, and makes an update $T_N[P][i] \leftarrow R$.

(3) If the announcement is a withdrawal of a route that $N$ has stored, *i.e.,* $W = 1$ and $R \in T_N[P]$, then $N$:

   (a) if $R = T_N[P][1]$, sends a withdrawal announcement ($N$, $N'$, $(N, R)$, $P, 1, \varepsilon$) to every neighbor as per $\mathsf{policy}_N$,

   (b) if $R = T_N[P][1]$ and $T_N[P][2] \neq \varepsilon$, sends a route advertisement $(N, N',$ $(N, T_N[P][2]),\ P, 0, \varepsilon)$ to every neighbor as per $\mathsf{policy}_N$,

(c) removes $R$ from $T_N[P]$ and decrements the rank of every route in $T_N[P]$ ranked higher than $R$.

$N$ ignores new announcements in all other cases. In the absence of adversaries and errors, no message in BGP should be rejected, so BGP should be correct for various interesting classes of networks believed to closely capture how routing is done on the Internet of today, such as the ones presented in [38, 47, 41] in Chapter 4 of this thesis. Although BGP route announcements in practice may contain more information that could be stored, for instance, in the *Aux* field, than what we present above, this information is not essential for our analysis.

### 3.4.2 The Secure Border Gateway Protocol

The Secure Border Gateway Protocol (S-BGP) [60] is an extension to BGP that relies on the full deployment of PKI such that each AS should know authentic and valid public keys of other ASes. In S-BGP, public-key cryptography is used to bind prefixes to their origins with certificates, called *address attestations*, issued by a third trusted party as well as to generate *route attestations*—certificates generated by intermediate ASes on route announcements they propagate. A recipient of a route announcement verifies the origin of the prefix in that announcement and the certificates of the ASes on the route that announcement has traversed. Let us now present the essential operations of S-BGP more concretely.

**Construction 3.4.1.** *Let* $\mathcal{I} = (\mathbf{G} = (\mathsf{ASes}, \mathsf{link}), \mathsf{Prefixes}, \mathsf{OrforPr}, \mathsf{relation}, \mathsf{preferto}, \mathsf{policy})$ *be an interdomain network, let* $\mathcal{SS} = (\mathsf{Kg}, \mathsf{Sign}, \mathsf{Ver})$ *be a signature scheme with* $\mathsf{MsgSp} = \{0,1\}^*$*, and let* $\mathcal{CP}_s = (\mathsf{Kg}_{\mathsf{CA}}, (\mathsf{CA}, \mathsf{U}), \mathsf{Vercert})$ *be the corresponding certification protocol as per Construction 3.2.1. In S-BGP* $= (\mathsf{Init}, \mathsf{An})$*, as part of* Init *the CA runs* $\mathsf{Kg}_{\mathsf{CA}}(1^k)$ *to generate* $(pk_{\mathsf{CA}}, K_{\mathsf{CA}})$ *and each AS runs* $\mathsf{Kg}(1^k)$ *to generate* $(pk, K)$*.* An *is defined as follows.*

*If AS $N_j$'s input $\mathbf{P}_{N_j}$ is nonempty (*i.e., $N_j \in \mathsf{Origins}$*), then for every prefix $P \in \mathbf{P}_{N_j}$, $N_j$ does the following:*

- *$CA$ and $N_j$ interact according to $(\mathsf{CA}, \mathsf{U})$, $N_j$ being $\mathsf{U}$. The input to $\mathsf{U}$ is $(pk_{\mathsf{CA}}, N_j, P)$, the input to $\mathsf{CA}$ is $(K_{\mathsf{CA}}, N_j, P)$ and the outputs of both parties are $(N_j, P, \mathsf{cert})$. Address attestation $AA^P_{N_j} \equiv \mathsf{cert}$ is $N_j$'s certificate of ownership of $P$.*

- *Next, for every $N_i \in \mathsf{policy}(N_j, \epsilon)$, $N_j$ runs $\mathsf{Sign}(K_{N_j}, (N_i, N_j, P))$ to produce a route attestation, $RA^i_{R_j}$, and sends $(N_j, N_i, R = (N_j), P, 0, Aux = (RA^i_{R_j}, AA^P_{N_j}))$ to $N_i$; here $R_j$ is $R$'s subroute authorized by $N_j$ for $N_i$ to use and propagate in its own route announcements.*

*For every new route announcement $(N_{j-1}, N_j, R = (N_{j-1}, \ldots, N_1), P, W, Aux = (RA^j_{R_{j-1}}, \ldots, RA^2_{R_1}, AA^P_{N_1}))$ that $N_j$ receives from some neighbor $N_{j-1}$, $N_j$ first performs address attestation and route attestation verification steps as follows. $N_j$ runs $\mathsf{Vercert}$ $(pk_{\mathsf{CA}}, N_1, P, AA^P_{N_1})$ and outputs $\perp$ if the output of this computation is $0$. Otherwise, $N_j$ runs $\mathsf{Ver}(pk_{N_i}, (N_{i+1}, \ldots N_1, P), RA^{i+1}_{R_i})$ for every $1 \leq i \leq j-1$ and outputs $\perp$ if at least one such computation outputs $0$. If none of the verification steps above results in $\perp$, then $N_j$ performs the same operations as $N_j$ would do in BGP upon receipt of $(N_{j-1}, N_j, R, P, W, \varepsilon)$, as per rules (1)-(3) specified in Section 3.4.1. Then, for every announcement $(N_j, N_{j+1}, R', P, W', \varepsilon)$ that $N_j$ would send to $N_{j+1}$ in BGP, $N_j$ now runs $\mathsf{Sign}(K_{N_j}, (N_{j+1}, R', P))$ to get $RA^{j+1}_{R'_j}$ and sends $(N_j, N_{j+1}, R', P, W', Aux')$ to $N_{j+1}$ instead, where $R' = (N_j, R)$ and $Aux' = (RA^{j+1}_{R'_j}, Aux)$.*

If the underlying signature scheme $\mathcal{SS}$ is correct, the execution of S-BGP is the same as that of BGP in terms of how ASes update their routing tables and how they decide which routes to announce to their neighbors. Therefore, S-BGP is correct for the same classes of networks as BGP if the underlying signature scheme $\mathcal{SS}$ used to generate address and route attestations is correct.

Note that in our description of S-BGP, ASes do not sign the withdrawal flag $W$. We do not consider attacks that involve modification of this field in this thesis because, as suggested in [62], IPSec [61, 35] could be used to prevent such attacks on announcements exchanged between neighboring ASes. This is why we primarily focus on authentication attacks that involve manipulation of the routing information being announced.

Note that in BGPSEC, functionality of PKI and generation of origin attestations are provided by the Resource Public Key Infrastructure (RPKI) while the rest of the protocol is essentially the same S-BGP as per Construction 3.4.1.

## *3.5 Security of Path-Vector Protocols*

In this section we present a security definition for path-vector protocols, show how it captures their security vulnerabilities, and discuss the attacks not captured in our model because they cannot be solved with cryptographic tools.

### 3.5.1 Intuition for the Formal Security Model

In our model, we do not consider malicious CA's, but we do consider malicious ASes. We consider an adversary which is given the CA's public key and the description of the network $\mathcal{I}$ with at least two ASes. The adversary also specifies which ASes will not have public keys and which ASes it wants to corrupt. The adversary is allowed to adaptively corrupt as many ASes as it wants at any point of its attack. In practice, it is unlikely that a malicious party knows the complete configuration of the network including the relations, and can corrupt as many ASes as it wants, but in the definition we target a very strong adversary. We allow the adversary to corrupt multiple ASes to capture collusion. On the Internet, collusion is certainly a plausible scenario, given that multiple ASes could be managed by a single administration with presence in different geographic locations. The adversary is given all the public and secret keys of the corrupted ASes. We assume that the adversary is stateful, *i.e.,* it

can preserve state in between stages. All ASes and the CA can interact: the honest ASes and the CA follow the protocol, while the adversary can act arbitrarily on behalf of the corrupted ASes. It can observe and modify all communication.

The adversary wins if it sends a route announcement to an honest AS, the AS accepts it and either (1) the prefix in the announcement does not belong to the corresponding origin, (2) there is an honest AS on the route that never sent the corresponding announcement for the same prefix, and (3) the route is invalid. The latter includes the possibilities of a non-existing (not-connected) route and a route that does not satisfy the export policies of at least one AS on that route. We now present our security definition for path-vector protocols concretely.

### 3.5.2 Path-Vector Protocol Security Definition

Let $k \in \mathbb{N}$ be the security parameter, $\mathcal{I} = (\mathbf{G} = (\mathsf{ASes}, \mathsf{link}), \mathsf{Prefixes}, \mathsf{OrforPr}, \mathsf{relation}, \mathsf{preferto}, \mathsf{policy})$ be an interdomain network, of size polynomial in $k$, such that $|\mathsf{ASes}| \geq 2$, and let $\mathcal{PV} = (\mathsf{Init}, \mathsf{An})$ be a path-vector protocol that is correct for $\mathcal{I}$. We define the experiment $\mathbf{Exp}^{\text{sec-rout-m}}_{\mathcal{I},\mathcal{PV}}(A)$, for $0 \leq m \leq |\mathsf{ASes}|$, involving a stateful adversary $A$ as follows.

Given the description of $\mathcal{I}$, $A$ selects the set $\mathsf{nopubk} \subsetneq \mathsf{ASes}$ of ASes that will not have public keys, such that $|\mathsf{nopubk}| = m$. Then, the public–secret key pairs for the CA and all ASes in $\mathsf{ASes} \setminus \mathsf{nopubk}$ are generated via $\mathsf{Init}(1^k)$. Here and further in this chapter $\mathbf{pk}$ denotes the vector of public keys of ASes in $\mathsf{ASes}\setminus\mathsf{nopubk}$ and $\mathbf{pk}[i]$ denotes its $i$'th component. Given all public keys, $A$ can output the initial sets of corrupted and honest ASes which form a partition of $\mathbf{G}$: $(\mathsf{Honest}, \mathsf{Corrupted}) \xleftarrow{\$} A(\mathcal{I}, pk_{\mathsf{CA}}, \mathbf{pk})$, so that $\mathsf{Honest} \cup \mathsf{Corrupted} = \mathsf{ASes}$ and $\mathsf{Honest} \cap \mathsf{Corrupted} = \emptyset$.

Next, $A$ is given all the secret keys of the corrupted ASes $\{\mathbf{sk}[i]\colon\ \mathbf{sk}[i]$ belongs to a corrupted AS$\}$, and it starts the execution of $\mathsf{An}$ on behalf of all ASes in $\mathsf{Corrupted}$ with the CA and also with the ASes in $\mathsf{Honest}$. The CA and the honest ASes follow

the protocol legitimately, while the adversary can act arbitrarily. In particular, $A$ is allowed to intercept and modify announcements exchanged between neighboring honest ASes as well as send messages on behalf of any honest AS. $A$ is given transcripts of all communication as it happens. $A$ is also allowed to adaptively corrupt more honest ASes, thereby reducing $\mathsf{Honest}$ and increasing $\mathsf{Corrupted}$, as it wishes during this stage of the experiment.

$A$'s goal is to have an honest AS, say $N_\ell \in \mathsf{Honest}$, accept an announcement of the form $(N_{\ell-1}, N_\ell,\ R = (N_{\ell-1}, \ldots, N_1),\ P, W, Aux)$, so that at least one of the following conditions is true. Note that the indexing of the ASes on the route is not essential for the definition and is done for simplicity only.

1. *Unauthentic origin:* $\mathsf{OrforPr}(P) \neq N_1$. In this case the experiment outputs 1.

2. *Unauthentic route:* there exists $1 \leq i \leq \ell - 1$ so that $N_i \in \mathsf{Honest}$ and $N_i$ never sent announcement $(N_i,\ N_{i+1}, R' = (N_i, \ldots, N_1), P,\ W', Aux')$ for any $W', Aux'$ to $N_{i+1}$. In this case the experiment outputs 2.

3. *Invalid route:* R is invalid. In this case the experiment outputs 3.

$\mathbf{Exp}^{\text{sec-rout-m}}_{\mathcal{I},\mathcal{PV}}(A)$ returns an output as soon as $A$ meets at least one of the winning conditions. If more than one condition above holds, $\mathbf{Exp}^{\text{sec-rout-m}}_{\mathcal{I},\mathcal{PV}}(A)$ outputs the smallest number. We define $A$'s advantage $\mathbf{Adv}^{\text{sec-rout-m-b}}_{\mathcal{I},\mathcal{PV}}(A)$ in this experiment as $\Pr\left[\,\mathbf{Exp}^{\text{sec-rout-m}}_{\mathcal{I},\mathcal{PV}}(A) = b\,\right]$, for $b \in \{1, 2, 3\}$.

We define $\mathcal{C}^{\mathcal{PV}}_m$ to be the class of all networks which have $m$ ASes without public keys and for which a path-vector protocol $\mathcal{PV}$ is correct, for $m \leq |\mathsf{ASes}|$. $\mathcal{PV}$ guarantees *origin authentication*, *route authentication*, and *route validity* with $m$-partial deployment ($m$-PD) for a class of networks $\mathcal{C}^{\mathcal{PV}}_m$, if for every $\mathcal{I} \in \mathcal{C}^{\mathcal{PV}}_m$, for every efficient adversary $A$, the probability that $\mathbf{Exp}^{\text{sec-rout-m}}_{\mathcal{I},\mathcal{PV}}(A)$ returns 1, 2 and 3 respectively is negligible in $k$. $\mathcal{PV}$ is *fully secure* with $m$-PD for a class of networks $\mathcal{C}^{\mathcal{PV}}_m$, if it guarantees origin authentication, route authentication and route validity with $m$-PD

for $\mathcal{C}_m^{\mathcal{PV}}$, *i.e.,* for every $\mathcal{I} \in \mathcal{C}_m^{\mathcal{PV}}$, for every efficient adversary $A$, the probability of $\mathbf{Exp}^{\text{sec-rout-m}}_{\mathcal{I},\mathcal{PV}}(A)$ returning 1, 2 or 3 is negligible in $k$. When $m = 0$, we omit the suffix 0-PD when qualifying security of protocols.

Note that, by definition, although $A$ is allowed to adaptively corrupt as many ASes as it desires at any point of the experiment, $A$ cannot be successful in $\mathbf{Exp}^{\text{sec-rout-m}}_{\mathcal{I},\mathcal{PV}}(A)$ if it is ever the case that $\mathsf{Honest} < 2$.

Our security definition does not consider rogue keys and replay attacks. This is very common as it is known that the standard measures, such as proofs of possession of secret keys during the key registration [6, 89] and the use of timestamps, can be used to provide protection against such attacks. To address rogue key attacks, we could require the adversary to output the public and secret keys of corrupted users to model the situation where users are required to perform proofs of knowledge of secret keys during key registration. However, all of our results would still trivially hold in this setting, so we do not complicate our model with this extension since rogue-key attacks are not essential to routing protocols and do not enhance the insights we get about the essential, routing-related attacks on BGP. Although it may be relevant to investigate whether simpler proofs of possession [89, 20] will suffice, we do not consider this point in this thesis. We discuss rogue key attacks with respect to RPKI in Section 3.8.3.

We also note that our security notion does not capture the goal of guaranteeing that the data that ASes send to those prefixes travels along the routes that they have learned and selected, or whether it reaches those prefixes at all. As discussed in [44], path-vector protocols cannot and were not intended to provide such guarantees. These are not goals of path-vector protocols, but of data-plane accountability and verification which we do not consider in this thesis.

Although our security model does not capture all complexities of routing protocols, in Sections 3.6-3.8 we show that even our simplified model can point out what is

necessary, not just sufficient, to achieve security with respect to essential, fundamental path-vector protocol vulnerabilities in full and partial PKI deployment scenarios.

### 3.5.3 Known Captured Attacks

In this section we discuss how our compact model captures many known vulnerabilities of path-vector protocols. For all figures in this section, a directed edge from $N$ to $N'$ indicates that $N$ is $N$'s customer, *i.e.,* $N$ pays $N'$ for all traffic exchanged on their link.

#### 3.5.3.1 Unauthentic Origin

The *Unauthentic origin* condition captures the prefix hijacking attack on BGP, where a corrupt AS claims to own a prefix or announces a more specific prefix, say $P$, that is owned by another AS. As a result, the corrupt AS could attract potentially all traffic destined to $P$. Such attacks happen almost on weekly basis and are believed to be the cause of most routing outages on the Internet. Some of the famous examples of prefix hijacks on the Internet that made it to the news include, but are not limited to, the Pakistan Telecom hijacking YouTube's prefix in February 2008 [24] and more recently Turk Telekom hijacking prefixes of public DNS servers of Google and Level3 in March 2014 [19]. The purpose of these attacks was to deny access to particular websites, such as YouTube and Twitter, for censorship purposes, either by creating a *black hole*—a locale where all traffic destined to $P$ disappears, or redirecting traffic to a bogus page. In addition, the attacker could intercept sensitive, government-related traffic to analyze it for malicious reasons, as speculated by some with regards to China Telecom diverting approximately 15% of Internet's traffic in April, 2010 for about 20 minutes [34]. Prefix deaggration attacks, in which an attacker deaggregates a prefix into more specific prefixes to attract traffic, are also captured by the unauthentic origin condition. This works because routers on the Internet select more specific prefixes over less specific ones by default. Figure 2(a) presents an example of such an

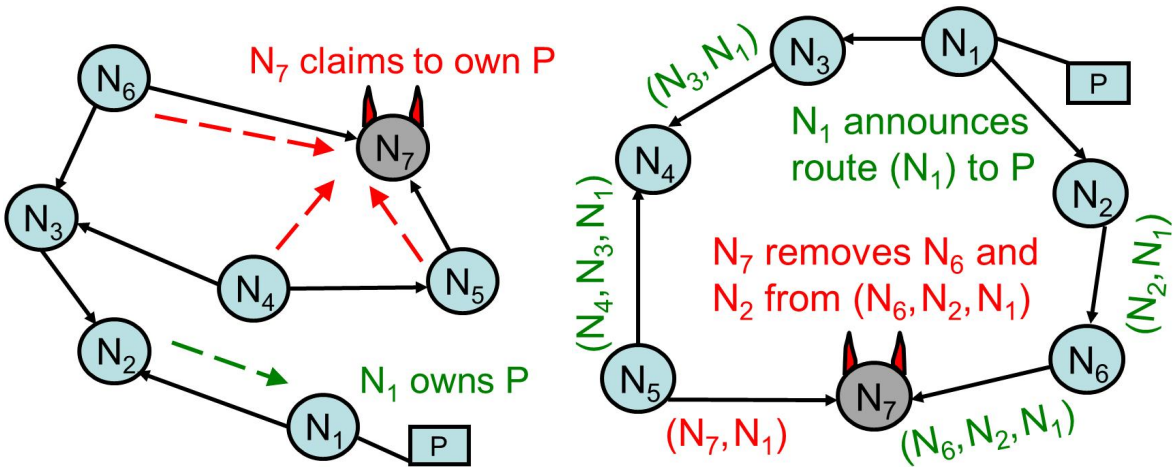


(a) Origin Authentication Attack (b) Route Authentication Attack

**Figure 2:** In (a) $N_7$ claims to own prefix $P$ and becomes a black hole by attracting majority of traffic destined to $P$ and dropping it. In (b) $N_7$ attracts $N_5$'s traffic by advertising a fake short route and then forwarding along a longer route via $N_6$.

attack, where AS $N_7$ announces to its neighbors ownership of prefix $P$, whose actual owner is $N_1$. As a result, $N_7$ is able to attract traffic from $N_4$, $N_5$, and $N_6$, because $N_7$ is closer to them than $N_1$. This traffic never reaches $N_1$ because, other than through ASes $N_5$ and $N_6$, $N_7$ does not have an alternative route to $N_1$.

Note that the implied requirement in this condition that every prefix should belong to the designated AS in the routing announcement is very strict because it disallows non-corrupted origin ASes to delegate the announcements of some or all of their prefixes to other non-delegated ASes. This type of delegation could be done for various purposes such as measuring risk or performance, but we ignore any such applications in this thesis as they are not essential to our analysis. Also note that RPKI [3] is a major, current effort by ARIN [1] to address origin authentication attacks, but by itself RPKI is not intended to address any other types of attacks that we capture in the next condition.

#### 3.5.3.2 Unauthentic Route

The *Unauthentic route* condition captures known attacks on BGP where an adversarial AS modifies the path attribute of a route announcement by adding and/or taking ASes out of this attribute as well as pretending to be a different AS altogether. By taking ASes out of the path attribute, the attacker could attract more traffic as the advertised route would seem shorter (and thus more preferred). and/or the advertised route may no longer contain AS(es) that the receiver of the advertisement wants to avoid for business/political reasons. Adding ASes to a route may make a route less attractive if it makes it seem longer, or contains the receiver of the announcement, which would present a loop and cause the receiver to ignore the announcement. This is how an attacker could force an AS not to select certain routes. Figure 2(b) presents an example of such an attack, where AS $N_7$ removes $N_6$ and $N_2$ from the shortest route that $N_7$ has to $P$, which is owned by $N_1$. This makes $N_5$ believe that $N_7$ is providing a shorter route to $N_1$ than the one through $N_4$, and hence $N_5$ picks the route through $N_7$. Thus, $N_5$ selects a suboptimal route to $P$, since the route to $P$ through $N_7$ is actually longer than that through $N_4$. The attacker benefits not only from intercepting $N_5$'s traffic but also from receiving $N_5$'s payment, since $N_5$ is $N_7$'s customer. We call this type of attack, where the attacker pretends to be a neighbor of the origin, *one-hop hijacks*, and we will focus on this type of attack in Chapter 5.

*Connection authentication* between adjacent ASes is a special case of route authentication in our security definition. $\mathcal{PV}$ guarantees connection authentication for some network $\mathcal{I}$, whose size is polynomial in $k$ and for which $\mathcal{PV}$ is correct, when the probability of the following event is negligible in $k$: the adversary $A$ in $\mathbf{Exp}^{\text{sec-rout-m}}_{\mathcal{I},\mathcal{PV}}(A)$ succeeds in having some honest AS $N_\ell$ accept an announcement of the form $(N_{\ell-1}, N_\ell,\ R = (N_{\ell-1}, \ldots,\ N_1),\ P, W, Aux)$, while $N_{\ell-1}$ is in Honest and has never output announcement $(N_{\ell-1},\ N_\ell, R, P,\ W', Aux')$ to $N_\ell$, for any $W'$ and $Aux'$. This event captures any attack in which an attacker actively modifies route

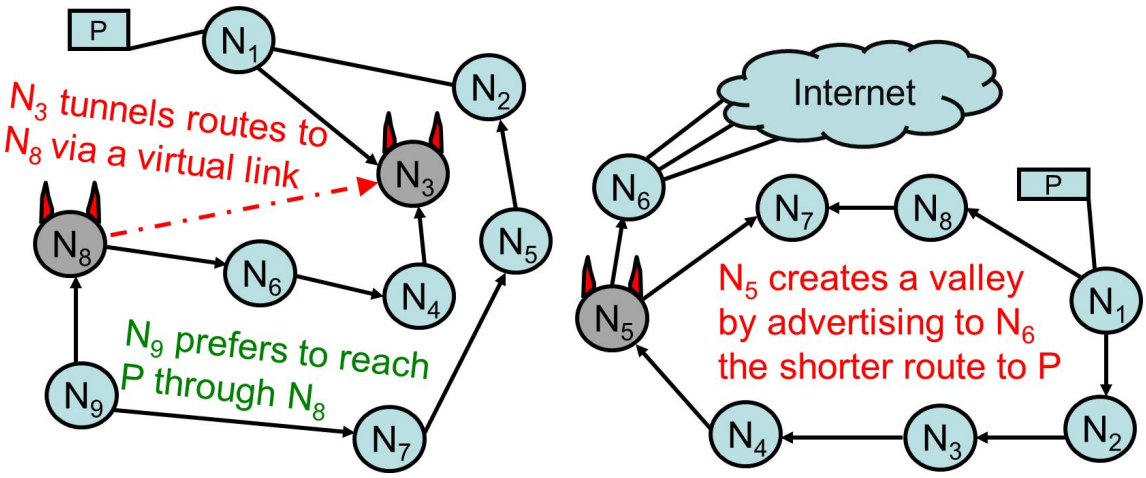


(a) Fake Link Attack (b) Export Policy Attack

**Figure 3:** In (a) colluders $N_8$ and $N_3$ create a fake link between each other and attract $N_9$'s traffic. In (b) $N_5$ attracts traffic of its provider $N_6$ by violating an export policy rule.

announcements traveling on a link between two non-corrupted ASes and/or impersonates a different AS. A good example of this is the withdrawal attack, where the attacker attempts to make a AS withdraw a route that the attacker had never advertised to that AS but that was previously advertised to that AS by another AS. Since connection authentication is a special case of route authentication, we do not analyze it separately. Furthermore, provable solutions have already been proposed to address them. As suggested in [62], IPSec [61, 35] could be used to prevent attacks on privacy, authenticity and integrity of route announcements exchanged between two neighboring, non-corrupted ASes. This is why we do not consider withdrawal attacks explicitly in our analysis, but focus on authentication attacks that involve manipulation of the route being announced.

#### 3.5.3.3 Invalid Route

The *Invalid route* condition captures two known types of attacks on S-BGP, both of which can be used to increase revenue as well as intercept and analyze possibly sensitive traffic. The wormhole attack consists of non-neighboring, colluding ASes

attracting traffic by creating a fake (virtual) link between themselves, by tunneling announcements between each other, *e.g.,* via IPSec, thereby announcing infeasible routes [105]. When tunneling announcements, they can essentially *skip* intermediate ASes. Figure 3(a) shows how two ASes, $N_3$ and $N_8$, create a fake (virtual) link between each other, although there is no direct route between them. They provide a seemingly shorter route to $P$, so $N_9$ selects a route through $N_8$ and $N_3$, which is actually longer than the route through $N_7$ that $N_9$ would have selected otherwise. The export policy attack consists of an attacker attracting traffic by violating export policy rules. In the example of Figure 3(b), both $N_6$ and $N_7$ are $N_5$'s providers. By announcing to $N_6$ the shorter route to $P$ through $N_7$ instead of the longer route through $N_4$, $N_5$ creates a *valley*, *i.e.,* a route through two of its providers, thereby violating a common export policy rule used on the Internet [38]. When forwarding traffic, however, $N_5$ can use the longer route through its customer $N_4$, thereby causing $N_6$ and other ASes in the Internet to use a route longer than they have intended. The same attack can be carried out when either $N_6$ or $N_7$ or both are $N_5$'s peers.

Note that in route validity attacks, the adversary introduces routes that are malicious to other users even though they are legitimate from the perspective of S-BGP, *i.e.,* a route that is authentic does not have to be valid. In both types of route validity attacks, the attackers could benefit from intercepting a victim's traffic as well as receiving extra payment from their customers for forwarding it. For networks with more sophisticated export policy rules, more complicated export policy attacks are possible. Route validity attacks have been studied in [98] and [47, 99] respectively, but no provably secure solution has yet been proposed. Also, such situations may be caused by route leaks or non-malicious, unintentional misconfigurations [83, 74, 76] that could still result in responsible ASes suffering from poor performance and potentially substantial, financial losses.

### 3.5.4 Attacks Crypto Cannot Prevent

Here we discuss several attacks not captured by our security model for the reason that such attacks cannot be prevented with strictly cryptographic methods.

Path-vector protocol divergence cannot be prevented with only cryptographic tools since the adversary could keep on withdrawing and then re-announcing the same set of routes ad infinitum. However, since the number of total routes to every prefix is finite, when a protocol diverges, some routes must be periodically withdrawn and then re-announced again resulting in what is called route flapping. Therefore, protocol divergence could be mitigated with tools that prevent route-flapping, *e.g.,* route dampening [25]. Convergence of path-vector protocols to suboptimal routes, *i.e.,* routes that are not the most preferred, also cannot be prevented with only cryptographic tools since the adversary could just make sure that some ASes never receive announcements of the most preferred routes.

Bellovin and Gansner have studied link cutting attacks which involve physically (*e.g.,* with a DDoS attack) taking out edges out of a topology so that certain route announcements fail to propagate [18]. These attacks do not involve the adversary listening and intercepting data without being noticed. Although in our security model the adversary, having access to all communication, can prevent any link from being operational, we do not capture this attack in our security model because, in general, cryptographic tools cannot resolve these attacks due to their physical nature.

Finally, contrary to common intuition, path-vector protocols cannot guarantee that a particular route announcement was propagated along the route shown in that announcement. More concretely, no path-vector protocol $\mathcal{PV}$ can guarantee that for every network $\mathcal{I} \in \mathcal{C}_m^{PV}$, for every efficient adversary $A$, for any $m \in \mathbb{N}$, the following event occurs with negligible probability in $\mathbf{Exp}^{\text{sec-rout-m}}_{\mathcal{I},\mathcal{PV}}(A)$: $N_\ell \in \mathsf{Honest}$ accepts an announcement $(N_{\ell-1}, N_\ell, R = (N_{\ell-1}, \ldots, N_1), P, W, Aux)$ such that there exists $1 \leq i \leq \ell - 1$ so that $N_i$ has never output announcement $(N_i, N_{i+1}, R' = (N_i, \ldots, N_1), P,$

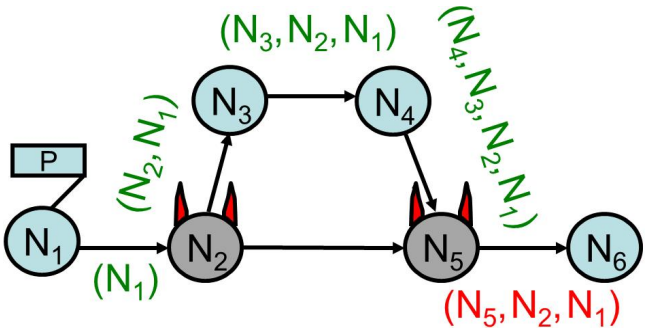


**Figure 4:** $N_5$ announces route $(N_5, N_2, N_1)$ to $N_6$ by signing on behalf of its colluding partner $N_2$, who never announced route $(N_2, N_1)$ to $N_5$.

$W', Aux'$) for any $W', Aux'$ to $N_{i+1}$. Notice that here $N_i$ is not required to be honest as it is in the unauthentic route condition in Section 3.5.2.

In Figure 4, we show this attack on S-BGP, where colluding corrupted ASes avoid using their expensive link by sending a route announcement through a route of honest ASes between them, and then taking these honest ASes out of the route announcement. Colluding ASes can do that because they can sign on behalf of each other. In this figure, colluding corrupted ASes $N_2$ and $N_5$ avoid using their expensive link $(N_2, N_5)$ by sending an announcement of a route to $P$ through honest ASes $N_3$ and $N_4$. After receiving this announcement from $N_4$, $N_5$ presents a route $(N_5, N_2, N_1)$ to $N_6$ by signing on behalf of its colluding partner $N_2$. $N_6$ accepts this announcement, even though $N_2$ has never announced route $(N_2, N_1)$ to $N_5$. Note that in real-life scenarios, ASes $N_2$ and $N_5$ could belong to a single administration with presence in different geographical locations and multiple distinct AS numbers.

## *3.6 How Secure is S-BGP?*

In this section we show that S-BGP guarantees *origin* and *route authentication*, assuming security of the building blocks, but that it is not fully secure because it does not guarantee *route validity*.

Let $\mathcal{SS} = (\mathsf{Kg}, \mathsf{Sign}, \mathsf{Ver})$ be a signature scheme, let $\mathcal{CP}_s = (\mathsf{Kg}, (\mathsf{CA}, \mathsf{U}), \mathsf{Vercert})$

be the corresponding straight-forward certification scheme as per Construction 3.2.1. Theorems 3.6.1-3.6.3 below state our results. While the first two are positive, the last result is negative.

**Theorem 3.6.1.** *S-BGP per Construction 3.4.1 guarantees* origin authentication *for* $\mathcal{C}_0^{S\text{-}BGP}$ *if the underlying* $\mathcal{SS}$ *is uf-cma-secure.*

*Proof.* The proof follows from Theorem 3.2.2 and Lemma 1 below. The latter is in fact more general than the above theorem. □

**Lemma 1.** *Construction 3.4.1 guarantees* origin authentication *for* $\mathcal{C}_0^{S\text{-}BGP}$ *if the underlying* $\mathcal{CP}$ *is uf-cda-secure.*

*Proof.* We show that for every adversary $A$ attacking origin authentication of S-BGP, there exists adversary $B$ attacking unforgeability of $\mathcal{CP}$ such that $\mathbf{Adv}_{\mathcal{CP}}^{\text{uf-cda}}(B) = \mathbf{Adv}_{\mathcal{I},\text{S-BGP}}^{\text{sec-rout-m-1}}(A)$ and the resources of $B$ are that of $A$.

Let $A$ be an efficient adversary attacking origin authentication of S-BGP for a network $\mathcal{I} = (\mathbf{G} = (\mathsf{ASes}, \mathsf{link}), \mathsf{Prefixes}, \mathsf{OrforPr}, \mathsf{relation}, \mathsf{preferto}, \mathsf{policy}) \in \mathcal{C}_0^{\text{S-BGP}}$, with $|\mathsf{ASes}| \geq 2$. We construct an adversary $B$ attacking $\mathcal{CP}$ as follows.

$B$ is given $pk_{\mathsf{CA}}$. $B$ generates $(pk_{N_j}, K_{N_j})$ for every $N_j \in \mathsf{ASes}$ by running $\mathsf{Kg}(1^k)$. $B$ then gives the description of $\mathcal{I}$ and all public keys to $A$, and the latter outputs the initial partition $(\mathsf{Honest}, \mathsf{Corrupted})$ of $\mathsf{ASes}$. Next, $B$ gives $A$ all the secret keys of the corrupted ASes, and then $A$ starts the execution of S-BGP on behalf of all ASes in $\mathsf{Corrupted}$ together with $B$ who executes S-BGP on behalf of all ASes in $\mathsf{Honest}$ and CA. $B$ follows S-BGP legitimately, whereas $A$ is allowed to act arbitrarily while observing all communication between all ASes in $\mathsf{ASes}$. $A$ is allowed to increase $\mathsf{Corrupted}$ by corrupting more ASes adaptively during its attack.

For each AS $N_j$ and prefix $P$ such that $N_j$ owns $P$ ($B$ can check this via $\mathsf{OrforPr}$) and either $N_j \in \mathsf{Honest}$ or $N_j \in \mathsf{Corrupted}$ and $A$ has requested address attestation $AA_{N_j}^P$ of $P$ for $N_j$, $B$ interacts with the CA via $(\mathsf{CA}(K_{\mathsf{CA}}, N_j, P),\ B(pk_{\mathsf{CA}},\ N_j,\ P))$

to get $(N_j, P, AA_j^P)$. $B$ stores all such certificates $AA_{N_j}^P$. This information together with all honest ASes' secret keys, allows $B$ to follow the computations according to the interactive algorithm $\mathsf{An}$.

Whenever $N_\ell \in \mathsf{Honest}$ accepts an announcement $(N_{\ell-1}, N_\ell, R = (N_{\ell-1}, \ldots, N_1), P, W, Aux = (RA_{R_{\ell-1}}^\ell, \ldots RA_{R_1}^2, AA_1^P))$ such that $\mathsf{OrforPr}(P) \neq N_1$, $B$ outputs $(N_1, P, AA_1^P)$.

Observe that $A$'s view in the simulated experiment has the same distribution as that in $\mathbf{Exp}_{\mathcal{I},\text{SBGP}}^{\text{sec-rout-m}}(A)$. Observer also that in accordance with S-BGP, $N_\ell$ accepts this announcement only if $\mathsf{Vercert}(pk_{\mathsf{CA}}, N_1, P, AA_{N_1}^P) = 1$, and, since $\mathsf{OrforPr}(P) \neq N_1$, this means that $B$ has not output $(N_1, P, AA_{N_1}^P)$ as a result of running $(\mathsf{CA}(K_{\mathsf{CA}}, N_1, P), B(pk_{\mathsf{CA}}, N_1, P))$ before. Thus, $\mathbf{Adv}_{\mathcal{CP}}^{\text{uf-cda}}(B) = \mathbf{Adv}_{\mathcal{I},\text{S-BGP}}^{\text{sec-rout-m-1}}(A)$. Finally, note that $B$'s running time is the same as that of $A$. □

**Theorem 3.6.2.** *S-BGP per Construction 3.4.1 guarantees* route authentication *for* $\mathcal{C}_0^{S\text{-}BGP}$ *if the underlying* $\mathcal{SS}$ *is uf-cma-secure.*

*Proof.* We show that for every adversary $A$ attacking route authentication of S-BGP, there exists adversary $B$ attacking unforgeability of $\mathcal{SS}$ such that

$$\mathbf{Adv}_{\mathcal{SS}}^{\text{uf-cma}}(B) = \frac{1}{|\mathsf{ASes}|}\mathbf{Adv}_{\mathcal{I},\text{S-BGP}}^{\text{sec-rout-m-2}}(A)$$

and the resources of $B$ are that of $A$ plus some overhead upper bounded by the size of network using S-BGP that $A$ is attacking.

Let $A$ be an efficient adversary attacking route authentication of S-BGP for a network $\mathcal{I} = (\mathbf{G} = (\mathsf{ASes}, \mathsf{link}), \mathsf{Prefixes}, \mathsf{OrforPr}, \mathsf{relation}, \mathsf{preferto}, \mathsf{policy}) \in \mathcal{C}_0^{\text{S-BGP}}$, with $|\mathsf{ASes}| \geq 2$. Let us now construct an adversary $B$ attacking $\mathcal{SS}$.

$B$ is given a public key $pk$ and the signing oracle $\mathsf{Sign}(K, \cdot)$. Let $n$ be the size of $\mathsf{ASes}$. $B$ first picks an AS's index at random $j \xleftarrow{\$} \{1, \ldots, n\}$ for AS $N_j \in \mathsf{ASes}$ and then generates public and secret keys for the CA and all ASes except $N_j$:

$(pk_{\mathsf{CA}}, K_{\mathsf{CA}}) \xleftarrow{\$} \mathsf{Kg}_{\mathsf{CA}}(1^k)$, $(\mathbf{pk}[1], \mathbf{sk}[1]), \ldots, (\mathbf{pk}[j-1], \mathbf{sk}[j-1]), (\mathbf{pk}[j+1], \mathbf{sk}[j+1])\mathbf{pk}[n], \mathbf{sk}[n]) \xleftarrow{\$} \mathsf{Kg}(1^k)$. $B$ sets $\mathbf{pk}[j] \leftarrow pk$.

Next, $B$ gives the description of $\mathcal{I}$ and all public keys to $A$ and the latter outputs its initial partition $(\mathsf{Honest}, \mathsf{Corrupted})$ of $\mathbf{G}$. If $N_j \in \mathsf{Corrupted}$, then $B$ aborts. Otherwise $B$ gives $A$ all the secret keys of the corrupted ASes.

Now $A$ starts the execution of S-BGP on behalf of all ASes in $\mathsf{Corrupted}$ together with $B$, who executes S-BGP on behalf of all ASes in $\mathsf{Honest}$ and CA. $B$ follows S-BGP legitimately, whereas $A$ can act arbitrarily. $B$ stores all the communication and also provides $A$ with all communication between all ASes. $B$ has all secret keys to simulate the execution of the protocol except for AS $N_j$. Whenever a secret-key operation is required from it, such as a route attestation for route $R$ destined to $N_j$'s neighbor, $B$ invokes its signing oracle to compute a signature on the corresponding data. $A$ is allowed to continue to corrupt more ASes adaptively as it wishes, and $B$'s simulation would change accordingly with the increase of $\mathsf{Corrupted}$ and the decrease of $\mathsf{Honest}$. $B$ aborts if $N_j$ ever becomes a member of $\mathsf{Corrupted}$.

Whenever $N_\ell \in \mathsf{Honest}$ accepts an announcement $(N_{\ell-1}, N_\ell, R = (N_{\ell-1}, \ldots, N_1), P, W, Aux = (RA^{\ell}_{R_{\ell-1}}, \ldots RA^{2}_{R_1}, AA^{P}_{1}))$ such that there exists $1 \leq i \leq \ell - 1$ so that $N_i \in \mathsf{Honest}$ but $N_i$ has never announced $(N_i, N_{i+1}, R' = (N_i, \ldots, N_1), P, W', Aux')$ for arbitrary $W', Aux'$ (we refer to this event by $A$ frames $i$), $B$ aborts if $N_i \neq N_j$. Otherwise (if $i = j$), $B$ outputs $((N_{i+1}, N_i, \ldots, N_1, P), RA_{R_i})$.

We see that if $B$ does not abort, then its simulation for $A$ is perfect, *i.e.,* $A$'s view has the same distribution as that in $\mathbf{Exp}^{\text{sec-rout-m}}_{\mathcal{I},\text{S-BGP}}(A)$.

Observe that, in accordance with S-BGP, $N_\ell$ accepts such an announcement only if $\mathsf{Ver}(pk, (N_{i+1}, N_i, \ldots, N_1, P), RA_{R_i}) = 1$, so $B$'s forgery is also valid. Similarly, $B$'s message $R'' = (N_{i+1}, N_i, \ldots, N_1, P)$ is new, *i.e.,* has not been queried to the signing oracle, because $\mathbf{Exp}^{\text{sec-rout-m}}_{\mathcal{I},\text{S-BGP}}(A) = 2$ only if $R''$ was not part of any announcement by $N_i$. This is true because, in S-BGP, the ID of the AS that is supposed to receive a route

announcement is always part of the message that is being signed to produce a route attestation. Therefore, $\Pr\left[\mathbf{Exp}^{\text{uf-cma}}_{\mathcal{SS}}(B)=1\right] = \frac{1}{|\mathsf{ASes}|}\Pr\left[\mathbf{Exp}^{\text{sec-rout-m}}_{\mathcal{I},\text{S-BGP}}(A)=2\right]$ which we justify as follows. Probability that $B$ wins by outputting $((N_{i+1}, N_i, \ldots, N_1, P), RA_{R_i})$, over all $i \in \mathsf{ASes}$, is

$$\begin{aligned}
\mathbf{Adv}^{\text{uf-cma}}_{\mathcal{SS}}(B) &= \textstyle\sum_{i\in\mathsf{ASes}} \frac{1}{|\mathsf{ASes}|}\Pr[\,A \text{ frames } i \mid i \notin \mathsf{Corrupted}\,]\Pr[\,i \notin \mathsf{Corrupted}\,] \\
&= \textstyle\frac{1}{|\mathsf{ASes}|}\sum_{i\in\mathsf{ASes}} \frac{\Pr[\,i\notin\mathsf{Corrupted} \mid A \text{ frames } i\,]\Pr[\,i\notin\mathsf{Corrupted}\,]\Pr[\,A \text{ frames } i\,]}{\Pr[\,i\notin\mathsf{Corrupted}\,]} \\
&= \textstyle\frac{1}{|\mathsf{ASes}|}\sum_{i\in\mathsf{ASes}} \Pr[\,A \text{ frames } i\,]\Pr[\,i \notin \mathsf{Corrupted} \mid A \text{ frames } i\,] \\
&= \textstyle\frac{1}{|\mathsf{ASes}|}\sum_{i\in\mathsf{ASes}} \Pr[\,A \text{ frames } i\,] \\
&= \textstyle\frac{1}{|\mathsf{ASes}|}\Pr\left[\mathbf{Exp}^{\text{sec-rout-m}}_{\mathcal{I},\text{S-BGP}}(A)=2\right] \\
&= \textstyle\frac{1}{|\mathsf{ASes}|}\mathbf{Adv}^{\text{sec-rout-m-2}}_{\mathcal{I},\text{S-BGP}}(A).
\end{aligned}$$

$B$ is efficient since, to simulate S-BGP, the number of queries it makes to $\mathsf{Sign}(K,\cdot)$ and their length are upper-bounded by the number of queries that $A$ makes and the size of $\mathcal{I}$ respectively. □

**Theorem 3.6.3.** *S-BGP as defined in Construction 3.4.1 does not guarantee route validity for* $\mathcal{C}_0^{S\text{-}BGP}$.

The proof formalizes the aforementioned attacks on S-BGP pointed out in [98, 47]. One attack deals with an adversary forging a connection that does not really exist in the network, and the other presents an adversary forging a route that violates the export policy of an intermediate AS. Either attack is sufficient to validate Theorem 3.6.3, and here we formalize just the former for simplicity. Note that the AS-level graph of the Internet is not a complete graph, so it is definitely vulnerable to this kind of attack.

*Proof.* We present an efficient adversary $A$ attacking route validity of S-BGP, by succeeding in having an honest AS accept an infeasible route, such that $\mathbf{Adv}^{\text{sec-rout-m-3}}_{\mathcal{I},\text{S-BGP}}(A)$

$= 1$.

We present a general attack in which the adversary corrupts two non-neighboring ASes that are on a valid route, obtains the corresponding route announcement from the corrupted AS closer to the origin, and then propagates the corresponding route announcement on behalf of the other corrupted AS. The route in the latter announcement is infeasible because the corrupted ASes are not neighbors, but there is no way to verify this fact by honest ASes down-stream from the corrupted AS that is farther away from the origin.

Consider an arbitrary network $\mathcal{I} \in \mathcal{C}_0^{\text{S-BGP}}$ that has at least one valid route $R$ that contains at least one pair of two non-neighboring ASes $N_i$ and $N_j$. $A$ is given the description of $\mathcal{I}$ and the public keys of the CA and all ASes. $A$ picks ASes $N_i$ and $N_j$ on a valid route $R = (N_\ell, \ldots, N_1)$, such that $1 < i < j < \ell$, which are not neighbors ($\mathsf{link}(N_i, N_j) = 0$). $A$ selects ($\mathsf{Honest} = \mathsf{ASes} \setminus \{N_i, N_j\}, \mathsf{Corrupted} = \{N_i, N_j\}$). $A$ gets the secret keys for the corrupted ASes and begins the execution of the interactive protocol $\mathsf{An}$ on their behalf. $A$ follows the protocol honestly. At some point of the protocol's execution, on behalf of $N_i$, $A$ receives an announcement of $R$'s subroute $\tilde{R}$, $(N_{i-1}, N_i, \tilde{R} = (N_{i-1}, \ldots, N_1), P, 0, Aux = (RA^i_{\tilde{R}_{i-1}}, \ldots, RA^2_{\tilde{R}_1}, AA_1^P))$ from $N_{i-1} \neq N_i$, where the components of $Aux$ are computed according to S-BGP's description in Section 3.4.2.

Then, $N_j$ sends the announcement $(N_j, N_{j+1}, R' = (N_j, N_i, \ldots, N_1), P, 0, Aux' = (RA^{j+1}_{R'_j}, RA^j_{R'_i}, \ldots, RA^2_{R'_1}, AA_1^P))$ to $N_{j+1}$. Note that in this announcement ASes $N_{i+1}, \ldots N_{j-1}$ are removed from $R$, so, since $\mathsf{link}(N_i, N_j) = 0$, $R'$ is infeasible. $N_{j+1} \in \mathsf{Honest}$ will not reject this announcement, because it will pass the verification process according to S-BGP, as all the signatures in $Aux'$ are valid, and there is no way in general for $N_{j+1}$ to verify whether $N_i$ and $N_j$ are neighbors or not.

$A$ is clearly efficient, and, for $m = 0$, $\mathbf{Exp}^{\text{sec-rout-m}}_{\mathcal{I},\text{S-BGP}}(A)$ will return 3 with probability 1, so $\mathbf{Adv}^{\text{sec-rout-m-3}}_{\mathcal{I},\text{S-BGP}}(A) = 1$. □

We now present a special case of the general attack described in the proof of Theorem 3.6.3 with the network depicted in Figure 3(a). Note that this network has at least one valid route $R = (N_8, N_6, N_4, N_3, N_1)$). The adversary $A$ is given the description of this network and the public keys of the CA and all ASes. $A$ corrupts two non-neighboring ASes $N_3$ and $N_8$ that are on a valid route $R$. $A$ gets the secret keys for the corrupted ASes and begins the execution of the interactive protocol $\mathsf{An}$ on their behalf. $A$ follows the protocol honestly. At some point of the protocol's execution, on behalf of $N_3$, $A$ receives an announcement from $N_1$, $(N_1, N_3, \tilde{R} = (N_1),\ P, 0, Aux = (RA^3_{\tilde{R}_1},\ AA^P_1))$, where the components of $Aux$ are computed according to S-BGP's description in Section 3.4.2. Then, on behalf of $N_8$, $A$ sends the announcement $(N_8,\ N_9, R' = (N_8,\ N_3, N_1, P, 0, Aux' = (RA^9_{R'_8}, RA^8_{R'_3}, RA^3_{R'_1},\ AA^P_1))$ to $N_9$. Note that in this announcement ASes $N_4$ and $N_6$ are removed from $R$, so, since $\mathsf{link}(N_3, N_8) = 0$, $R'$ is infeasible. Honest AS $N_9$ will not reject this announcement, because the latter will pass the verification process according to S-BGP, because all the signatures in $Aux'$ are valid, and there is no way for $N_9$ to verify whether $N_3$ and $N_8$ are neighbors or not.

## *3.7 Fully Secure BGP*

The attack on S-BGP from the proof of Theorem 3.6.3 exploits ASes' inability to verify whether remote ASes in the route announcements are neighbors. To address this attack, in this section we suggest a modification to S-BGP and show that the resulting protocol *provably* guarantees route validity assuming the underlying signature scheme is secure. We then argue that this modification is necessary. The modified protocol is fully secure with respect to our security definition from Section 3.5 under the same assumption, so we call it *Fully Secure* BGP (FS-BGP).

**Construction 3.7.1.** *Let* $\mathcal{I} = (\mathbf{G} = (\mathsf{ASes}, \mathsf{link}),\ \mathsf{Prefixes},\ \mathsf{OrforPr},\ \mathsf{relation},\ \mathsf{preferto},\ \mathsf{policy})$ *be an interdomain network, let* $\mathcal{SS} = (\mathsf{Kg},\ \mathsf{Sign}, \mathsf{Ver})$ *be a signature scheme,*

*and let* $\mathcal{CP}_s = (\mathsf{Kg}_{\mathsf{CA}}, (\mathsf{CA}, \mathsf{U}), \mathsf{Vercert})$ *be the corresponding certification protocol as per Construction 3.2.1. Let S-BGP* $= (\mathsf{Init}, \mathsf{An})$ *be the construction from Section 3.4.2. FS-BGP* $= (\mathsf{Init}, \mathsf{An}')$ *is defined exactly like S-BGP, but* $\mathsf{An}'$ *requires a few extra operations.*

*After all address attestations are generated and before any announcement is sent, each AS* $N_j$ *interacts with the CA via* $(\mathsf{CA}, \mathsf{U})$*. In what follows, smaller input is always on the left corresponding to any link* $(N_j, N_i)$*, and for convenience only, suppose that* $N_j = \min(N_j, N_i)$*, for every* $N_i \in \mathsf{Neighbors}(N_j)$*. For this interaction, the input to* $\mathsf{U}$ *is* $(pk_{\mathsf{CA}}, N_j, ((N_j, N_i), \mathsf{relation}(N_j, N_i)))$*, the input to* $\mathsf{CA}$ *is* $(K_{\mathsf{CA}}, N_j, ((N_j, N_i), \mathsf{relation}(N_j, N_i)))$ *and the outputs of both parties are* $(N_j, ((N_j, N_i), \mathsf{relation}(N_j, N_i)), \mathsf{cert})$*. We define* link attestation *to be* $LA_{N_jN_i} \equiv \mathsf{cert}$*. If* $N_j$ *owns prefix* $P \in \mathsf{Prefixes}$*, for every* $N_i \in \mathsf{policy}_{N_j}((N_j), \varepsilon)$*,* $N_j$ *generates a route attestation* $RA^i_{R_j}$ *just as in S-BGP and sends* $(N_j, N_i, R = (N_j), P, 0, Aux = ((\mathsf{relation}(N_j, N_i), LA_{N_jN_i}), RA^i_{R_j}, AA^P_{N_j}))$ *to* $N_i$*.*

*For every new route announcement* $(N_{j-1}, N_j, R = (N_{j-1}, \ldots, N_1), P, W, Aux = (\mathsf{relation}(N_{j-1}, N_j), LA_{N_{j-1}N_j}, RA^j_{R_{j-1}}, \ldots \mathsf{relation}(N_1, N_2), LA_{N_1N_2}, RA^2_{R_1}, AA^P_{N_1}))$ *that* $N_j$ *receives,* $N_j$ *first performs address and route attestation verification just as in S-BGP, and, if these steps do not result in* $\perp$*, then* $N_j$ *performs link attestation verification as follows.* $N_j$ *runs* $\mathsf{Vercert}$ $(pk_{\mathsf{CA}}, N_i, ((N_i, N_{i+1}), \mathsf{relation}(N_i, N_{i+1})), LA_{N_iN_{i+1}})$*, for every* $1 \leq i \leq j-1$*, and outputs* $\perp$ *if at least one such computation outputs* $0$*. Otherwise,* $N_j$ *outputs* $\perp$ *if there is at least one* $N_i$*, for* $1 \leq i \leq j-1$*, such that* $N_{i+1} \notin \mathsf{policy}_{N_i}((N_i, \ldots, N_1), \mathsf{relation}(N_i, N_{i-1}))$*.*

*If none of the verification steps above results in* $\perp$*, then* $N_j$ *performs the same operations as* $N_j$ *would do in S-BGP upon receipt of* $(N_{j-1}, N_j, R, P, W, RA^j_{R_{j-1}}, \ldots RA^2_{R_1}, AA^P_{N_1})))$*, and then, for every message* $(N_j, N_{j+1}, R', P, W', Aux')$ *that* $N_j$ *would send to* $N_{j+1}$ *in S-BGP,* $N_j$ *now sends* $(N_j, N_{j+1}, R', P, W', Aux'')$ *to* $N_{j+1}$ *instead, where* $R' = (N_j, R)$ *and* $Aux'' = (\mathsf{relation}(N_j, N_{j+1}), LA_{N_jN_{j+1}}, RA^{j+1}_{R_j}, \ldots$

$\mathsf{relation}(N_1, N_2)$, $LA_{N_1N_2}$, $RA^2_{R_1}$, $AA^P_{N_1}$)).

Note that FS-BGP is correct for the same classes of networks that BGP is correct for, if the underlying signature scheme $\mathcal{SS}$ used to generate address, route attestations and link attestations is correct.

**Theorem 3.7.2.** *FS-BGP as defined in Construction 3.7.1 is fully secure for* $\mathcal{C}_0^{FS\text{-}BGP}$ *if the underlying* $\mathcal{SS}$ *is uf-cma.*

*Proof.* The proof follows from Theorems 3.2.2, 3.6.1, 3.6.2 and Lemma 2 stated below. □

**Lemma 2.** *FS-BGP, as defined above, guarantees route validity for any network* $\mathcal{I} \in \mathcal{C}_0^{FS\text{-}BGP}$ *if the underlying* $\mathcal{CP}$ *is uf-cda-secure.*

*Proof.* The proof is very similar to the proof of Lemma 1. We show that for every adversary $A$ attacking route validity of S-BGP, there exists adversary $B$ attacking unforgeability of $\mathcal{CP}$ such that $\mathbf{Adv}^{\text{uf-cda}}_{\mathcal{CP}}(B) = \mathbf{Adv}^{\text{sec-rout-m-3}}_{\mathcal{I},\text{S-BGP}}(A)$ and the resources of $B$ are that of $A$.

Let $A$ be an efficient adversary attacking route validity of S-BGP for a network $\mathcal{I} = (\mathbf{G} = (\mathsf{ASes}, \mathsf{link}), \mathsf{Prefixes}, \mathsf{OrforPr}, \mathsf{relation}, \mathsf{preferto}, \mathsf{policy}) \in \mathcal{C}_0^{\text{S-BGP}}$, with $|\mathsf{ASes}| \geq 2$. We construct an adversary $B$ attacking unforgeability of $\mathcal{CP}$ as follows.

We use $\mathsf{CA}$ and CA interchangeably. Given the CA's public key $pk_{\mathsf{CA}}$, $B$ generates the public and secret keys for all ASes, and gives $A$ all public keys, including that of the CA. After $A$ outputs the initial sets of honest and corrupted ASes, $B$ interacts with $A$ according to the interactive protocol $\mathsf{An}'$. We observe that the only information that $B$ cannot initially compute are the address attestation certificates for any of the prefixes and link attestation certificates for any of the links. To obtain these certificates, $B$ can sequentially interact with the CA via $(\mathsf{CA}, B)$ on the appropriate inputs. Note that $A$'s forgery, *i.e.,* an announcement that contains an invalid route

that passes the verification, can be converted into $B$'s forgery. This is because an invalid route implies that at least one link attestation in the announcements field $Aux$ contains a valid signature on the data that the CA never signed. This is because an invalid route requires that at least one pair of subsequent ASes on a route advertised in that announcement are not neighbors and/or one of them violated the export policy rule. In our case, the policy only depends on the relationships between neighboring ASes. Therefore,

$$\mathbf{Adv}_{\mathcal{CP}}^{\text{uf-cda}}(B) = \mathbf{Adv}_{\mathcal{I},\text{S-BGP}}^{\text{sec-rout-m-3}}(A).$$

Note that $B$'s running time is the same as that of $A$. □

Assigning link attestations for every link in the Internet may seem impractical because the Internet contains many more edges than ASes (possibly over 200K versus 40K [30, 7]), their management is harder due to periodic reconfiguration, and ASes may be unwilling to make their connections, business relationships and export policies public. However, some ASes already post their policies on the Internet Routing Registry (IRR) [80]. Also, in principle, just as with address attestations, such certificates could be downloaded and verified off-line instead of being passed along with announcement on-line. The results of such verification could then be cached, which would significantly speed up the process of verifying origin authentication and route validity when processing routing announcements.

Furthermore, we argue that link attestations are necessary to prevent route feasibility attacks in general. This is because if a path-vector protocol guarantees route validity, every announcement received as part of this protocol can itself serve the role of a certificate for the links between the ASes in the route of that announcement. Since in our model arbitrary ASes on any route could be corrupted, such certificates would have to be generated independently by trusted parties. Analogously, to guarantee route validity when export policies of ASes are not publicly known and/or are not next-hop, more sophisticated certificates and in greater amounts (potentially one

for every route of every AS, and to every prefix of every origin) would have to be issued by a trusted authority to ensure that honest ASes can check for export policy violations of remote ASes.

Several plausible solutions to route leaks–unintentional export policy violations—and route validity attacks have been suggested without provable security analysis in [99, 76]. Although these solutions are more practical than FS-BGP because they are mostly based on restricted models of ASes business relationships and export policies, *e.g.,* models presented in [38], as our analysis in Section 3.10 shows, they do provide strictly weaker security guarantees and require on-line verification of route validity. Also, because business relationships and export policies of ASes on the Internet may be more complicated than in the model of [38], as we argued above, more sophisticated solutions than the ones proposed in [99, 76] would be necessary.

Link attestations are similar to AS Policy Certificates in SoBGP, which we will analyze later in this thesis in Section 3.9. FS-BGP is also similar to TASRS [81] that makes use of the reverse DNS and DNSSEC to address route leaks.

## *3.8 Partial Deployment of PKI*

In this section we study what happens to security guarantees when PKI is only partially deployed. We first show that neither S-BGP nor FS-BGP can guarantee route authenticity for networks in which there is at least a single AS without a public key, and then present variants of these protocols with which full security can be guaranteed in partial PKI scenarios.

### 3.8.1 Achieving Security in Partial PKI Deployment is Difficult

Before stating our main introductory result in Theorem 3.8.4, to develop intuition as to why providing security guarantees in scenarios with partial PKI deployment is a very difficult problem, we present a simple example of an attack where only one AS has no public key and only one AS is corrupted. First, we formalize the modification

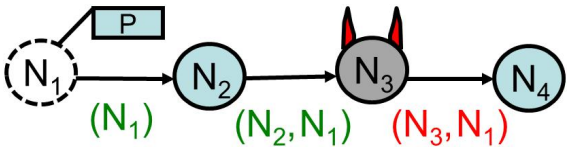


**Figure 5:** $N_1$ does not have a public key, and the adversary corrupts only $N_3$. In this route authentication attack $N_3$ takes $N_2$ out of the route and announces a shorter, infeasible route to $N_4$.

of allowing some ASes not to have public keys in S-BGP as follows.

**Construction 3.8.1.** *Let* $\mathcal{I} = (\mathbf{G} = (\mathsf{ASes}, \mathsf{link}), \mathsf{Prefixes}, \mathsf{OrforPr}, \mathsf{relation}, \mathsf{preferto}, \mathsf{policy})$ *be an interdomain network and* $k$ *a security parameter. We define S-BGP with partial deployment (S-BGP-PD)* $= (\mathsf{Init}', \mathsf{An}')$ *as a path-vector protocol identical to S-BGP* $= (\mathsf{Init}, \mathsf{An})$ *but with the following modifications. During execution of* $\mathsf{Init}(1^k)$ *not every AS has to generate a public key. During execution of* $\mathsf{An}'$*, ASes that do not have public keys do not generate route attestations. Also, route announcements that contain ASes without public keys are not checked for contents of route attestations corresponding to those keyless ASes during the route attestation verification.*

Notice that S-BGP is just a special case of S-BGP-PD when all ASes have keys. In Figure 5 we present a simple route authentication attack that shows that S-BGP-PD does not guarantee route authentication when $m = 1$, for $\mathcal{C}_1^{\text{S-BGP-PD}}$. The attack consists of the adversary taking an intermediate AS $N_2$ out of the route announcement during the execution of S-BGP-PD. This results in $N_4$ accepting a shorter, infeasible route, so in this scenario $\mathbf{Exp}^{\text{sec-rout-m}}_{\mathcal{I},\text{S-BGP-PD}}(A)$ returns 2 with probability 1.

**Remark 3.8.2.** *Given an attack in a network with* $m$ *ASes without public keys one can always construct an attack in a network with* $m'$ *ASes without public keys, for any* $m' > m$*, by making more ASes keyless in the same network.*

This is because increasing the number of keyless ASes cannot make a plausible attack implausible. Without affecting the attack, the number of ASes with public keys in

the network can be increased by adding neighbors to an origin. Thus, the attack in Figure 5 shows that for no $m \geq 1$ does S-BGP-PD guarantee route authentication with $m$-PD for $\mathcal{C}_m^{\text{FS-BGP-PD}}$.

Providing security guarantees in scenarios with partial PKI deployment is a difficult problem because ASes that do not have public keys cannot generate route attestations. The attack in Figure 5 works because S-BGP-PD does not guarantee route feasibility since ASes cannot find out using this protocol whether some remote ASes are neighbors or not. When not all ASes have public keys, providing ASes with the capability of verifying neighborship of remote ASes ultimately would require a certificate from a third trusted party, such as link attestations in FS-BGP. Let us define FS-BGP-PD to account for partial PKI deployment similarly to Construction 3.8.1. Notice that FS-BGP is just a special case of FS-BGP-PD when all ASes have keys. It can be easily shown that the route authentication attack in Figure 5 would not be possible if ASes were to use FS-BGP-PD to establish a route to $P$.

**Remark 3.8.3.** *FS-BGP-PD guarantees origin authentication and route validity with $m$-PD for any network in $\mathcal{C}_m^{FS\text{-}BGP\text{-}PD}$, for any $m \leq |\mathsf{ASes}|$, if the underlying $\mathcal{CP}$ is uf-cda-secure.*

This is because for networks in $\mathcal{C}_m^{\text{FS-BGP-PD}}$, in FS-BGP-PD origin authentication and route validity do not depend on whether ASes have public keys or not. However, we now show with the following result that even when origin authentication and route validity are guaranteed, route authentication cannot be guaranteed when $|\mathsf{nopubk}| > 0$.

**Theorem 3.8.4.** *For no $m \geq 1$ does FS-BGP-PD guarantee route authentication with $m$-PD for $\mathcal{C}_m^{FS\text{-}BGP\text{-}PD}$.*

*Proof.* Consider a network which contains at least one AS $N_i$ which does not have a public key, has a choice of at least two routes to the same prefix, and has a neighbor

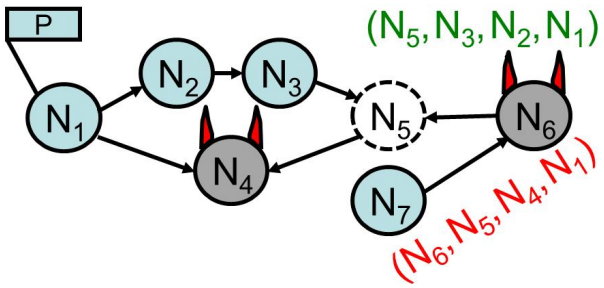


**Figure 6:** Only $N_5$ does not have a public key, and the adversary corrupts $N_4$ and $N_6$. In this Valid-Route Switching (VRS) route authentication attack $N_6$ announces to $N_7$ a valid route to $P$ that $N_5$ did not authorize $N_6$ to announce.

$N_j$ whose only access to that prefix is through $N_i$ and to whom $N_i$ is willing to export at least two different routes to that prefix.

We construct an efficient adversary $A$ such that $\mathbf{Adv}^{\text{sec-rout-m-2}}_{\mathcal{I},\text{S-BGP-PD}}(A) = 1$ as follows. Since $A$ can observe all communication, it can learn of all the routes announced to $N_i$ from $N_i$'s neighbors in addition to the routes $N_i$ announces to $N_j$. $A$ can intercept $N_i$'s announcement to $N_j$ and switch the route in that announcement to another valid route, available to $N_i$, that $N_i$ is willing to export to $N_j$. Since $N_i$ does not have a public key, it cannot generate a route attestation for its original announcement, so $N_j$ is bound to accept this false announcement that contains a valid route. Therefore, $\mathbf{Exp}^{\text{sec-rout-m}}_{\mathcal{I},\text{S-BGP-PD}}(A)$ returns 2 with probability 1 in this scenario, so $\mathbf{Adv}^{\text{sec-rout-m-2}}_{\mathcal{I},\text{S-BGP-PD}}(A) = 1$. The theorem then follows due to Remark 3.8.2. □

To show a pictorial example of the proof of Theorem 3.8.4, in Figure 6 we present an attack where the adversary switches a valid route announced by AS $N_5$ without a public key for another valid route that $N_5$ never announced. The adversary corrupts two ASes, $N_4$ and $N_6$. In this network, $N_5$ prefers the longer customer route through $N_3$ to the provider route through $N_4$ (recall that a directed edge from one AS to another indicates that the former pays the latter for all traffic exchanged on their link). However, $N_6$ switches $N_5$'s more preferred route to the one through

$N_4$ in its announcement to $N_7$, who accepts this route as authentic since $N_5$ does not have a public key (and thus cannot generate a route attestation). Therefore, $\mathbf{Exp}^{\text{sec-rout-m}}_{\mathcal{I},\text{S-BGP-PD}}(A)$ returns 2 with probability 1 in this scenario. By Remark 3.8.2 the same attack can be carried out for any $m > 1$.

The attack in the proof of Theorem 3.8.4, deserves a special name because we later show it to be the only type of attacks that can prevent FS-BGP-PD from being fully secure later in this section. A similar type of attack was known in the networking community to prevent SoBGP from guaranteeing route authentication.

**Definition 3.8.5** (The Valid-Route-Switching Attack)**.** *Let* $\mathcal{I} = (\mathbf{G} = (\mathsf{ASes}, \mathsf{link}),$ $\mathsf{Prefixes}, \mathsf{OrforPr}, \mathsf{relation}, \mathsf{preferto}, \mathsf{policy})$ *be a network in* $\mathcal{C}_m^{\mathcal{PV}}$*, for any* $1 \leq m \leq |\mathsf{ASes}|$*, such that* $|\mathsf{ASes}| \geq 2$*, let* $\mathcal{PV} = (\mathsf{Init}, \mathsf{An})$ *be a path-vector protocol correct for* $\mathcal{I}$ *and let* $k$ *be the security parameter such that the size of the description of* $\mathcal{I}$ *is polynomial in* $k$*. We consider the experiment* $\mathbf{Exp}^{\text{sec-rout-m}}_{\mathcal{I},\mathcal{PV}}(A)$*, involving an adversary* $A$*.*

*When* $\mathbf{Exp}^{\text{sec-rout-m}}_{\mathcal{I},\mathcal{PV}}(A)$ *outputs* 2, i.e., *when* $N_\ell \in \mathsf{Honest}$ *accepts announcement* $(N_{\ell-1}, N_\ell,\ R = (N_{\ell-1}, \ldots, N_1),\ P, W, Aux)$*, such that* $\exists 1 \leq i \leq \ell - 1$ *so that* $N_i \in \mathsf{Honest}$ *has never output announcement* $(N_i,\ N_{i+1}, R' = (N_i, \ldots, N_1), P,\ W', Aux')$ *for any* $W', Aux'$ *to* $N_{i+1}$*, if in addition* $N_i \in \mathsf{nopubk}$ *and* $R'$ *is a valid route to* $P$*, then this event is called a Valid-Route-Switching (VRS) attack.*

In the definition of the Valid-Route-Switching (VRS) attack, an honest AS $N_i$ may never announce to $N_{i+1}$ a valid route $R'$ to a particular prefix $P$ because $N_i$ may have never received any route announcements to $P$ from its neighbors or because $R'$ is not $N_i$'s most preferred route to $P$. Notice that VRS attacks can cause subroute inconsistency, so they can cause FS-BGP-PD to come to a stable but inconsistent state as in the example in the proof of 3.8.4, where FS-BGP-PD does not diverge since all ASes select their most preferred routes after a finite number of transmitted route announcements, but there is an inconsistency between the preferred routes of

$N_7$ and $N_5$.

### 3.8.2 The Relaxed Path-Vector Protocol Security Definition

In this section we first motivate two relaxations to our security definition, and we then justify that these relaxations are in fact reasonable on the Internet of today due to physical security of communication links and the trust relationship that neighboring ASes can establish when they agree to form business relationships. We then formalize and integrate these relaxations into our security model, in a form of restrictions on the adversary, to present a new security definition for path-vector protocols adequate for scenarios with partial PKI deployment. Finally, we present refinements to S-BGP-PD and FS-BGP-PD that address the weakness pointed out in the proof of Theorem 3.8.4, and prove that the refined protocols meet our new definition.

#### 3.8.2.1 Security Relaxations

Currently available technology allows honest neighboring ASes, whether with public keys or not, to establish communication channels that guarantee authentication and integrity. ASes could establish communication channels with their neighbors via IPSec that could guarantee integrity and authenticity, for which they may not need public keys as they could establish pre-shared keys off line. BGP TTL security hack [100] could also be used for this purpose. Although most of the time ASes establish connections at Internet Exchange Points (IXP), sometimes connections between ASes are established via fiber-optic cables outside of IXP's. Such cables mostly run underground and may be closely monitored for performance deviations. The transmitted data along such cables is transformed into optical signals that are impossible to interpret without expensive equipment. Thus, although attacking such cables is feasible in principle, as has been shown with recent revelations about the US National Security Agency (NSA) surveillance programs, it would be impractical for operationally limited adversaries in real life.

Sender authenticity could be added with appropriate gateway configurations, which associate neighboring ASes to specific outgoing and incoming ports, such that announcements get dropped when they come to the port not associated with the neighbor claiming to have sent them.

Note that, although other types of physical attacks on links between ASes are possible and have been studied before [18], these types of attacks do not involve listening and intercepting data without being noticed. The only purpose of these attacks is to take out links out of a topology so that certain route announcements are never made.

To establish a business relationship between themselves, neighboring ASes must be able to establish some level of trust between each other. Many ASes on the Internet are now multi-homed, so framing AS business partners on the Internet could lead to unwanted consequences such as the tearing down of their business contracts and possibly physical links connecting them, which could result in significant financial losses. Having established trust with their neighbors, ASes that do not have public keys could rely on their trusted down-stream neighbors with public keys to vouch for the former with their signatures.

As mentioned above, on the Internet, most connections between ASes are made at public or private IXP's which, intuitively, serve the role of rendez-vous points for ASes to exchange traffic. ASes that wish to connect at a particular IXP have to establish a physical connection at that IXP. Thus, since IXP's make a profit by providing basic infrastructure for ASes to make connections and become neighbors, it would be in their interest to facilitate the establishment of physically secure communication channels and trust between the participating ASes, as this would guarantee longer lasting business relationships for those ASes (which would imply longer lasting profits for the IXP connecting them).

We formally present these two main points in the following two relaxations.

Security Relaxations

1. (Physical-Link-Security Relaxation) $A$ is not allowed to (i) send announcements on behalf of honest neighboring ASes and (ii) intercept and modify announcements exchanged between neighboring honest ASes.

2. (Trusted-Next-Neighbor Relaxation) Whenever experiment $\mathbf{Exp}^{\text{sec-rout-m}}_{\mathcal{I},\mathcal{PV}}(A)$ outputs 2, *i.e.,* $N_\ell \in \mathsf{Honest}$ accepts announcement $(N_{\ell-1}, N_\ell, R = (N_{\ell-1}, \ldots, N_1), P, W, Aux)$, and there exists $1 \leq i \leq \ell - 1$ such that $N_i \in \mathsf{Honest}$ never output $(N_i,\ N_{i+1}, R' = (N_i, \ldots, N_1), P,\ W', Aux')$ for any $W', Aux'$ to $N_{i+1}$, $N_{i+1} \in \mathsf{Honest}$ if $N_i \in \mathsf{nopubk}$.

#### 3.8.2.2 The Relaxed Security Definition

In what follows, we incorporate Relaxations 1 and 2 described above into a new security definition for path-vector protocols where adversary's behavior is restricted according to these relaxations.

The Relaxed Security Definition We relax the security definition from Section 3.5 as follows.

**Definition 1.** *Let* $\mathcal{I} = (\mathbf{G} = (\mathsf{ASes}, \mathsf{link}),\ \mathsf{Prefixes},\ \mathsf{OrforPr},\ \mathsf{relation},\ \mathsf{preferto},\ \mathsf{policy})$ *be a network in* $\mathcal{C}^{\mathcal{PV}}_m$*, for any* $1 \leq m \leq |\mathsf{ASes}|$*, such that* $|\mathsf{ASes}| \geq 2$*, let* $\mathcal{PV} = (\mathsf{Init}, \mathsf{An})$ *be a path-vector protocol and let* $k$ *be the security parameter such that the size of the description of* $\mathcal{I}$ *is polynomial in* $k$*. We define experiment* $\mathbf{Exp}^{\text{r-sec-rout-m}}_{\mathcal{I},\mathcal{PV}}(A)$ *involving adversary* $A$ *to be identical to the experiment* $\mathbf{Exp}^{\text{sec-rout-m}}_{\mathcal{I},\mathcal{PV}}(A)$ *involving an adversary* $A$ *from the definition from Section 3.5 except that Relaxations 1-2 must hold.*

We define $A$'s advantage $\mathbf{Adv}^{\text{r-sec-rout-m-b}}_{\mathcal{I},\mathcal{PV}}(A)$ in this experiment as
$\Pr\left[\mathbf{Exp}^{\text{r-sec-rout-m}}_{\mathcal{I},\mathcal{PV}}(A) = b\right]$, for $b \in \{1, 2, 3\}$. We say that $\mathcal{PV}$ guarantees *relaxed origin authentication*, *route authentication*, and *route validity* with $m$-PD for a class of networks $\mathcal{C}^{\mathcal{PV}}_m$, if for every network $\mathcal{I} \in \mathcal{C}^{\mathcal{PV}}_m$, for every efficient adversary $A$ the

probability that experiment $\mathbf{Exp}^{\text{r-sec-rout-m}}_{\mathcal{I},\mathcal{PV}}(A)$ returns 1, 2 and 3 respectively, while Relaxations 1-2 hold, is negligible in $k$. The *relaxed full security* is defined analogously to security definition in Section 3.5.

#### 3.8.2.3 Secure Constructions

We now slightly modify S-BGP-PD and then show that it meets the above definition.

**Construction 3.8.6.** *Let* $\mathcal{I} = (\mathbf{G} = (\mathsf{ASes}, \mathsf{link}), \mathsf{Prefixes}, \mathsf{OrforPr}, \mathsf{relation}, \mathsf{preferto}, \mathsf{policy})$ *be an interdomain network. We define S-BGP-PD with a restriction (S-BGP-PDR)* $= (\mathsf{Init}, \mathsf{An}')$ *as a path-vector protocol identical to S-BGP-PD* $= (\mathsf{Init}, \mathsf{An})$ *but with the following restrictions in* $\mathsf{An}'$*. When an AS receives an announcement of a route, that AS rejects the announcement if that route contains more than one AS without public keys in a row at any part of that route. Also, an AS without a public key does not propagate a route that was announced by its neighbor who also does not have a public key.*

We define FS-BGP-PD with a restriction (FS-BGP-PDR) analogously. Note that in S-BGP-PDR and FS-BGP-PDR, the last two ASes on a route could be without public keys. This new restriction implicitly requires that ASes reject announcements that are missing a signature for at least one AS in that route who has a public key. Although checking whether an AS has a public key or not may be difficult in practice, this is in fact necessary, otherwise an adversarial AS could simply strip an honest AS's signature and send a bogus route on its behalf.

**Theorem 3.8.7.** *S-BGP-PDR as defined in Construction 3.8.6 guarantees relaxed route authentication with m-PD for* $\mathcal{C}_m^{\text{S-BGP-PDR}}$*, for any* $m \leq |\mathsf{ASes}|$*, if the underlying* $\mathcal{SS}$ *is uf-cma-secure.*

*Proof.* We show that for every adversary $A$ attacking route authentication of S-BGP-PDR, there exist adversaries $B$ and $C$ attacking unforgeability of $\mathcal{SS}$ such that

$$\mathbf{Adv}^{\text{uf-cma}}_{\mathcal{SS}}(B) + \mathbf{Adv}^{\text{uf-cma}}_{\mathcal{SS}}(C) \geq \frac{1}{|\mathsf{ASes}|}\mathbf{Adv}^{\text{r-sec-rout-m-2}}_{\mathcal{I},\text{S-BGP-PDR}}(A),$$

and the resources of each are that of $A$ plus some overhead upper bounded by the size of the network using S-BGP-PDR that $A$ is attacking.

Suppose $\mathcal{C}_m^{\text{S-BGP-PDR}} \neq \emptyset$ and let $A$ be an efficient adversary attacking route authentication of S-BGP-PDR for a network $\mathcal{I} = (\mathbf{G} = (\mathsf{ASes}, \mathsf{link}), \mathsf{Prefixes}, \mathsf{OrforPr}, \mathsf{relation}, \mathsf{preferto}, \mathsf{policy}) \in \mathcal{C}_m^{\text{S-BGP-PDR}}$, $1 \leq m \leq |\mathsf{ASes}|$ and $|\mathsf{ASes}| \geq 2$, whose description is polynomial in $k$.

As a result of $A$'s attack, $N_\ell \in \mathsf{Honest}$ accepts an announcement $(N_{\ell-1}, N_\ell, R = (N_{\ell-1}, \ldots, N_1), P, W, Aux)$, such that $\exists\ 1 \leq i \leq \ell - 1$ so that $N_i \in \mathsf{Honest}$, $N_i$ has never sent announcement $(N_i, N_{i+1}, R' = (N_i, N_{i-1}, \ldots, N_1), P, W', Aux')$, for some $W'$, $Aux'$, and $N_{i+1} \in \mathsf{Honest}$ if $N_i \in \mathsf{nopubk}$, while Relaxations 1-2 hold. We refer to this event by $A$ frames $i$.

Notice that either $N_i$ has a public key or it does not. $N_\ell \in \mathsf{Honest}$ could not have accepted a route announcement with two ASes without a public key in a row, so, by construction of S-BGP-PDR and Relaxation 2, $N_{i+1}$ must have a public key and be honest if $N_i \in \mathsf{nopubk}$ and $i < \ell - 1$. Since $N_i \in \mathsf{Honest}$, $N_i$ would not send an announcement to an AS that is not its neighbor, so there must be a link between $N_i$ and $N_{i+1}$. If $N_{i+1}$ has never accepted the announcement that $N_i$ has never actually sent, it must be that $i < \ell - 1$, since, by definition of the attack, $N_\ell$ did accept $(N_{\ell-1}, N_\ell, R = (N_{\ell-1}, \ldots, N_1), P, W, Aux)$. If $N_{i+1}$ did accept it (which means that it must have received it), $A$ must have either generated that announcement and sent it on behalf of $N_i$ or intercepted and modified some other $N_i$'s announcement. However, this cannot happen as it would violate Relaxation 1, in which case $A$ would not win (note that this also includes the case when $i = \ell - 1$). More concretely, exactly one of the following two conditions must hold when $A$ frames $i$:

(1) $N_i \in \mathsf{nopubk}$, $i < \ell - 1$, and $N_{i+1}$ never accepted announcement $(N_i, N_{i+1}, R' = (N_i, N_{i-1}, \ldots, N_1), P, W', Aux')$, for any $W'$, $Aux'$ or

(2) $N_i \notin \mathsf{nopubk}$.

When condition (1) holds, we construct adversary $B$ attacking unforgeability of $\mathcal{SS}$ as follows. $B$ is given a public key $pk$ and the signing oracle $\mathsf{Sign}(K, \cdot)$ in $\mathbf{Exp}^{\text{uf-cma}}_{\mathcal{SS}}(B)$. After giving $A$ the description of $\mathcal{I}$, who then selects the set $\mathsf{nopubk} \subsetneq \mathsf{ASes}$ of ASes who will not have public keys, $B$ picks an AS at random $N_x \xleftarrow{\$} \mathsf{ASes}$ and then generates public-private key pairs for all ASes not in $\mathsf{nopubk} \cup \{N_x\}$ using $\mathsf{Kg}(1^k)$. $B$ then sets $\mathbf{pk}[x] \leftarrow pk$ and gives $A$ all the public keys. $A$ outputs initial partition $(\mathsf{Honest}, \mathsf{Corrupted})$ of $\mathbf{G}$. If $N_x \in \mathsf{Corrupted}$, then $B$ aborts its attack. Otherwise, $B$ gives $A$ all the secret keys of the corrupted ASes. The rest of the proof for this condition is identical to that of Theorem 3.6.2. Therefore, $\Pr\left[\,\mathbf{Exp}^{\text{uf-cma}}_{\mathcal{SS}}(B) = 1 \mid \text{cond } 1\,\right]$ that $B$ wins by outputting $((N_{i+1}, N_i, \ldots, N_1, P), RA_{R_i^{i+1}})$, when condition (1) holds, over all $i \in \mathsf{ASes}$, is $\frac{1}{|\mathsf{ASes}|}\sum_{i \in \mathsf{ASes}} \Pr\left[\,A \text{ frames } i \mid \text{cond } 2\,\right]$. $B$ is efficient since, to simulate S-BGP-PDR, the number of queries it makes to $\mathsf{Sign}(K, \cdot)$ and their length are upper-bounded by the number of queries that $A$ makes and the size of $\mathcal{I}$ respectively.

When condition (2) is true, we construct adversary $C$ attacking unforgeability of $\mathcal{SS}$ the same way as adversary $B$ when condition (1) is true, only in this case, at the end of $A$'s attack, $C$ would output $((N_{i+2}, N_{i+1}, \ldots, N_1, P), RA_{R_{i+1}})$. Note that since $N_{i+1}$ never accepted $(N_i, N_{i+1}, R' = (N_i, N_{i-1}, \ldots, N_1), P, W', Aux')$, $N_{i+1}$ could not have sent $(N_{i+1}, N_{i+2}, (N_{i+1}, R'),\ P, W'', Aux'')$ to $N_{i+2}$ because $N_{i+1} \in \mathsf{Honest}$ due to Relaxation 2. Therefore, $C$'s output is "new" in the sense that $C$ never queried $((N_{i+2}, N_{i+1}, \ldots, N_1, P)$ to the signing oracle. Thus, $\Pr\left[\,\mathbf{Exp}^{\text{uf-cma}}_{\mathcal{SS}}(C) = 1 \mid \text{cond } 2\,\right]$ is also $\frac{1}{|\mathsf{ASes}|}\sum_{i \in \mathsf{ASes}} \Pr\left[\,A \text{ frames } i \mid \text{cond } 2\,\right]$. $C$ is efficient since, to simulate S-BGP-PDR, the number of queries it makes to $\mathsf{Sign}(K, \cdot)$ and their length are upper-bounded by the number of queries that $A$ makes and the size of $\mathcal{I}$ respectively.

We thus have that

$$
\begin{aligned}
\mathbf{Adv}^{\text{r-sec-rout-m-2}}_{\mathcal{I},\text{S-BGP-PDR}}(A) &= \Pr\left[\,\mathbf{Exp}^{\text{r-sec-rout-m}}_{\mathcal{I},\text{S-BGP-PDR}}(A) = 2\,\right] \\
&= \sum_{j=1}^{2}\left(\sum_{i\in\mathsf{ASes}}\Pr\left[\,A \text{ frames } i \mid \text{cond j}\,\right]\Pr\left[\,\text{cond j}\,\right]\right) \\
&= |\mathsf{ASes}|\Pr\left[\,\mathbf{Exp}^{\text{uf-cma}}_{\mathcal{SS}}(B) = 1 \mid \text{cond 1}\,\right]\Pr\left[\,\text{cond 1}\,\right] \\
&+ |\mathsf{ASes}|\Pr\left[\,\mathbf{Exp}^{\text{uf-cma}}_{\mathcal{SS}}(C) = 1 \mid \text{cond 2}\,\right]\Pr\left[\,\text{cond 2}\,\right] \\
&\leq \sum_{j=1}^{2}|\mathsf{ASes}|\Pr\left[\,\mathbf{Exp}^{\text{uf-cma}}_{\mathcal{SS}}(B) = 1 \mid \text{cond j}\,\right]\Pr\left[\,\text{cond j}\,\right] \\
&+ \sum_{j=1}^{2}|\mathsf{ASes}|\Pr\left[\,\mathbf{Exp}^{\text{uf-cma}}_{\mathcal{SS}}(C) = 1 \mid \text{cond j}\,\right]\Pr\left[\,\text{cond j}\,\right] \\
&= |\mathsf{ASes}|\left(\Pr\left[\,\mathbf{Exp}^{\text{uf-cma}}_{\mathcal{SS}}(B) = 1\,\right] + \Pr\left[\,\mathbf{Exp}^{\text{uf-cma}}_{\mathcal{SS}}(C) = 1\,\right]\right)\ .
\end{aligned}
$$

□

**Corollary 3.8.8.** *FS-BGP-PDR is relaxed fully secure with m-PD for* $\mathcal{C}_m^{FS\text{-}BGP\text{-}PDR}$*, for* $m \leq |\mathsf{ASes}|$*, if the underlying* $\mathcal{SS}$ *and* $\mathcal{CP}$ *are uf-cma-secure and uf-cda-secure respectively.*

*Proof.* The proof follows from Theorems 3.7.2 and 3.8.7 and Remark 3.8.3. □

A significant practical implication of Theorem 3.8.7 and Corollary 3.8.8 is that new ASes who have just joined the Internet but do not have public keys, do not have to get a public key as long as they establish a trust relationship with their neighbors in the sense that for any route announcement that they make, they are sure that their neighbors who have public keys will vouch for them.

The following results emphasize that the restrictions in the relaxed path-vector protocol security definition posed by Relaxations 1-2 and the requirement to ignore routes that have more than one AS without a public key in a row, as is done in S-BGP-PDR and FS-BGP-PDR, are in fact necessary. The latter restriction, in the worst case, could cause some parts of the network to become disconnected as many routes may be ignored.

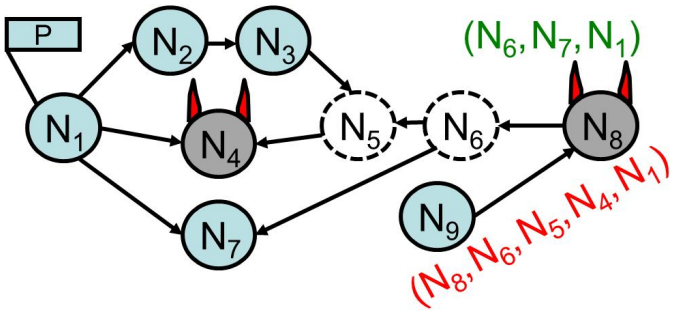


**Figure 7:** $N_5$ and $N_6$ do not have public keys, and the adversary corrupts $N_4$ and $N_8$. In this VRS route authentication attack $N_8$ announces to $N_9$ a valid route to $P$ that $N_5$ never authorized $N_6$ to announce. Note that Relaxations 1-2 are satisfied since $N_6$ is honest and the adversary does not need to intercept and modify communication between honest ASes.

**Theorem 3.8.9.** *For the statements in Theorem 3.8.7 and Corollary 3.8.8 to hold, each relaxation (Physical-Link-Security or Trusted-Next-Neighbor) is necessary given the other one.*

*Proof.* The proof is demonstrated in Figure 6. If the Trusted-Next-Neighbor relaxation does not hold, then the adversary can perform the same attack as in the proof of Theorem 3.8.4. If the Physical-Link-Security relaxation does not hold, then the adversary can do the same by intercepting and modifying the route announcement on a link between $N_5$ and $N_6$, while corrupting no AS. In either case, $\mathbf{Adv}^{\text{sec-rout-m-2}}_{\mathcal{I},\text{S-BGP-PDR}}(A) = 1$ and $A$ is efficient. □

**Theorem 3.8.10.** *Even when the underlying $\mathcal{SS}$ is uf-cma-secure, S-BGP-PD as per Construction 3.8.1 and FS-BGP-PD do not guarantee relaxed route authentication with m-PD for $\mathcal{C}_m^{\text{S-BGP-PDR}}$ and $\mathcal{C}_m^{\text{FS-BGP-PDR}}$ respectively, for any $m \geq 2$.*

*Proof.* This is shown in Figure 7, for $m = 2$. While $N_5$ prefers a customer route through $N_3$, $N_6$ prefers a shorter route through $N_7$. On behalf of $N_8$, the adversary announces a valid route to $N_9$ that goes through $N_4$. $N_9$ accepts this route, because there is no way for $N_9$ to find out that $N_5$ never announced to $N_6$ a route through

$N_4$. This is because neither $N_5$ nor $N_6$ has a public key, although both are honest. Observe that in this case $\mathbf{Adv}^{\text{r-sec-rout-m-2}}_{\mathcal{I},\text{S-BGP-PD}}(A) = 1$ and $A$ is efficient. □

In Section 3.8.4, we show that it is possible to guarantee route authentication even without relying on Relaxation 2 but with a very restricted version of S-BGP-PD.

### 3.8.3 What If There Is No PKI?

We show that if all prefixes and links are certified by a trusted certification authority, even when no AS has a public key, ASes are guaranteed to discover valid routes with authentic origins, and that VRS attacks are the *only* attacks that prevent FS-BGP-PD from guaranteeing route authentication. In light of this result, we then discuss the feasibility of achieving reasonable security without PKI.

**Theorem 3.8.11.** *If the underlying $\mathcal{SS}$ is uf-cma-secure and the underlying $\mathcal{CP}$ is uf-cda-secure, for any $1 \leq m \leq |\mathsf{ASes}|$, if $\mathbf{Exp}^{\text{sec-rout-m}}_{\mathcal{I},FS\text{-}BGP\text{-}PD}(A) = 2$ (see security definition in Section 3.5), then $A$ must have carried out a VRS attack.*

*Proof.* Let us define advantage $\mathbf{Adv}^{\text{r-sec-rout-m-2}}_{\mathcal{I},\text{S-BGP-PD,no-VRS}}(A)$ of any adversary $A$ attacking route authentication of FS-BGP-PD to be the probability that $A$ wins without performing a VRS attack $\Pr\left[\,\mathbf{Exp}^{\text{r-sec-rout-m}}_{\mathcal{I},\text{FS-BGP-PD}}(A) = 2 \mid \text{no VRS}\,\right]$. We show that for every adversary $A$ attacking route authentication of FS-BGP-PD, there exist adversaries $B$ and $C$ attacking unforgeability of $\mathcal{SS}$ and $\mathcal{CP}$ respectively such that

$$\mathbf{Adv}^{\text{uf-cma}}_{\mathcal{SS}}(B) + \frac{1}{|\mathsf{ASes}|}\mathbf{Adv}^{\text{uf-cda}}_{\mathcal{CP}}(C) \geq \frac{1}{|\mathsf{ASes}|}\mathbf{Adv}^{\text{r-sec-rout-m-2}}_{\mathcal{I},\text{S-BGP-PD,no-VRS}}(A),$$

and the resources of each are at most that of $A$ plus some overhead upper bounded by the size network using FS-BGP-PD that $A$ is attacking.

Suppose $\mathcal{C}^{\text{FS-BGP-PD}}_m \neq \emptyset$ and let $A$ be an efficient adversary attacking route authentication of FS-BGP-PD for a network $\mathcal{I} \in \mathcal{C}^{\text{FS-BGP-PD}}_m$ with $m \geq 1$ and $|\mathsf{ASes}| \geq 2$, whose description is polynomial in $k$.

As a result of $A$'s attack, $N_\ell \in \mathsf{Honest}$ accepts an announcement $(N_{\ell-1},\ N_\ell,\ R = (N_{\ell-1}, \ldots, N_1),\ P, W, Aux)$, such that $\exists\ 1 \leq i \leq \ell - 1$ and $N_i \in \mathsf{Honest}$ has never sent announcement $(N_i, N_{i+1}, R' = (N_i, N_{i-1}, \ldots, N_1), P,\ W', Aux')$, for some $W'$, $Aux'$, to $N_{i+1}$. Without loss of generality, let us consider the closest such $N_i$ to the origin $N_1$. The following are all the possible reasons for why $N_i$ would not send that announcement to $N_{i+1}$.

1. $(N_{i+1}, N_i, \ldots, N_1)$ is a valid route to $P$, but $N_i$ has never received announcement $(N_{i-1}, N_i, (N_{i-1}, N_{i-2}, \ldots, N_1),\ P, \ldots)$;

2. $N_i$ has received announcement $(N_{i-1}, N_i, (N_{i-1}, N_{i-2}, \ldots,\ N_1), P, \ldots)$ but rejected it;

3. $N_i$ has received and accepted announcement $(N_{i-1}, N_i, (N_{i-1}, N_{i-2}, \ldots, N_1), P, \ldots)$, but $N_i$ did not announce $R'$ to $N_{i+1}$ because

   (a) $N_{i+1} \notin \mathsf{Neighbors}(N_i)$,

   (b) $N_{i+1} \notin \mathsf{policy}_{N_i}((N_i, N_{i-1}, N_{i-2}, \ldots, N_1), \mathsf{relation}(N_i, N_{i-1}))$,

   (c) $(N_{i-1}, \ldots, N_1)$ is not $N_i$'s preferred route to $P$.

If $N_i$ has a public key, then $N_\ell$ would accept $(N_{\ell-1},\ N_\ell,\ R = (N_{\ell-1}, \ldots, N_1),\ P, W, Aux)$ only after checking the validity of $N_i$'s route attestation of $R'$, so in this case we could construct an adversary $B$ attacking the unforgeability of the underlying $\mathcal{SS}$ the same way as in proof of condition (1) of Theorem 3.8.7. $B$ is given a public key $pk$ and the signing oracle $\mathsf{Sign}(K, \cdot)$ in $\mathbf{Exp}^{\text{uf-cma}}_{\mathcal{SS}}(B)$. After giving $A$ the description of $\mathcal{I}$, who then selects the set $\mathsf{nopubk} \subsetneq \mathsf{ASes}$ of ASes without public keys, $B$ picks an AS at random $N_x \xleftarrow{\$} \mathsf{ASes}$ and then generates public-private key pairs for all ASes not in $\mathsf{nopubk} \cup \{N_x\}$ using $\mathsf{Kg}(1^k)$. $B$ then sets $\mathbf{pk}[x] \leftarrow pk$ and gives $A$ all the public keys. $A$ outputs initial partition $(\mathsf{Honest}, \mathsf{Corrupted})$ of $\mathbf{G}$. If $N_x \in \mathsf{Corrupted}$, then $B$ aborts its attack. Otherwise $B$ gives $A$ all the secret

keys of the corrupted ASes. The rest of the proof for this condition is identical to that of Theorem 3.6.2. Therefore, $\Pr\left[\mathbf{Exp}^{\text{uf-cma}}_{\mathcal{SS}}(B) = 1 \mid N_i \notin \mathsf{nopubk}\right]$ that $B$ wins by outputting $((N_{i+1}, N_i, \ldots, N_1, P),\ RA_{R_i^{i+1}})$ when $N_i$ has a public key, over all $i \in \mathsf{ASes}$, is $\frac{1}{|\mathsf{ASes}|}\sum_{i\in\mathsf{ASes}} \Pr[\, A \text{ frames } i \mid N_i \notin \mathsf{nopubk}\,]$. $B$ is efficient since, to simulate FS-BGP-PD, the number of queries it makes to $\mathsf{Sign}(K, \cdot)$ and their length are upper-bounded by the number of queries that $A$ makes and the size of $\mathcal{I}$ respectively.

Because FS-BGP-PD guarantees origin authentication and route validity (see Theorem 3.7.2 and Remark 3.8.3), if $N_i \in \mathsf{nopubk}$, then $A$ must succeed in a VRS attack only. This is because if reason 2 holds, then either $(N_{i-1}, N_{i-2}, \ldots, N_1)$ is invalid or $\mathsf{OrforPr}(P) \neq N_1$ (recall that we have chosen $N_i$ to be the closest framed AS to the origin). Note that reason 2 also contains less interesting issues such as lack of a route attestation from some intermediate AS $N_j \notin \mathsf{nopubk}$, for $1 \leq j < i$, in the announcement or a bogus route/address attestation that does not verify during S-BGP's attestation verification steps (see Construction 3.4.1). However, if $N_i$ would not accept this announcement due to any of these issues, then neither would $N_\ell$, since both $N_i$ and $N_\ell$ are honest. If either of the reasons 3(a) or 3(b) holds, then $(N_{i+1}, R')$ is invalid. $R$ must be invalid if at least one of its subroutes, in this case $(N_{i+1}, R')$, is invalid. Thus, if any one of reasons 2, 3(a) or 3(b) is true, since $N_\ell \in \mathsf{Honest}$, it could not have accepted announcement $(N_{\ell-1},\ N_\ell,\ R = (N_{\ell-1}, \ldots, N_1),\ P, W, Aux)$, since FS-BGP-PD guarantees origin authentication and route validity. If $N_i \in \mathsf{nopubk}$ and only one of reasons 1 or 3(c) is true, then $(N_{i+1}, R')$ is a valid route to $P$ and $A$ succeeds in a VRS attack. Note that reason 3(c) also captures the scenario in which $R$ contains $N_i$, in which case $N_i$ would ignore $R$ as a loop-preventative measure.

Thus, if $N_i \in \mathsf{nopubk}$ and $A$ does not succeed in a VRS attack, then we can construct adversary $C$ attacking the unforgeability of the underlying $\mathcal{CP}$ as follows. $C$ is given CA's public key $pk_{\mathsf{CA}}$ in $\mathbf{Exp}^{\text{uf-cda}}_{\mathcal{CP}}(B)$. $C$ first gives $A$ the description of $\mathcal{I}$, who then selects the set $\mathsf{nopubk} \subseteq \mathsf{ASes}$ of ASes who will not have public keys. $C$

then generates public-private key pairs for all ASes not in $\mathsf{nopubk}$ using $\mathsf{Kg}(1^k)$. $C$ then gives $A$ all the public keys, including that of the CA. $A$ outputs initial partition $(\mathsf{Honest}, \mathsf{Corrupted})$ of $\mathbf{G}$. $C$ gives $A$ all the secret keys of the corrupted ASes. $A$ starts the execution of FS-BGP-PD on behalf of all ASes in $\mathsf{Corrupted}$ together with $C$ who executes FS-BGP-PD on behalf of all ASes in $\mathsf{Honest}$ and the CA. $C$ follows FS-BGP-PD legitimately, whereas $A$ is allowed to act arbitrarily. $C$ stores all the communication and provides it to $A$.

For each AS $N_i$ and prefix $P$, such that $N_i$ owns $P$ ($C$ can check this with $\mathsf{OrforPr}$), where either $N_i \in \mathsf{Honest}$ or $N_i \in \mathsf{Corrupted}$ and $A$ has requested address attestation $AA^P_{N_i}$ of $P$ on behalf of $N_i$, $C$ sequentially interacts with the CA via ($\mathsf{CA}(K_{\mathsf{CA}}, N_i, P)$, $B(pk_{\mathsf{CA}}, N_i, P)$) to get $(N_i, P, AA^P_i)$. Similarly, for each AS $N_i$ and its neighbor $N_j$ ($C$ can check this with $\mathsf{link}$), where either $N_i \in \mathsf{Honest}$ or $N_i \in \mathsf{Corrupted}$ and $A$ has requested link attestation $LA^P_{N_iN_j}$ on behalf of $N_i$, $C$ sequentially interacts with the CA via ($\mathsf{CA}(K_{\mathsf{CA}}, N_i, ((\min(N_i, N_j), \max(N_i, N_j)), \mathsf{relation}(\min(N_i, N_j), \max(N_i, N_j))))$, $B(pk_{\mathsf{CA}}, N_i, ((\min(N_i, N_j), \max(N_i, N_j)),$ $\mathsf{relation}(\min(N_i, N_j), \max(N_i, N_j)))))$ to get $(N_j, ((\min(N_j, N_i), \max(N_j, N_i)),$ $\mathsf{relation}(\min(N_j, N_i), \max(N_j, N_i))), LA_{N_iN_j})$. $C$ stores all address and link attestations. This information together with all honest ASes' secret keys, allows $C$ to follow the computations according to the interactive algorithm $\mathsf{An}$. Observe that $C$'s simulation for $A$ is perfect.

When $A$ outputs $(N_{\ell-1}, N_\ell, R = (N_{\ell-1}, \ldots, N_1), P, W, Aux)$, such that $N_\ell \in \mathsf{Honest}$ accepts it and $\exists\ 1 \leq i \leq \ell - 1$ an $N_i \in \mathsf{Honest}$ has never sent announcement $(N_i, N_{i+1}, R' = (N_i, N_{i-1}, \ldots, N_1), P, W', Aux')$, for some $W'$, $Aux'$, $C$ proceeds as follows. If $R'$ is invalid, then there must be $1 < j < i$ such that either $\mathsf{link}(N_j, N_{j+1}) = 0$, in which case $C$ outputs $(N_j, (N_j, N_{j+1}, \text{rel}, LA_{N_jN_{j+1}})$, where rel is a fake relationship between $N_j$ and $N_{j+1}$ since they are not neighbors but is presented in $Aux'$, or $\mathsf{link}(N_j, N_{j+1}) = 1$ but $N_{j+1} \notin \mathsf{policy}_{N_j}((N_j, N_{j-1}, \ldots, N_1), \mathsf{relation}(N_j, N_{j-1}))$ in

which case $C$ outputs $(N_{j-1}, (N_{j-1}, N_j, \text{rel}, LA_{N_{j-1}N_j}))$, where rel $\neq \mathsf{relation}(N_{j-1}, N_j)$ but is presented in $Aux'$. Note that these are the only reasons why $R'$ would not be valid. This is because $\mathsf{policy}$ is publicly available in our model, so if $N_i$ would not accept the announcement with $R'$, then neither would $N_\ell$ accept the announcement with $R$ for the same reason. Otherwise, if $N_1$ does not own $P$, then $C$ outputs $(N_1, P, AA_1^P)$, where $AA_1^P$ must be the last entry of $Aux'$. Otherwise, if reason 3(a) is true, then $C$ outputs $(N_i, (N_i, N_{i+1}, \text{rel}, LA_{N_iN_{i+1}})$, where rel is a fake relationship between $N_i$ and $N_{i+1}$ but is presented in $Aux'$. Otherwise, if reason 3(b) is true, then $C$ outputs $(N_{i-1}, (N_{i-1}, N_i, \text{rel}, LA_{N_{i-1}N_i}))$ where rel $\neq \mathsf{relation}(N_{i-1}, N_i)$ but is presented in $Aux'$.

Since reasons 2-3(b) above cover all possible non-VRS-attack events that could occur when $N_i \in \mathsf{nopubk}$, $C$'s probability of breaking uf-cda security of $\mathcal{CP}$ is the same as that of $A$ breaking route authenticity of $\mathcal{PV}$:

$$\Pr\left[\,\mathbf{Exp}^{\text{uf-cda}}_{\mathcal{CP}}(C) = 1\,\right] = \sum_{i \in \mathsf{ASes}} \Pr\left[\,A \text{ frames } i \mid N_i \in \mathsf{nopubk}\,\right].$$

Note that $C$ is as efficient as $A$.

Thus we have that if $A$ does not succeed in a VRS attack, then

$$\begin{aligned}
&\mathbf{Adv}^{\text{r-sec-rout-m-2}}_{\mathcal{I},\text{S-BGP-PD,no-VRS}}(A) \\
=\;& \textstyle\sum_{i \in \mathsf{ASes}} \Pr\left[\,A \text{ frames } i \mid N_i \notin \mathsf{nopubk},\ \text{no VRS}\,\right]\Pr\left[\,N_i \notin \mathsf{nopubk},\ \text{no VRS}\,\right] \\
+\;& \textstyle\sum_{i \in \mathsf{ASes}} \Pr\left[\,A \text{ frames } i \mid N_i \in \mathsf{nopubk},\ \text{no VRS}\,\right]\Pr\left[\,N_i \in \mathsf{nopubk},\ \text{no VRS}\,\right] \\
=\;& |\mathsf{ASes}|\Pr\left[\,N_i \notin \mathsf{nopubk},\ \text{no VRS} \;\middle|\; \mathbf{Exp}^{\text{uf-cma}}_{\mathcal{SS}}(B) = 1\,\right]\Pr\left[\,N_i \notin \mathsf{nopubk},\ \text{no VRS}\,\right] \\
+\;& \Pr\left[\,\mathbf{Exp}^{\text{uf-cda}}_{\mathcal{CP}}(C) = 1 \mid N_i \in \mathsf{nopubk},\ \text{no VRS}\,\right]\Pr\left[\,N_i \in \mathsf{nopubk},\ \text{no VRS}\,\right] \\
\leq\;& |\mathsf{ASes}|\Pr\left[\,\mathbf{Exp}^{\text{uf-cma}}_{\mathcal{SS}}(B) = 1\,\right] + \Pr\left[\,\mathbf{Exp}^{\text{uf-cda}}_{\mathcal{CP}}(C) = 1\,\right].
\end{aligned}$$

□

### 3.8.4 S-BGP Security without Relaxation 2 in Partial PKI Deployment

In this section we show that if we do not rely on security Relaxation 2, it is still possible to guarantee route authentication but with a very restricted version of S-BGP-PD, where only the last two ASes on any route are allowed not to have public keys.

**Construction 3.8.12.** *Let* $\mathcal{I} = (\mathbf{G} = (\mathsf{ASes}, \mathsf{link})$, $\mathsf{Prefixes}$, $\mathsf{OrforPr}$, $\mathsf{relation}$, $\mathsf{preferto}$, $\mathsf{policy})$ *be an interdomain network and* $k$ *a security parameter. We define S-BGP-PD with an extra restriction (S-BGP-PDxR)* $= (\mathsf{Init}, \mathsf{An}')$ *as a path-vector protocol identical to S-BGP-PD* $= (\mathsf{Init}, \mathsf{An})$ *but with the following restrictions in* $\mathsf{An}'$*. When an AS receives an announcement of a route from a neighbor, that AS rejects that announcement if that route contains at least one AS without a public key other than the neighbor sending the announcement. An AS does not announce a route if it contains at least one AS without a public key other than itself.*

We define FS-BGP-PD with an extra restriction (FS-BGP-PDxR) analogously.

**Theorem 3.8.13.** *S-BGP-PDxR guarantees route authentication with* $m$*-PD for networks in* $\mathcal{C}_m^{S\text{-}BGP\text{-}PDxR}$*, for* $m \geq 1$*, if the underlying* $\mathcal{SS}$ *is uf-cma-secure and the Physical-Link-Security Relaxation holds (see Security Relaxation 1 in Section 3.8).*

*Proof.* This theorem trivially holds if $\mathcal{C}_m^{\text{S-BGP-PDxR}} = \emptyset$. Otherwise, let $A$ be an efficient adversary attacking route authentication of S-BGP-PDxR for a network $\mathcal{I} \in \mathcal{C}_m^{\text{S-BGP-PDxR}}$ with $m \geq 1$ and $|\mathsf{ASes}| \geq 2$, whose description is polynomial in $k$. As a result of $A$'s attack, $N_\ell \in \mathsf{Honest}$ accepts an announcement $(N_{\ell-1}, N_\ell, R = (N_{\ell-1}, \ldots, N_1), P, W, Aux)$, such that $\exists\ 1 \leq i \leq \ell - 1$ so that $N_i \in \mathsf{Honest}$ has never sent announcement $(N_i, N_{i+1}, R' = (N_i, N_{i-1}, \ldots, N_1), P, W', Aux')$, for some $W'$, $Aux'$, while Relaxation 1 holds.

Either $N_i$ has a public key or it does not. If $N_i \in \mathsf{nopubk}$ and $i = \ell - 1$, then $A$ must have either generated that announcement and sent it on behalf of $N_i$ or intercepted

and modified some other $N_i$'s announcement. This cannot happen as it would violate the Physical-Link-Security Relaxation (see Security Relaxation 1 in Section 3.8), in which case $A$ would not win. If $N_i \in \mathsf{nopubk}$ and $i < \ell - 1$, then $R'$ must contain at least one AS $N_{i+1}$ between $N_i$ and $N_\ell$. However, since $N_\ell \in \mathsf{Honest}$, by construction of S-BGP-PDxR, it could not have accepted a route announcement with at least one AS without a public key other than $N_{\ell-1}$, so we exclude this case from the proof. The only remaining option is that $N_i \notin \mathsf{nopubk}$. The rest of the proof is identical to that of condition (2) in Theorem 3.8.7. □

**Corollary 3.8.14.** *FS-BGP-PDxR is guarantees origin authentication as well as route authentication and validity m-PD for* $\mathcal{C}_m^{FS\text{-}BGP\text{-}PDxR}$*, for* $m \leq |\mathsf{ASes}|$*, if the underlying* $\mathcal{SS}$ *and* $\mathcal{CP}$ *are uf-cma-secure and uf-cda-secure respectively and the Physical-Link-Security Relaxation holds.*

*Proof.* The proof follows from Theorems 3.7.2 and 3.8.7 and Remark 3.8.3. □

These results show that guarantees analogous to full security can be provided as long as the adversary is not capable of controlling communication between honest parties and only the last couple of ASes on routes are allowed not to have public keys. As was already pointed out in [41], these are the smaller networks that are likely to be at the edge of the Internet, *i.e.,* stub networks. Stub networks do not have any customers of their own, *e.g.,* small university and corporate networks. Results in this section are significant because stub networks make up over 85% of the Internet [30].

### 3.8.5 Discussion of Practical Implications in Partial PKI Deployment

Ultimately, our partial PKI deployment results with S-BGP-PDR and S-BGP-PDxR (together with FS-BGP-PDR and FS-BGP-PDxR) show that it is possible to achieve well-defined provable security guarantees when some ASes, *e.g.,* stubs, do not possess public keys. As we will discuss in Chapter 5, such approach is similar to the *simplex S*BGP* suggested in [68, 40]. Although in simplex S*BGP stub networks do not

verify routing signatures in routing announcements, they are allowed to delegate the operation of signing of their announcements to their ISP's. If stubs were also to delegate the verification of routing announcements to their ISP's, then simplex S*BGP would reduce the cost of deploying S-BGP at over 85% of ASes on the Internet while still providing meaningful security guarantees.

On the other hand, the goal of path-vector protocols is for ASes to learn of routes in the network to all prefixes, so the importance of Theorem 3.8.11 is that FS-BGP-PD guarantees that ASes learn of valid routes with authentic origins and that, even without PKI, the worst thing that can happen compared to when FS-BGP is deployed, is that due to a VRS attack, at least one honest AS $N_\ell$ accepts at least one route $R = (N_{\ell-1}, \ldots, N_1)$ to some prefix $P$, such that for at least one honest intermediate AS $N_i$ in $R$, subroute $(N_{i-1}, \ldots, N_1)$ is not $N_i$'s the most preferred route to $P$, which would mean that the protocol does not converge due to a subroute consistency violation. Although requiring link-attestations diminishes the practical gains of having no PKI, having no PKI is still very practical and facilitates gradual, Internet-wide deployment of FS-BGP-PD as it relieves ASes of storing public keys of all other ASes and generating signatures for their every announcement. It also reduces communication overhead by getting rid of ASes' signatures in ASes' route announcements.

With respect to adversarial control of the flow of traffic on the Internet, Theorem 3.8.11 is a major milestone in understanding the security and efficiency tradeoffs that can be achieved in full versus no PKI deployment. Although with a VRS attack an adversary could cause an honest AS to send traffic along an unintended route without that AS's knowledge, the adversary could do the same without a VRS attack by simply diverting traffic to an unintended route of its choosing without that source's knowledge. The latter is an issue of data-plane accountability, and if the Internet does not deploy a provably secure accountability protocol, *e.g.,* [15, 48], then FS-BGP-PD

with no PKI is just as good as with fully deployed PKI with respect to such an adversary. On the other hand, the only provably secure accountability protocols that are known to date require ASes to deploy a PKI or have shared keys, so having no PKI for FS-BGP-PD would yield no practical gains if the Internet does deploy a provably secure accountability protocol. Thus, in the beginning stages of partial deployment of secure path-vector protocols, when there is no PKI, it may be more beneficial to deploy link certificates rather than have some ASes possess public keys but deploy no link certificates at all.

BGPSEC is being currently standardized to run on top of the RPKI, which is being deployed. The keys provided by the RPKI would be used in BGPSEC, and the results in this section are relevant to settings when either RPKI is partially deployed (*i.e.,* not every AS gets a certificate for a prefix and a key) or RPKI is fully deployed but some ASes choose not to use their private keys to generate route attestations. Also, if the Internet were to be divided into some ASes that use S-BGP while the rest stick to BGP (partial deployment scenario we consider in Chapters 4 and 5, then our results with respect to S-BGP's and FS-BGP's security guarantees would apply only to each of the connected subgraphs of the Internet that choose to use S-BGP separately. To maintain overall Internet connectivity, ASes running S-BGP would have to use BGP when communicating with ASes that do not use S-BGP.

If origin authentication could be guaranteed with RPKI, then it is plausible that a similar system could be used to establish link certificates as is done in FS-BGP. We note, however, that if an adversary is allowed to corrupt various ASes in the RPKI and/or an analogous hierarchy for certifying communication links (*i.e.,* entities that generate and/or certify keys, AS numbers, and communication links may be rogue), as we suggested in Section 3.5, to have well-defined, provable security guarantees in such scenarios, more sophisticated models and protocols would be needed to address rogue key and certificate attacks.

## *3.9 SoBGP Definition and Security Analysis*

In SoBGP [104], Origin Authorization Certificates are used to bind prefixes to certain ASes (just like address attestations in S-BGP) while AS Policy Certificates are used to allow ASes to learn of links and policies of remote ASes. Although similar to link attestations, AS Policy Certificates are not generated for communication links by a third trusted party but by ASes (possibly corrupted) themselves who then disseminate their policies to their neighbors. Note that as with link attestations in FS-BGP, such Policy Certificates in SoBGP could in principle be downloaded and verified off-line. There is no equivalent of S-BGP Route Attestations in SoBGP, and this together with AS Policy Certificates are the most essential differences between SoBGP and S-BGP. In this section we formally define SoBGP and show that, although it guarantees origin authentication, it does not guarantee route authentication and route validity.

**Construction 3.9.1.** *Let* $\mathcal{I} = (\mathbf{G} = (\mathsf{ASes}, \mathsf{link})$, $\mathsf{Prefixes}$, $\mathsf{OrforPr}$, $\mathsf{relation}$, $\mathsf{preferto}$, $\mathsf{policy})$ *be an interdomain network, let* $\mathcal{SS} = (\mathsf{Kg}, \mathsf{Sign}, \mathsf{Ver})$ *be a signature scheme with* $\mathsf{MsgSp} = \{0,1\}^*$*, and let* $\mathcal{CP}_s = (\mathsf{Kg}_{\mathsf{CA}}, (\mathsf{CA}, \mathsf{U}), \mathsf{Vercert})$ *be the corresponding certification protocol as per Construction 3.2.1. In SoBGP* $= (\mathsf{Init}, \mathsf{An})$*, as part of* $\mathsf{Init}$ *the CA runs* $\mathsf{Kg}_{\mathsf{CA}}(1^k)$ *to generate* $(pk_{\mathsf{CA}}, K_{\mathsf{CA}})$ *and each AS runs* $\mathsf{Kg}(1^k)$ *to generate* $(pk, K)$*.* $\mathsf{An}$ *is defined as follows.*

*If AS* $N_j$*'s input* $\mathbf{P}_{N_j}$ *is nonempty (*i.e., $N_j \in \mathsf{Origins}$*), then for every prefix* $P \in \mathbf{P}_{N_j}$*,* $N_j$ *does the following (note that this is just like in S-BGP as per Construction 3.4.1):*

- *CA and* $N_j$ *interact according to* $(\mathsf{CA}, \mathsf{U})$*,* $N_j$ *being* $\mathsf{U}$*. The input to* $\mathsf{U}$ *is* $(pk_{\mathsf{CA}}, N_j, P)$*, the input to* $\mathsf{CA}$ *is* $(K_{\mathsf{CA}}, N_j, P)$ *and the outputs of both parties are* $(N_j, P, \mathsf{cert})$*. Origin Authorization* $OA^P_{N_j} \equiv \mathsf{cert}$ *is* $N_j$*'s certificate of ownership of* $P$*.*

- *Next, every* $N \in \mathsf{ASes}$ *runs* $\mathsf{Sign}(K_N(N, N', \mathsf{relation}(N, N')) = \sigma$*, for every* $N' \in$

$\mathsf{Neighbors}(N)$, *to produce Policy Certificates* $PC_{NN'} \equiv (N, N', \mathsf{relation}(N, N'), \sigma)$ *that* $N$ *makes publicly available to all other ASes in* ASes. *Note that in our model of an interdomain network* policy *is publicly available, so AS Policy Certificates would need to be more sophisticated in scenarios where this is not true. In SoBGP, ASes may also be able to specify which other ASes their neighbors are allowed to export their routes, but we omit this detail as it is not essential to our security analysis of this protocol.*

*For every new route announcement* $(N_{j-1}, N_j, R = (N_{j-1}, \ldots, N_1), P, W, Aux = OA^P_{N_1})$ *that* $N_j$ *receives from some neighbor* $N_{j-1}$, $N_j$ *first performs origin authorization and policy certificate verification steps as follows.* $N_j$ *runs* $\mathsf{Vercert}\ (pk_{\mathsf{CA}}, N_1, P, OA^P_{N_1})$ *and outputs* $\perp$ *if the output of this computation is* 0. *Otherwise,* $N_j$ *runs* $\mathsf{Ver}(pk_{N_i}, PC_{N_iN_{i+1}})$, *for every* $1 \leq i \leq j-1$, *and outputs* $\perp$ *if at least one such computation outputs* 0 *or if there is at least one* $N_i$, *for* $1 \leq i \leq j-1$, *such that* $N_{i+1} \notin \mathsf{policy}_{N_i}((N_i, \ldots, N_1), \mathsf{relation}(N_i, N_{i-1}))$. *If none of the verification steps above results in* $\perp$, *then* $N_j$ *performs the same operations as* $N_j$ *would do in BGP upon receipt of* $(N_{j-1}, N_j, R, P, W, \varepsilon)$, *as per rules (1)-(3) specified in Section 3.4.1. Then, for every announcement* $(N_j, N_{j+1}, R', P, W', \varepsilon)$ *that* $N_j$ *would send to* $N_{j+1}$ *in BGP,* $N_j$ *sends* $(N_j, N_{j+1}, R', P, W', Aux)$ *to* $N_{j+1}$ *instead, where* $R' = (N_j, R)$.

If the underlying signature scheme $\mathcal{SS}$ is correct, the execution of SoBGP is the same as that of BGP in terms of how ASes update their routing tables and how they decide which routes to announce to their neighbors. Therefore, SoBGP is correct for the same classes of networks as BGP if the underlying signature scheme $\mathcal{SS}$ used to generate origin authorizations and policy certificates.

**Theorem 3.9.2.** *SoBGP per Construction 3.9.1 guarantees* origin authentication *for* $\mathcal{C}_0^{SoBGP}$ *if the underlying* $\mathcal{SS}$ *is uf-cma-secure.*

*Proof.* The proof follows from Theorem 3.2.2 and Lemma 3 stated below. □

**Lemma 3.** *Construction 3.9.1 guarantees* origin authentication *for* $\mathcal{C}_0^{SoBGP}$ *if the underlying* $\mathcal{C}P$ *is uf-cda-secure.*

*Proof.* The proof is identical to that of Theorem 1 because the mechanism for achieving origin authentication in SoBGP with Origin Authorization Certificates is essentially identical to that in S-BGP with Address Attestations. □

**Theorem 3.9.3.** *SoBGP per Construction 3.9.1 does not guarantee* route authentication *for* $\mathcal{C}_0^{SoBGP}$.

*Proof.* The proof is essentially the same as that of Theorem 3.8.4, where the adversary causes subroute inconsistency with a VRS attack. This is because AS Policy Certificates in SoBGP (certifying physical communication links) do not guarantee route authentication by themselves for the same reason link attestations do not guarantee route authentication in S-BGP-PD by themselves when not every AS has a public key. □

The following result points out the fact that link certification is not enough to guarantee route validity in general due to collusion when the certification is done by the ASes themselves.

**Theorem 3.9.4.** *SoBGP as defined in Construction 3.9.1 does not guarantee route validity for* $\mathcal{C}_0^{SoBGP}$.

*Proof.* (Sketch) The proof is very similar to that of 3.6.3. Here we present a specific example of an attack for the network in $\mathcal{C}_0^{\text{SoBGP}}$ depicted in Figure 3(a), where there is at least one valid route of length greater than three ASes. The adversary corrupts two non-neighboring ASes $N_3$ and $N_8$ that are on a valid route and creates a policy certificate for the fake link between them. There is nothing in SoBGP to prevent this from happening since both ASes are corrupted. The adversary then obtains the corresponding route announcement from the corrupted AS closer to the origin $N_3$,

and then propagates the corresponding route announcement on behalf of the other corrupted AS $N_8$ to $N_9$. The route in the latter announcement is infeasible because the corrupted ASes are not actually neighbors, but there is no way to verify this fact by the honest AS $N_9$ that is down-stream from the corrupted AS farther away from the origin $N_8$. This is because the policy certificate of the fake link passes verification since it was created in a legitimate manner from the perspective of SoBGP. □

The results in this section demonstrate that SoBGP is not as good of a candidate for securing the Internet's routing infrastructure as S-BGP, because SoBGP meets only a single security goal, namely origin authentication, with respect to our security model, while S-BGP meets two security goals, namely origin and route authentication. In the next section, however, we will discuss how SoBGP in fact can guarantee route validity, albeit not route authentication, with respect to a weaker security model.

## *3.10 Alternative Solutions to Route Validity Attacks*

Recall that in a route validity attack the adversary succeeds when a route that is either infeasible and/or invalid is accepted by a non-corrupted AS, and in the previous section we have shown that SoBGP does not guarantee route validity. Also, recall that when a route $R = (N_\ell, \ldots, N_1)$ is feasible, $R$ is invalid if $N_{i+1} \notin \mathsf{policy}(N_i, (N_i, \ldots, N_1))$ for at least one $1 \leq i \leq \ell - 1$. In this section we investigate different solutions to route validity attacks, including SoBGP, with respect to weaker attackers and networks that satisfy certain routing conditions. Note that although we show in Chapter 4 that convergence under these conditions can be guaranteed, we do not concern ourselves specifically with convergence in this investigation. Rather, we are interested in exploiting these conditions to facilitate development of more efficient solutions to route validity attacks than link attestations used in FS-BGP.

### 3.10.1 Commercial Routing Conditions

BGP convergence has been shown to hold for various classes of networks [38, 47, 41], with commercial routing (*i.e.,* network models that capture commercial routing policies of ASes on the Internet of today), and we focus specifically on the following variant of these conditions, which will refer to by *CC*.

1. For a particular network $\mathcal{I}$, for every prefix in $\mathsf{Prefixes}$, every AS $N \in \mathsf{ASes}$ prefers customer routes to peer and provider routes to that prefix, where $R = (N', \ldots)$ is called a customer, peer, or provider route if $N'$ is $N$'s customer, peer or provider respectively. In case of ties, shorter routes are preferred over longer routes, and further ties are resolved via a consistent tie-breaking rule.

2. For an origin AS $N \in \mathsf{ASes}$, $N$'s every neighbor is in $\mathsf{policy}(N, (N))$. For any route announcement advertising some route $R = (N', \ldots)$ that any $N \in \mathsf{ASes}$ receives from its any neighbor $N'$, if $\mathsf{relation}(N', N) \neq (\mathsf{cust}, \mathsf{prov})$, then $N$'s neighbor is in $\mathsf{policy}(N, (N, R))$ if and only if that neighbor is $N$'s customer. Otherwise, $N$'s every neighbor is in $\mathsf{policy}(N, (N, R))$.

For convergence to be guaranteed, it is not required for every origin to advertise its prefixes to its every neighbor and for every AS to advertise its every non-customer and customer routes to its every customer and neighbor respectively [38]. In practice, some origins may choose to not advertise some of their prefixes to some of their neighbors for traffic-load balancing. In this case, origins would notify their neighbors about their prefixes that those neighbors are not going to have direct access to. However, for simplicity of presentation, throughout our analysis we only focus on networks that satisfy *CC*. We comment on how our security results apply to SoBGP at the end of this section.

### 3.10.2 S-BGP-XB

According to *CC*, a feasible route $R = (N_\ell, \ldots, N_1)$ is invalid if for at least one $1 \leq i \leq \ell - 1$, at least one of the following is true about $R$ [99]: (1) the valley attack— $N_i$ exports a provider route to another provider, i.e. $\mathsf{relation}(N_i, N_{i-1}) = \mathsf{relation}(N_i, N_{i+1}) = (\mathsf{cust}, \mathsf{prov})$; (2) the step attack—$N_i$ exports a provider or a peer route to a peer or to a provider respectively, i.e. $\mathsf{relation}(N_i, N_{i-1}) = (\mathsf{cust}, \mathsf{prov})$ and $\mathsf{relation}(N_i, N_{i+1}) = (\mathsf{peer}, \mathsf{peer})$ or $\mathsf{relation}(N_i, N_{i-1}) = (\mathsf{peer}, \mathsf{peer})$ and $\mathsf{relation}(N_i, N_{i+1}) = (\mathsf{cust}, \mathsf{prov})$; and (3) the mirror attack—$N_i$ exports a peer route to a peer, i.e. $\mathsf{relation}(N_i, N_{i-1}) = \mathsf{relation}(N_i, N_{i+1}) = (\mathsf{peer}, \mathsf{peer})$.

The solution to such policy violation attacks proposed in [99], let us call it S-BGP-XB, is to augment S-BGP by a 1-bit flag. In summary, for any announcement, an AS sets this flag to 0 when advertising a route to its provider, and it sets this flag to 1 otherwise. Every AS rejects an announcement if it comes from a customer and this flag is set to 1 or if it comes from a non-customer and this flag was set to 1 by an AS other than the sender of this announcement. Let us now describe S-BGP-XB more concretely.

**Construction 3.10.1.** *Let* $\mathcal{I} = (\mathbf{G} = (\mathsf{ASes}, \mathsf{link}), \mathsf{Prefixes}, \mathsf{OrforPr}, \mathsf{relation}, \mathsf{preferto}, \mathsf{policy})$ *be an interdomain network, let* $\mathcal{SS} = (\mathsf{Kg}, \mathsf{Sign}, \mathsf{Ver})$ *be a signature scheme, and let* $\mathcal{CP}_s = (\mathsf{Kg}_{\mathsf{CA}}, (\mathsf{CA}, \mathsf{U}), \mathsf{Vercert})$ *be the corresponding certification protocol as per Construction 3.2.1. In S-BGP* $= (\mathsf{Init}, \mathsf{An})$, $\mathsf{Init}$ *is defined exactly as in S-BGP.*

$\mathsf{An}$ *is defined exactly as in S-BGP with the following additional steps. For every origin AS* $N_j$*, for every prefix* $P \in \mathsf{Prefixes}$ *such that* $\mathsf{OrforPr}(P) = N_j$*, for every* $N_i \in \mathsf{policy}(N_j, (N_j))$*, to generate a route attestation* $RA^i_{R_j}$*,* $N_j$ *runs* $\mathsf{Sign}(K_{N_j}, (N_i, N_j, P, X))$*, where* $X = \varepsilon$ *if* $\mathsf{relation}(N_j, N_i) = (\mathsf{cust}, \mathsf{prov})$ *and* $X = 1$ *otherwise. To* $N_i$*,* $N_j$ *sends announcement* $(N_j, N_i, R = (N_j), P, 0, Aux = ((|X|, |X|), (RA^i_{R_j}, X), AA^P_{N_j}))$.

*For every new route announcement* $(N_{j-1}, N_j, R = (N_{j-1}, \ldots, N_1), P, W, Aux =$

$((X,k), RA^{j}_{\tilde{R}_{j-1}}, \ldots\ RA^{2}_{\tilde{R}_1}, AA^{\tilde{P}}_{N_1}))$ *that* $N_j$ *receives from* $N_{j-1}$, $N_j$ *performs address attestation verification the same way as in S-BGP. When* $N_j$ *performs route attestation verification,* $N_j$ *first checks if* $X = 0$ *or if* $X = 1$. *In the former case, route attestation verification is then executed exactly as it would be done in S-BGP. In the latter case,* $N_j$ *outputs* $\perp$ *if* $k < 1$ *or* $k > j - 1$. *Otherwise,* $N_j$ *runs* $\mathsf{Ver}(pk_{N_i}, (N_{i+1}, \ldots N_1, P, X_i), RA^{i+1}_{\tilde{R}_i})$ *for every* $1 \leq i \leq j-1$ *and outputs* $\perp$ *if at least one such computation outputs* 0, *where* $X_i = 1$ *if* $k = i$ *and* $X_i = \varepsilon$ *otherwise. Otherwise,* $N_j$ *outputs* $\perp$ *if* $\mathsf{relation}(N_{j-1}, N_j) = (\mathsf{cust}, \mathsf{prov})$ *or if* $\mathsf{relation}(N_{j-1}, N_j) = (\mathsf{peer}, \mathsf{peer})$ *and* $k < j - 1$.

*If none of the verification steps results in* $\perp$, *then* $N_j$ *performs the same decision process as* $N_j$ *would do in S-BGP. For every announcement* $(N_j, N_\kappa, R', P, W', Aux')$ *that* $N_j$ *would send to* $N_\kappa$ *as a result of this decision process in S-BGP,* $N_j$ *now sends* $(N_j, N_\kappa, R', P, W', Aux'')$ *to* $N_\kappa$ *instead, where* $R' = (N_j, R)$, $Aux'' = ((X', k'), (RA^{\kappa}_{R'_j}, RA^{j}_{\tilde{R}_{j-1}}, \ldots\ RA^{2}_{\tilde{R}_1}, AA^{\tilde{P}}_{N_1}))$, $RA^{\kappa}_{R'_j} = \mathsf{Sign}(K_{N_j}, (N_\kappa, N_j, \ldots, N_1, P, X'))$, $X' = X$ *and* $k' = k$ *if* $X = 1$, $X' = 1$ *and* $k' = j$ *if* $\mathsf{relation}(N_j, N_k) \neq (\mathsf{cust}, \mathsf{prov})$, *and* $X' = \varepsilon$ *and* $k' = 0$ *otherwise.*

Although in [99] the use of back-up routes in S-BGP-XB is required, for simplicity of presentation, we exclude back-up routes from the description of S-BGP-XB here. We only work with networks which have at least one valid route between any pair of ASes in that network. Let $\mathcal{C}$ be the class of networks that satisfy *CC*.

Note that S-BGP-XB is not fully secure for every network in $\mathcal{C}$ because it does not guarantee route feasibility. However, we can show this to be true even for networks for which route infeasibility is never an issue, e.g. for complete graphs. Let $\mathcal{C}_c \subset \mathcal{C}$ be the class of all networks in $\mathcal{C}$ that form complete graphs.

**Theorem 3.10.2.** *S-BGP-XB does not guarantee route validity for every network in* $\mathcal{C}_c$.

*Proof.* Consider network $\mathcal{I} = (\mathbf{G} = (\mathsf{ASes}, \mathsf{link}), \{P\}, \mathsf{OrforPr}, \mathsf{relation}, \mathsf{preferto}, \mathsf{policy})$, where $\mathsf{ASes} = \{N_1, N_2, N_3, N_4\}$, $\mathsf{link}(N_i, N_j) = 1$ for all $N_i, N_j \in \mathsf{ASes}$ such that $i \neq j$, $\mathsf{OrforPr}(P) = N_1$, $\mathsf{relation}(N_2, N_3) \in \{(\mathsf{prov}, \mathsf{cust}), (\mathsf{peer}, \mathsf{peer})\}$, $\mathsf{relation}(N_1, N_4) = \mathsf{relation}(N_1, N_3) = \mathsf{relation}(N_2, N_1) = \mathsf{relation}(N_2, N_4) = \mathsf{relation}(N_4, N_3) = (\mathsf{prov}, \mathsf{cust})$. Given $\mathcal{I}$, $A$ sets $\mathsf{Honest} \equiv \{N_1, N_4\}$. $N_1$ sends $(N_1, N_2, R = (N_1), P, 0, (0, 0), RA^i_{R_1}, AA^P_{N_1})$ and $(N_1, N_i, R = (N_1), P, 0, (1, 1), RA^i_{R_1}, AA^P_{N_1})$, for $i \in \{3, 4\}$, and these announcements are accepted by every AS because they are authentic. At this point, every AS in $\mathsf{ASes}$ has a one-hop route to $P$, but exchange of announcements is not yet over. Since $N_2$ and $N_3$ are both corrupted, we do not need to consider messages that they exchange. At the end, $N_3$ sends $(N_3, N_4, \tilde{R} = (N_3, N_2, R), P, 0, (0, 0), RA^4_{\tilde{R}_3}, RA^3_{\tilde{R}_2}, RA^2_{R_1}, AA^P_{N_1})$ and $N_4$ accepts this announcement, even though $\tilde{R}$ is not valid. This is because there is no way for $N_4$ to check whether $\mathsf{relation}(N_2, N_3) = (\mathsf{prov}, \mathsf{cust})$ or not. Furthermore, $N_4$ selects this invalid customer route over its one-hop provider route to $P$ because customer routes are preferred over provider routes due *CC*. Thus, we have that $\mathbf{Adv}^{\text{sec-rout-m-3}}_{\mathcal{I},\text{S-BGP-XB}}(A) = 1$, and $A$ is efficient. □

Notice that the attack in the proof of Theorem 3.10.2 requires only that two corrupted ASes be strategically positioned by the adversary on the route such that an export policy violation can be hidden from the remaining ASes on the route. What if adversary were required to meet the following condition?

**Definition 2** (Break-Point Condition)**.** *For any three consecutive ASes* $(\ldots, N_{i-1}, N_i, N_{i+1}, \ldots)$*, for* $1 < i < \ell - 1$*, in any route being advertised such that* $\mathsf{relation}(N_{i-1}, N_i) \neq (\mathsf{cust}, \mathsf{prov})$ *and* $\mathsf{relation}(N_{i+1}, N_i) \neq (\mathsf{cust}, \mathsf{prov})$*,* $N_{i-1} \notin \mathsf{Corrupted}$*.*

Observer that if the Break-Point Condition holds, then the attack in the proof of Theorem 3.10.2 would not work. $N_{i-1}$ is the point where the flow of route announcements for a particular prefix changes its direction in the sense that it is no longer flowing from customers to their providers only. We now show that, together with

other necessary conditions, the break-point condition is sufficient for S-BGP-XB to guarantee route validity for any network in $\mathcal{C}_c$.

**Definition 3.** *We introduce a new experiment, called* $\mathbf{Exp}^{\text{sec-rout-bp}}_{\mathcal{I},\mathcal{PV}}(A)$*, which is identical to* $\mathbf{Exp}^{\text{sec-rout-m}}_{\mathcal{I},\mathcal{PV}}(A)$ *with the restrictions that* $m = 0$ *and for* $\mathbf{Exp}^{\text{sec-rout-bp}}_{\mathcal{I},\mathcal{PV}}(A)$ *to return 2, in addition to the necessary conditions in* $\mathbf{Exp}^{\text{sec-rout-m}}_{\mathcal{I},\mathcal{PV}}(A)$*, the break-point condition must be met. We say that* $\mathcal{PV}$ *guarantees origin authentication, route authentication, and route validity* with the break-point condition *(BP), if probability of* $\mathbf{Exp}^{\text{sec-rout-bp}}_{\mathcal{I},\mathcal{PV}}(A)$ *returning* 1, 2, *or* 3 *respectively is negligible in* $k$*, where* $k$ *is the security parameter such that the size of* $\mathcal{I}$ *is polynomial in* $k$*. We say that* $\mathcal{PV}$ *is fully secure for a network* $\mathcal{I}$ *with BP, if it guarantees origin authentication, route authentication and route validity with BP. We say that* $\mathcal{PV}$ *is fully secure with BP for a class of networks* $\mathcal{C}$ *for which* $\mathcal{PV}$ *is correct, if* $\mathcal{PV}$ *is fully secure with BP for any network* $\mathcal{I} \in \mathcal{C}$*.*

**Lemma 3.10.3.** *S-BGP-XB guarantees route validity with BP for any network in* $\mathcal{C}_c$ *if the underlying* $\mathcal{SS}$ *is uf-cma-secure.*

*Proof.* (Sketch) Since route feasibility is not an issue in $\mathcal{C}_c$, consider an efficient adversary $A$ who can succeed in having some $N_\ell \in \mathsf{Honest}$ accept an announcement advertising a feasible route $R = (N_{\ell-1}, \ldots, N_1)$ such that for at least one $1 \leq i \leq \ell - 1$, either (a) $\mathsf{policy}(N_i, N_{i-1}) = \mathsf{policy}(N_i, N_{i+1}) = (\mathsf{cust}, \mathsf{prov})$, or (b) $\mathsf{policy}(N_i, N_{i-1}) = (\mathsf{cust}, \mathsf{prov})$ and $\mathsf{policy}(N_i, N_{i+1}) = (\mathsf{peer}, \mathsf{peer})$, or (c) $\mathsf{policy}(N_i, N_{i-1}) = (\mathsf{peer}, \mathsf{peer})$ and $\mathsf{policy}(N_i, N_{i+1}) = (\mathsf{cust}, \mathsf{prov})$, or (d) $\mathsf{policy}(N_i, N_{i-1}) = \mathsf{policy}(N_i, N_{i+1}) = (\mathsf{peer}, \mathsf{peer})$. In any case, $N_{i-1} \in \mathsf{Honest}$ because the break-point condition is satisfied, so $N_{i-1}$ must have set $X = 1$, $k = i - 1$ and must have generated route attestation $RA^i_{\tilde{R}_{i-1}} = \mathsf{Sign}(K_{N_{i-1}}, (N_i, N_{i-1} \ldots, N_1, P, 1))$. According to S-BGP-XB, however, $N_\ell$ must have accepted an announcement $(N_{\ell-1}, N_\ell, R = (N_{\ell-1}, \ldots, N_1), P, W, Aux = ((0,0), RA^\ell_{\tilde{R}_{\ell-1}}, \ldots\ RA^2_{\tilde{R}_1}, AA^{\tilde{P}}_{N_1}))$, which means that $\mathsf{Ver}(pk_{N_{i-1}}, (N_i, \ldots N_1, P, \varepsilon)$,

$RA^i_{\tilde{R}_{i-1}}$) resulted in 1. Given that, similarly to the proof of Theorem 3.6.2, we can construct an efficient adversary $B$ that can create a forgery for $\mathcal{SS}$ with a chosen-message attack. □

Note that S-BGP-XB requires on-line verification of the extra flag and is not fully secure with BP for $\mathcal{C}$ because it does not guarantee route feasibility. Introduction of FS-BGP-like link attestations but without explicit relationship indicators may resolve this issue, although it would also greatly diminish the efficiency advantage of introducing only a single flag to prevent route validity attacks. However, the main practical implications of this result concern the scenarios where the adversary is restricted to corrupting only a single AS, an acceptable threat model [47, 99] that we also consider in Chapter 5, which we present in the following observations.

**Observation 3.10.4.** *For networks in $\mathcal{C}$, S-BGP-XB provides guarantees analogous to origin authentication as well as route authentication and validity with respect to an adversary that is allowed to corrupt only a single AS. This is because when only a single AS is corrupted, route feasibility is guaranteed and the break-point condition is trivially satisfied.*

**Observation 3.10.5.** *SoBGP provides guarantees analogous to origin authentication route validity with respect to an adversary that is allowed to corrupt only a single AS. This is because when only a single AS is corrupted, route feasibility is guaranteed. SoBGP cannot, however, guarantee route authentication even in this setting due to VRS attacks.*

## *3.11 Concluding Remarks*

In this chapter we developed a framework for the provable-security treatment of path-vector routing protocols. We defined an interdomain network, a path-vector protocol and designed a formal security model for such protocols, which incorporates

three general security requirements and is strong in terms of adversarial capabilities. Using our framework we analyzed security of the S-BGP, and we proved that S-BGP meets two out of the security definition's three requirements, namely origin and route authentication, assuming the underlying signature scheme is secure. We showed that S-BGP does not guarantee route validity and then studied how the protocol can be modified to meet all three security requirements at the same time with FS-BGP, as well as the more light-weight S-BGP-XB with respect to a weaker but well-defined threat model. We showed that SoBGP fails to meet the goals of route authentication and route validity, making S-BGP a better candidate for securing the Internet's routing infrastructure, but we also studied conditions under which SoBGP can guarantee route validity with respect to a weaker but still well-defined threat model. Whether using FS-BGP, S-BGP-XB, or SoBGP, our results suggest that network operators would have to be willing to make their routing policies and business relationships known to remote ASes in order to have any security guarantees against route validity attacks. This is because without revealing such information, there is no way in general for remote ASes to verify validity of routes.

Finally, we investigated the possibility of relaxing the PKI requirement, such that not all ASes have certified keys, while relying on non-traditional, non-cryptographic security assumptions, and presented the necessary and sufficient conditions to achieve weaker but still well-defined security guarantees in this setting. These results facilitate our understanding of how gradual deployment, as well as full deployment but where, for efficiency reasons, not all parties want to execute parts of the protocol that require the use of their private keys, of security-enhanced BGP variants on the Internet could be made possible. We show that if all prefixes and links are certified by a trusted certification authority, even when no AS has a public key, ASes are guaranteed to discover valid routes with authentic origins, and the worst thing that can happen is that an honest AS may accept a route to some prefix such that for at least one

honest AS on that route, the latter does not prefer its part of that route the most. We have then discussed that in this setting, due the Internet's lack of any provably secure accountability mechanism, the Internet as a whole may be just as protected against adversaries whose primary goal is to divert traffic onto unwanted routes as when PKI is fully deployed.

Thus, the framework we developed in this chapter and our results should be useful for protocol developers, standards bodies, and government agencies not only with respect to verifying security guarantees of previous as well as future routing protocols, but also in understanding the issue of gradual deployment secure variants of BGP. While on the one hand our results suggest how to achieve well-defined security guarantees with partial PKI deployment, on the other hand our results also suggest that in the initial stages of partial deployment of security-enhanced BGP variants, it may be more beneficial to deploy link certificates rather than have some ASes possess public keys while deploying no link certificates at all. Our results highlight importance of considering the trade offs between operational complexities and whether route validity verification could be done off line, *e.g.,* with link attestations in FS-BGP-PD even with no PKI deployment, or may be required to be done on-line, *e.g.,* with a more light-weight S-BGP-XB that requires full PKI deployment. Our results also highlight the crucial role that RPKI may play in the deployment evolution of security-enhanced BGP protocols on the Internet, especially if RPKI may also be used to deploy link attestations.

# CHAPTER IV

# EVALUATING NETWORK STABILITY GUARANTEES

## *4.1 Introduction*

Almost every observed BGP attack and misconfiguration to date [77, 86, 24, 34, 83, 85, 19] shares a common characteristic: the wrongdoer announces the same bogus information throughout the duration of the incident. We refer to this class of attacks by *fixed-route attacks*. Prefix hijacks as well as general route authentication and validity attacks we consider in this thesis fall into this class. It is known that routing policies of ASes can interact in ways that lead to persistent routing oscillations, where some ASes endlessly change the routes they select to reach a particular destination [51]. BGP oscillations render the network unpredictable and can significantly harm network performance, causing traffic to be mis-ordered, delayed, and even dropped. We have shown in Figure 1(c) that a stable network running BGP can be destabilized by a single fixed-route attacker, and in this chapter we investigate the conditions required to avoid such instabilities in presence of fixed-route attackers.

Due to Internet's scale and complexity, S*BGP is likely to coexist with BGP for possibly a very long time. For example, IPv6 and DNSSEC have been in deployment since at least 1999 and 2007 respectively. Thus, in this chapter we address the main question of network stability when S*BGP is only partially deployed. This is challenging because it is expected that, for backwards compatibility, in a partial deployment scenario secure ASes may have to use legacy, insecure BGP to exchange routing information with ASes that do not deploy S*BGP [68]. This would help in preventing losing connectivity to certain parts of the Internet that do not deploy S*BGP. This means, that ASes that become secure may still have to accept many insecure routes

via BGP. This issue has been mostly ignored by the research community before, either by assuming that ASes will never accept insecure routes [12, 28], by studying only the full deployment scenario where every AS has already deployed S*BGP [46, 25], or by focusing on other challenges such as creating incentives for ASes to adopt S*BGP in the first place [40, 28]. Recall that every AS on the Internet uses its routing policy to select a single, most preferred route to every destination, so in partial S*BGP deployment, secure ASes may be forced to make a choice between secure and insecure routes. However, there seems to be no consensus between network operators on how secure routes should be prioritized with respect to insecure routes [42].

While BGP routing policies differ between ASes and are often kept private, many commercial routing models, believed to capture routing policies of ASes on the Internet, have been studied in the community [38, 39, 51, 40, 14, 56, 57]. In Section 4.2 we present a class of routing models, a variation of previously studied models, that captures ASes preferences for secure routes when S*BGP is only partially deployed.

In Section 4.3 we demonstrate that stable networks running BGP can become unstable resulting from partial deployments of S*BGP. We then identify conditions under which stability is maintained in the presence of fixed-route attackers and partially deployed S*BGP with respect to our routing models in Section 4.4. Our results also allow us to quantify the convergence rate in terms of asynchronous rounds [37, 92], *i.e.,* periods of time in which each AS gets at least one update message from each neighbor, and then processes and sends updates to its neighbors according to its policies at least once after receiving these updates.

Finally, we conclude this chapter with a summary and discussion of our results' practical implications in Section 4.6.

## *4.2 Routing Model*

In this chapter we do not consider colluding attackers, and we do not allow them to have control of communication of other ASes. However, we allow attackers to announce different routing information to different neighbors, as long they consistently advertise the same information to the same neighbors for the duration of their attacks. We have shown in Chapter 3 that S-BGP guarantees origin and route authentication, and the same result can be shown to hold for BGPSEC assuming RPKI's functionality as a black box. Thus, these protocols allow an AS to verify the correctness of the AS-level route information it learns from its neighbors. S-BGP and BGPSEC verify that every AS on a route sent a routing announcement for that route. We also showed in Chapter 3 that SoBGP does not guarantee route authentication, so we do not consider it in this chapter and use S*BGP to denote only S-BGP and BGPSEC protocols. Note that the notation and BGP routing models that we use in this chapter are for convenience different from the model and notation we considered in Chapter 3. We omit any notation and routing model details that may be useful for analyzing provable security guarantees but not needed for analyzing stability guarantees.

For S*BGP protocols to prevent routing attacks, verification of routes alone is not sufficient. ASes also need to use this information to make their routing decisions. When S*BGP is only partially deployed, an AS that adopts S*BGP must be able to process and react to insecure routing information, so that it can still route to destination ASes that have not yet adopted S*BGP. In accordance with the current standard [68], in this chapter we will assume that AS learns a route via S*BGP only if every AS on that route has deployed S*BGP. Otherwise, the route is propagated via legacy BGP [97].

We call an AS that has adopted S*BGP a *secure AS*, and a route learned via S*BGP (*i.e.,* a route where every AS is secure) a *secure route*. All other routes are called *insecure*. If a secure AS learns of both secure and insecure routes, what

role should security play in route selection? To blunt routing attacks, secure routes should be preferred over insecure routes, but how should expensive or long secure routes be ranked relative to revenue-generating or short insecure routes? We address this question next.

### 4.2.1 Secure Routing Models

Recall from Chapter 3 that we represent AS-level topology with a graph $G = (V, E)$, where the set of vertices $V$ represents ASes and the set of links (edges) $E$ represents direct BGP links between neighboring ASes. The class of commercial routing models presented in this section is based on previous research in this domain [38, 39, 51, 40, 14, 56, 57, 42].

#### 4.2.1.1 Routing Policies Without Security Considerations

- Neighboring ASes have one of two business relationships: *customer-provider*, in which the customer AS purchases connectivity from the provider AS, and *peering*, in which the two neighboring ASes agree to carry transit traffic between their customers for free.

- To select a route from multiple available routes to every destination AS $d$, each AS considers the following (in order):

    1. Local pref (**LP**): prefer revenue-generating routes through customer neighbors to routes through its peer neighbors, and prefer the latter to routes through provider neighbors.

    2. AS routes (**SP**): prefer shorter routes to longer routes.

    3. Tiebreak (**TB**): use a consistent rule (*e.g.,* geographic location, device ID) to break ties among remaining routes.

- After selecting a single route as above, an AS announces that route to a subset of its neighbors according to its Export policy (**Ex**): if a route is through a customer, the route is exported to all neighbors, and it is exported to customers only otherwise.

We say that an AS $s$ *learns a route* or *has a route* $R$ if $R$ was announced to $s$ by one of its neighbors, and we say that AS $s$ *uses a route* $R$ if it chooses $R$ from its set of available routes. Recall that AS $s$ has a customer, peer, or provider route if its neighbor on that route is a customer, peer, or provider respectively. Recall from Chapter 3 that customer-to-provider relationships are denoted with directed edges from the customers to their providers, while peer-to-peer relationships are denoted with undirected edges.

The relative ranking of the **LP**, **SP**, and **TB** are standard in most router implementations [33]. The **LP** and **Ex** steps are based on the classical economic models of BGP routing studied in [38, 39, 56, 57]. **LP** captures ASes' incentives to send traffic along revenue-generating customer routes, as opposed to routing through peers (which does not increase revenue), or routing through providers (which comes at a monetary cost). **Ex** captures ASes's willingness to transit traffic only when paid to do so by a customer.

#### 4.2.1.2 Routing Policies with Security Considerations

To model routing in scenarios where S*BGP is partially deployed and ASes that run S*BGP may have to make a route-selection decision between secure and insecure routes, we incorporate an additional consideration, Secure Paths (**SecP**), into its routing policy: prefer a secure route over an insecure route. While the security $1^{st}$ model is the most idealistic from the security perspective, it is likely the least realistic. During incremental deployment, network operators are expected to cautiously

incorporate S*BGP into routing policies, placing security $2^{nd}$ or $3^{rd}$, to avoid disruptions due to (1) changes to traffic engineering, and (2) revenue lost when expensive secure routes are chosen instead of revenue-generating customer routes. The security $1^{st}$ model might be used only once these disruptions are absent (*e.g.,* when most ASes have transitioned to S*BGP), or to protect specific, highly-sensitive IP prefixes. A survey of 100 network operators [42] found that 10% would rank security $1^{st}$, 20% would rank security $2^{nd}$ and 41% would rank security $3^{rd}$. The remaining operators opted not to answer this question.

In our investigation, we use all three different models for incorporating this consideration into the above routing model.

- **Security** $1^{st}$ : the **SecP** is placed above the **LP** consideration. In this model security is ASes's most important consideration.
- **Security** $2^{nd}$ : the **SecP** comes between the **LP** and **SP** considerations. In this model all ASes place economic considerations above security concerns.
- **Security** $3^{rd}$ : the **SecP** comes after both business considerations and AS-route length. In this model, also used in [40], the **SecP** serves the role of a tie-breaker and comes between **SP** and **TB** steps.

#### 4.2.1.3 Stubborn ASes: Generalized Model of Local Preferences

Although a survey of network operators [42] suggests that about 80% of network operators prefer customer routes over peer and provider routes, this survey also points out that some network operators, especially those of content providers, prefer shorter peer routes over longer customer routes. Let us consider a generalized model of local preferences $\mathbf{LP}_k$, for various values of $k$, in which routes are ranked as follows:

- customer routes of length 1
- peer routes of length 1

- $\cdots$
- customer routes of length $k$
- peer routes of length $k$
- customer routes of length $> k$
- peer routes of length $> k$
- provider routes

Observer that $\mathbf{LP}_\infty$ is equivalent to a routing policy where ASes almost equally prefer customer and peer routes over provider routes:

- prefer peer and customer routes over provider routes (selecting customer routes over peer routes if they are of the same length)
- prefer shorter routes over longer routes
- break ties in favor of customer routes
- use intradomain criteria (*e.g.,* geographic location, device ID) to break ties among remaining routes

Additionally, observe that $\mathbf{LP}_0$ yields the local preference model **LP** described in Section 4.2.1.1, where customer routes are preferred to peer routes, which are preferred to provider routes.

The **SP** and **TB** steps follow the $\mathbf{LP}_k$ analogous to the secure routing model described above. For any $k \in \mathbb{Z}^+$, the security $1^{st}$ model prioritizes the **SecP** step above $\mathbf{LP}_k$, while the security $2^{nd}$ model prioritizes the **SecP** step between $\mathbf{LP}_k$ and **SP**, and in the security $3^{rd}$ model **SecP** step is prioritized below **SP** and serves the role of a tie-breaker before **TB**.

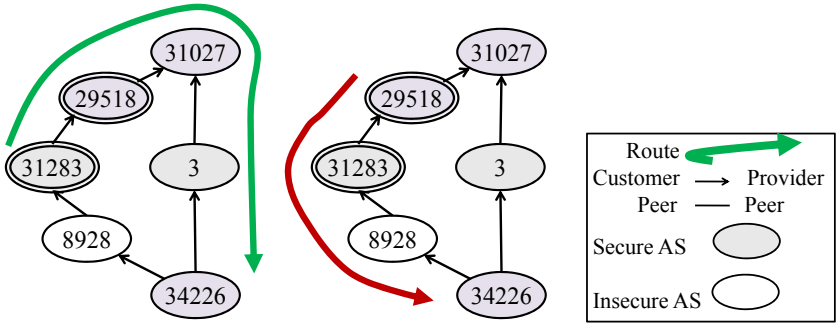


**Figure 8:** Example of a S*BGP Wedgie.

## *4.3 Security-Ranking Disagreements Can Be Bad*

It is important to note that in each of our S*BGP routing models, the ranking of the **SecP** step in the route selection process is consistent across ASes. The alternative—lack of consensus amongst network operators as to where to place security in the route selection process—can lead to more than just confusion. It can result in undesirable phenomena that we discuss next.

Consider the network in Figure 8, taken from the UCLA AS-level topology from 24 September 2012 [31], and suppose that all ASes in this network, except AS 8928, have deployed S*BGP. The Swedish ISP AS 29518 places security below **LP** in its route selection process, while the Norwegian ISP AS 31283 prioritizes security above all else (including **LP**). Thus, while AS 29518 prefers the customer route through AS 31283, AS 31283 prefers the secure route through its provider AS 29518.

The following undesirable scenario, known as a BGP Wedgie [50] can occur. Initially, the network is in an intended stable routing configuration in which AS 31283 uses the secure route through its provider AS 29518 (left). Now suppose the link between AS 31027 and AS 3 fails. Routing now converges to a different stable configuration, where AS 29518 prefers the customer route through AS 31283 (right). When the link comes back up, BGP does not revert to the original stable configuration, and the system is stuck in an unintended routing outcome.

BGP Wedgies [50] can cause unpredictable network behavior that is difficult to debug. Furthermore, Sami *et al.* [92] have shown that the existence of multiple stable configurations, as in Figure 8, implies that persistent routing oscillations are possible. We address this problem in the next section.

## *4.4 When Can Stability Be Guaranteed?*

In this section we show that convergence to unique stable routing state is guaranteed for any S*BGP deployment scenario, as long as all ASes prioritize secure routes the same way. Furthermore, we bound the convergence rate in terms of asynchronous rounds. An asynchronous round is a period of time in which each AS gets at least one update message from each of its neighbors, and is activated at least once after receiving these updates. We say that a non-attacking AS gets activated when it processes the most recent update messages received from neighboring ASes, select its most preferred available, loop-free route according to its routing policies, and then propagates this route to its neighbor in accordance with **Ex**.

**Theorem 4.4.1.** *S*BGP convergence to a unique stable routing state is guaranteed in security* $1^{st}$, $2^{nd}$, *and* $3^{rd}$ *models, for any local preference model* $\boldsymbol{LP}_{k\in\mathbb{N}}$, *for any S*BGP deployment scenario, in presence of any number of fixed route attackers. Moreover, convergence to a stable state is guaranteed within* $(2\mathcal{X}+1)$ *asynchronous rounds, where* $\mathcal{X}$ *is the height of the customer-provider hierarchy.*

*Proof.* To prove Theorem 4.4.1, we first consider algorithms for computing routing outcomes for security $1^{st}$, $2^{nd}$, and $3^{rd}$ models in Section 4.5 and then show that these algorithms correctly output the routing S*BGP outcomes in Lemmas 5-21 in Section 4.5.5. The proof of the theorem then follows. □

Intuitively, customer-provider hierarchy is the analog of a diameter on the Internet graph. On the Internet pf today, the average depth of the customer-provider hierarchy

is approximately five levels [31], so the above theorem implies that even in presence of fixed-route attackers, convergence in S*BGP routing with commercial routing policies is quite fast.

## *4.5 Computing Routing Outcomes*

Below we present algorithms for computing S*BGP routing outcomes in the presence of fixed-route attackers, for security $1^{st}$, $2^{nd}$, and $3^{rd}$ S*BGP routing models with respect to $\mathbf{LP}_k$ local preference model. These algorithms receive as input a set of pairs $(M, d)$, where $M \subseteq V \setminus \{d\}$ is a set containing fixed-route attackers, non-negative integer $k$, and the set of secure ASes $S$. These algorithms output the S*BGP routing outcome for each of our three S*BGP security prioritization models. We emphasize out that our algorithms can also be used to compute routes during normal conditions, when there is no attacker $M = \emptyset$, and when no AS is secure $S = \emptyset$. In these algorithms, which extend the algorithmic approach used in [46, 40, 43] to handle partial S*BGP deployment in the presence an adversary described in Section 5.4.2.1, we carefully construct a partial multi-rooted routing tree by performing multi-stage breadth-first-search (BFS) computations with $d$ and attackers in $M$ as the roots.

In Section 4.5.5, we prove that our algorithms are correct, in the sense that they compute the appropriate S*BGP routing outcomes. .

### 4.5.1 Notation and Preliminaries

Since BGP (and S*BGP) sets up routes to each destination independently, we focus on routing to a unique destination $d$. We say that a route is legitimate if it does not contain an attacker $m$ in $M$; this could be either because there is no attacker, *i.e.,* $M = \emptyset$, or because no attacker is on that route. We say that a route is attacked or bogus otherwise. Observe that in the presence of at least one fixed route attacker $m \in M$, all attacked routes must contain $m$.

We consider $\mathbb{N}$ to be the set of non-negative integers. We say that a route $R =$

$\{v_i, v_{i-1}, ..., v_1, d\}$ contains AS $x$, if at least one AS in $\{v_i, v_{i-1}, ..., v_1, d\}$ is $x$. We borrow the following definition of *perceivable routes* from [72].

**Definition 4** (Perceivable routes)**.** *A simple, loop-free route $R = \{v_{i-1}, \ldots, v_1, d\}$ is* perceivable *at AS $v_i$ if one of the two following conditions holds:*

1. *$R$ is legitimate and for every $0 < j < i$ it follows that $v_j$ announcing the route $(v_j, \ldots, d)$ to $v_{j+1}$ does not violate **Ex**.*
2. *$R$ is attacked, so $R$ contains the closest to $v_i$ attacker $m = v_j$ that chooses to announce $(v_j, \ldots, v_1, d)$ to $v_{j+1}$ throughout the attack, and for every $j < j' < i$, it follows that $v_{j'}$ announcing the route $(v_{j'}, \ldots, d)$ to $v_{j+1}$ does not violate **Ex**.*

Intuitively, an AS's set of perceivable routes captures all the routes this AS could potentially learn during the S*BGP convergence process. All non-perceivable routes from an AS can safely be removed from consideration as the **Ex** condition ensures that they will not propagate from the destination and attacker(s) to that AS.

**PR and BPR Sets** Let $\mathsf{PR}(v_i, M, d)$ be the set of perceivable routes from $v_i$ for when fixed-route attacker(s) in $M$ attack destination $d$. Given a set of secure ASes $S$, for every AS $v_i$ we define the $\mathsf{BPR}(v_i, S, M, d)$ to be the set of all perceivable routes in $\mathsf{PR}(v_i, M, d)$ that are preferred by $v_i$ over all other perceivable routes, before the arbitrary tiebreak step **TB**, according to the secure routing policy model (*i.e.,* security $1^{st}$, $2^{nd}$, or $3^{rd}$) under consideration. Recall that we set $M = \emptyset$ when there is no attacker and $S = \emptyset$ when no AS is secure. We define $\mathsf{Nxt}(v_i, S, M, d)$ to be the set of all neighbors of $v_i$ that are next hops of all routes in $\mathsf{BPR}(v_i, S, M, d)$. For simplification, we will just use $\mathsf{Nxt}(v_i)$ when it is clear what $S$, $M$ and $d$ are.

Observe that in each of our models, all routes in $\mathsf{BPR}(v_i, S, M, d)$ must (1) belong to the same type—customer routes, peer routes, or provider routes, (2) be of the same length, and (3) either be all secure or all insecure.

### 4.5.2 Algorithm for Security $3^{rd}$

We now present our algorithm for computing the S*BGP routing outcome in the security $3^{rd}$ model in the presence of a set of secure ASes $S$ and fixed-route attackers $M$, for any $\mathbf{LP}_{k\in\mathbb{N}}$. We note that this algorithm can also be used to compute the routing outcome when no AS is secure, *i.e.,* $S = \emptyset$ and/or there is no attacker *i.e.,* $M = \emptyset$. As in [72] (which studies a somewhat different BGP routing model and does not consider S*BGP) we exhibit an iterative algorithm **Fix-Routes** (FR) that, informally, at each iteration fixes a single AS's route and adds that AS to a set $\mathcal{F} \subseteq V$. This goes on until all ASes are in $\mathcal{F}$ (that is, all ASes' routes are fixed). We will later prove that FR outputs the BGP routing outcome.

FR consists of four subroutines: **Fix Stubborn Routes** (FStuBB), **Fix Customer Routes** (FCR), **Fix Peer Routes** (FPeeR), and **Fix Provider Routes** (FPrvR), that FR executes in that order. Note that at the very beginning of this algorithm, $\mathcal{F}$ contains only the legitimate destination $d$ and all the fixed-route attackers in $M$. We let $r$ be the FR iteration and initialize it to 0. We initially set $\mathsf{PR}^r(v_i) = \mathsf{PR}(v_i, M, d)$ and $\mathsf{BPR}^r(v_i) = \mathsf{BPR}(v_i, S, M, d)$ for every AS $v_i$. FR then executes:

1. **Run FStuBB($k$)**;
2. **Run FCR($\infty$)**;
3. **Run FPeeR($\infty$)**;
4. **Run FPrvR**;

We now describe each subroutine in detail.

**FStuBB($k$) Subroutine:**

FStuBB takes parameter $k$, initializes $i = 1$, and executes the following while $i \leq k$.

1. Run FCR($i$);

2. Run FPeeR($i$);
3. i++;

**The FCR($K$) Subroutine:**

FCR takes parameter $K$ as input. Intuitively, FCR constructs a partial multi-rooted tree (rooted at $d$ and all attackers in $M$) of height at most $K$ on the graph, using a BFS computation in which only customer-to-provider edges are traversed. While there is at least one AS $s \notin \mathcal{F}$ such that $\mathsf{PR}^{r-1}(s)$ contains at least one customer route of length at most $K$, we fix the route of (at least) one AS by executing the following steps:

1. r++;
2. Select the AS $v_i \notin \mathcal{F}$ that has the shortest **customer** route in its set $\mathsf{BPR}^{r-1}(v_i)$ of length at most $K$ (if there are multiple such ASes, choose one arbitrarily);
3. Add $v_i$ to $\mathcal{F}$; set $\mathsf{Nxt}(v_i)$ to be $v_i$'s next-hop on the route in $\mathsf{BPR}^{r-1}(v_i)$ selected according to its tie-breaking rule **TB** preferring secure routes (if any) over insecure routes;
4. Remove, for every AS $v_j$, all routes in $\mathsf{PR}^{r-1}(v_j)$ that contain $v_i$ but whose suffix at $v_i$ is not in $\mathsf{BPR}^{r-1}(v_i)$ to obtain the new set $\mathsf{PR}^{r}(v_j)$; set $\mathsf{BPR}^{r}(v_j)$ to be $v_j$'s most preferred routes in $\mathsf{PR}^{r}(v_j)$;
5. Add all ASes $v_j$ such that $\mathsf{PR}^{r}(v_j) = \emptyset$ to $\mathcal{F}$.

**The FPeeR($K$) Subroutine:**

FPeerR takes parameter $K$ as input. At this point, $\mathcal{F}$ contains only $d$, $M$, and ASes with either empty, customer or peer routes. We now use only single peer-to-peer edges to connect new yet-unexplored ASes to the ASes that were locked in the partial routing tree in the previous stages of the algorithm. While there is at least one AS $s \notin \mathcal{F}$ such that $\mathsf{PR}^{r-1}(s)$ contains at least one peer route of length at most $K$, the following steps are executed:

1. r++;
2. Select an AS $v_i \notin \mathcal{F}$ with a peer route of length at most $K$ in $\mathsf{PR}^{r-1}(s)$ (if there are multiple such ASes, choose one arbitrarily);
3. Add $v_i$ to $\mathcal{F}$; set $\mathsf{Nxt}(v_i)$ to be $v_i$'s next-hop on the route in $\mathsf{BPR}^{r-1}(v_i)$ selected according to its tie-breaking rule **TB** preferring secure routes (if any) over insecure routes;
4. Remove, for every AS $v_j$, all routes in $\mathsf{PR}^{r-1}(v_j)$ that contain $v_i$ but whose suffix at $v_i$ is not in $\mathsf{BPR}^{r-1}(v_i)$ to obtain the new set $\mathsf{PR}^r(v_j)$; set $\mathsf{BPR}^r(v_j)$ to be $v_j$'s most preferred routes in $\mathsf{PR}^r(v_j)$
5. add all ASes $v_j$ such that $\mathsf{PR}^r(v_j) = \emptyset$ to $\mathcal{F}$.

**The FPrvR Subroutine:**

We now run a BFS computation in which only provider-to-customer edges are traversed, that is, only ASes who are direct customer of those ASes that have already been added to the partial two-rooted tree are explored. This step starts with $\mathcal{F}$ and the configuration of the routing system and the $\mathsf{PR}$ and $\mathsf{BPR}$ sets the way it is after the consecutive execution of FStuBB, FCR ,and FPeeR.

While there is an AS $s \notin \mathcal{F}$ such that $\mathsf{PR}^{r-1}(s)$ contains at least one provider route, we execute the identical steps as in FCR($\infty$), with the exception that we look for the $v_i$ that has the shortest **provider** route in its set $\mathsf{BPR}^{r-1}(v_i)$.

### 4.5.3 Algorithm for security $2^{nd}$

Our algorithm for the security $2^{nd}$ model is a refinement of the iterative algorithm **Fix Routes** (FR) presented above for the security $3^{rd}$ model. This new algorithm is also based on four stages of BFS, but in each stage we are careful to prioritize ASes with *secure* routes over ASes with insecure routes.

We present the following three new subroutines. (1) **Fix Secure Customer Routes** (FSCR): FSCR is identical to FCR, with the sole exception that for the

AS chosen at each iteration $r$ has a $\mathsf{BPR}^{r-1}$ that contains a *secure* customer route; (2) **Fix Secure Provider Routes** (FSPrvR): FSPrvR is identical to FPrvR, with the sole exception that for the AS chosen at each iteration $r$ has a $\mathsf{BPR}^{r-1}$ that contains a *secure* provider route; (3) **Fix Secure and Insecure Stubborn Routes** (FSIStuBB): FSIStuBB is identical to FStuBB, with the addition that it executes FSCR($i$) prior to executing FCR($i$), for each iteration $i$. The variant of FR for the security $2^{nd}$ model executes the subroutines in the following order:

1. **Run FSIStuBB($k$)**;
2. **Run FSCR($\infty$)**;
3. **Run FCR($\infty$)**;
4. **Run FPeeR($\infty$)**;
5. **Run FSPrvR**;
6. **Run FPrvR**;

### 4.5.4 Algorithm for Security $1^{st}$

Once again, we present a variant of the Fix Routes (FR) algorithm. This multi-stage BFS computation first discovers all ASes that can reach the destination $d$ via secure routes and only then discovers all other ASes (as in our algorithm for the security $3^{rd}$ model).

We present the following two new subroutines. (1) **Fix Secure Peer Routes** (FSPeeR): FSPeeR is identical to FSPeeR, except that the AS chosen at each iteration $r$ has a *secure* peer route in its $\mathsf{BPR}^{r-1}$ set. (3) **Fix Secure Stubborn Routes** (FSStuBB): FSStuBB is identical to FStuBB, with the exception that it executes FSCR($i$) instead of FCR($i$) and FSPeeR($i$) instead of FPeeR($i$), for each iteration $i$. This variant of FR executes the subroutines in the following order:

1. **Run FSStuBB($k$)**;
2. **Run FSCR($\infty$)**;

3. **Run FSPeeR($\infty$)**;
4. **Run FSPrvR**;
5. **Run FStuBB($k$)**;
6. **Run FCR($\infty$)**;
7. **Run FPeeR($\infty$)**;
8. **Run FPrvR**;

### 4.5.5 Correctness of Algorithms

We now prove that our algorithms for computing the S*BGP routing outcomes indeed output the desired outcomes.

#### 4.5.5.1 Correctness of Algorithm for Security $3^{rd}$

The proof that our algorithm for the security $3^{rd}$ model outputs the correct S*BGP routing outcome in this model for any local preference $\mathbf{LP}_{k\in\mathbb{N}}$ follows from the combination of the lemmas below. Recall that each of our algorithms computes, for every AS $v_i$, a next-hop AS $\mathsf{Nxt}(v_i)$. Let $R_{v_i}$ be the route from $v_i$ induced by these computed next-hops.

**Lemma 4.** *Under S*BGP routing, for any $\boldsymbol{LP}_{k\in\mathbb{N}}$, the route of every AS added to $\mathcal{F}$ in FStuBB(k) is guaranteed to stabilize to the route $R_{v_i}$.*

*Proof.* If $k = 0$, then the Lemma trivially holds because then the only AS(es) added to $\mathcal{F}$ in FStuBB would be $d$ and $M$. Otherwise, we prove the lemma by induction on the FStuBB iteration. Consider the first iteration, *i.e.,* FStuBB runs FCR(1) followed by FPeeR(1). Observe that the ASes chosen as a result of FCR(1) must be direct providers of $d$ and/or attackers in $M$ announcing routes of length 0, *i.e.,* have customer routes of length 1 to at least one of these ASes (note that at this point only prefix-hijackers in $M$ can be considered). Hence, in the security $3^{rd}$ model, once these ASes learn of $d$ and/or attackers in $M$ announcing routes of length 0, each will

select a direct customer route to one of these ASes and never choose a different route thereafter because this would be its most preferred route. Similarly, observe that the ASes chosen as a result of FPeeR(1) must be direct peers of $d$ and/or attackers in $M$ announcing routes of length 0, *i.e.,* have a peer route of length 1 to one of these ASes (again, note that at this point only prefix-hijackers in $M$ can be considered). Hence, in the security $3^{rd}$ model, once these ASes learn of $d$ and/or attackers in $M$ announcing routes of length 0, each will select a direct peer route to one of these ASes and never choose a different route thereafter because this would be its most preferred route.

Now, let us assume that for every AS chosen in FStuBB iterations up to $k$ the statement of this lemma holds. Let $v_i$ be the AS chosen at the first iteration $r$ of the execution of FCR($k$) and consider $v_i$'s $\mathsf{BPR}$ set at that time. By definition, every route in the $\mathsf{BPR}$ set is perceivable, so it must comply with **Ex** at each and every hop along the route starting at $v_{i-1}$ and down to $d$ or the closest to $v_i$ fixed-route attacker. Observe, that this, combined with the fact that all routes in $v_i$'s $\mathsf{BPR}$ set are customer routes of length $k$ implies that the suffix of every such route must be a perceivable customer route of length $k-1$. Consider some AS $v_j$ that is $v_i$'s next-hop on some route in $v_i$'s $\mathsf{BPR}$ set. Notice that $v_j$'s route is fixed prior to or at some iteration in FCR($k-1 \geq 1$) because $v_j$ must either be a fixed route attacker or have a perceivable customer route of length $k-1$. Hence, by the induction hypothesis, at some point in the S*BGP convergence process, $v_j$'s route converges to $R_{v_j}$ for every such AS $v_j$. Observe that, from that point in time onward, $v_i$'s best available routes are precisely those capture by $\mathsf{BPR}$ in iteration $r$ of the execution of FCR($k$). Hence, from that moment onwards, $v_i$ will repeatedly select the route $R_{v_i}$ according to the tiebreak step **TB** and never select a different route thereafter.

The argument for any $v_i$ being chosen at iteration $r$ of the execution of FPeeR($k$) is identical to that for FCR($k$) with the exception that all routes in $v_i$'s $\mathsf{BPR}$ set must

be peer routes of length $k$, which, combined with **Ex**, must mean that the suffix of each such route must also be a perceivable customer route of length $k-1$. □

**Lemma 5.** *Under S*BGP routing, for any $\boldsymbol{LP}_{k\in\mathbb{N}}$, the route of every AS added to $\mathcal{F}$ in FCR($\infty$) is guaranteed to stabilize to the route $R_{v_i}$.*

*Proof.* We prove the lemma by induction on the number of FCR iterations. Consider an AS $v_i$ chosen at the very first iteration of execution of FCR($\infty$). Observe that, according to **Ex**, the AS chosen at this iteration must be a provider of some AS $v_j$ in $\mathcal{F}$ that either is $d$, is a fixed-route attacker, or has a customer route of length $k$, *i.e.,* have a customer route of length $k+1$ via some AS $v_j \in \mathcal{F}$ fixed prior to or in FStuBB($k$). By Lemma 4, at some point in the S*BGP convergence process, $v_j$ converges to $R_{v_j}$, for every such AS $v_j$. Observe that once all such ASes' routes have converged and onwards, $v_i$'s best available routes are precisely those captured by BPR in the first iteration of execution of FCR($\infty$), so it will select a customer route of length $k+1$ via an AS in $\mathcal{F}$ at that time and never choose a different route thereafter because this is its most preferred route.

Now, let us assume that for every AS chosen in iterations up to $r$ the statement of this lemma holds. Let $v_i$ be the AS chosen at iteration $r+1$ of the execution of FCR($\infty$) and consider $v_i$'s BPR set at that time. By definition, every route in the BPR set is perceivable and so must comply with **Ex** at each and every hop along the route starting at $v_{i-1}$ and down to $d$ or the closest to $v_i$ fixed-route attacker. Observe, that this, combined with the fact that all routes in $v_i$'s BPR set are customer routes, implies that the suffix of every such route is also a perceivable customer route. Consider some AS $v_j$ that is $v_i$'s next-hop on some route in $v_i$'s BPR set. Notice that $v_j$'s route is fixed at some iteration up to $r$ because $v_j$ has a shorter perceivable customer route than $v_i$. Hence, by Lemma 4 and the induction hypothesis, at some point in the S*BGP convergence process, $v_j$'s route converges to $R_{v_j}$ for every such AS $v_j$. Observe that, from that point in time onward, $v_i$'s best available routes are precisely those capture

by BPR in the $(r+1)^{\text{th}}$ iteration of FCR. Hence, from that moment onwards, $v_i$ will repeatedly select the route $R_{v_i}$ according to the tiebreak step **TB** and never select a different route thereafter. □

**Lemma 6.** *Under S*BGP routing, for any $\boldsymbol{LP}_{k\in\mathbb{N}}$, the route of every AS added to $\mathcal{F}$ in FPeeR($\infty$) is guaranteed to stabilize to the route $R_{v_i}$.*

*Proof.* Consider an AS $v_i$ chosen at some iteration of the execution of FPeeR($\infty$). Observe that, due to **Ex**, if $(v_i, v_{i-1}, \ldots, d)$ is a perceivable peer route then $(v_{i-1}, \ldots, d)$ must either be $d$, be a fixed route attacker, or have a perceivable customer route. Hence, for every such route in $v_i$'s BPR set, it must be the case that the route of $v_i$'s next-hop on this route $v_j$ was fixed prior to or in either FStuBB($k$) or FCR($\infty$). By Lemmas 4 and 5, at some point in the S*BGP convergence process, $v_j$'s route converges to $R_{v_j}$ for every such AS $v_j$. Observe that, from that point in time onward, $v_i$'s best available routes are precisely those capture by its BPR set at the iteration of FPeeR in which $v_i$ is chosen. Hence, $v_i$ will select the route $R_{v_i}$ according to the tiebreak step **TB** and never select a different route thereafter. □

**Lemma 7.** *Under S*BGP routing, for any $\boldsymbol{LP}_{k\in\mathbb{N}}$, the route of every AS added to $\mathcal{F}$ in FPrvR is guaranteed to stabilize to the route $R_{v_i}$.*

*Proof.* We prove the lemma by induction on the number of FPrvR iterations. Consider the first iteration of the execution of FPrvR. Let $v_i$ be the AS chosen at this iteration, let $v_j$ be a next-hop of $v_i$ on some route $R$ in $v_i$'s BPR set, and let $Q$ be the suffix of $R$ at $v_j$. Observe that $Q$ cannot possibly be a provider route, for otherwise $v_j$ would have been chosen in FPrvR before $v_i$. Hence, $Q$ must be either empty, an attacked route, a customer route, or a peer route, and so $v_j$'s route must have been fixed prior to or in either FStuBB($k$), FCR($\infty$), or FPeeR($\infty$). Hence, by Lemmas 4-6, every such $v_j$'s route will eventually converge to $R_{v_j}$. Observe that once all such ASes' routes have converged and onwards, $v_i$'s best available routes are precisely those

captured by BPR in the first iteration of FPrvR. Hence, $v_i$ will select the route $R_{v_i}$ according to the tiebreak step **TB** and never select a different route thereafter.

Now, let us assume that for every AS chosen in iterations up to $r$ the statement of the lemma holds. Let $v_i$ be the AS chosen at iteration $r+1$ of FPrvR and consider $v_i$'s BPR set at this time. Let $v_j$ again be a next-hop of $v_i$ on some route $R$ in $v_i$'s BPR set, and let $Q$ be the suffix of $R$ at $v_j$. Observe that if $Q$ is a provider route then $v_j$'s route must have been fixed in FPrvR at some point in previous iterations. If, however, $Q$ is either an attacked route, customer route or a peer route, then $v_j$'s route must have been fixed prior to or in either FStuBB($k$), FCR($\infty$), or FPeeR($\infty$). Hence, by Lemmas 4-6 and the induction hypothesis, under S*BGP convergence, every such $v_j$'s route will eventually converge to $R_{v_j}$. From that moment onwards, $v_i$'s best available routes are precisely those captured by BPR in the $(r+1)^{\text{th}}$ iteration of FprvR. Hence, $v_i$ will select the route $R_{v_i}$ according to the tiebreak step **TB** and never select a different route thereafter. □

#### 4.5.5.2 Correctness of Algorithm for Security $2^{nd}$

The proof that our algorithm for the security $2^{nd}$ model outputs the S*BGP routing outcome in this model for any local preference $\mathbf{LP}_{k\in\mathbb{N}}$ follows from the combination of the lemmas below. Let $R_{v_i}$ be the route from $v_i$ induced by the algorithm's computed next-hops.

**Lemma 8.** *Under S*BGP routing, for any $\boldsymbol{LP}_{k\in\mathbb{N}}$, the route of every AS added to $\mathcal{F}$ in FSIStuBB(k) is guaranteed to stabilize to the route $R_{v_i}$.*

*Proof.* The proof is essentially the same as the proof of Lemma 4. The main difference is that in FSIStuBB($k$), secure customer routes get fixed prior to insecure customer routes. □

**Lemma 9.** *Under S*BGP routing, for any $\boldsymbol{LP}_{k\in\mathbb{N}}$, the route of every AS added to $\mathcal{F}$ in FSCR($\infty$) is guaranteed to stabilize to the route $R_{v_i}$.*

*Proof.* The proof is essentially the same as the proof of Lemma 5, but where all routes must be secure. □

**Lemma 10.** *Under S*BGP routing, for any $\boldsymbol{LP}_{k\in\mathbb{N}}$, the route of every AS added to $\mathcal{F}$ in FCR($\infty$) is guaranteed to stabilize to the route $R_{v_i}$.*

*Proof.* As in proof of Lemma 5, we prove this lemma by induction on the number of FCR iteration. Consider the first iteration of the execution of FCR($\infty$). Let $v_i$ be the AS chosen at this iteration and let $v_j$ be a next-hop on a route in $v_i$'s $\mathsf{BPR}$ set. Observe that, due to **Ex**, $v_j = d$, or $v_j$ is a fixed route attacker, or $v_j$'s has a customer route that was fixed in FSIStuBB($k$) or FSCR($\infty$). Otherwise, $v_j$ would have been selected in FCR before $v_i$. Hence, by Lemmas 8 and 9, it holds that under each such $v_j$'s route will stabilize at some point, and from that point onwards $v_i$ will repeatedly select $R_{v_i}$. Recall that we made a similar argument in Lemma 5.

Now, let us assume that for every AS chosen in iterations up to $r$ the statement of the lemma holds. Let $v_i$ be the AS chosen at iteration $r+1$ of the execution of FCR($\infty$). Consider $v_i$'s $\mathsf{BPR}$ set at that time and consider an AS $v_j$ that is $v_i$'s next-hop on some route in $v_i$'s $\mathsf{BPR}$ set. Notice that, due to **Ex**, either $v_j$ is in $d \cup M$, or $v_j$ must have a customer route that must have been fixed at some iteration prior to $r+1$ in either FCR($\infty$), if $v_j$ has a shorter perceivable customer route than $v_i$, in FSCR($\infty$), if $v_j$ has a secure customer route to $d$ longer than $k$, or in FSIStuBB($k$) if $v_j$ has a secure or insecure customer route to $d$ of length at most $k$. Hence, by Lemmas 8 and 9 and the induction hypothesis, at some point in the S*BGP convergence process, $v_j$'s route converges to $R_{v_j}$ for every such AS $v_j$. From that point in time onward, $v_i$ will repeatedly select $R_{v_i}$. □

**Lemma 11.** *Under S*BGP routing, for any $\boldsymbol{LP}_{k\in\mathbb{N}}$, the route of every AS added to $\mathcal{F}$ in FPeeR is guaranteed to stabilize to the route $R_{v_i}$.*

*Proof.* The proof is identical to the proof of Lemma 6. □

**Lemma 12.** *Under S*BGP routing, for any* $\boldsymbol{LP}_{k\in\mathbb{N}}$*, the route of every AS added to* $\mathcal{F}$ *in FSPrvR is guaranteed to stabilize to the route* $R_{v_i}$*.*

*Proof.* The proof is essentially the proof of Lemma 7, but where all routes must be secure. □

**Lemma 13.** *Under S*BGP routing, for any* $\boldsymbol{LP}_{k\in\mathbb{N}}$*, the route of every AS added to* $\mathcal{F}$ *in FPrvR is guaranteed to stabilize to the route* $R_{v_i}$*.*

*Proof.* As in the proof of Lemma 5, we prove this lemma by induction on the number of FPrvR iterations. Consider the first iteration. Let $v_i$ be the AS chosen at this iteration, let $v_j$ be a next-hop of $v_i$ on some route $R$ in $v_i$'s $\mathsf{BPR}$ set, and let $Q$ be the suffix of $R$ at $v_j$. Note that $Q$ must be either empty, an attacked route, a customer route, or a peer route, in which case $v_j$'s route must have been fixed prior to fixed prior to FSPrvR or $Q$ is a secure provider route, in which case $v_j$'s route was fixed in FSPrvR. With Lemmas 8-12 and an argument similar to that in the proof of Lemma 12, we can argue that $v_i$'s route will eventually converge to $R_{v_i}$ at some point in the S*BGP routing process.

Now, let us assume that for every AS chosen in iterations up to $r$ the statement of the lemma holds. Let $v_i$ be the AS chosen at iteration $r+1$ of FPrvR and consider $v_i$'s $\mathsf{BPR}$ set at this time. Let $v_j$ again be a next-hop of $v_i$ on some route $R$ in $v_i$'s $\mathsf{BPR}$ set, and let $Q$ be the suffix of $R$ at $v_j$. Observe that if $Q$ is a provider route then $v_j$'s route must have been fixed in either FSPrvR or in FPrvR at some point in iterations up to $r$. If, however, $Q$ is either a customer route or a peer route, then $v_j$'s route must have been fixed in either FSIStuBB($k$), FSCR($\infty$), FCR($\infty$) or FPeeR($\infty$). Hence, by Lemmas 8-12 and the induction hypothesis, every such $v_j$'s route will eventually converge to $R_{v_j}$. Thus we can conclude that $v_i$'s route too will converge to $R_{v_i}$. □

#### 4.5.5.3 Correctness of Algorithm for Security $1^{st}$

The proof that our algorithm for the security $1^{st}$ model outputs the S*BGP routing outcome in this model follows from the combination of the lemmas below. The proofs of these lemmas are almost identical to the proofs for the other two secure routing models, so we omit the details. Again, let $R_{v_i}$ be the route from $v_i$ induced by the algorithm's computed next-hops.

**Lemma 14.** *Under S*BGP routing, for any $\boldsymbol{LP}_{k\in\mathbb{N}}$, the route of every AS added to $\mathcal{F}$ in FSStuBB is guaranteed to stabilize to the route $R_{v_i}$.*

*Proof.* The proof is essentially the same as the proof of Lemma 4, but where all routes must be secure. □

**Lemma 15.** *Under S*BGP routing, for any $\boldsymbol{LP}_{k\in\mathbb{N}}$, the route of every AS added to $\mathcal{F}$ in FSCR is guaranteed to stabilize to the route $R_{v_i}$.*

*Proof.* The proof is essentially the same as the proof of Lemma 5, but where all routes must be secure. □

**Lemma 16.** *Under S*BGP routing, for any $\boldsymbol{LP}_{k\in\mathbb{N}}$, the route of every AS added to $\mathcal{F}$ in FSPeeR is guaranteed to stabilize to the route $R_{v_i}$.*

*Proof.* The proof is essentially the same as the proof of Lemma 6, but where all routes must be secure. □

**Lemma 17.** *Under S*BGP routing, for any $\boldsymbol{LP}_{k\in\mathbb{N}}$, the route of every AS added to $\mathcal{F}$ in FSPrvR is guaranteed to stabilize to the route $R_{v_i}$.*

*Proof.* The proof is essentially the same as the proof of Lemma 7, but where all routes must be secure. □

**Lemma 18.** *Under S*BGP routing, for any $\boldsymbol{LP}_{k\in\mathbb{N}}$, the route of every AS added to $\mathcal{F}$ in FStuBB is guaranteed to stabilize to the route $R_{v_i}$.*

*Proof.* The proof is essentially the same as the proof of Lemma 4, but where all routes must not be secure because all the ASes with secure customer and peer routes of length at most $k$ have converged to their favorite routes added to $\mathcal{F}$ in FSStuBB($k$), and therefore to those routes during S*BGP convergence process by Lemma 14. □

**Lemma 19.** *Under S*BGP routing, for any $\boldsymbol{LP}_{k\in\mathbb{N}}$, the route of every AS added to $\mathcal{F}$ in FCR is guaranteed to stabilize to the route $R_{v_i}$.*

*Proof.* The proof is identical to the proof of Lemma 10. □

**Lemma 20.** *Under S*BGP routing, for any $\boldsymbol{LP}_{k\in\mathbb{N}}$, the route of every AS added to $\mathcal{F}$ in FPeeR is guaranteed to stabilize to the route $R_{v_i}$.*

*Proof.* The proof is identical to the proof of Lemma 11. □

**Lemma 21.** *Under S*BGP routing, for any $\boldsymbol{LP}_{k\in\mathbb{N}}$, the route of every AS added to $\mathcal{F}$ in FPrvR is guaranteed to stabilize to the route $R_{v_i}$.*

*Proof.* The proof is identical to the proof of Lemma 13. □

#### 4.5.5.4 Rate of Convergence

Here we provide intuition for why convergence to a stable state in S*BGP routing is guaranteed to take place in $2\mathcal{X}+1$ asynchronous rounds, where $\mathcal{X}$ is the height of the customer-provider hierarchy.

Among routes of the same local preference, due to **SP**, shorter routes are preferred to longer routes. Thus, with the algorithms described above that we use to compute routes for all three secure routing models, for all $\mathbf{LP}_{k\in\mathbb{N}}$, all computed legitimate customer routes must be of length at most $\mathcal{X}$. Due to **Ex**, this implies that all computed legitimate peer routes must be of length at most $\mathcal{X}+1$, and all computed legitimate provider routes must be of length at most $2\mathcal{X}+1$. Supposing that $\mathcal{Y}$ is maximal length of any bogus route announce by an attacker in $M$, this implies that

the longest computed customer, peer, or provider route must be no longer than $\mathcal{X}+\mathcal{Y}$, $\mathcal{X}+\mathcal{Y}+1$, and $2\mathcal{X}+\mathcal{Y}+1$ respectively.

Since we have shown that our algorithms are correct, the longest route that any AS can converge to in S*BGP execution must be of length at most $2\mathcal{X}+\mathcal{Y}+1$, for all three secure routing models and for all $\mathbf{LP}_{k\in\mathbb{N}}$. Note that because due to **SP** shorter routes are preferred to longer routes, provided they are of the same local preference, our algorithms for computing routes mimic the propagation of routing announcements that happen during S*BGP execution at each asynchronous round. Also, note that $d$ and all fixed-route attackers in $M$ announce their routes at the beginning of the S*BGP's execution, so $\mathcal{Y}$ plays no role in how long it takes non-attacking ASes to learn of such bogus routes in terms of asynchronous rounds. Thus, the longest it could take for any AS to converge to any route during S*BGP execution for any of our three secure routing models and for any $\mathbf{LP}_{k\in\mathbb{N}}$, is $2\mathcal{X}+1$ asynchronous rounds.

## *4.6 Concluding Remarks*

In this chapter we have developed a family of S*BGP routing models that are general enough to capture commercial routing policies with various security prioritizations and local preferences of ASes. Using our routing models, we have studied the conditions under which network stability could be guaranteed for any S*BGP deployment scenarios and in presence of an arbitrary number and type of fixed-route attackers.

The main results in this chapter suggest that, whether S*BGP is deployed or not, major routing outages, such as [77, 86, 24, 34, 83, 85, 19], should not be expected to result in Internet-wide routing oscillations because they constitute fixed-route attacks in a commercial routing setting. Additionally, our results suggest that network operators would have to agree on how to prioritize security in their routing policies to avoid such undesirable routing anomalies such as BGP Wedgies and persistent routing oscillations.

# CHAPTER V

# QUANTIFYING SECURITY BENEFITS AND HARM IN FULL AND PARTIAL DEPLOYMENT

## *5.1 Introduction*

In Chapters 3 and 4 we have studied various S*BGP candidates to remedy BGP security vulnerabilities, and we have characterized the conditions under which they provide sufficient security and stability guarantees. In this chapter we study another crucial deployment property of S*BGP protocols. RPKI is currently being deployed, and it has been shown that fully-deployed RPKI could do much to improve routing security on the Internet [46]. RPKI requires neither changes to the BGP message structure nor on-line cryptographic computations, and, although S*BGP protocols provide protection from more sophisticated attacks than RPKI, S*BGP protocols require both [68]. The deployment of RPKI is already a significant challenge [5], so in this chapter we ask the question: Is the juice even worth the squeeze? In other words, we would like to find out if the security benefits from S*BGP protocols over those provided with RPKI be worth the extra effort required to deploy it.

To address this question we study the interplay between local policies of ASes' and the impact that S*BGP deployments could have on the Internet in terms of their benefits and harm. S*BGP's impact on the Internet may very much depend on the routing policies used by individual ASes, the AS-level topology, and the set of ASes deploying S*BGP. For instance, suppose a secure AS is faced with a choice between a *secure route* (learned via S*BGP) and an *insecure route* (learned via legacy, insecure BGP) to the same destination. In an ideal scenarios, we might expect that AS to always select the secure route over the insecure one, but a network operator of that AS

must balance security against economic and performance concerns. As such, a *long* secure route through a *costly* provider might be less preferred than a *short* insecure route through a *revenue-generating* customer. In fact, the BGPSEC standard is very careful to provide maximum flexibility in this regard, allowing network operators to weigh security and other route properties according to their own local policies [68]. However, such local policy decisions may have negative global implications. For instance, if a secure AS selects an insecure route, this could lead its secure neighbors to also use an insecure route, regardless of their own routing policies, for the lack of other alternatives. This means that an AS that has deployed S*BGP may not learn of any secure routes to certain destinations simply because its secure neighbors prefer insecure routes to those destinations. We formalize this particular issue with respect to fixed-route attackers and show when it can be avoided in Section 5.2.

In Section 5.3 we present the metric that we use to quantify benefits and harm of various S*BGP deployments with respect to fixed-route attackers in this chapter. This metric provides us with a framework for quantifying security benefits and harm of any S*BGP deployment for any security prioritization (*i.e.,* $1^{st}$, $2^{nd}$, and $3^{rd}$) and local preference model $\mathbf{LP}_k$ described in Chapter 4.

To deal with the large space of parameters that we explore, such as attackers, destinations, S*BGP deployment scenarios, and different routing policies, we have designed parallel simulation algorithms. All simulations and examples described in this chapter were run over variants of the UCLA AS-level topology from 24 September 2012 [31]. We describe our experimental set up in more detail in Section 5.4.

The vast number of choices for the set of ASes that could adopt S*BGP makes evaluating security benefits very challenging, and to deal with this intractability we have designed a novel methodology for efficiently computing bounds on the maximum and minimum security improvements for any deployment scenario. We present our

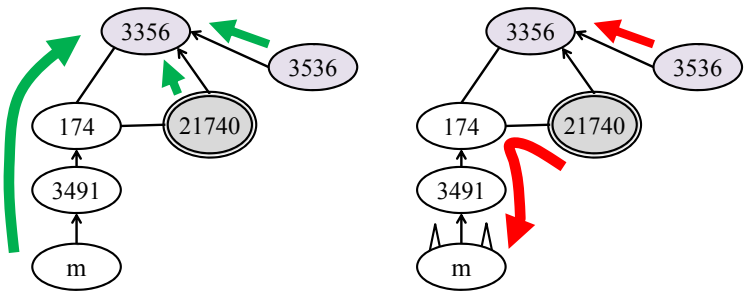


**Figure 9:** Example of a protocol downgrade attack when security is $2^{nd}$ or $3^{rd}$.

empirical results on these bounds for $\mathbf{LP}_0$ and all three security prioritization models in Section 5.5, and we then show in Section 5.6 how close we can get to these bounds with many different partial deployment scenarios. To verify robustness of our results with respect to other local preference models, namely $\mathbf{LP}_{k\in\{1,2,50\}}$, we have also studied the bounds on metric improvements as well as impact of various partial deployment scenarios with respect to those bounds in Sections 5.7 and 5.8.

We conclude this chapter with a high-level discussion of our results and their practical implications in Section 5.9.

## *5.2 Are Secure ASes Protected From Attacks?*

It seems natural to expect that any AS not deploying S*BGP would fall victim to a fixed-route attack (*e.g.,* prefix hijack). It also seems natural to expect that any AS deploying S*BGP with a secure route to be protected from such an attack, but, unfortunately, this cannot be guaranteed. In this section we discuss this troubling aspect of S*BGP in partial deployment [58].

### 5.2.1 Protocol Downgrade Attack

In a protocol downgrade attack, a source AS that uses a secure route to the legitimate destination under normal conditions, downgrades to an insecure bogus route during an attack as a result of its own routing policies [58].

Consider an example of this phenomenon in Figure 9. This figure depicts how AS

21740, a webhosting company, suffers a protocol downgrade attack, in the security $2^{nd}$ or $3^{rd}$ models. Under normal conditions (left), AS 21740 has a secure provider route directly to the destination Level 3 AS 3356, a Tier 1 ISP. AS 21740 does not have a peer route via AS 174 due to **Ex**. During the attack (right), $m$ announces that it is directly connected to Level 3, and so AS 21740 sees a bogus, insecure 4-hop peer route, via his peer AS 174. Note that AS 21740 has no idea that this route is bogus because it looks just like any other route that might be announced with legacy BGP. In the security $2^{nd}$ and $3^{rd}$ models, AS 21740 prefers an insecure peer route over a secure provider route, and will therefore downgrade to the bogus route during the attack.

Although such attacks have been considered before in [58], their significance with respect to partial S*BGP deployments has not been quantified. In Section 5.6, we show that protocol-downgrade attacks can be a serious problem, rendering even large partial deployments of S*BGP ineffective against attacks.

In the example of Figure 9 AS 21740 was too eager to give up its secure route in favor of the shorter insecure route during the attack. To mitigate this issue one could in principle apply an idea similar to route-flap damping to prevent such an attack, where AS 21740 might always prefer an *old* secure route over an *new* insecure route. Such a policy can be very tricky in partial S*BGP deployments, however, as network operators would have to damp out potentially many legitimate insecure routes that may be learned alongside secure routes under normal conditions, *i.e.,* not due to a routing attack.

### 5.2.2 When Can Protocol Downgrades Be Avoided?

We have demonstrated in Figure 9 that protocol downgrade attacks can happen in the security $2^{nd}$ and $3^{rd}$ models. We now show that they do not happen when security is $1^{st}$, The following result confirms that every AS $s$ that uses a secure route that

contains no attackers under normal conditions, will continue to use that secure route when attackers launch their attacks.

**Theorem 5.2.1.** *In the security $1^{st}$ model, for any local preference model $\boldsymbol{LP}_{k\in\mathbb{N}}$, destination AS $d$, set of fixed-route attackers $M$, and AS $s$ that, under normal conditions, has a secure route to $d$ that does not go through any attacker in $M$, $s$ will use a secure route to $d$ even when any attackers in $M$ run their attack.*

*Proof.* The theorem follows from the correctness of the algorithm in Section 4.5.4 for computing routes when security is $1^{st}$. Suppose the set of secure routes is $S$. Consider a source AS $s$ who has its secure route $R_s$ fixed during the FSStuBB, FSCR, FSPeerR, or FSPrvR subroutine of the algorithm in Section 4.5.4, when the set of secure ASes is $S$ during normal conditions when there is no attack. If $R_s$ does not contain any attacker in $M$, then $s$ will have its route fixed to exactly the same secure route $R_s$ during the FSStuBB, FSCR, FSPeeR, or FSPrvR subroutine of the algorithm in Section 4.5.4 when the set of secure ASes is $S$ and any attackers in $M$ run their attack. This follows because all routes that contain any attacker in $M$ must be fixed after the FSStuBB, FSCR, FSPeeR, or FSPrvR portions of the algorithm. This is because, by definition, all routes containing an attacker in $M$ that runs its attack must be insecure during the attack. With an inductive argument we can then show that all ASes on route $R_s$ will therefore be fixed to the same route that they used under normal conditions, and the theorem follows. □

While the theorem holds only if an attacker $m \in M$ is on AS $s$'s route, this is not a significant restriction because, otherwise, $m$ would not need to attack to attract traffic from $s$ to $d$ in the first place. Hence, in the security $1^{st}$ model the attacker's best hope is to attract a large fraction of the ASes that cannot route to the destination along a secure route.

## 5.3 How to Quantify Security Benefits?

To quantify the impact that a particular S*BGP deployment could have in terms of security benefits, we focus on the scenario where a single, fixed-route attacker AS $m$ attacks a single destination AS $d$. All ASes except $m$ use the routing policies we described in Chapter 4. The attacker $m$ wants to convince ASes to route to $m$, instead of the legitimate destination AS $d$ that is authorized to originate the prefix(es), that $m$ is trying to attack. It will do this by sending bogus AS-route information using legacy BGP.

We define a security metric as the average fraction of all ASes using legitimate routes to a destination being attacked by a single fixed-route attacker. The average is taken over all sources, all destinations and all attackers. Let us define it now more concretely.

Suppose $S$ is the the set of secure ASes deploying S*BGP and consider a fixed-route attacker $m$ that attacks a destination $d$ by announcing a bogus route to $d$. Let $H(m, d, S)$ be the number of *happy* source ASes that choose a legitimate route to $d$ instead of a bogus route to $m$. Our metric is:

$$H_{M,D}(S) = \frac{1}{|D|(|M|-1)(|V|-2)} \sum_{m \in M} \sum_{d \in D \setminus \{m\}} H(m, d, S)$$

The goal of attackers we consider in this chapter is to attract traffic from as many ASes as possible, and our metric therefore measures the average fraction of ASes that do not choose a route to the attacker. Observe that the value of this metric for any deployment scenario must always be a number between 0 and 1, and that it allows us to compare any deployment scenarios with respect to any secure routing model and any local preference model $\mathbf{LP}_{k \in \mathbb{N}}$ by averaging over fixed sets $D$ and $M$ (that are independent of $S$).

Additionally, this simple definition provides us with extra flexibility. Since we cannot predict where an attack will come from, or which ASes the attacker will target,

the metric averages over all attackers in a set $M$ and destinations in a set $D$. and we can choose $M$ and $D$ to be any subset of the ASes in the graph, depending on (i) where we expect attacks to come from, and (ii) which destinations we are particularly interested in protecting. When we want to capture the idea that all destinations are of equal importance, we average over all destinations (note that the "China's 18 minute mystery" of 2010 [34] fits into this framework well, since the hijacker targeted prefixes originated by a large number of (seemingly random) destination ASes) However, we can also zoom in on important destinations $D$ (*e.g.,* content providers [67, 24, 83]) by averaging strictly over those destinations. We can, analogously, zoom in on certain types of attackers $M$ by averaging over them only.

Our metric definition naturally leads to the following two questions with respect to any S*BGP deployment scenario:

1. How hard is it computationally to maximize the security metric, provided that only a fixed number of ASes could participate in the deployment?
2. As more ASes participate in the deployment, can the security metric be guaranteed to not decrease?

We address these questions next.

### 5.3.1 How Hard Is It to Decide Whom to Secure?

To address this question, consider the following computational problem, that we call *Max-$\kappa$-Security*: Given an AS graph, $G = (V, E)$, a specific attacker-destination pair $(m, d)$, where $m$ is a fixed-route attacker, and a parameter $\kappa > 0$, find a set $S$ of secure ASes of size $\kappa$ that maximizes the total number of happy ASes. Ideally, we would be able to efficiently (*e.g.,* in polynomial time in terms of the size of the problem) select the smallest set of ASes that maximizes the value of the metric. However, we next show that this cannot be guaranteed in general.

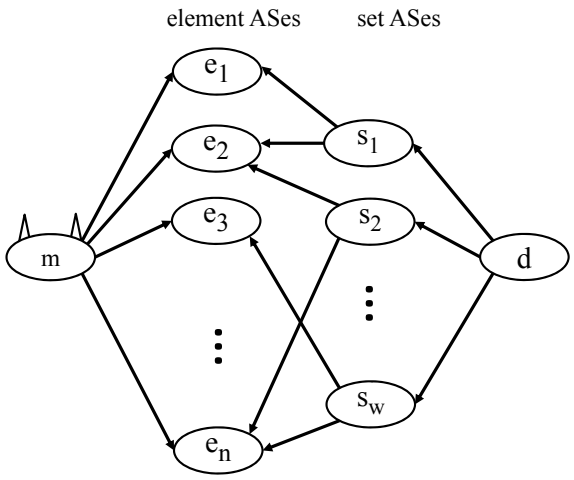


**Figure 10:** Gadget for proof of Theorem 5.3.2

**Theorem 5.3.1.** *Max-$\kappa$-Security is NP-hard for any security model, $1^{st}$, $2^{nd}$, or $3^{rd}$, and any local preference model $\boldsymbol{LP}_{k\in\mathbb{N}}$,*

To prove this result, let us consider a slightly different problem that we will call the *Decisional-$\kappa$-$\ell$-Security* problem (D$\kappa\ell$SP). Given an AS graph, a specific attacker-destination pair $(m, d)$, , where $m$ is a fixed-route attacker, a and parameters $\kappa > 0$ and $1 \leq \ell \leq |V|$, determine if there is a set of secure ASes $S$ of size $\kappa$ that results in at least $\ell$ happy ASes. Notice that this problem is in NP since we can check the number of happy ASes in polynomial time given the algorithms discussed in Section 4.5, and is certainly poly-time reducible to "Max-$\kappa$-Security". Therefore, the following theorem implies Theorem 5.3.1:

**Theorem 5.3.2.** *D$\kappa\ell$SP is NP-Complete for any security model $1^{st}$, $2^{nd}$, or $3^{rd}$, and any local preference model $\boldsymbol{LP}_{k\in\mathbb{N}}$.*

*Proof.* We present a poly-time reduction from the Set Cover Decisional Problem (SCDP). In SCDP, we are given a set $N$ with $n$ elements, a family $F$ of $w$ subsets of $N$ and an integer $\gamma \leq w$, and we must decide if there exist $\gamma$ subsets in the family $F$ that can cover all the elements in $N$.

We prove this result with respect to single fixed-route attacker announcing an edge $(m, d)$ (*i.e.,* a bogus route of length 1 to $d$) and then comment on how this proof could be extended for any fixed-route attackers. Our reduction is shown in Figure 10. For each element $e_i \in N$ in the SCDP instance, we create an AS $e_i$ in our D$k\ell$SP instance and connect it to the attacker via a provider-to-customer edge. For each subset $s_j \in F$, we create an AS $s_j$ in our our D$k\ell$SP instance and connect it to the destination $d$ via a provider-to-customer edge. We connect AS $e_i$ to AS $s_j$ via a provider-to-customer edge if $e_i \in s_j$ in the SCDP problem. Moreover, we require that every $e_i$'s has a tiebreak criteria **TB** that prefers the route through $m$ over any route through any $s_j$. Notice that the perceivable routes at every $e_i$ are of the same length and type, namely, two-hop customer routes. Finally, we let $\ell = n + w + 1$, and let $\kappa = n + \gamma + 1$.

Suppose that our SCDP instance has a $\gamma$-cover. We argue that this implies that our corresponding D$\kappa\ell$SP should be able to choose a set $S$ of $\kappa$ secure ASes that ensure that at least $\ell$ ASes are happy. The following set $S$ of secure ASes suffice: $S = \{d, e_1, ..., e_n\} \cup \{s_j | s_j \text{ is in the } \gamma \text{ cover}\}$. Notice that $S$ is of size $\kappa = n + \gamma + 1$, and results in exactly $\ell = n + w + 1$ happy ASes. This follows because $d$ is happy by definition, all the set ASes $s_1, ..., s_w$ are happy regardless of the choice of $S$, and all the element ASes $e_1, ..., e_n$ choose legitimate routes to the destination because they have secure routes to $d$ by construction.

On the other hand, suppose we are able to secure exactly $\kappa$ ASes while ensuring that $\ell$ ASes are happy. First, note that all the set ASes $s_1, ..., s_w$ and the destination AS are happy regardless of which ASes are secure. Next, note that if any of the $n$ element ASes $e_1, ..., e_n$ are insecure, then by construction it will choose a route to the attacker and be unhappy, and we will have less than $\ell$ happy ASes. Similarly, if the destination $d$ is insecure, by construction all of the element ASes will choose an insecure route to the attacker. Thus, if we secure all the element ASes and the

destination, we have $\kappa - 1 - n = \gamma$ remaining ASes to secure. By construction, these must be distributed amongst the set ASes, and thus we will have a $\gamma$-cover by construction.

Observe that this result holds in all three secure routing models and any local preference model $\mathbf{LP}_{k\in\mathbb{N}}$ because our reduction is agnostic to how ASes rank security as well as length of peer versus customer routes in their route preference decisions, since the perceivable routes at every element AS $e_i$ have the same length and type.

Finally, observer that to extend this proof for any fixed route attacker $m$, announcing a bogus route of length $x > 1$, we just need to modify the proof as follows. We replace $d$ with a sequence of $x - 1$ dummy ASes $(d_{x-1}, \cdots, d_1, d)$ such that all ASes $s_1, ..., s_w$ connect to $d_{x-1}$ via provider-to-customer edges instead of $d$ and each dummy AS in this sequence has a customer route to $d$. We then let $\ell = n + w + x$, and let $\kappa = n + \gamma + x$. □

To extend this result to multiple destinations $D$ and attackers $M$, we can show the hardness of the following variant of the "Max-k-Security" problem: given $G(V, E)$, sets $M, D \subseteq V$ and an integer $k$, the objective is to maximize the average number of happy ASes across all $(m, d)$ pairs in $M \times D$. The argument is the same as the above, except that now we create multiple copies of the $m$ and $d$ ASes (and their adjacent edges) in Figure 10, and let $M$ be the copies of the $m$ ASes and $D$ be the copies of the $d$ ASes.

The main practical implication of this result is that there is no guarantee that selecting the set of ASes on the Internet could be done efficiently. However, we provide a framework for binding the maximum and minimum possible metric values for any S*BGP deployment later in this chapter in Section 5.5.

### 5.3.2 Is Security Monotonic?

The most obvious expectation from S*BGP deployment is that the Internet should become more protected against attackers as more ASes adopt S*BGP. Unfortunately, however, this is not always the case. In this section we demonstrate that security in fact is not monotonic, in the sense that securing more ASes can actually make other ASes unhappy.

To explain this, we use an example toy network taken from the UCLA AS graph, where the destination (victim) AS $d$ is Pandora's AS40426 (a content provider) and the attacker $m$ is an anonymized Tier 2 network. We consider the network before and after a partial deployment of S*BGP $S$ and see how the set of happy ASes changes. Here $S$ consists of all 100 Tier 2s, all 17 content providers, and all of their stubs.

#### 5.3.2.1 Collateral Damages

In Figure 11 we show how AS 52142, a Polish ISP, suffers from collateral damage when security is $2^{nd}$. On the left, we show the network prior to S*BGP deployment. AS 52142 is offered two routes that are both insecure: a 3-hop route through his provider AS 5617 to the legitimate destination AS 40426, and a 5-hop bogus route to the attacker.

Attacker $m$ advertises a fake edge $(m, d)$ to all its neighbors. Note that although the route to $m$ is really 4 hops long, $m$ falsely claims to have a link to AS 40426 so AS 52142 thinks it is 5 hops long to $d$ via $m$. AS 52142 will choose the legitimate route because it is shorter. On the right, we show the network after S*BGP deployment. AS 5617 has become secure and now prefers the secure route through its neighbor Cogent AS 174. However, AS 5617's secure route is 5 hops long (right), significantly longer than the 2 hop route AS 5617 used prior to S*BGP deployment (left). Thus, after S*BGP deployment AS 52142 learns a 6-hop legitimate route through AS 5617, and a 5-hop bogus route. Since AS 52142 is insecure, it chooses the shorter route,

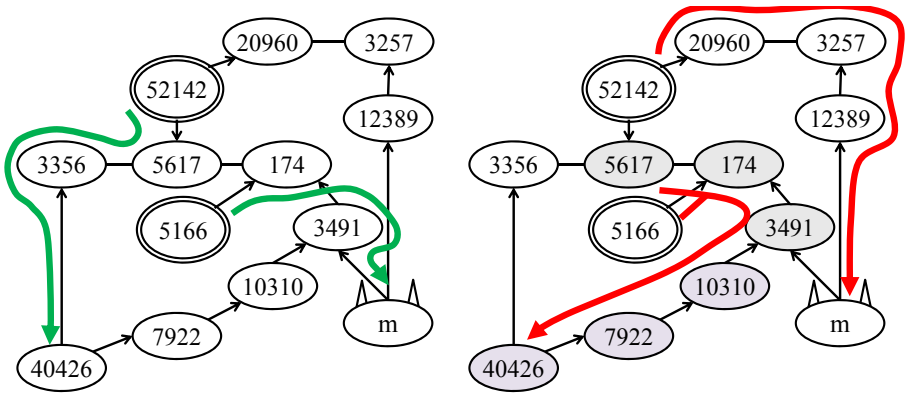


**Figure 11:** Example of collateral benefits and damages when security is $2^{nd}$.

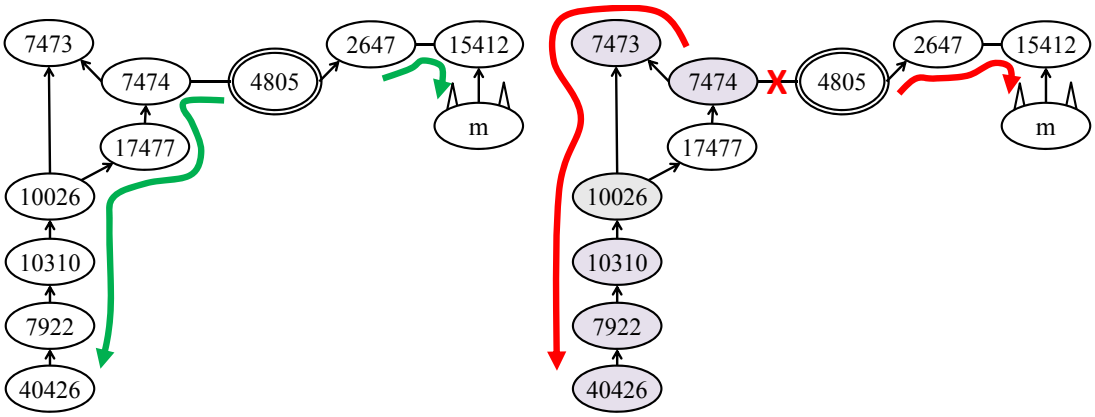


**Figure 12:** Example of collateral damages when security is $1^{st}$.

and becomes unhappy as collateral damage, which we define as follows.

**Definition 5** (Collateral Damage)**.** *A source AS $s \notin T$ experiences collateral damage from S*BGP deployment $T$ with respect to an attacker $m$ and destination $d$ if $s$ was happy when the ASes in $S$ are secure, but $s$ is unhappy when the ASes in $T$ are secure, and $T \supset S$.*

Figure 11 revealed that collateral damages can be caused by secure ASes choosing long secure routes over shorter insecure ones when security is $2^{nd}$. When security is $1^{st}$, collateral damages can also be caused by secure ASes choosing expensive secure routes over cheaper insecure ones.

Consider the network shown in Figure 12. We show how AS 4805, Orange Business in Oceania, suffers from collateral damage when security is $1^{st}$. On the left, we show the network prior to S*BGP deployment. Orange Business AS 4805 learns two routes: a legitimate route through its peer Optus Communications AS 7474, and a bogus route through its provider AS 2647. Again, attacker $m$ advertises a fake edge $(m, d)$ to all its neighbors. Since AS 4805 prefers peer routes over provider routes per our **LP** rule, it will choose the legitimate route and avoid the attack. On the right, we show what happens after S*BGP deployment. Now, Optus Communications AS 7474 has started using a secure route. However, this secure route is through its provider AS 7473. Observe that AS 7474 is no longer willing to announce a route to its peer AS 4805 as this would violate the export policy **Ex**. AS 4805 is now left with the bogus provider route through AS 2647, and becomes unhappy as collateral damage.

#### 5.3.2.2 Can Collateral Damage Be Avoided?

Observer that the collateral damage in Figure 11 above occurs because AS 5617 prefers a longer secure route over a shorter insecure route. Such travesty could not have occurred were security $3^{rd}$, and, in fact, we can show that collateral damages do not happen in the security $3^{rd}$ model in general.

**Theorem 5.3.3.** *In the security $3^{rd}$ model, for any $\boldsymbol{LP}_{k\in\mathbb{N}}$, if an AS $s$ has a route to a destination $d$ that avoids a fixed-route attacker $m$ when the set of secure ASes is $S$, then $s$ has a route to a destination $d$ that avoids attacker $m$ for every set of secure ASes in $T \supset S$.*

*Proof.* The theorem follows from the correctness of our algorithm for computing routing outcomes when security is $3^{rd}$ described in Section 4.5.2. First, with an inductive argument we can show that every AS $s$ that the algorithm fixes to a secure route in deployment $S$ is also fixed to a secure route in $T$. It then follows that all such ASes stabilize to a legitimate route in both $S$ and $T$. We next argue that every AS $s$ that

the algorithm fixes to an insecure legitimate route in $S$ is also fixed to a legitimate route in $T$. There are two cases: (a) if $s$ is fixed to a secure route in $T$, it uses a legitimate route, (b) otherwise, with an inductive argument we can show that the algorithm computes the same next hop $\mathsf{Nxt}(s)$ for $s$ in both deployments $T$ and $S$, and since the route was legitimate in $S$, it will be legitimate in $T$ as well. □

The security $3^{rd}$ model is our only monotone model, in the sense that more secure ASes cannot result in fewer happy ASes, so the metric $H_{M,D}(S)$ grows monotonically in $S$. On this positive note, we next show that collateral benefits can happen in all three secure routing models.

#### 5.3.2.3 Collateral Benefits

In this section we show that insecure ASes can also become happy as a collateral benefit, because other ASes obtained secure routes.

In Figure 11 we show how AS 5166, with the Department of Defense Network Information Center, obtains collateral benefits when its provider AS 174, Cogent, deploys S*BGP. On the left, we show the network prior to the deployment of S*BGP. Focusing on Cogent AS 174, we see that it falls victim to the attack, choosing a bogus route through its customer AS 3491. As a result, AS 5166 routes to the attacker as well. On the right, we show the network after S*BGP deployment. Now, both AS 174 and AS 3491 are secure, and choose a longer secure customer route to the legitimate destination. As a result, AS 5166, which remains insecure, becomes happy as a collateral benefit.

**Definition 6** (Collateral Benefit)**.** *A source AS $s \notin S, T$ experiences collateral benefit from an S*BGP deployment $T$ with respect to an attacker $m$ and destination $d$ if $s$ is unhappy when the ASes in $S$ are secure, but $s$ is happy when the ASes in $T$ are secure, and $T \supset S$.*

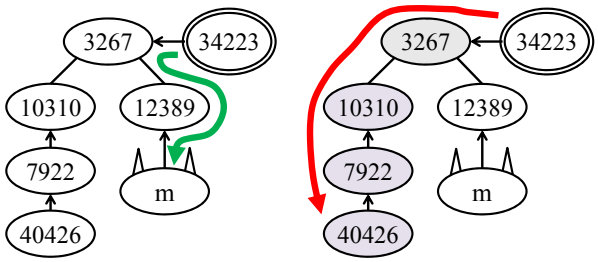


**Figure 13:** Example of collateral benefits when security is $3^{rd}$.

**Table 1:** S*BGP partial deployment phenomena in different security models.

| Security model | $1^{st}$ | $2^{nd}$ | $3^{rd}$ |
|---|---|---|---|
| **Protocol Downgrade Attacks** | $X$ | ✓ | ✓ |
| **Collateral Benefits** | ✓ | ✓ | ✓ |
| **Collateral Damages** | ✓ | ✓ | $X$ |

Collateral benefits are possible in all three routing policy models. We now show an example of collateral benefits when security is $3^{rd}$. Consider the example of Figure 13. We show how AS34223, a Russian ISP, obtains collateral benefits in the security $3^{rd}$ model. The left subfigure shows how AS34223 and its provider AS3267 react to the attack before S*BGP deployment. Attacker $m$ again advertises a fake edge $(m, d)$ to all its neighbors. AS3267 learns two peer routes of equal length – one bogus route to the attacker $m$ and one legitimate route to Pandora's AS 40426. AS3267 then tiebreaks in favor of the attacker, so both AS3267 and his customer AS34223 become unhappy. On the right, we show what happens after partial S*BGP deployment. AS3267 has a secure route to Pandora of equal length and type as the insecure route to $m$, so AS3267 chooses the secure route, and his insecure customer AS34223 becomes happy as a collateral benefit.

The reader is invited to refer to Table 1 for a summary of the various tricky phenomena that could result from partial deployments of S*BGP.

## *5.4 Empirical Methods*

To perform empirical evaluations of various S*BGP deployments with respect to our metric, we simulated routing outcomes, each requiring time $O(|V|)$, over all possible $|M||D|$ attacker and destination pairs. We sometimes take $M = D = V$ so that our computations approach $O(|V|^3)$. The algorithms that we used to compute routing outcomes were presented in the previous chapter in Section 4.5. We explain the details of their implementation as well as our threat model and the empirical AS-level Internet topologies that we used in our studies below.

### 5.4.1 Simulations Explained

For each destination $d$, our simulations compute the following.

1. The S*BGP routing outcome for various S*BGP routing models and for every deployment set $S$ considered in this thesis to enable computations that quantify protocol downgrade attacks;
2. The BGP routing outcome with respect to every possible $(m, d)$ pair when $S = \emptyset$ to determine which ASes are happy and unhappy in the baseline scenario;
3. The S*BGP routing outcome for every possible $(m, d)$ pair for various S*BGP routing models and for every deployment set $S$ considered in this thesis to compute the metric improvements, to detect phenomena like collateral benefits and damages, and to quantify protocol downgrade attacks;

To do this, we use the algorithms in Sections 4.5.2-4.5.4, where the we execute the FStuBB, FSIStuBB, FSStuBB, FCR, FSCR, FPeeR, FSPeeR, FPrvR, and FSPrvR subroutines using breath-first searches. The overall complexity of our simulations is therefore $O(|M||D|(|V| + |E|)$ for each deployment $S$. We optimize the running time of our simulations in two ways.

**Reusing Information:** Instead of running multiple computations from scratch our simulations re-use information and pass it on from one computation to the next.

**Parallelization:** We run these computations in parallel across all destinations $d$. Our code was written in C++ and parallelization was achieved with MPI on a BlueGene, Blacklight, and Gordon supercomputing platforms.

### 5.4.2 Threat Model

For our empirical evaluations we consider a future scenario where RPKI and origin authentication are deployed, and the challenge is in engineering global S*BGP adoption. We therefore disregard attacks that are prevented by origin authentication, such as prefix- and subprefix-hijacks [25, 14, 24, 34, 77, 19] and instead focus on attacks that are effective even in the presence of origin authentication, because these are the attacks that S*BGP was designed to prevent.

Previous studies on S*BGP security [46, 12, 28] focused the scenario with full S*BGP deployment assuming that any secure AS would reject any insecure route that it receives. As we emphasized earlier in this chapter, this assumption is invalid in the context of a partial deployment of S*BGP, where S*BGP has to coexist alongside BGP. In this setting, some destinations may only be reachable via insecure routes. Moreover, even a secure AS may prefer to use an insecure route that was announced via BGP for economic or performance reasons [85, 46]. This is possible when security is $2^{nd}$ or $3^{rd}$, and we have shown that it can result in protocol downgrades in Section 5.2.

#### 5.4.2.1 The One-Hop Hijack Attack

Recall that we focus on the scenario where a single attacker AS $m$ attacks a single destination AS $d$. The attacker $m$'s objective is to maximize the number of source ASes that send traffic via $m$, rather than to the legitimate destination $d$ avoiding $m$. This goal captures $m$'s incentive to attract traffic from as many source ASes as possible, for the purposes of eavesdropping, tampering and even dropping it [14, 46, 45].

It can be shown that it is NP-hard for $m$ to determine a bogus route to export to each neighbor that maximizes the number of source ASes it attracts, based on a similar result shown in [46]. Thus, in our empirical evaluations we consider the next simplest, yet still very disruptive [14, 46], attack, in which $m$ that is not neighbors with $d$, pretends to be directly connected to $d$. Since there is no need to explicitly include IP prefixes for our studies in this chapter, this results in a single attacker AS $m$ announcing the bogus AS-level route $(m, d)$ using legacy, insecure BGP to all of its neighbor ASes. We call this attack the *one-hop hijack attack*. Recall that examples of this attack were shown in Figures 9, 11 and 13.

Observe that using our provable security framework developed in Chapter 3, it is easy to argue that this attack can be prevented by fully deployed SoBGP, S-BGP and BGPSEC. With SoBGP, the attacker claims to have an edge to $d$ that does not exist in the graph. With S-BGP or BGPSEC the attacker claims to have learned a route $(m, d)$ that $d$ never announced. Hence, throughout this chapter we will refer to all these protocols with S*BGP in our analysis. Note that the bogus route is announced via legacy BGP, so the recipient ASes cannot validate it with S*BGP, and thus will accept it without suspicions.

### 5.4.3 Empirical AS-level Internet Topologies

All of our simulations were performed over the UCLA AS-level topology from 24 September 2012 [31]. We preprocessed the empirical topology by (1) renaming all 4-byte ASNs in more convenient way, and (2) recursively removing all ASes that had no providers that had low degree and were not Tier 1 ISP's. The resulting graph had 39056 ASes, 73442 customer-provider links and 62129 peer-to-peer links.

We will sometimes in this chapter refer to the tiers of ASes [36] in Table 2. The list of 17 content providers (CPs) in Table 2 was obtained from recent empirical work on interdomain traffic volumes [64, 66, 67, 93, 10].

**Table 2:** Tier Classification of ASes on the Internet.

| | |
|---|---|
| **Tier 1** | 13 ASes with high customer degree & no providers |
| **Tier 2** | 100 top ASes by customer degree & with providers |
| **Tier 3** | Next 100 ASes by customer degree & with providers |
| **CPs** | 17 Content providers: AS 15169, 8075, 20940, 22822, 32934, 15133, 16265, 16509, 2906, 23286, 40428, 714, 10310, 38365, 14907 13414, 4837 |
| **Small CPs** | Top 300 ASes by peering degree (other than Tier 1, 2, 3, and CP) |
| **Stubs-x** | ASes with peers but no customers |
| **Stubs** | ASes with no customers & no peers |
| **SMDG** | Remaining ASes |

Because empirical AS graphs often miss many of peer-to-peer links in Internet eXchange Points (IXP) [91, 11, 8], we constructed a second graph where we augmented the UCLA graph with over 550K peer-to-peer edges between ASes listed as members of the same IXP, on September 24, 2012, on voluntary on-line sources, *e.g.,* IXPs websites, EuroIX, Peering DB, Packet Clearing House, *etc.*. Our list contained 332 IXPs and 10,835 mappings of member ASes to IXPs. After connecting every pair of ASes that are present in the same IXP, and were not already connected in our original UCLA AS graph, with a peer-to-peer edge, our graph was augmented with 552,933 extra peering links. As it may be the case that not all ASes at an IXP peer with each other [8], our augmented graph is an upper bound on the number of missing links in this empirical AS topology. For robustness, we also tested our empirical results with respect to this augmented topology.

## *5.5 Invariants to Deployment*

Given the vast number of possible configurations for a partial deployment of S*BGP, we present a framework for exploring the security benefits of S*BGP vis-a-vis origin authentication, while making no assumptions about which ASes are secure. We show how to determine an upper bound on security benefits available with any S*BGP deployment for any routing model in Section 5.5.3.1 and then compare it to the

**Table 3:** Status of source $s$ when $m$ attacks $d$.

| | |
|---|---|
| **Happy** | Chooses a legitimate secure/insecure route to $d$. |
| **Unhappy** | Chooses a bogus insecure route to $m$. |
| **Immune** | Happy *regardless of which ASes are secure.* |
| **Doomed** | Unhappy *regardless of which ASes are secure.* |
| **Protectable** | Neither immune nor doomed. |

security benefits available with origin authentication in Sections 5.5.2 and 5.5.4. All the results shown in this section correspond to the $\mathbf{LP}_0$ local preference model, and we show corresponding plots for other local preference models in Section 5.7

### 5.5.1 Tiebreaking and Computing Bounds on The Metric

Recall from Chapter 4 that our model fully determines an AS's routing decision up to the tiebreak step **TB** of its routing policy. Since computing $H_{M,D}(S)$ only requires us to distinguish between *happy* and *unhappy* ASes, the tiebreak step matters only when a source AS $s$ has to choose between (1) an insecure route(s) to the legitimate destination $d$ (that makes it happy), and (2) an insecure bogus route(s) to $m$ that makes it unhappy. Importantly, $s$ has no idea which route is bogus and which is legitimate, as both of them are insecure. Therefore, to avoid making uninformed guesses about how ASes choose between equally-good insecure routes, we will compute upper and lower bounds on our metric. To get a lower bound, we assume that every AS $s$ in the aforementioned situation will always choose to be unhappy (*i.e.,* option (2)). The upper bound is obtained by assuming $s$ always chooses to be happy (*i.e.,* option (1)).

To compute upper and lower bounds on the set of happy ASes. We use the three algorithms in Sections 4.5.2-4.5.4 for a given attacker-destination pair $(m, d)$, set of secure ASes $S$ and routing model. To do this, each algorithm records, for every AS discovered in the BFS computation, whether (1) all routes in its $\mathsf{BPR}$ at that iteration

lead to the destination, or (2) all these routes lead to the attacker or (3) some of these routes lead to the destination and others to the attacker. The number of ASes in the $1^{st}$ category is then set to be a lower bound on the number of happy ASes. The total number of ASes in the $1^{st}$ and $3^{rd}$ category is set to be an upper bound on the number of happy ASes.

The correctness of this approach follows from the correctness of our algorithms, shown in Section 4.5.5, and the fact that all the routes in the $\mathsf{BPR}^r(v_i)$ of an AS $v_i$ at iteration $r$ have the same length, type, and are either all secure or insecure, so the **TB** criteria completely determines which of these routes are chosen. As such, ASes in the $1^{st}$ category choose legitimate routes (and are happy) regardless of the **TB** criteria, ASes in the $2^{nd}$ category choose attacked routes (and are unhappy) regardless of the **TB** criteria, and whether ASes in the $3^{rd}$ category are happy completely depends on the **TB** criteria.

### 5.5.2 Origin Authentication Gives Good Security

At this point, we could compute the metric for various S*BGP deployment scenarios, show that most source ASes are happy, argue that S*BGP has improved security, and conclude our analysis. This, however, would not give us the full picture, because it is possible that most of the happy ASes would have been happy even if S*BGP had not been deployed. Thus, to understand if the juice is worth the squeeze, we need to ask how many more attacks are prevented by a particular S*BGP deployment scenario, relative to those already prevented by RPKI with origin authentication. More concretely, we need to compare the fraction of happy ASes before and after the ASes in $S$ deploy S*BGP. To do this, we compare the metric for a deployment scenario $S$ against the baseline scenario, where RPKI and origin authentication are in place, but no AS has adopted S*BGP, so that the set of secure ASes is $S = \emptyset$.

In [46], the authors evaluated the efficacy of origin authentication against attacks

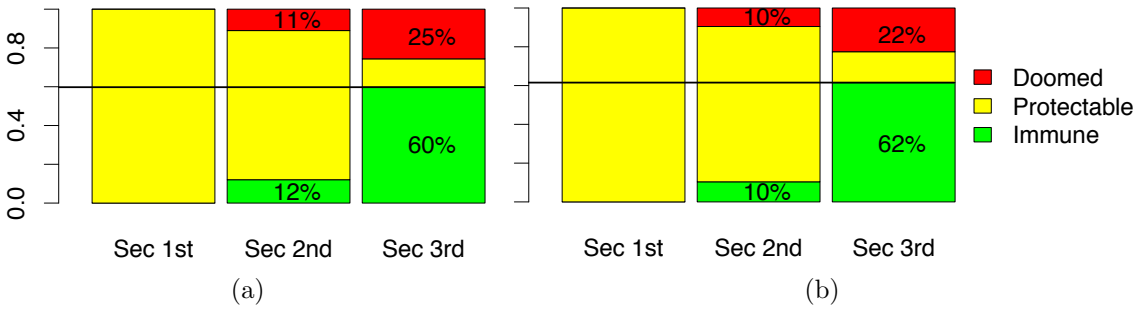


**Figure 14:** Partitions for the (a) UCLA (b) and IXP-augmented topologies.

that it was not designed to prevent — namely, the one-hop hijack attack of Section 5.4.2.1. They randomly sampled pairs of attackers and destinations and plotted the distribution of the fraction of unhappy source ASes (ASes that route through the attacker, see Table 3). Figure 3 of [46] shows that attacker is able to attract traffic from less than half of the source ASes in the AS graph, on average. We now perform a computation and obtain a result that is similar in spirit. Rather than randomly sampling pairs of attackers and destinations as in [46], we instead compute a lower bound on our metric over all possible attackers and destinations. We find that $H_{V,V}(\emptyset) \geq 60\%$ on the basic UCLA graph, and $H_{V,V}(\emptyset) \geq 62\%$ on our IXP-augmented graph, shown in Figure 14.

It is striking that both our and [46]'s result indicate more than half of the AS graph is already happy even before S*BGP is deployed. To understand why this is the case, recall that with origin authentication, an attacking AS $m$ must announce a bogus route $(m, d)$ that is one hop longer than the route$(d)$ announced by the legitimate destination AS $d$. When we average over all $(m, d)$ pairs and all the source ASes, bogus routes through $m$ will appear longer, on average, than legitimate routes through $d$. Since route length plays an important role in route selection, on average, more source ASes choose the legitimate route.

### 5.5.3 Does S*BGP Provide Better Security Than Origin Authentication?

How much further can we get with a partial deployment of S*BGP? We now obtain bounds on the improvements in security that are possible for a given routing policy model, but for any set $S$ of secure ASes.

We can obtain these bounds thanks to the following crucial observation. ASes can be partitioned into three distinct categories with respect to each attacker-destination pair $(m, d)$. Some ASes are *doomed* to route through the attacker regardless of which ASes are secure. Others are *immune* to the attack regardless of which ASes are secure. Only the remaining ASes are *protectable*, in the sense that whether or not they route through the attacker depends on which ASes are secure (see Table 3).

To bound our metric $H_{M,D}(S)$ for a given routing policy model (*i.e.,* security $1^{st}$, $2^{nd}$, or $3^{rd}$) and across all partial-deployment scenarios $S$, we first partition source ASes into categories — doomed, immune, and protectable — for each $(m, d)$ pair and each routing policy model. By computing the average fraction of immune ASes across all $(m, d) \in M \times D$ for a given routing model, we get a lower bound on $H_{M,D}(S)$ $\forall S$ and that routing model. We similarly get an upper bound on $H_{M,D}(S)$ by computing the average fraction of ASes that are not doomed.

#### 5.5.3.1 Partitions: Doomed, Protectable and Immune

Let us return to Figure 9 to explain our partitioning.

**Definition 7** (Doomed)**.** *A source AS $s$ is doomed with respect to pair $(m, d)$ if $s$ routes through $m$ no matter which set $S$ of ASes is secure.*

For example, AS 174 in Figure 9 is doomed when security is $2^{nd}$ (or $3^{rd}$). If security is $2^{nd}$ (or $3^{rd}$), AS 174 always prefers the bogus customer route to the attacker over a (possibly secure) peer route to the destination AS 3356, for every $S$.

**Definition 8** (Immune)**.** *A source AS $s$ is immune with respect to pair $(m, d)$ if $s$*

*will route through d no matter which set S of ASes is secure.*

For example, AS 3536 in Figure 9 is one example. This single-homed stub customer of the destination AS 3356 can never learn a bogus route in any of our security models. When security is $2^{nd}$ or $3^{rd}$, another example of an immune AS is AS 10310 in Figure 11. Its customer route to the legitimate destination AS 40426 is always more attractive than its provider route to the attacker in these models.

**Definition 9** (Protectable)**.** *AS s is protectable with respect to pair* $(m, d)$ *if it can either choose the legitimate route to d, or the bogus one to m, depending on S (*i.e., *s is neither doomed nor immune).*

For example, with security $1^{st}$, AS 174 in Figure 9 becomes protectable. If it has a secure route to the destination AS 3356, AS 174 will choose it and be happy, but otherwise, it will choose the bogus route to $m$.

#### 5.5.3.2 Which ASes are Protectable?

The intuition behind the following partitioning of ASes is straightforward. We discuss the subtleties involved in proving whether an AS is doomed or immune in Section 5.5.9.

**Security $1^{st}$:** Here, we suppose that all ASes are protectable. The few exceptions (*e.g.,* the single-homed stub of Figure 9) have little impact on the count of protectable ASes.

**Security $2^{nd}$:** Here, an AS is doomed if it has a route to the attacker with better local preference **LP** than every available route to the legitimate destination. For example, the bogus customer route offered to AS 174 in Figure 9 has higher **LP** than the legitimate peer route. An immune AS has a route to the destination that has higher **LP** than every route to the attacker. For protectable AS, its best available routes to the attacker and destination have exactly the same **LP**.

**Security $3^{rd}$:** Here, a doomed AS has a route to $m$ with (1) better **LP** OR (2) equal **LP** and shorter length **SP**, than every available route to $d$. The opposite holds for an immune AS. A protectable AS has best available routes to $m$ and $d$ with equal **LP** and route length **SP**.

### 5.5.4 Bounding Security for All Deployments

For each routing model, we found the fraction of doomed, protectable, and immune source ASes for each attacker destination pair $(m, d)$, and took the average over all $(m, d) \in V \times V$. We used these values to get upper and lower bounds on $H_{V,V}(S)$ for all deployments $S$, for each routing model.

The colored parts of each bar in Figure 14 represent the average fraction of immune, protectable, and doomed source ASes, averaged over all $O(|V|^2)$ possible pairs of attackers and destinations. Since $H_{V,V}(S)$ is an average of the fraction of happy source ASes over all pairs of attackers and destinations, the upper bound on the metric $H_{V,V}(S)$ $\forall S$ is the average fraction of source ASes that are *not* doomed. The upper bound on the metric $H_{V,V}(S)$ $\forall S$ is therefore: $\approx 100\%$ with security $1^{st}$, 89% with security $2^{nd}$, and 75% with security $3^{rd}$. Figure 14(b), the same figure computed on our IXP-augmented topology, looks almost exactly the same, with the proportions being $\approx$ 100%, 90% and 78%. Meanwhile, the heavy solid line is the lower bound on the metric $H_{V,V}(\emptyset)$ in the baseline setting where $S = \emptyset$ and there is only origin authentication. In Section 5.5.2 we found that $H_{V,V}(\emptyset) = 60\%$ (and 62% for the IXP-edge-augmented graph shown in Figure 14(b)). Therefore, we can bound the maximum change in our security metric $H_{V,V}(S)$ $\forall S$ for each routing policy model by computing the distance between the solid line and the boundary between the fraction of doomed and protectable ASes. Below we discuss our findings.

Figure 14 demonstrates that the maximum gains over origin authentication that

are provided by the security $3^{rd}$ model are quite slim — at most 15% — regardless of which ASes are secure. This follows because the upper bound on the metric $H_{V,V}(S) \leq 75\%$ for any $S$ while the lower bound on the baseline setting is $H_{V,V}(\emptyset) \geq 60\%$. Moreover, these are the *maximum* gains $\forall S$. Similar exercise shows that the maximum improvement for the IXP-augmented graph is 16%. Note that in a realistic S*BGP deployment, however, the gains are likely to be much smaller. This result is disappointing, since the security $3^{rd}$ model is likely to be the most preferred by network operators, but it is not especially surprising. S*BGP is designed to prevent route shortening attacks; however, in the security $3^{rd}$ model ASes prefer short (possibly bogus) insecure routes over a long secure routes, so it is natural that this model realizes only minimal security benefits.

Figure 14 confirms, on the other hand that the maximum gains over origin authentication are better when security is $2^{nd}$: $89 - 60 = 29\%$ for the UCLA graph and $90 - 62 = 28\%$ for the IXP-augmented graph. We discuss whether these gains can be realized in realistic partial-deployment scenarios in Section 5.6.

Note that the fraction of immune ASes in the security $2^{nd}$ (12%) and $1^{st}$ ($\approx 0\%$) models is strangely lower than the fraction of happy ASes in the baseline scenario (60%). Recall that in Section 5.3.2.1 we explained this counterintuitive observation by showing that more secure ASes can sometimes result in fewer happy ASes, a phenomenon we dubbed collateral damage. We also showed that such phenomena occur only in the security $1^{st}$ and $2^{nd}$ models. It is this negative side effect that leads to the decrease in the number of immune ASes.

### 5.5.5 Robustness to Destination Tier

Thus far, we have been averaging our results over all possible attacker-destination pairs in the graph. However, some destination ASes might be particularly important to secure, perhaps because they source important content (*e.g.,* the content provider

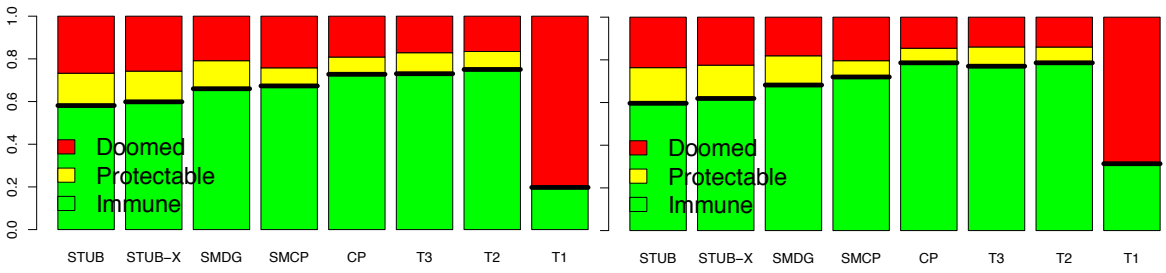


**Figure 15:** Partitions by destination tier when security is $3^{rd}$ for the (a) UCLA and (b) IXP-augmented topologies.

ASes (CPs)) or transit large volumes of traffic (the Tier 1 ASes). As such, we broke down the metric over destinations in each tier in Table 2.

In Figure 15 we show the partitioning into immune, protectable and doomed ASes in the security $3^{rd}$ model, but this time averaged individually over all destinations in each tier, and all possible attackers $V$. The thick horizontal line over each vertical bar again shows the corresponding lower bound on our metric $H_{V,\text{Tier}}(\emptyset)$ when no AS is secure. Apart from the Tier 1s, that we will discuss next, we observe similar trends as in Section 5.5.4, with the improvement in security ranging from $8-15\%$ for all tiers. The same holds for the security $2^{nd}$ model, shown in Figure 16(a). We observe a similar trend in our partition results for the IXP-augmented topology shown in Figures 15(b) and 16(b).

### 5.5.6 It is Difficult to Protect Tier 1 Destinations

Strangely enough, Figure 15(a) shows that when Tier 1 destinations are attacked in the security $3^{rd}$ model, the vast majority ($\approx 80\%$) of ASes are doomed, and only a tiny fraction are protectable. The same holds when security is $2^{nd}$ in Figure 16(a). Therefore, in these models, S*BGP can do little to blunt attacks on Tier 1 destinations. We observer a similar result for the IXP-augmented topology.

How can it be that Tier 1s, the largest and best connected (at least in terms of customer-provider edges) ASes in our AS graph, are the most vulnerable to attacks?

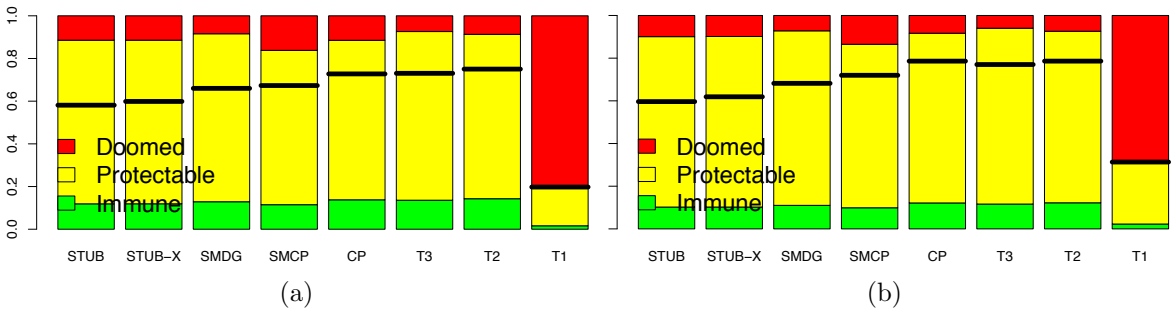


**Figure 16:** Partitions by destination tier when security is $2^{nd}$ for the (a) UCLA and (b) IXP-augmented topologies.

Ironically, it is the Tier 1s' very connectivity that harms their security. Because the Tier 1s are so well-connected, they can charge most of their neighbors for Internet service. As a result, most ASes reach the Tier 1s via costly provider routes that are the least preferred type of route according to the **LP** step in our routing policy models. Meanwhile, it turns out that when a Tier 1 destination is attacked, most source ASes will learn a bogus route to the attacker that is *not* through a provider, and is therefore preferred over the (possibly secure) provider route to the T1 destination in the security $2^{nd}$ or $3^{rd}$ models. In fact, this is exactly what lead to the protocol downgrade attack on the Tier 1 destination AS 3356 in Figure 9. In Section 5.6 we will confirm that this is a serious hurdle to protecting Tier 1 destinations.

### 5.5.7 Which Attackers Cause the Most Damage?

Let us now break things down by the type of the attacker, to get a sense of type of attackers that S*BGP is best equipped to defend against.

In Figure 17 we bucket our counts of doomed, protectable, and immune ASes for the security $3^{rd}$ model by the attacker type in Table 2, for all $|V|^2$ possible attacker-destination pairs. As the degree of the attacker increases, its attack becomes more effective. The number of immune ASes steadily decreases, and the number of doomed ASes correspondingly increases, as the tier of the attacker grows from stub to Tier 2. Meanwhile, the number of protectable ASes remains roughly constant across tiers.

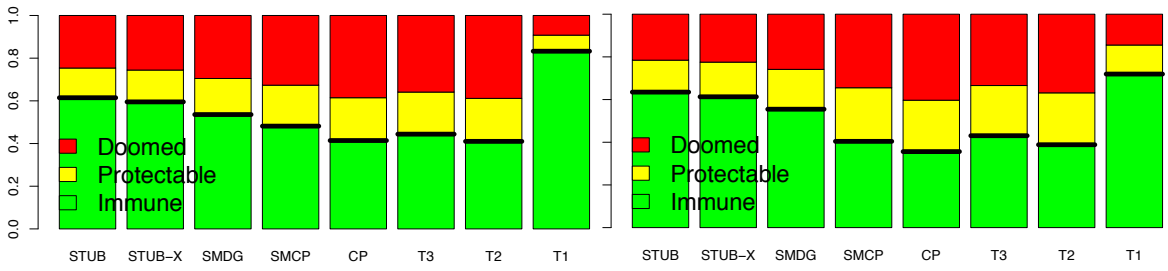


**Figure 17:** Partitions by attacker tier when security is $3^{rd}$ for the (a) UCLA and (b) IXP-augmented topologies.

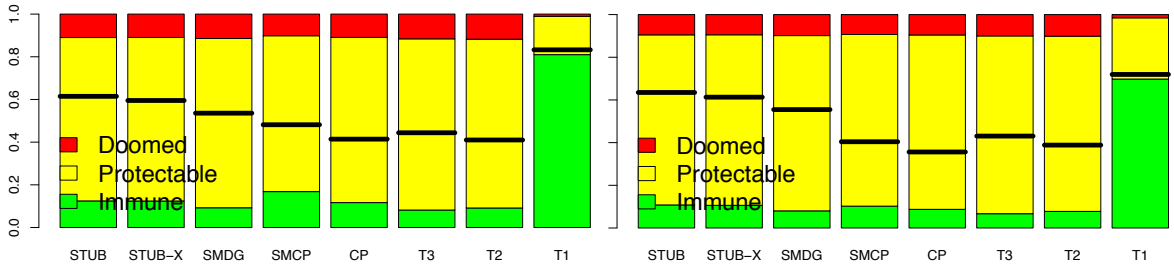


**Figure 18:** Partitions by attacker tier when security is $2^{nd}$ for the (a) UCLA and (b) IXP-augmented topologies.

The striking exception to this trend is that the Tier 1 attackers are significantly less effective than even the lowest degree (stub) attackers. While this observation might seem unnatural at first, there is a perfectly reasonable explanation. When a Tier 1 attacks, its bogus route will look like a provider route from the perspective of most other source ASes in the graph. Because the **LP** step of our routing model depreferences provider routes relative to peer and customer routes, the Tier 1 attacker's bogus route will be less attractive than any legitimate route through a peer or provider, and as such, most ASes will be immune to the attack. The same observations hold when security is $2^{nd}$ (see Figure 18), and we also observer the same trend for the IXP-augmented topology.

### 5.5.8 Which Sources Benefit the Most From S*BGP?

However, before we completely give up on the Tier 1s obtaining any benefit from S*BGP, we reproduced Figures 15-18 but this time, bucketing the results by the tier of source. We found that each source tier, including the Tier 1s, has roughly the same average number of doomed (25%), immune (60%), and protectable (15%) ASes. It follows that, while S*BGP cannot protect Tier 1 destinations from attack, S*BGP still has the potential to prevent a Tier 1 sources from choosing a bogus route. We also observer the same trend for the IXP-augmented topology.

### 5.5.9 Computing Partitions

In this section we describe how we compute the sets of immune, doomed, and protectable ASes for any attacker-destination pair $(m, d)$ with respect to all of our routing models. To do this, we set $S = \emptyset$ and compute the BGP routing outcome for that $(m, d)$ pair using the algorithm described in Section 4.5.2.

#### 5.5.9.1 Computing Partitions: Security $3^{rd}$

To determine the partitions for the security $3^{rd}$ model, this algorithm records, for every AS discovered in the BFS computation whether (1) all routes in its BPR set at that iteration lead to the destination, or (2) all these routes lead to the attacker or (3) some of these routes lead to the destination and others to the attacker. We classify ASes in the $1^{st}$ category as immune, ASes in the $2^{nd}$ category as doomed, and ASes in the $3^{rd}$ category as protectable. The following result shows that this is consistent with our definitions of immune, doomed, and protectable ASes in Section 5.5.3.1 for the security $3^{rd}$ model.

**Corollary 5.5.1.** *In the security $3^{rd}$ routing model, for any $\boldsymbol{LP}_{k\in\mathbb{N}}$, for any destination $d$, attacker $m$, source $s$, and deployment $S \subseteq V$, $s$ will stabilize to a route of the same type and length as any route in $\mathsf{BPR}(s, \emptyset, m, d)$.*

*Proof.* This follows from the correctness of our algorithm for computing routes in the security $3^{rd}$ model described in Section 4.5.2. Note that because in the security $3^{rd}$ model route security is prioritized below route length, all routes in $\mathsf{BPR}^r(s)$ must be contained in $\mathsf{BPR}(s, \emptyset, m, d)$, where $\mathsf{BPR}^r(s)$ is the set of best perceivable routes of $s$ during iteration $r$ of the subroutine FStuBB($k$), FCR($\infty$), FPeeR($\infty$) or FPrvR of our algorithm, when $\mathsf{BPR}(s, S, m, d)$ contains customer or peer routes of length at most $k$, customer routes of length greater than $k$, peer routes of length greater than $k$, or provider routes respectively. Recognize that by the correctness of our algorithm, $s$ must stabilize to a route in $\mathsf{BPR}^r(s)$ for some iteration $r$ of exactly one of these subroutines.

Therefore, any $s$ that has customer routes in $\mathsf{BPR}(s, \emptyset, m, d)$ will be fixed to a customer route of length at most or greater $k$ in the FStuBB($k$) or FCR($\infty$) subroutine respectively, for any choice of $S$. Similarly, if $s$ has peer routes in $\mathsf{BPR}(s, \emptyset, m, d)$, it will be fixed to a peer route of length at most or greater than $k$ in the FStuBB($k$) or FPeeR($k$) subroutine respectively, for any choice of $S$. Finally, if $s$ has provider routes in $\mathsf{BPR}(s, \emptyset, m, d)$, it will be fixed to a provider route FPrvR subroutine, for any choice of $S$. Therefore, the type of the route will be fixed to the same type as that of the $\mathsf{BPR}(s, \emptyset, m, d)$ for all $S$. Moreover, when we choose to fix the route of $s$ in the appropriate subroutine, we do so by selecting $s$ with shortest routes out of all the sources that have not been fixed, and regardless of $S$, so it follows that the length of the route will be the same for all $S$. □

Corollary 5.5.1 tells us that for determining whether $s$ is immune, doomed or protectable in security $3^{rd}$ model, it is sufficient to keep track of all the routes of the best type and shortest length of $s$, *i.e.,* all the routes in $\mathsf{BPR}(s, \emptyset, m, d)$, because $s$ is guaranteed to stabilize to one of these routes. Therefore, if all such routes are legitimate or attacked, then $s$ will always stabilize to a legitimate or attacked route, under any S*BGP deployment $S$, so $s$ must be immune or doomed respectively.

However, if some of these routes are legitimate and some are attacked, then whether $s$ stabilizes to a route to $m$ or $d$ depends on deployment $S$, so $s$ must be protectable.

#### 5.5.9.2 Computing Partitions: Security $2^{nd}$

The algorithm for determining partitions for the security $2^{nd}$ model is slightly different from that used when security is third. We still use the algorithm from Section 4.5.2, except that now, for every AS discovered in the BFS computation we need to keep track of all perceivable routes in its PR set that are of the same type as the routes in its BPR set. We keep track of whether (1) all such routes lead to the destination, or (2) all such routes lead to the attacker or (3) some of these routes lead to the destination and others to the attacker. We classify ASes in the first category as immune, ASes in the second category as doomed, and ASes in the third category as protectable. The following result shows that our algorithm for computing partitions when security is $2^{nd}$ is correct.

**Corollary 5.5.2.** *In the security $2^{nd}$ routing model, for any $\boldsymbol{LP}_{k\in\mathbb{N}}$, for any destination $d$, attacker $m$, source $s$, and deployment $S \subseteq V$, $s$ will stabilize to a route of the same type as any route in $\mathsf{BPR}(s, \emptyset, m, d)$.*

*Proof.* This follows from the correctness of our algorithm for computing routes in the security $2^{nd}$ model described in Section 4.5.3. Because in the security $2^{nd}$ model security is prioritized above route length, but below route type, all the routes in $\mathsf{BPR}(s)^r$ must be contained in the set of routes in $\mathsf{PR}(s, m, d)$ that are of the same type as routes in $\mathsf{BPR}(s, \emptyset, m, d)$. Recall that $\mathsf{BPR}^r(s)$ is the set of best perceivable routes of $s$ during iteration $r$ of the appropriate subroutines FSIStuBB($k$), FSCR($\infty$) and FCR($\infty$), FPeeR($\infty$), or FSPrvR and FPrvR of our algorithm, if $\mathsf{BPR}(s, S, m, d)$ contains customer or peer routes of length at most $k$, customer routes of length greater than $k$, peer routes of length greater than $k$, or provider routes respectively. Also, note that by the correctness of our algorithm, $s$ must stabilize to a route in $\mathsf{BPR}^r(s)$

for some iteration $r$ of exactly one of these subroutines.

Therefore, if $s$ has customer routes in $\mathsf{BPR}(s, \emptyset, m, d)$, it will be fixed to a customer route of length at most or greater than $k$ during FSIStuBB($k$) or either of the FSCR($\infty$) and FCR($\infty$) subroutines of this algorithm respectively, for any choice of $S$. Similarly, if $s$ has peer routes in $\mathsf{BPR}(s, \emptyset, m, d)$, it will be fixed to a customer route of length at most or greater than $k$ during FSIStuBB($k$) or FPeeR($\infty$) subroutine of this algorithm respectively, for any choice of $S$. Finally, if $s$ has provider routes in $\mathsf{BPR}(s, \emptyset, m, d)$, it will be fixed to a route in either FSPrvR or FPrvR subroutines for any choice of $S$. □

Corollary 5.5.2 tells us that to determine if $s$ is immune, doomed or protectable in security $2^{nd}$ model, it is sufficient to keep track of all the routes of the best type of $s$, *i.e.,* . all $s$'s perceivable routes of the same type as routes in $\mathsf{BPR}(s, \emptyset, m, d)$, because $s$ is guaranteed to stabilize to one of these routes. Therefore, if all such perceivable routes are legitimate or attacked, then $s$ must stabilize to a legitimate or attacked route under any S*BGP deployment $S$, so $s$ must be immune or doomed respectively. However, if some of these routes are legitimate and some are attacked, then whether $s$ stabilizes to a route to $m$ or $d$ depends on deployment $S$, so $s$ must be protectable.

#### 5.5.9.3 Computing Partitions: Security $1^{st}$

Recall that in our empirical evaluations we have assumed that all source ASes are protectable in security $1^{st}$ model (see Figure 14 for example). Technically, however, there can be doomed and immune ASes in the security $1^{st}$ model, in a few exceptional cases. In this section we argue that the number of such ASes is negligible.

We can characterize doomed ASes when security if $1^{st}$ as follows.

**Observation 5.5.3.** *In the security $1^{st}$ model, for any $\boldsymbol{LP}_{k\in\mathbb{N}}$, for a particular destination-attacker pair $(d, m)$, a source AS $v_i$ is doomed if and only if every one of its perceivable routes $\mathsf{PR}(v_i, m, d)$ contains $m$.*

If every perceivable route from $v_i$ to $d$ contains $m$, then there is no S*BGP deployment scenario that could result in $v_i$ being happy. On the other hand, if $v_i$ is not doomed, then there must be at least one S*BGP deployment scenario that results in $v_i$ being happy, in which case $v_i$ must select a route to $d$ that does not contain $m$.

ASes that are single-homed to the attacking AS $m$ are certainly doomed, per Observation 5.5.3. There are $11,953$ and $11,585$ single-homed stub ASes (without peers) for the regular and the IXP-augmented graphs respectively. As an upper bound, we consider only the former number. Recall from Section 5.3 that our security metric is defined as the average of happy sources, where the average is taken over all sources and all appropriate destination-attacker pairs. It follows that for any one destination, there can be at most $11,953$ doomed single-homed ASes when summed over all attackers and all sources. Therefore, the fraction of doomed sources does not exceed .001% and .01% when considering all and only non-stub attackers respectively. While Observation 5.5.3 suggests there could be other doomed ASes, other than the just the single-homed stub ASes, the Internet graph is sufficiently well-connected to ensure that the number of such ASes is small enough.

Characterizing immune ASes is more tricky than doomed ASes. If for every possible S*BGP deployment scenario there exists a route to $d$ that is available to $v_i$ and that is preferred by $v_i$ to any route through $m$, then $v_i$ must be immune. On the other hand, if $\forall$ $v_i$'s exportable routes $R$, $\exists$ at least one intermediate AS $v_j \in R$ such that either one of $v_j$'s routes to $d$ is not exportable to $v_{j+1}$ or the resulting route $R^* v_j R^{v_j}$ is not preferred by $v_i$ to one of its available attacked routes $R^A$, then we can cause $v_i$ to become unhappy by securing only the ASes contained in $R^{v_j}$, possibly for all such routes $R$. It's important for $R^A$ to be available when $S = \emptyset$, so that by securing $R^{v_j}$ we cannot secure an AS in $R^A$ making $R^A$ unavailable, otherwise $R^{v_j}$ would be unavailable. Such complicated scenarios are not expected to happen frequently, so we opt out for a more simplified characterization.

**Observation 5.5.4.** *In the security $1^{st}$ model, for any $\boldsymbol{LP}_{k\in\mathbb{N}}$, for a particular destination-attacker pair $(d, m)$, a source AS $v_i$ is immune if every one of its perceivable routes $\mathsf{PR}(v_i, \emptyset, m)$ contains $d$ and does not contain $m$.*

Estimating the fraction of immune ASes is slightly more complicated than the fraction of doomed ASes. In addition to single-homed networks, ASes that are the direct provider of the legitimate destination $d$ will be immune. This follows because in our threat model, attackers will always announce that they are directly connected to $d$. Therefore, ASes that are direct providers of $d$ will have a one-hop customer route to $d$, which will always be preferred over any two-hop route offered by $m$. Note that such scenarios occur exactly as many times as there are customer-provider links, namely $73,442$. Thus, recognizing that for any one attacker, there can be at most $11,953 + 73,442 = 85,395$ immune single-homed networks and/or direct providers of $d$ when summed over all destinations and all sources, it follows that the fraction of immune sources does not exceed .006% and .04% when considering all and only non-stub attackers respectively.

## *5.6 How Close Can We Get to the Upper Bounds?*

In Section 5.5.4 we presented upper bounds on the improvements in security from S*BGP deployment for choice of secure ASes $S$. We found that while only meager improvements over origin authentication are possible in the security $3^{rd}$ model, better results are possible in the security $2^{nd}$ and $1^{st}$ models. However, achieving the bounds in Section 5.5.4 could require full S*BGP deployment at every AS. In this section we consider the question of what happens in the more realistic deployment scenarios when S*BGP is only partially deployed. We also present prescriptive guidelines for partial S*BGP deployment. In Section 5.6.2.2 we explain the error bars for some of the plots presented in this section. All the results shown in this section correspond to the $\mathbf{LP}_0$ local preference model, and we show corresponding plots for other local

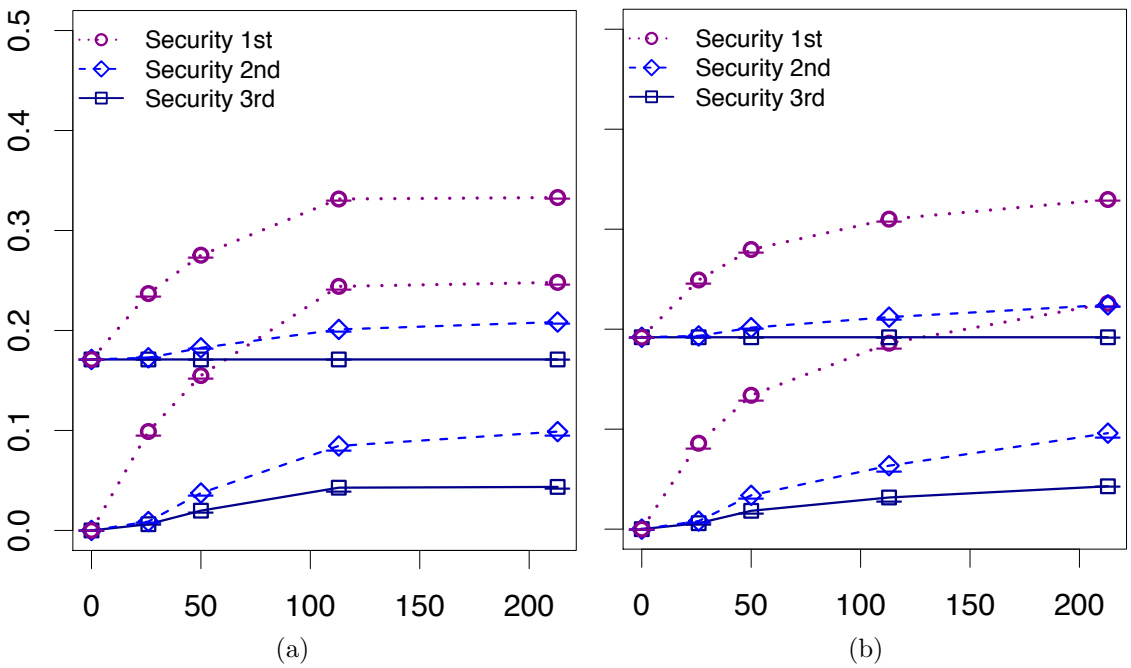


**Figure 19:** Tier 1+2+3 rollout. For each step $S$ in the rollout, upper and lower bounds on the metric improvement $H_{M',V}(S) - H_{M',V}(\emptyset)$ are presented for (a) the UCLA and (b) IXP-augmented topologies. The $x$-axis is the number of non-stub ASes in $S$.

preference models in Section 5.8

### 5.6.1 Looking at Large Partial S*BGP Deployments

Instead of focusing on choosing the optimum set $S$ of ASes to secure (an intractable feat), we will instead consider a few partial deployment scenarios among high-degree ASes $S$, as suggested in practice [87] and in the literature [40, 12, 28].

We will focus on set of attackers $M'$ that consists of all non-stub ASe in our graph, *i.e.,* all ASes except Stubs or Stubs-x as per Table 2. Ruling out stub ASes is consistent with the idea that stubs cannot launch attacks if their providers perform prefix filtering [46, 25], a functionality that can be achieved via IRRs [4] or even the RPKI [82], and does not require S*BGP.

#### 5.6.1.1 Security Across All Destinations

Gill *et al.* [40] have suggested bootstrapping S*BGP deployment by having secure ISPs deploy S*BGP in their customers that are stub ASes, and with that in mind we consider the following rollout.

**Tier 1, Tier 2 and Teir 3 Rollout:** Other than the empty set, we consider four different secure sets. We secure $X$ Tier 1s, $Y$ Tier 2s, and $Z$ Tier 3s together with all their stubs, where $(X, Y, Z) \in \{(13, 13, 0), (13, 37, 0), (13, 100, 0), (13, 100, 100)\}$. This rollout corresponds to securing about 33%, 40%, 50% and 57% of the AS graph. With such large deployments, we would hope to see very large improvements in security.

The results are shown in Figure 19, which plot, for each secure routing policy model, the increase in the upper and lower bound on $H_{M',V}(S)$ for each set $S$ of secure ASes in the rollout ($y$-axis), versus the number of non-stub ASes in $S$ ($x$-axis). We make a few important observations:

We notice that even with a large S*BGP deployment, the improvement in security benefit is highly dependent on the vagarities of the intradomain tiebreaking criteria used to decide between insecure routes. Recall the discussion on tiebreaking in Section 5.5.1. Even when we secure 57% of ASes in the security $1^{st}$ model (the last step of our rollout), there is still a gap of almost 10% between the lower and upper bounds of our metric. Thus, in a partial S*BGP deployment, there is a large fraction of ASes that are balanced on a knife's edge between an insecure legitimate route and an insecure bogus route. Only the unknown-to-us intradomain routing policies of these ASes can save them from attack. This is inherent to any partial deployment of S*BGP, even in the security $1^{st}$ model.

As expected, the biggest improvements come in the security $1^{st}$ model, where ASes make security their highest priority and deprecate all economic and operational considerations. When security is $1^{st}$ and 57% of the AS graph deploy S*BGP (at the last step in the rollout), the improvement over the baseline scenario is approximately

25%, which is significant. While we might hope that the security $2^{nd}$ model would present improvements that are similar to those achieved when security is $1^{st}$, this is unfortunately not the case. In both the security $2^{nd}$ and $3^{rd}$ models we see similarly disappointing increases in our metric. We explain this observation in Section **??**.

Inspired by a real-life prefix hijacking event that impacted many destination ASes [34], we evaluated the metric $H_{m,V}(S)$ when only that specific AS $m$ had attacked all destinations $D = V$ for our Tier 1 + 2 + 3 rollout. The results are similar to Figure 19.

#### 5.6.1.2 Focusing Strictly On the Content Providers

Since much of the Internet's traffic originates at the content providers (CPs) [67], let us consider the impact of S*BGP deployment on CPs only. We analyze the Tier 1+2+3 rollout as above, but with all 17 CPs secure, and computed the metric over CP destinations only, *i.e.,* $H_{M',CP}(S)$. However, the results of this rollout, shown in Figure 20, are similar to those in shown in Figure 19. For the last point in the rollout, we observe improvements of at least 26% 9.4%, and 4% for security $1^{st}$, $2^{nd}$, and $3^{rd}$ respectively. We note, however, that CP destinations have a higher fraction of happy sources than other destinations on average, as can also be seen in Figure 15.

#### 5.6.1.3 Different Destinations See Different Benefits

Thus far, we have looked at the impact of S*BGP in aggregate across all destinations $d \in V$ (or $d \in CP$). Because secure routes can only exist to secure destinations, in this section we look at the impact S*BGP may have on each secure destination by considering $H_{M',d}(S)$ for each secure destination $d \in S$ separately.

Let us zoom in on the second-to-last step in our first rollout, where 13 Tier 1s, 100 Tier 2s, and all of their stubs are secure (50% of ASes). For this step, in Figure 21 we plot upper and lower bounds on the change in the metric, *i.e.,* $H_{M',d}(S) - H_{M',d}(\emptyset)$,

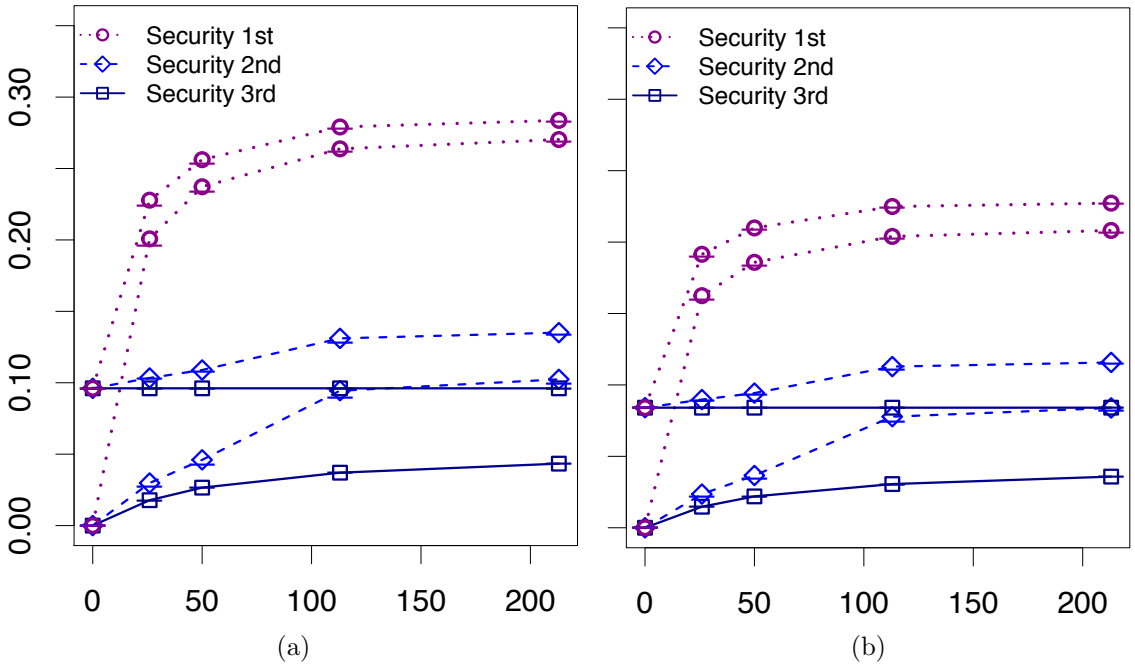


**Figure 20:** Tier 1+2+3+CP rollout. For each step $S$ in the rollout, upper and lower bounds on the metric improvement $H_{M',V}(S) - H_{M',V}(\emptyset)$, for strictly CP destinations, are presented for (a) the UCLA and (b) IXP-augmented topologies. The $x$-axis is the number of non-stub, non-CP ASes in $S$, and the averaging is done with respect to CP destinations only.

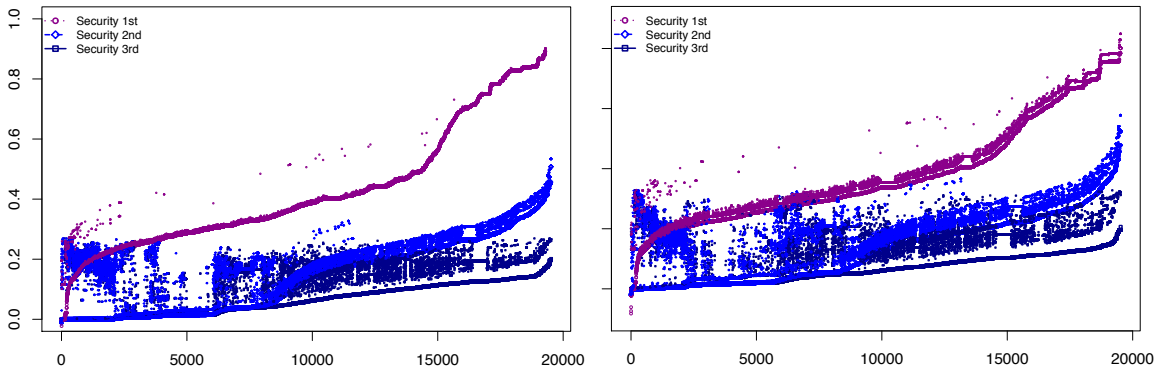


**Figure 21:** Non-decreasing sequence of $H_{M',d}(S) - H_{M',d}(\emptyset)\ \forall\ d \in S$ for (a) the UCLA and (b) the IXP-augmented topologies. $S$ is all T1s, T2s, and their stubs.

for each individual secure destination $d \in S$. For each of our three models, the lower bound for each $d \in S$ is plotted as a non-decreasing sequence. These are the three smooth lines. The corresponding upper bound for each $d \in S$ was plotted as well. For security $1^{st}$, we see that the upper and lower bounds are almost identical, and for security $2^{nd}$ and $3^{rd}$, the upper bounds are the clouds that hover over the lower bounds. We make the following observations.

We find that when security is $1^{st}$, a secure destination can reap the full benefits of S*BGP even in partial deployment, albeit a large one. To see this, we computed the true value of $H_{M',d}(S)$ for all secure destinations $d \in S$, and found that it was between $96.8 - 97.9\%$ on average (across all $d \in S$).

Figure 21 also reveals that many destinations obtain roughly the same benefits from S*BGP when security is $2^{nd}$ as when security is $3^{rd}$. 93% of 7500 secure destinations that see $< 4\%$ of improvement in Figure 21 when security is $3^{rd}$, see the same improvement when security is $2^{nd}$. The reason for this is that there are certain types of protocol downgrade attacks that succeed both when security is $2^{nd}$ and when security is $3^{rd}$, *i.e.,* when the bogus route has better **LP** than the legitimate route (*e.g.,* Figure 9). In Section 5.6.3 we will show that protocol downgrade attacks is one of the major causes for such low metric improvements. Therefore, for destinations

where these **LP**-based protocol downgrade attacks are most common, the security $2^{nd}$ model looks much like the security $3^{rd}$ model.

Figure 21 also reveals that when security is $1^{st}$, secure destinations that obtain the largest ($> 40\%$) increases in their security metric $H_{M',d}(S)$ (relative to the baseline setting $H_{M',d}(\emptyset)$) include: (a) all 13 Tier 1s, and (b) $\geq 99\%$ of Tier 1 stub destinations (*i.e.,* stub ASes such that all their providers are Tier 1 ASes). On the other hand, these same destinations experience the worst improvements when security is $2^{nd}$ or $3^{rd}$ (*i.e.,* at most 3%).

To explain this, recall from Section 5.5.6 that when security is $2^{nd}$ or $3^{rd}$, most source ASes that want to reach a Tier 1 destination are *doomed*, because of protocol downgrade attacks like the one shown in Figure 9. This explains the meager benefits these destinations obtain when security is $2^{nd}$ or $3^{rd}$. On the other hand, protocol downgrade attacks fail when security is $1^{st}$. Therefore, in the security $1^{st}$ model, the Tier 1 destinations (and by extension, Tier 1 stub destinations) obtain excellent security when S*BGP is partially deployed. Furthermore, they see most significant gains simply because they were so highly vulnerable to attacks in the absence of S*BGP (see Figure 15).

Finally, we observe that when security is $2^{nd}$, about half of the secure destinations $d \in S$ see benefits that are discernibly better than what is possible when security is $3^{rd}$, though not quite as impressive as those when security is $1^{st}$. These destinations include some Tier 2s and their stubs, but never any Tier 1s.

Note that similar observations hold for earlier steps in the rollout for the UCLA and IXP-augmented topologies.

#### 5.6.1.4 Other Partial Deployments

The results of the previous section motivate considering deployments that exclude securing the Tier 1 ISPs. We considered a number of other deployment scenarios

that exclude Tier 1s in this section, but our results suggest that it will be difficult to find a small and simple deployment scenario $S$ where the security benefits obtained when security is $2^{nd}$ are much better than those when security is $3^{rd}$.

**Securing just the Tier 2s:** We reproduce the analysis of Section 5.6.1.1 with a rollout among only the Tier 2s and theirs stubs. There are 100 Tier 2 ISPs in our AS graph (see Table 2), and our Tier 2 rollout secures $Y$ Tier 2 ASes, and all of their stubs, where $Y \in \{13, 26, 50, 100\}$. This amounts to securing about 18%, 24%, 30%, and 38% of ASes.

The results shown in Figure 22 are similar to those in Figure 19, except that the metric grows even more slowly, and we see smaller improvements when security is $1^{st}$. We see this also in Figure 23 which reproduce the results of Figure 21 for the last step of the Tier 2 rollout (amounting to 38% of the AS graph). Not that this is consistent with the observation that most dramatic improvements observed when security is $1^{st}$ are for Tier 1 destinations discussed in Section 5.6.1.3. The improvements for Tier 2 destinations and their stubs are much smaller when security is $1^{st}$. This causes the gap between the security $2^{nd}$ and $1^{st}$ models to become smaller for the Tier 2 rollout, relative to the Tier 1+2 rollout. However, the gap between security $2^{nd}$ and $1^{st}$ is smaller not only because Tier 2s see bigger improvements when security is $2^{nd}$. This is also because they see worse improvements when security is $1^{st}$.

**Securing all non-stub ASes:** Perhaps the fact that our rollouts include so many stubs can explain the meager improvements in our metric? However, we also found this to be false. Securing all $\approx 6K$ non-stubs in the AS graph did not result in large improvements to the security metric. Specifically, we see a 6.2%, 4.7% and 2.2% worst-case improvement in the metric $H_{M',D}(S)$ when security is $1^{st}$, $2^{nd}$, and $3^{rd}$ respectively. This scenario therefore is similar to the last step in our Tier 2 rollout, with exception that the gap between the security $2^{nd}$ and $1^{st}$ model is even smaller. This is corroborated by Figure 24, which reproduces the results of Figure 21 for the

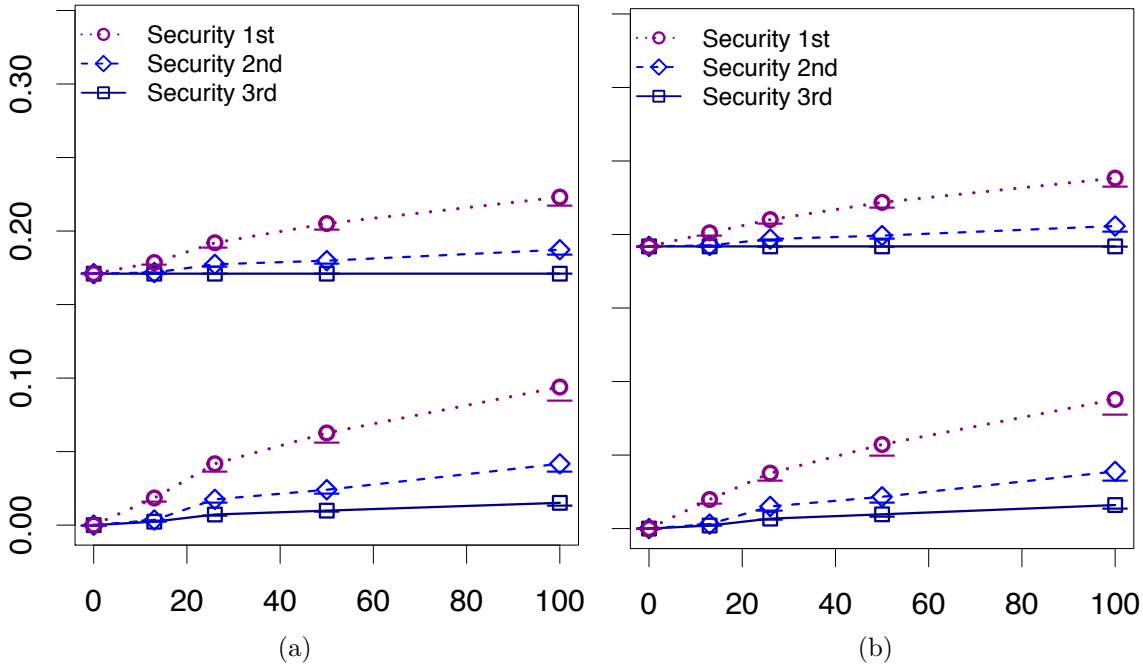


**Figure 22:** Tier 2 rollout. For each step $S$ in the rollout, upper and lower bounds on the metric improvement $H_{M',V}(S) - H_{M',V}(\emptyset)$ are presented for (a) the UCLA and (b) IXP-augmented topologies. The $x$-axis is the number of non-stub ASes in $S$.

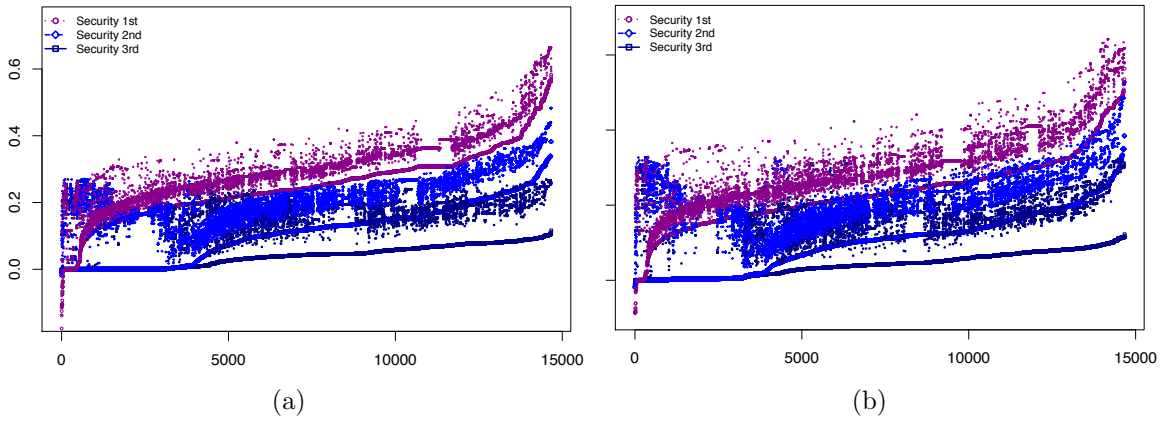


**Figure 23:** Non-decreasing sequence of $H_{M',d}(S) - H_{M',d}(\emptyset)\ \forall\ d \in S$ for (a) the UCLA and (b) IXP-augmented topologies. $S$ is all T2s, and their stubs.

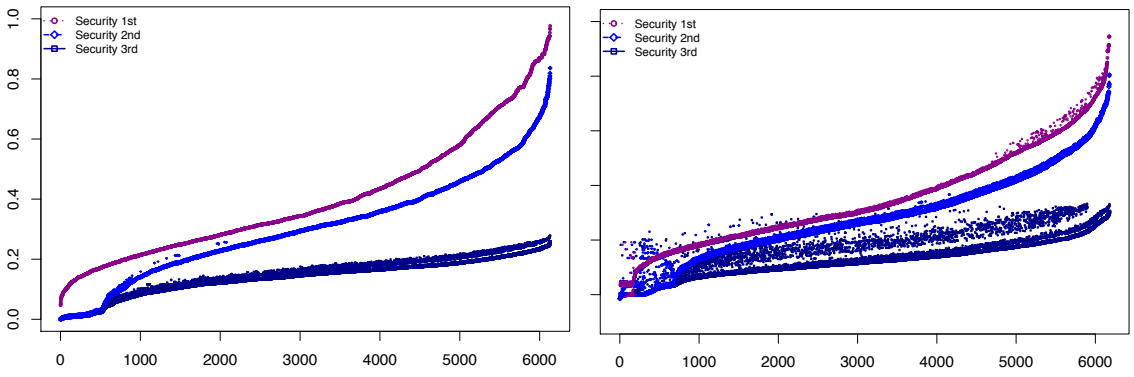


**Figure 24:** Non-decreasing sequence of $H_{M',d}(S) - H_{M',d}(\emptyset)\ \forall\ d \in S$ for (a) the UCLA and (b) IXP-augmented topologies. $S$ is all non stubs.

scenario where only non-stub ASes are secure. We see that the benefits available when security is $2^{nd}$ almost reach those that are possible when security is $1^{st}$.

Taken together, our results suggest that in the security $1^{st}$ model, destinations that are Tier 1s or their stubs see the largest improvements in security. In such cases, the security $2^{nd}$ model behaves much like the security $3^{rd}$ model. However, in cases where Tier 1s and their stubs are not secure, the gap between the security $2^{nd}$ and $1^{st}$ model diminishes, in exchange for smaller gains when security is $1^{st}$.

### 5.6.2 Prescriptive Deployment Guidelines

In this section we suggest a few S*BGP deployment guidelines.

#### 5.6.2.1 On the Choice of Early Adopters

Previous work [40, 12, 28] has suggested that Tier 1s should be the earliest adopters of S*BGP due to their centrality and the high volumes of traffic they transit. However, our discussion in Sections 5.5.6 and 5.6.1.3 suggests that securing Tier 1s might not lead to good security benefits at the early adoption stage, when ASes are most likely to rank security $2^{nd}$ or $3^{rd}$. We confirm this even further.

Even in a deployment that includes *all* 13 Tier 1 ASes and their stubs (*i.e.,* 7872 ASes or $\approx$ 20% of the AS graph), security benefit improvements were almost

imperceptible. With security $2^{nd}$ or $3^{rd}$, the average change in $H_{M',d}(S) - H_{M',d}(\emptyset)$ over secure destinations $d \in S$ causes the metric to increase by $< 0.2\%$.

Following [40, 87], we have considered securing the CPs, the Tier 1s and all of their stubs, and obtained similar results. We notice that deployment at more than 20% of the ASes in the AS graph, including the large and well-connected Tier 1s, still results in very little improvement in security. Recall that in Section 5.5.6 and Figure 15 we showed that when Tier 1 destinations are attacked, the vast majority of source ASes are doomed and almost none is protectable. It follows that if a source retains a secure route to a Tier 1 destination during an attack, that source is likely to be immune. The same argument also applies to other secure destinations, *i.e.,* such as CPs of stub customers of T1s,. because, in the deployment scenarios above, most secure routes traverse a Tier 1 as their first hop. Because almost every source AS that continued to use a secure route during an attack would have routed to the legitimate destination even if no AS deployed S*BGP, we see little improvements in our security metric.

In Figure 25 we confirm this by showing what happens to the secure routes to each CP destination when security is $3^{rd}$. The height of each bar is the fraction of routes to each CP destination that are secure under normal conditions. The lower part of the bar shows secure routes that were lost to protocol downgrade attacks (averaged over all attacks by non-stubs in $M'$), and the middle part shows the fraction of secure routes from immune source ASes to the destination. We clearly see that (1) most secure routes are lost to protocol downgrade attacks, and (2) almost all the secure routes that remain during attacks are from source ASes that are immune. As can be seen from Figure 26, similar observations hold when security is $2^{nd}$.

On the other hand, we found that early deployments at the Tier 2 ISPs actually fare better than those at the larger, and better connected Tier 1s. For example, securing the 13 largest Tier 2s (in terms of customer degree) and all their stubs (a

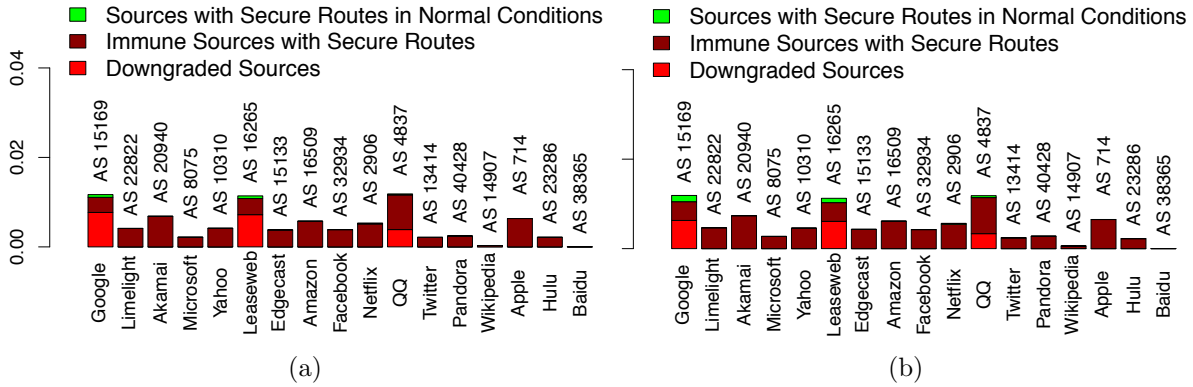


**Figure 25:** What happens to secure routes to each CP destination during attack for (a) the UCLA and (b) IXP-augmented topologies. $S$ is the Tier 1s, the CPs, and all their stubs when security is $3^{rd}$. Y-axis is the average fraction of sources.

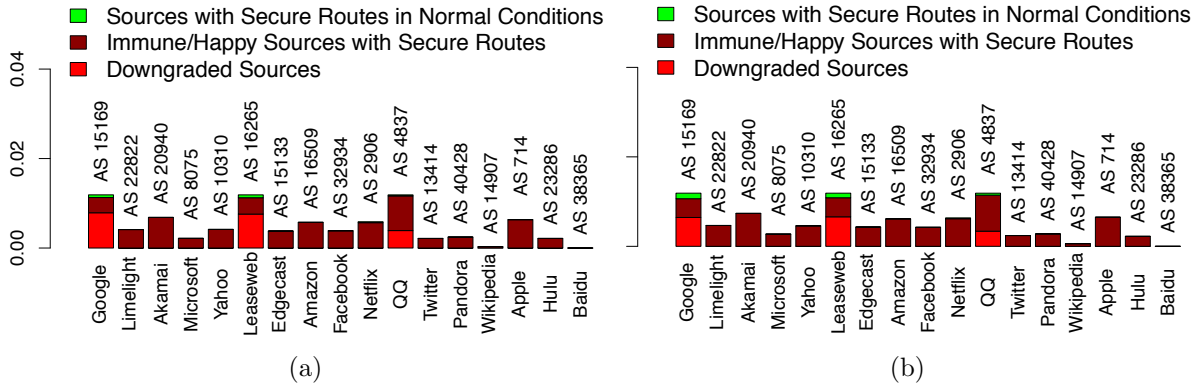


**Figure 26:** What happens to secure routes to each CP destination during attack for (a) the UCLA and 26(b) IXP-augmented topologies. $S$ is the Tier 1s, the CPs, and all their stubs when security is $2^{nd}$. Y-axis is the average fraction of sources.

total of 6918 ASes) provides a change in $H_{M',d}(S) - H_{M',d}(\emptyset)$, averaged over secure destinations $d \in S$, of $\approx 1\%$ when security is $2^{nd}$ or $3^{rd}$. This also agrees with our observations in Section 5.6.1.4.

#### 5.6.2.2 Use Simplex S*BGP at Stubs

We now discuss [68, 40]'s suggestion for reducing complexity by securing stubs with *simplex S*BGP* and explain the error bars in our figures corresponding to our S*BGP rollout experiments. Stub ASes have no customers of their own, and therefore (by **Ex**) they will never send S*BGP announcements for routes through other ASes. They will, however, announce routes to their own IP prefixes. As we have already mentioned in Chapter 3, in simplex S*BGP, either (1) ISPs are allowed to send S*BGP messages on behalf of their stub customers or (2) stubs are allowed to deploy S*BGP in a unidirectional manner, sending outgoing S*BGP messages but receiving legacy BGP messages. Since a stub propagates only outgoing BGP announcements for a very small number of IP prefixes, *i.e.,* the prefixes owned by that stub, simplex mode can decrease computational load, and make S*BGP adoption less costly.

However, given that over 85% of ASes are stubs, we ask the question if simplex S*BGP could severely harm security. We now show that this is not an issue. The error bars in Figures 19, 20, and 22 show what happens when we suppose that all stubs run simplex S*BGP, and, accounting for the worst case scenario, always fall victim to the attack if at least one of their learned routes is bogus. Recognize that in simplex S*BGP stubs are not able to verify route announcements. We observer that there is little change in the metric. To explain this, we note that (1) a stub's routing decision does not affect any other AS's routing decision, since by **Ex** stubs do not propagate BGP routes from one neighbor to another, and (2) a stub's routing decisions are limited by the decisions made by its providers, so if its providers avoid attacks, so will the stub, but (3) the stub acts like a secure destination, and therefore

(non-stub) ASes establishing routes to the stub still benefit from S*BGP. These results indicate that simplex S*BGP at stubs can lower the complexity of S*BGP deployment almost without having any serious impact on the overall security. Stub ASes that are concerned about their own security as sources, not only as destinations, can, of course, always choose to deploy full S*BGP.

### 5.6.3 Root-Cause Analysis

We now examine the reasons for the changes in our security metric as S*BGP is deployed. Our strategy is to investigate which of the phenomena in Table 1 discussed in Section 5.3.2 have the biggest impact on security. and then check how they play out with respect to the Tier 1 and Tier 2 rollout of Section 5.6.1.1. Recall that for this deployment scenario $S$ is all 13 Tier 1s, all 100 Tier 2s and all of their stubs, *i.e.,* roughly 50% of the AS graph.

Let us start with a root cause analysis for the security $3^{rd}$ model in Figure 27(a) (left). Recall that Theorem 5.3.3 showed that collateral damages do not occur in the security $3^{rd}$ model, and so we do not consider them here. The bottom three parts of the bar show the fraction of secure routes available in normal conditions, prior to any routing attacks. Averaging is done across all $V^2$ sources and destinations as well as over all attackers in $M'$. During routing attacks, these routes can be broken down into three types: (1) secure routes lost to protocol downgrade attacks (lowest part of the bar), (2) secure routes wasted on ASes that would have been happy even in the absence of S*BGP (second lowest part), and (3) secure routes of protected ASes that were unhappy in the absence of S*BGP (third lowest part). Note that improvements in our security metric can only result from the small fraction of secure routes in class (3), and the remaining secure routes either (1) disappear due to protocol downgrades, or (2) are wasted on ASes that would have avoided the attack even without S*BGP.

The top two parts of the bar show how (the lower bound on) the metric $H_{M',V}(S)$

grows relative to the baseline scenario $S = \emptyset$ due to: (a) secure routes in class (3), and (b) the lower bound on the fraction of insecure ASes that obtained collateral benefits. Notice how Figure 27(a) thus illustrates the importance of collateral benefits. Figure 28(a) shows similar results for the IXP-augmented topology.

We perform the same analysis for the security $1^{st}$ model in Figure 27(b). Recall that by Theorem 5.2.1, protocol downgrade attacks occur only rarely in this model, so these are not visible in the figure. However, we now have to account for collateral damages, which we depict with the smaller sliver on right of the figure. We obtain the change in the metric by subtracting the collateral damages from the gains resulting from (a) offering secure routes to unhappy ASes and (b) collateral benefits. Fortunately, we find collateral damages to be a relatively rare phenomenon.Figure 28(b) shows similar results for the IXP-augmented topology.

Our analysis reveals that changes in the metric can be computed via the following equation, which we call Source Conservation:

$$
\begin{aligned}
\text{changes in the metric} \;=\;& \text{secure routes created under normal conditions} \\
+\;& \text{collateral benefits} \\
-\;& \text{protocol downgrades} \\
-\;& \text{secure routes wasted on ASes that are already happy} \\
-\;& \text{collateral damages}
\end{aligned}
$$

Intuitively, the Source Conservation equation says that the fraction of ASes that change their routes from the attacker to the destination and vice versa must be the result of all the *good* events that can happen to the sources (*i.e.,* obtaining secure routes and collateral benefits) minus all the *bad* events that can happen to the sources (*i.e.,* protocol downgrades, collateral damages, wasted secure routes ). Note that in this equation we ignore any sources that loose connectivity (to the destination or

attacker) as a result of adding security, while sources that loose connectivity as a result of introducing an attacker are captured with collateral damages. Regardless, our empirical evaluation confirms that the fraction of ASes that lose connectivity for any reason is negligible. With the exception of collateral damage, we find that all of these phenomena have significant impact on the security metric.

The Source Conservation equation also drives home the point that the number of routes learned via S*BGP under normal conditions is a poor proxy for capturing security of the network. More sophisticated metrics, such as the one we used in our analysis, are required to evaluate how well a network is protected from attacks in general (not just the one-hop hijacks!) when S*BGP is partially deployed.

Results when security is $2^{nd}$ look very similar to the results when security is $3^{rd}$, with the addition of a small amount of collateral damage. To conclude, when security is $2^{nd}$ or $3^{rd}$, (1) protocol downgrade attacks cause many secure routes that were available under normal conditions to disappear, and (2) those ASes that retain their secure routes during the attack would have been happy even if S*BGP had not been deployed. The result is meager increases in the security metric. Meanwhile, when security is $1^{st}$, few downgrades occur, and the security metric is greatly improved.

Although *protocol upgrades*—the exact opposite phenomenon of a protocol downgrade that results in a source AS that uses an insecure route to the legitimate destination under normal conditions, to upgrade to a secure route during an attack as a result of its routing policies—our empirical evaluations confirm that protocol upgrades happen very rarely, so we do not show them in Figures 27-28. Note that for the same reason as with protocol downgrades, protocol upgrades cannot happen when security is $1^{st}$. Due to the routing policies we use in our analysis, an attack cannot result in the creation of more secure routes. This is why protocol upgrades can only happen when some AS on a route switches to a longer, bogus route with higher local preference as a result of an attack, causing some other AS to switch to a longer

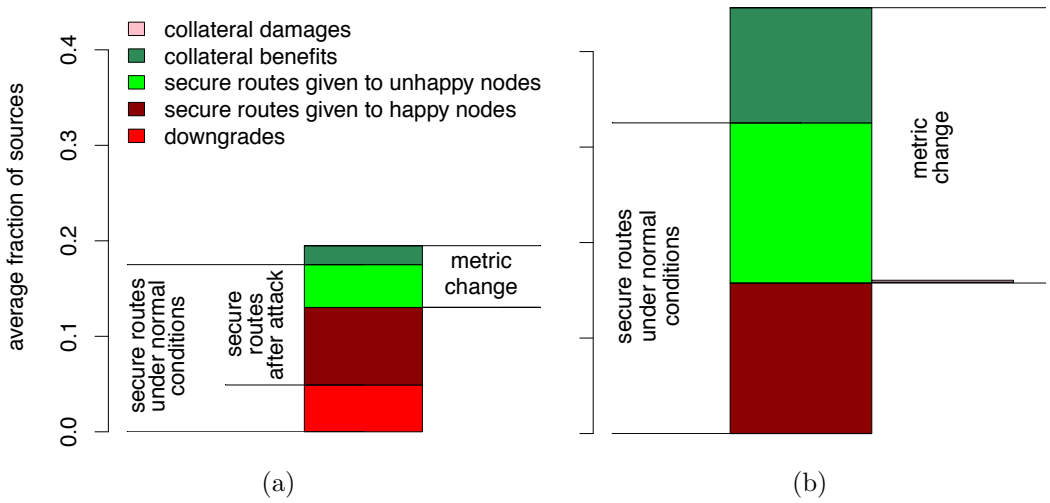


**Figure 27:** The break down of metric changes for deployment scenario with 13 T1's + 100 T2's + all their stubs when security is (a) $3^{rd}$ and (b) $1^{st}$, for the UCLA topology.

secure route with the same local preference. This is why protocol upgrades cannot happen when security is $2^{nd}$; they can only happen when security is $3^{rd}$. While ASes that experience protocol downgrades must be doomed, ASes that experience protocol upgrades must be immune (captured by the Source Conservation equation) or protectable (a very rare case).

### 5.6.4 Computing Protocol Downgrades

Since protocol downgrades play such an important role in our empirical results, we now briefly describe how we compute them in our experiments. To quantify the success of protocol downgrade attacks with respect to an attacker-destination pair $(m, d)$ and a set of secure ASes $S$, we need to first establish which ASes have a secure route to the destination under normal conditions, that is, when there is no attack. To do this, we compute the S*BGP routing outcome when there is no attacker, by setting $M = \emptyset$ for the set $S$, for the specific model under consideration. The algorithm records for every AS discovered in this BFS computation whether (1) all routes in its BPR set at that iteration are secure or (2) all these routes are insecure. We then compute the S*BGP routing outcome for the pair $(m, d)$ for the set $S$ (for the routing specific

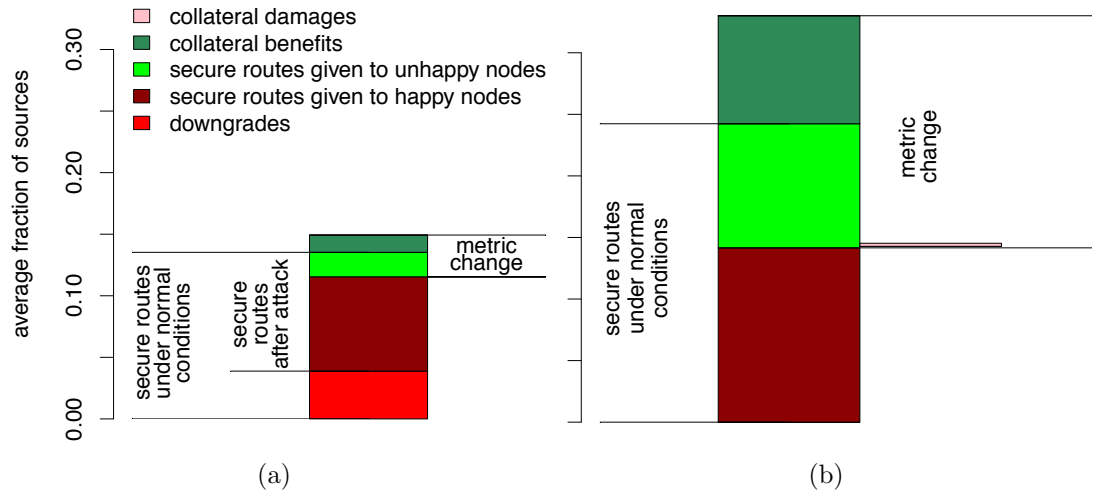


**Figure 28:** The break down of metric changes for deployment scenario with 13 T1's + 100 T2's + all their stubs when security is (a) $3^{rd}$ and (b) $1^{st}$, for the IXP-augmented topology.

model under consideration)). Again, the algorithm records for every AS discovered in this BFS computation whether (1) all routes in its BPR set at that iteration are secure or (2) all these routes are insecure. We conclude that a protocol-downgrade attack against an AS is successful if that AS falls in category (1) in the first of these computations and in category (2) in the second computation. The correctness of this approach follows from the correctness of our algorithms in Section 4.5 for any security prioritization and $\mathbf{LP}_{k\in\mathbb{N}}$.

## *5.7 Sensitivity to Routing Policy: Partitions*

Thus far, all our analysis in this chapter has been done with respect to the $\mathbf{LP}_0$ model of local preferences presented in Section 4.2.1.1. In this section we investigate $\mathbf{LP}_1$ $\mathbf{LP}_2$ and $\mathbf{LP}_{50}$ models of local preference, and consider how they impact the results we presented in Section 5.5. As we will confirm later in this section, $\mathbf{LP}_{50}$ plays a role of $\mathbf{LP}_\infty$ for the empirical AS-level topologies that we used.

### 5.7.1 Partition Results with $\mathbf{LP}_1$ Policy Variant

We begin our analysis with the $\mathbf{LP}_1$ policy variant. Here, a peer route of length less than or equal to 1 hops is preferred over a longer customer route.

#### 5.7.1.1 General Partitions

Figure 29 shows the partitions for the $\mathbf{LP}_1$ policy variant, for the UCLA graph and for the IXP-augmented graph. (*cf.,* Figure 14 in Section 5.5.4). Recall that the thick solid horizontal line shows the fraction of happy source ASes in the baseline scenario, when no AS is secure. As in Section 5.5.4, we find that with security $3^{rd}$ only limited improvements in the metric $H_{V,V}(S)$ are possible, relative to the baseline scenario $H_{V,V}(\emptyset)$: $75 - 60 = 15\%$ for the UCLA AS graph, and $78 - 62 = 16\%$ for the IXP-augmented graph, the latter being greater than what we saw for our original **LP** model. In the security $2^{nd}$ model, we see better improvements than security $3^{rd}$, and we get the same values as with the $\mathbf{LP}_0$ model: $89 - 60 = 29\%$ for the UCLA AS graph, and $90 - 62 = 28\%$ for the IXP-augmented graph.

#### 5.7.1.2 Partitions by Destination Tier

In Figure 30 we show the partitions broken down by destination tier (see Table 2) when security is $2^{nd}$ and $3^{rd}$ for the $\mathbf{LP}_1$ policy variants, for the UCLA graph and for the IXP augmented graph (*cf.,* Figures 15-16 and Section 5.5.5). The thick solid horizontal lines show the fraction of happy source ASes in the baseline scenario (where no AS is secure) for each destination tier. We immediately make the following observations.

We see that most of the protectable ASes are for stub and SMDG destinations. While in Section 5.5.5 (Figures 15-16) we found that most destination tiers have roughly the same number of protectable ASes, here we notice slightly different trends. The higher-degree AS destinations, , *i.e.,* Tier 2s, Tier 3s, and CPs, have fewer protectable ASes and more immune ASes than for the $\mathbf{LP}_0$ model. This is even more

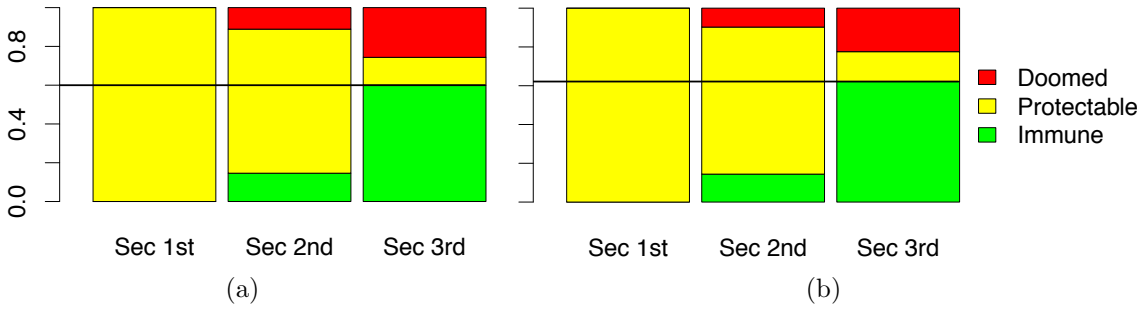


(a) (b)

**Figure 29:** Partitions for the $\mathbf{LP}_1$ policy variant, (a) UCLA graph (b) IXP-augmented graph.

apparent for the IXP augmented graph in the $\mathbf{LP}_2$ model.

While in Section 5.5.6 we found that most ASes that wish to reach Tier 1 destinations are doomed for the $\mathbf{LP}_0$ model, the vast majority of source ASes that wish to reach Tier 1 destinations are immune when security is $3^{rd}$ for the $\mathbf{LP}_1$ model. Although, Tier 1 destinations still do not have quite as many immune ASes as the Tier 2s do.

As we will discuss later, these trends will also be apparent for $\mathbf{LP}_2$ and $\mathbf{LP}_{50}$ models.

### 5.7.2 Partition Results with $\mathbf{LP}_2$ Policy Variant

We continue with analysis of the $\mathbf{LP}_2$ policy variant. Here, a peer route of length less than or equal to 2 hops is preferred over a longer customer route.

#### 5.7.2.1 General Partitions

In Figure 31 we show the partitions for the $\mathbf{LP}_2$ policy variants, for the UCLA and IXP-augmented topologies. We find that with security $3^{rd}$, only limited improvements in the metric $H_{V,V}(S)$ are possible, relative to the baseline scenario $H_{V,V}(\emptyset)$: $82-71 = 11\%$ for the UCLA AS graph, and $88-72 = 13\%$ for the IXP-augmented graph, both of which are less than what we saw for $\mathbf{LP}_0$ model. In the security $2^{nd}$ model, we again see better improvements than security $3^{rd}$, but not quite as much as we saw

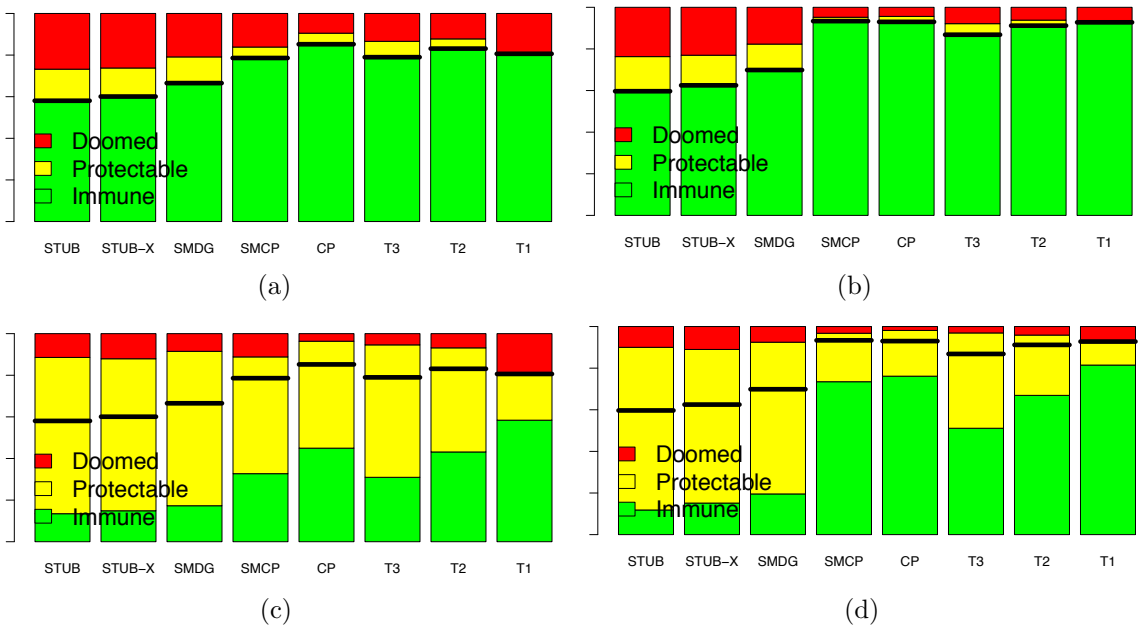


**Figure 30:** Partitions by destination tier for the $\mathbf{LP}_1$ policy variant. (a) UCLA graph, security $3^{rd}$. (b) IXP-augmented graph, security $3^{rd}$. (c) UCLA graph, security $2^{nd}$. (d) IXP-augmented graph, security $2^{nd}$. The Y-axis runs from 0 to 1.

with $\mathbf{LP}_0$ model: $92 - 71 = 21\%$ for the UCLA AS graph, and $94 - 72 = 22\%$ for the IXP augmented graph. Interestingly, however, we do see one major difference between the UCLA AS graph and the IXP-augmented graph in this model: there are many more immune ASes when security is $2^{nd}$ for the IXP-augmented graph (41% vs. 55%). We discuss the observation in more detail in Section 5.7.4.

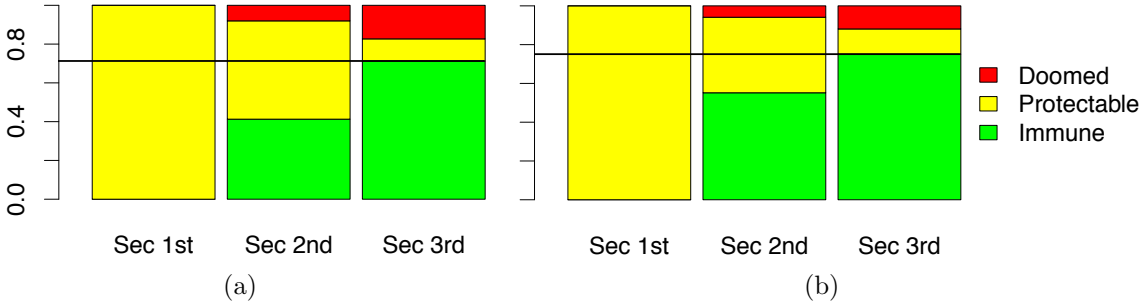


**Figure 31:** Partitions for the $\mathbf{LP}_2$ policy variant, (a) UCLA graph (b) IXP-augmented graph.

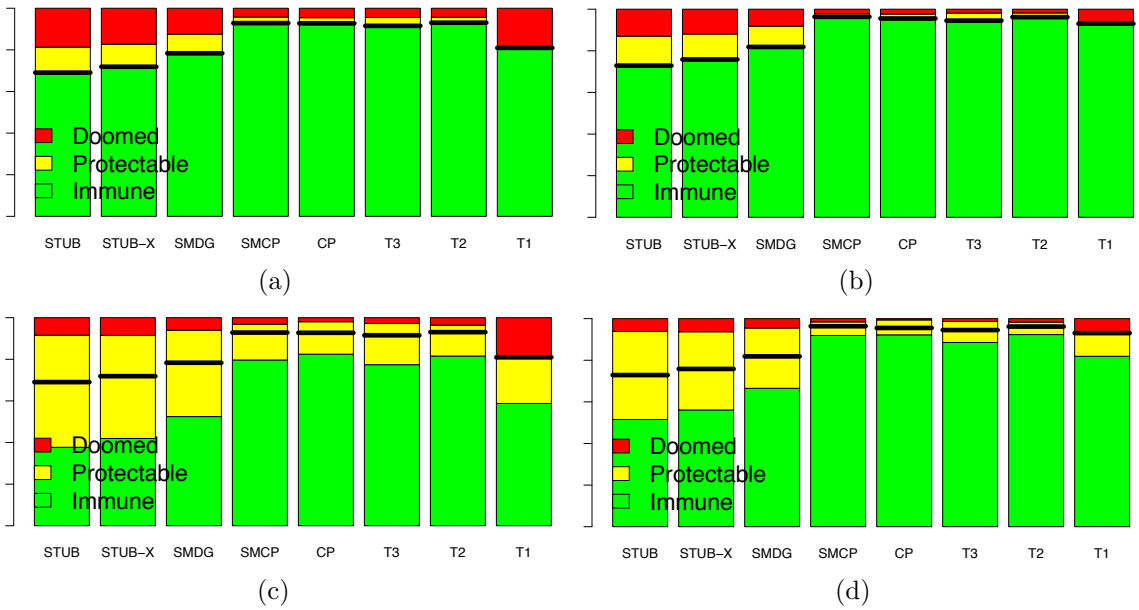


**Figure 32:** Partitions by destination tier for the $\mathbf{LP}_2$ policy variant. (a) UCLA graph, security $3^{rd}$. (b) IXP-augmented graph, security $3^{rd}$. (c) UCLA graph, security $2^{nd}$. (d) IXP-augmented graph, security $2^{nd}$. Y-axis runs from 0 to 1.

#### 5.7.2.2 Partitions by Destination Tier

In Figure 32 we show the partitions broken down by destination tier when security is $2^{nd}$ and $3^{rd}$ for the $\mathbf{LP}_2$ policy variant, for the UCLA graph and for the IXP augmented graph. We see that as with the $\mathbf{LP}_1$ model, here most of the protectable ASes are stub and SMDG (low-degree non-stub ASes) destinations. The higher-degree AS destinations have very few protectable ASes but many more immune ASes as compared to the results we obtained for the $\mathbf{LP}_0$ model. Similarly, Tier 1's are not such terrible destinations as for the $\mathbf{LP}_0$ model.

### 5.7.3 Partition Results with $\mathbf{LP}_{50}$ Policy Variant

We continue with the analysis of the $\mathbf{LP}_{50}$ policy variant, where a peer route of any length is preferred to any longer customer route. Throughout our experiments, we found that the maximal peer route length was 26, which confirms that $\mathbf{LP}_{50}$ is equivalent to $\mathbf{LP}_\infty$ in the context of our empirical evaluations.

#### 5.7.3.1 General Partitions

Figure 33 shows the partitions for the $\mathbf{LP}_{50}$ policy variant, for the UCLA and the IXP augmented topologies. As in the $\mathbf{LP}_2$ case, we find that with security $3^{rd}$, only limited improvements in the metric $H_{V,V}(S)$ are possible, relative to the baseline scenario $H_{V,V}(\emptyset)$: $89 - 72 = 17\%$ for the UCLA AS graph, and $92 - 74 = 18$ for the IXP-augmented graph. which is greater than what we saw for any other **LP** models that we looked at. On the other hand, the possible improvements in the security $2^{nd}$ model are lower than for any other **LP** model we looked at: $92 - 72 = 20\%$ for the UCLA AS graph, and $93 - 74 = 19\%$ for the IXP-augmented graph. Thus, the maximum improvements for the security $3^{rd}$ and $2^{nd}$ models are the closest when $k = \infty$. When security is $3^{rd}$, we observe a lower fraction of doomed ASes while the fraction of immune ASes is comparable to that of $\mathbf{LP}_2$, albeit still greater than for $\mathbf{LP}_{k\in\{0,1\}}$ models. This indicates that large $k$ in this case converts more doomed ASes into neutral ASes (protectable ASes when security is $3^{rd}$) that treat peer and customer routes almost equally (*cf.,* Section 4.2.1.3). At the same time, when security is $2^{nd}$, we observe a higher fraction of immune ASes while the fraction of doomed ASes is comparable to that of $\mathbf{LP}_2$, albeit still smaller than for $\mathbf{LP}_{k\in\{0,1\}}$ models. This indicates that large $k$ in this case converts more protectable ASes into immune ASes as they stubbornly stick to their customer and peer routes that they treat almost equally.

#### 5.7.3.2 Partitions by Destination Tier

In Figure 34 we show the partitions broken down by destination tier when security is $2^{nd}$ and $3^{rd}$ for the $\mathbf{LP}_{50}$ policy variants, for the UCLA graph and for the IXP-augmented graph. Similar to the $\mathbf{LP}_{k\in\{1,2\}}$ cases, most of the protectable ASes are stub and SMDG destinations, the higher-degree AS destinations have very few protectable ASes but many more immune ASes, Tier 1 destinations seem to be pretty

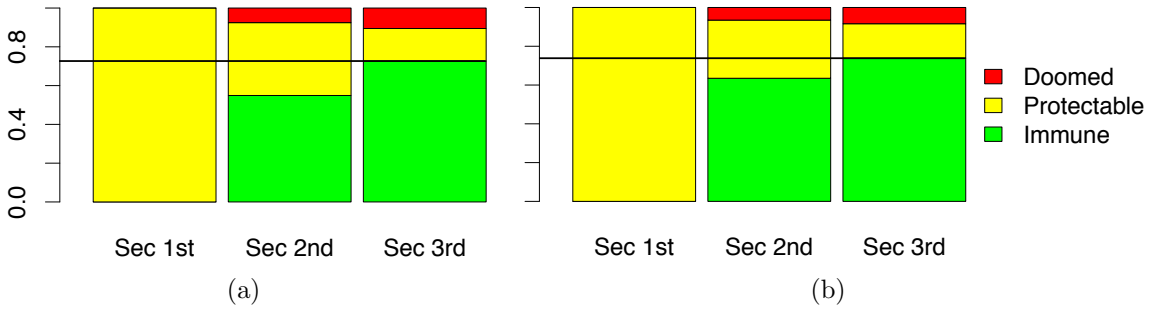


(a) (b)

**Figure 33:** Partitions for the $\mathbf{LP}_{50}$ policy variant, (a) UCLA graph (b) IXP-augmented graph.

good destinations.

### 5.7.4 Summary of Partitions Results for the LP Routing Policy Variants

In this section we summarize our study partitions with respect to various $\mathbf{LP}_k$ policy variants for $k \in \{0, 1, 2, 50\}$. Our analysis confirms that the following high-level observations are robust with respect to the $\mathbf{LP}_{k \in \mathbb{N}}$ local policy variants: (1) possible security metric improvements are meager when security is $3^{rd}$ but are better when security is $2^{nd}$ and (2) Tier 1 ASes are not good candidates for initial deployment. In the rest of the section we also comment on other trends that we observe as we increase $k$ from 0 to 50. We start by discussing our overall partitions analysis and then focus on high-degree destination ASes.

#### 5.7.4.1 General Partitions Analysis Summary

Figure 35 presents a general partitions summary for all the **LP** variants we have empirically evaluated in this thesis, for the security $3^{rd}$ and $2^{nd}$ models. We immediately notice that in terms of immune and happy sources, $\mathbf{LP}_1$ is very similar to $\mathbf{LP}_0$, while $\mathbf{LP}_2$ is very similar to $\mathbf{LP}_{50}$. As $k$ increases from 0 to 50, $k = 2$ plays the role of a breaking point at which there is a noticeable improvement, but after which not much change is observed in terms of immune and happy sources.

We also observe that as $k$ increases in $\mathbf{LP}_k$, for the UCLA and the IXP-augmented

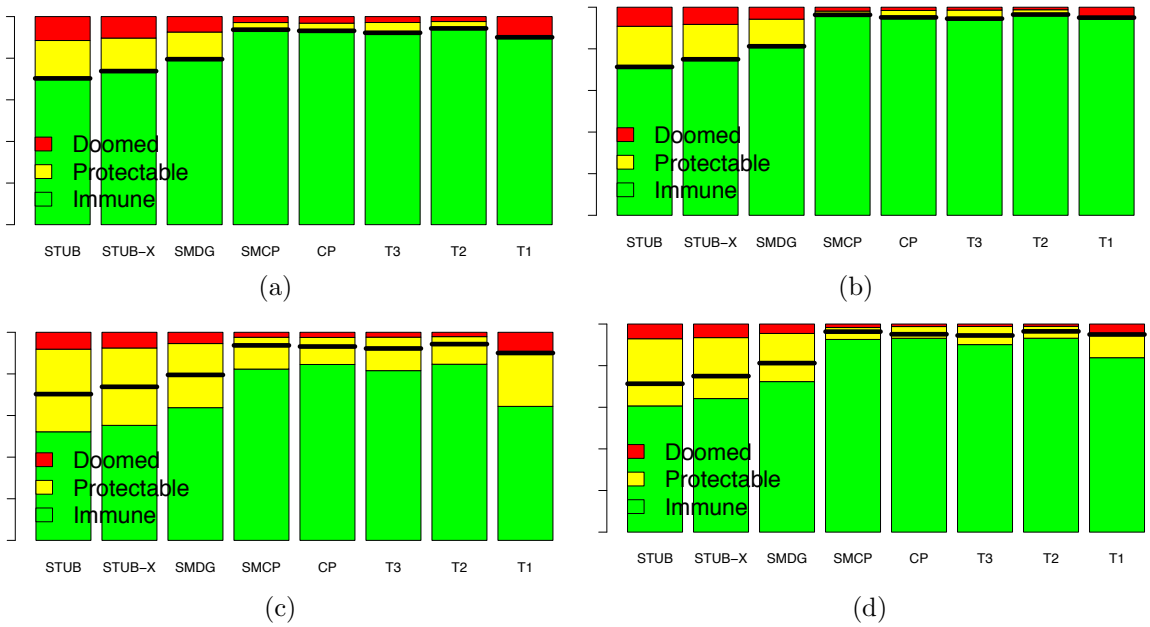


**Figure 34:** Partitions by destination tier for the $\mathbf{LP}_{50}$ policy variant. (a) UCLA graph, security $3^{rd}$. (b) IXP-augmented graph, security $3^{rd}$. (c) UCLA graph, security $2^{nd}$. (d) IXP-augmented graph, security $2^{nd}$. Y-axis runs from 0 to 1.

graphs, the average fraction of

1. doomed ASes decreases when security is $3^{rd}$
2. immune ASes increases when security is $2^{nd}$
3. protectable ASes decreases when security is $2^{nd}$
4. protectable ASes decreases until $k = 2$ and then increases for $k = 50$, by approximately 2% beyond what it is for $k = 0$, when security is $3^{rd}$

The main reason behind these trends is that the attacker on average advertises a route one hop longer than any legitimate destination. Thus, as ASes stubbornly select longer peer routes over (possibly longer) customer routes, either the average fraction of doomed ASes decreases at the expense of an increase in immune or protectable ASes when security is $3^{rd}$ or the average fraction of immune ASes increases at the expense of a decrease in protectable and doomed ASes. As we explained in Section 5.7.3, the sudden increase in the fraction of protectable ASes for $\mathbf{LP}_{50}$ when security is $3^{rd}$ is accompanied by a sudden drop in the fraction of the doomed ASes.

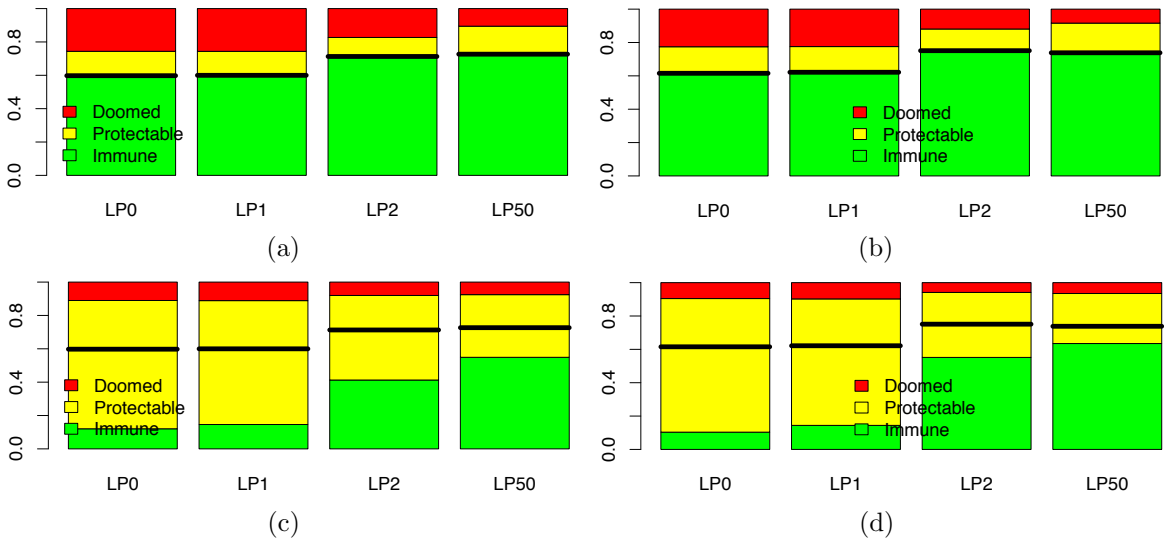


**Figure 35:** Partitions summary for the $\mathbf{LP}_{k\in\{0,1,2,50\}}$ local preference policy variants for (a) UCLA graph, security $3^{rd}$, (b) IXP-augmented graph, security $3^{rd}$, (c) UCLA graph, security $2^{nd}$, and (d) IXP-augmented graph, security $2^{nd}$.

For the UCLA and the IXP-augmented graphs, for any **LP** variant, there is little room for improvement in the security metric beyond the RPKI baseline (noted with horizontal bars) when security is $3^{rd}$, and there is noticeably more room for improvement over the RPKI baseline when security is $2^{nd}$.

When security is $3^{rd}$, for all **LP** policy variants, the fractions of immune and protectable ASes are greater for the IXP-augmented graph than for the UCLA graph, and the opposite holds for the fraction of doomed sources. Similarly, when security is $2^{nd}$, for all **LP** policy variants, the fraction of doomed sources is smaller for the IXP-augmented graph than for the UCLA graph. However, the fraction of immune and protectable ASes for the IXP-augmented graph is smaller and greater than for the UCLA graph respectively when $k < 2$, and vice versa when $k \geq 2$. This suggests that as the number peering edges increases in the AS-level graph (which results in an increase in the number of peering routes available to source ASes) and ASes value route length more than cost (preferring shorter peer routes over longer customer routs), more ASes are likely to be immune to attacks and less impacted by introduction of S*BGP.

#### 5.7.4.2 Focusing on High-Degree Destinations

Figure 36 presents a general partitions summary focusing only on the Tier 1 ASes for all the **LP** variants we have empirically evaluated in this thesis, for the security $3^{rd}$ and $2^{nd}$ models. Here we also observe that as $k$ increases, in $\mathbf{LP}_k$ local preference model the average fraction of immune and doomed ASes increases and decreases respectively for both security $3^{rd}$ and $2^{nd}$ models. We also see from this figure that, there is little room for improvement in the security metric beyond the RPKI baseline when security is $3^{rd}$, and there is noticeably more room for improvement over the RPKI baseline when security is $2^{nd}$. For both security $3^{rd}$ and $2^{nd}$ models, although when $k = 0$ most sources are doomed, most sources are immune for any $k \geq 1$. This makes sense because high-degree ASes, including Tier 1's, are fewer in numbers and have many peering edges. Thus, as ASes prefer shorter peer routes to longer customer routes, they are less likely to fall victim to the attack when trying to reach a high-degree destinations. Since the attacker on average advertises a route one hop longer than any legitimate destination, the effect of an increase of immune ASes at the expense of a decrease of doomed ASes is more pronounced among Tier 1 destinations in $\mathbf{LP}_k$ routing models for $k \geq 1$. This effect was also noticed when

When security is $2^{nd}$ and $3^{rd}$, for all **LP** variants, the fractions of immune and doomed ASes are greater and smaller for the IXP-augmented graph than for the UCLA graph respectively. This makes sense, because as there are more peering edges in the graph, more ASes are less likely to fall victim to the attack since the attacker on average announces a rout 1-hop longer than the destination.

## *5.8 Sensitivity to Routing Policy: Large Deployments*

In this section we recreate analysis of S*BGP deployment rollouts in Section 5.6 for the for $\mathbf{LP}_{k\in\{1,2,50\}}$ local preference models. Specifically, we look at the T1+T2+T3, T1+T2+T3+CP, and T2 rollouts shown in Figures 37, 38, and 39 respectively. Our

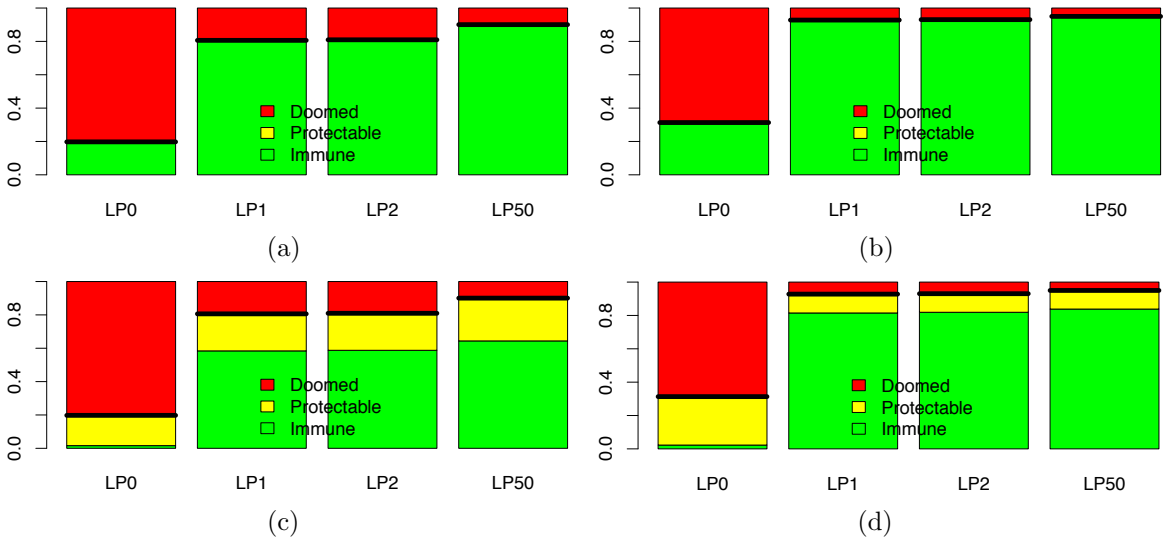


**Figure 36:** Partitions summary for the $\mathbf{LP}_{k\in\{0,1,2,50\}}$ local preference policy variants, strictly for Tier 1 destinations, for (a) UCLA graph, security $3^{rd}$, (b) IXP-augmented graph, security $3^{rd}$, (c) UCLA graph, security $2^{nd}$, and (d) IXP-augmented graph, security $2^{nd}$.

analysis in this section confirms that the high-level observations we made in Section 5.7 are robust with respect to the $\mathbf{LP}_{k\in\mathbb{N}}$ models: (1) security metric improvements are meager when security is $3^{rd}$ and (2) security metric improvements are only slightly better when security is $2^{nd}$.

For all of our rollout experiments with $\mathbf{LP}_{k\in\{0,1,2,50\}}$ routing models, we notice that as $k$ increases, although there appears to be no consistent trend with respect to metric improvements, the following quantities increase or stay approximately the same for all three secure routing models:

1. secure routes under normal conditions,
2. secure routes under attack,
3. happy ASes with secure routes.

We demonstrate these trends for the last step in our T1+T2+T3 rollout in Figure 40. This figure shows that when security is first, as $k$ increases, although there seems to little difference between the counts of ASes with secure routes under normal conditions

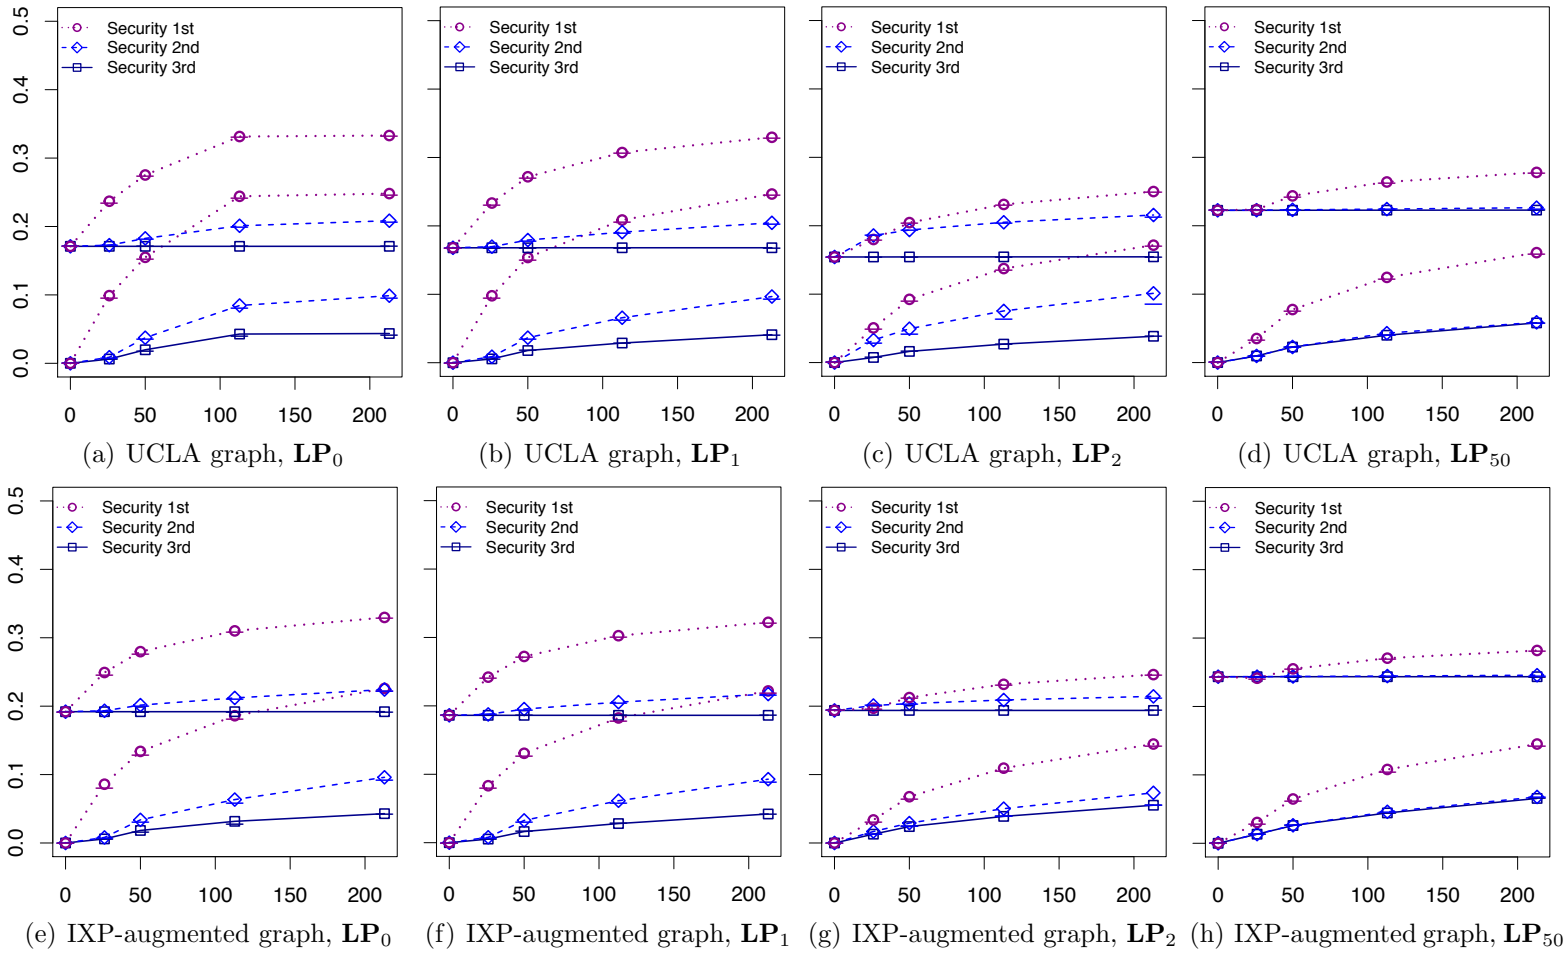


(a) UCLA graph, $\mathbf{LP}_0$ (b) UCLA graph, $\mathbf{LP}_1$ (c) UCLA graph, $\mathbf{LP}_2$ (d) UCLA graph, $\mathbf{LP}_{50}$

(e) IXP-augmented graph, $\mathbf{LP}_0$ (f) IXP-augmented graph, $\mathbf{LP}_1$ (g) IXP-augmented graph, $\mathbf{LP}_2$ (h) IXP-augmented graph, $\mathbf{LP}_{50}$

**Figure 37:** Metric improvements for the Tier 1+2+3 rollout. X-axis is the number of non-stub ASes for each step in the rollout.

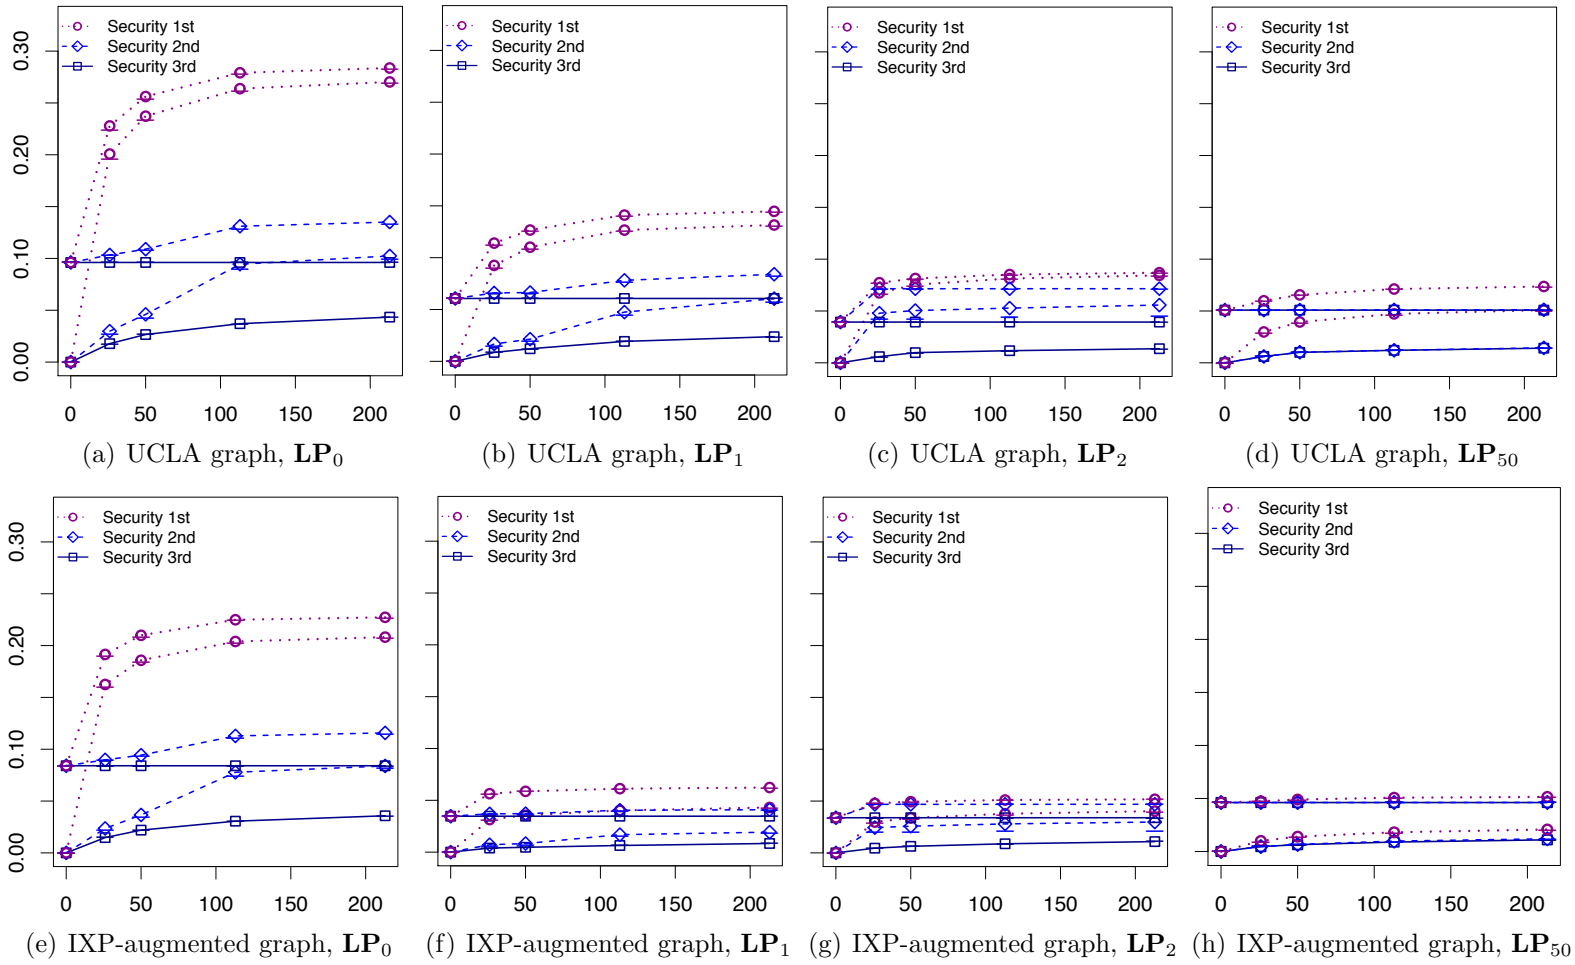


(a) UCLA graph, $\mathbf{LP}_0$ (b) UCLA graph, $\mathbf{LP}_1$ (c) UCLA graph, $\mathbf{LP}_2$ (d) UCLA graph, $\mathbf{LP}_{50}$

(e) IXP-augmented graph, $\mathbf{LP}_0$ (f) IXP-augmented graph, $\mathbf{LP}_1$ (g) IXP-augmented graph, $\mathbf{LP}_2$ (h) IXP-augmented graph, $\mathbf{LP}_{50}$

**Figure 38:** Metric improvements for the Tier 1+2+3+CP rollout. Averaging is done for CP destinations only. X-axis is the number of non-stub, non-CP ASes for each step in the rollout.

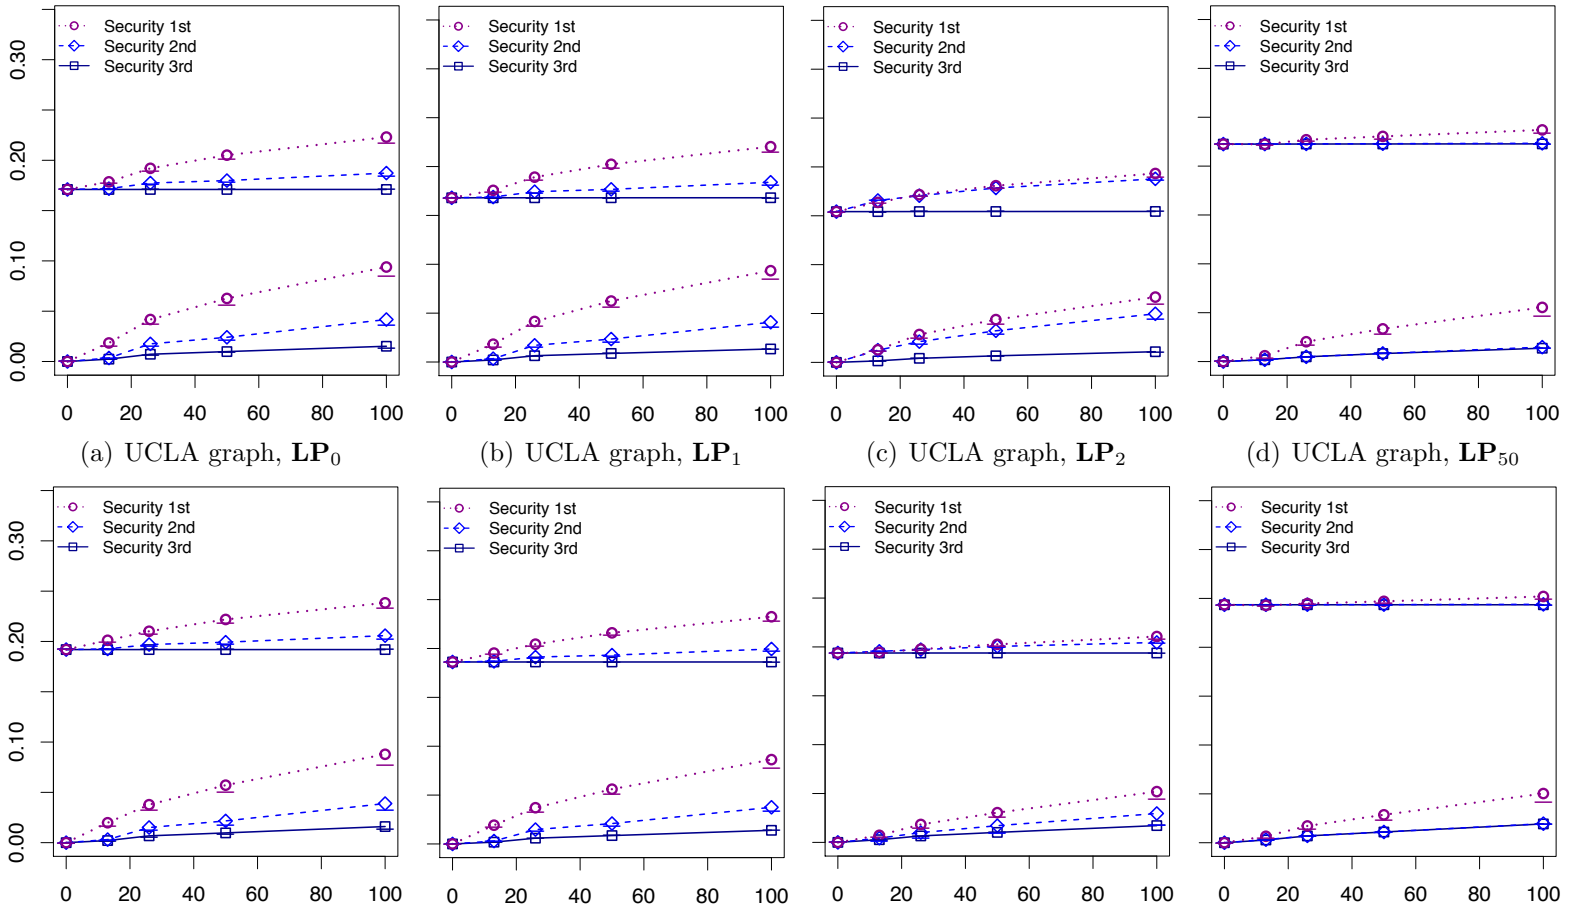


(a) UCLA graph, $\mathbf{LP}_0$ (b) UCLA graph, $\mathbf{LP}_1$ (c) UCLA graph, $\mathbf{LP}_2$ (d) UCLA graph, $\mathbf{LP}_{50}$

(e) IXP-augmented graph, $\mathbf{LP}_0$ (f) IXP-augmented graph, $\mathbf{LP}_1$ (g) IXP-augmented graph, $\mathbf{LP}_2$ (h) IXP-augmented graph, $\mathbf{LP}_{50}$

**Figure 39:** Metric improvements for the Tier 2 rollout. X-axis is the number of non-stub ASes for each step in the rollout.

(which, due to the lack of downgrades in this security prioritization model, implies the same for the counts of ASes with secure routes under attack) for $\mathbf{LP}_{k\in\{0,1,2,50\}}$ models, there is a clear increasing trend of happy ASes with secure routes.

From our empirical evaluations we also observe that although collateral damages do not play much of a role, collateral benefits do play a noticeable role in metric improvements. There also appears to be no consistent trend with respect to these phenomena as $k$ increases.

Although security benefits when security is $2^{nd}$ are noticeably, albeit only slightly, better than when security is $3^{rd}$ with respect to the $\mathbf{LP}_{k\in\{0,1,2\}}$ models, they are almost identical with respect to $\mathbf{LP}_{50}$ model. We also observe that the addition of extra peering links also causes the improvements to decrease. This is particularly evident with CP's in Figure 38.

Such trends are consistent with our partitions analysis from Section 5.7 that indicated an increase in immune ASes and a decrease in doomed ASes for security $2^{nd}$ and $3^{rd}$ with an increase in $k$. As was discussed in Section 5.7, the major reason for such trends is an increase of happy ASes as ASes value route length more in their routing policies. An increase in peering edges in the underlying topology only makes this effect stronger, especially for CP's and other high-degree ASes.

The bottom line is that RPKI forces BGP attackers to announce routes that are on average longer than the routes announced by the legitimate origins. We have seen that because of this, the RPKI baseline is reasonably high (60% and 62% of happy ASes with no S*BGP deployment for the UCLA and the IXP-augmented topologies respectively) for the $\mathbf{LP}_0$ model, and the effect becomes even stronger as ASes value route length higher in their routing policies.

Finally, although our experiments confirm that for the $\mathbf{LP}_1$ model 13 Tier 2's with the highest customer degree are better candidates for initial deployment than 13 Tier 1's (recall similar discussion in Section 5.6.2.1 with respect to $\mathbf{LP}_0$), they result

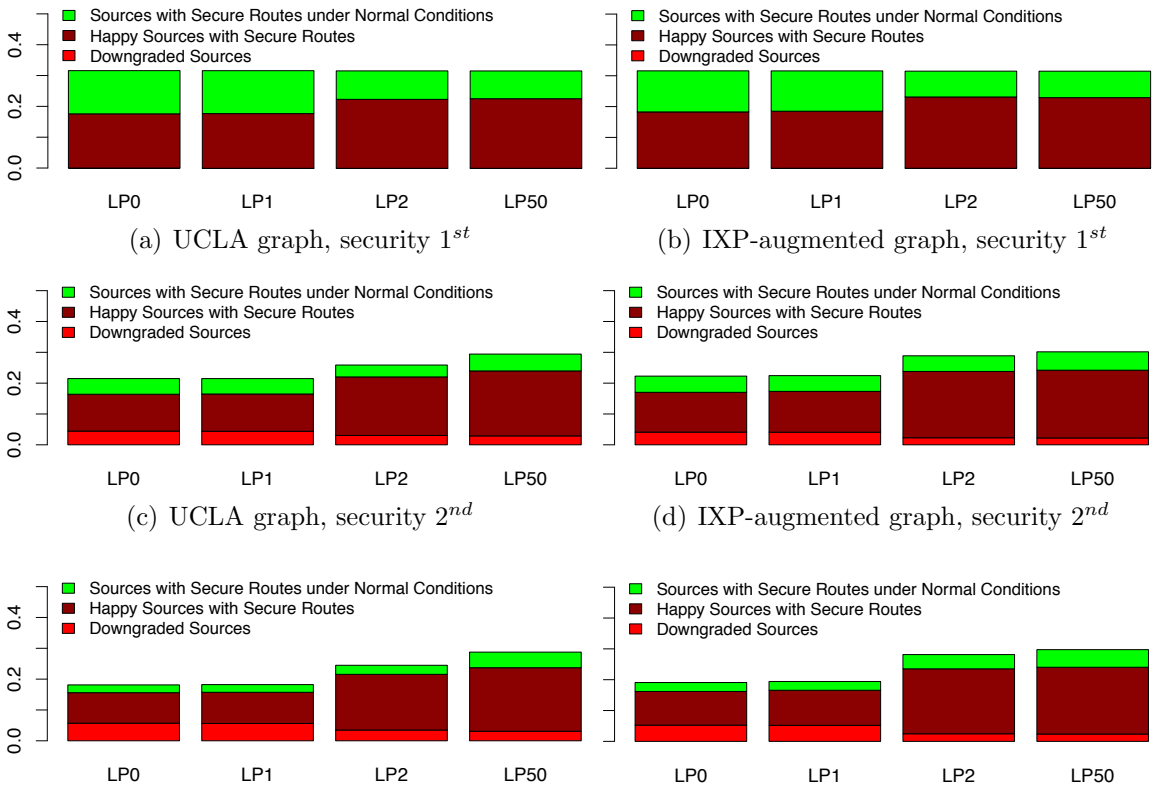


(a) UCLA graph, security $1^{st}$

(b) IXP-augmented graph, security $1^{st}$

(c) UCLA graph, security $2^{nd}$

(d) IXP-augmented graph, security $2^{nd}$

(e) UCLA graph, security $3^{rd}$

(f) IXP-augmented graph, security $3^{rd}$

**Figure 40:** The breakdown of secure routes for the last step in our Tier 1+2+3 rollout, for all security and local preference models.

in approximately equivalent security improvements for the $\mathbf{LP}_{k\in\{2,50\}}$ models when security is $3^{rd}$ and $2^{nd}$.

## *5.9 Concluding Remarks*

In this chapter we developed a framework for quantifying security benefits and possible harm that could result from S*BGP deployments on the Internet. With large scale simulations we quantified the significance of the various complications, such as protocol downgrades and collateral damages, provided upper and lower bounds on the possible S*BGP security benefits, and evaluated many S*BGP deployment scenarios with respect to those bounds, for many different routing models when S*BGP is only partially deployed.

On the one hand, our results give rise to guidelines for facilitating initial S*BGP

deployment on the Internet. Our findings suggest to deploy S*BGP at Tier 2 ISPs rather than at the more connected Tier 1s, because the latter result in only marginal security benefits unless ASes prioritize security above cost and performance considerations (a very unlikely scenario in the initial stages of deployment) in their routing policies. Our findings also suggest that deploying lightweight simplex S*BGP at stub ASes, instead of the full-fledged S*BGP, is a good idea because this reduces deployment complexity at the majority of ASes without compromising overall security benefits.

On the other hand, given the significant effort that may be required to deploy S*BGP, and in light of the possible complications we have highlighted in this chapter, our results cast some doubt about the value of deploying S*BGP. We find that partially-deployed S*BGP provides, on average, only limited security benefits over route origin authentication when ASes do not prioritize security $1^{st}$. The major reason for such disappointing results is that most ASes avoid the attacker even in the scenario of no S*BGP deployment. The second major reason for our results is that many ASes downgrade from secure routes to insecure bogus routes due to their local policies. Thus, the main practical implication of our results in this chapter is that either network operators would have to find a way to deal with protocol downgrades (*e.g.,* prioritize security above all other considerations in their routing policies) or we might have to wait for a very large deployment (*i.e.,* possibly over 60%) before obtaining significant benefits from partially deployed S*BGP.

For these reasons and because deploying S*BGP may require significant effort, route origin authentication, *e.g.,* with RPKI, may be good enough by itself. This is consistent with our suggestion in Chapter 3, in which we said that deployment of some type of link certificates to address route validity attacks may initially make more sense than partially deploying S*BGP, because, like RPKI objects, while providing a certain well-defined level of protection, link certificates could be verified off line.

# CHAPTER VI

# CONCLUSIONS AND FUTURE WORK

## *6.1 Summary of Contributions*

In this thesis we have presented an evaluation of the security benefits and complications of various security-enhanced interdomain routing protocols. Specifically we made the following contributions.

### 6.1.1 Provable Security

We have designed a general security definition for path-vector routing protocols, that captures many BGP security vulnerabilities. With our model we showed that S-BGP provides origin and route authentication but does not provide route validity. We also showed that although SoBGP provides origin authentication, it does not provide route authentication and validity. We have considered protocol modifications and relaxations to our adversarial model to address these weaknesses of S-BGP and SoBGP. We have also demonstrated that although in general it is impossible to have any security guarantees even if there is only one AS without cryptographic keys, it is possible to compensate for the loss of PKI-related security with two new, non-traditional assumptions on trust between neighboring ASes and attacker's operational capabilities of intercepting their communication.

### 6.1.2 Network Stability

We have demonstrated that a stable network can be destabilized by a single fixed-route attacker as well as due to disagreements ASes may have about how to prioritize secure routes to insecure ones when S-BGP or BGPSEC are only partially deployed. We have then presented sufficient conditions, for various routing models,

under which convergence to a unique stable state could be guaranteed irrespective of the number, locations and type of the fixed-route attackers and the characteristics of S-BGP/BGPSEC deployment (full, partial, or no deployment).

### 6.1.3 Quantifying Benefits in Deployment

Using theoretical analysis and large-scale simulations we have studied BGPSEC security benefits over RPKI with respect to multiple partial BGPSEC deployment scenarios, routing models, and underlying Internet AS-level topologies. We have discovered that if network operators do not prioritize security above all other considerations in their routing policies, partial BGPSEC deployments result in only marginal security benefits over those already by the RPKI. We have shown that partial BGPSEC deployments can result in new vulnerabilities such as protocol downgrades, collateral damages, and routing instabilities, and we have discussed strategies for mitigating them.

## *6.2 Discussion of Practical Implications*

The Internet is a very large and complicated system, and experience with the long and slow deployments of new network protocols such as IPv6 and DNSSEC emphasize how important it is to study *usefulness* of communication protocols in full as well as partial deployment scenarios before actually deploying them. For any protocol to survive a long deployment journey, there has to be belief and consensus among network operators on how it should be deployed and how to deal with its pitfalls. We discuss general practical implications that we hope the network operator communities would take away from our studies.

### 6.2.1 Consensus Is Crucial

As we have pointed out in Chapter 4, there seems to be no security prioritization that could avoid protocol downgrades as well as collateral damages. However, network

operators would have to figure out how to reach consensus on how to security should be prioritized to avoid potential routing instabilities and BGP Wedgies as well as on how to deal with protocol downgrades and collateral damages.

### 6.2.2 No Free Lunch

Our analysis in Chapter 3 suggests that there seems to be no easy way to deal with route validity attacks and route leaks in general. Although BGPSEC might be a good candidate solution for dealing with route authentication attacks on BGP, unless network operators are willing to make their routing policies publicly available, it may not be possible to prevent route validity issues, which may continue to cause regular routing outages on the Internet.

Our empirical studies of BGPSEC described in Chapter 5 suggest that if BGPSEC ever gets deployed, unless it is fully deployed right away (which is very unlikely), either network operators would have to prioritize security above all other considerations such as route length and cost, or we may have to wait for a very large deployment, *i.e.,* over 60%, before significant security benefits could be observed. On the other hand, if network operators are willing to make their routing policies publicly available, then, as we pointed out in Chapter 3, it may be more beneficial for ASes to address route validity issues than engage in small partial BGPSEC deployments.

## *6.3 Future Work*

We conclude this thesis by summarizing possible future direction for studying security in interdomain routing.

### 6.3.1 Surmounting Partial S*BGP Deployment Vulnerabilities

In Chapter 5 we have found that partially-deployed S*BGP results in only marginal security benefits over origin authentication if network operators do not prioritize security above all other considerations in their routing policies, while also introducing

new vulnerabilities. We have discovered that protocol downgrades play a major role in this result, so an important direction for future work would be to further investigate ways to limit the impact of this vulnerability without ranking security above all other routing considerations. For example, one could add hysteresis to S*BGP, so that an AS does not immediately drop a secure route when a better insecure route appears. Alternatively, one could investigate deployment scenarios that create connected components of secure ASes that agree to prioritize security $1^{st}$ for routes between ASes in that component. The main challenge of this approach would be to do this without disrupting existing traffic engineering or business arrangements.

### 6.3.2 More Efficient and Deployable Solutions

One of the main criticisms of S-BGP and BGPSEC is that it is very inefficient in terms of processing and communication overhead. There have been various proposals for more efficient route attestation mechanisms in S-BGP [54, 26, 23], and an important direction for future work would be to incorporate such proposals with our framework in order to design more practical and deployable secure routing protocols in a provably secure manner. It is also worth exploring non-cryptographic security solutions that could be deployed in addition to or instead of S*BGP. For example, origin authentication with RPKI combined with anomaly detection and prefix filtering could be easier to deploy, while being possibly as effective as partially-deployed S*BGP.

### 6.3.3 Quantifying Security Benefits in Partial RPKI Deployments

In our studies of BGPSEC deployments in Chapter 5 we have assumed that RPKI is fully deployed, so prefix hijacks are not possible. However, it would be interesting to see what security benefits and vulnerabilities could arise from partial deployments of BGPSEC when RPKI is only partially deployed. Additionally, as we discussed in Chapter 3, it would be interesting to investigate the necessary and sufficient conditions

under which RPKI could be modified to address route validity attacks, and then quantify possible security benefits and harm when it is only partially deployed.

# REFERENCES


[1] "American Registry for Internet Numbers (ARIN)." https://www.arin.net/.

[2] "The Internet Assigned Numbers Authority (IANA)." http://www.iana.org/.

[3] "Resource Public Key Infrastructure (RPKI)." https://www.arin.net/resources/rpki.html.

[4] "IRR power tools." http://sourceforge.net/projects/irrpt/, 2011.

[5] "Working group 6 secure bgp deployment report," tech. rep., FCC CSRIC http://transition.fcc.gov/bureaus/pshs/advisory/csric3/CSRICIII_9-12-12_WG6-Final-Report.pdf, 2012.

[6] Adams, C. and Farrell, S., "Internet X.509 Public Key Infrastructure: Certificate management protocols," 2004.

[7] Ager, B., Chatzis, N., Feldman, A., Sarrar, N., Uhlig, S., and Willinger, W., "Anatomy of a large European IXP," in *ACM SIGCOMM 2012*, Aug. 2012.

[8] Ager, B., Chatzis, N., Feldmann, A., Sarrar, N., Uhlig, S., and Willinger, W., "Anatomy of a large european IXP," in *ACM SIGCOMM*, 2012.

[9] Aiello, W., Ioannidis, J., and McDaniel, P., "Origin authentication in interdomain routing," in *CCS '03: Proceedings of the 10th ACM conference on Computer and communications security*, (New York, NY, USA), pp. 165–178, ACM Press, 2003.

[10] Alexa, "The top 500 sites on the web." http://www.alexa.com/topsites, October 1 2012.

[11] Augustin, B., Krishnamurthy, B., and Willinger, W., "IXPs: Mapped?," in *ACM IMC*, 2009.

[12] Avramopoulos, I., Suchara, M., and Rexford, J., "How small groups can secure interdomain routing," tech. rep., Princeton University Comp. Sci., 2007.

[13] Backes, M., Cervesato, I., Jaggard, A. D., A.Scedrov, and Tsay, J.-K., "Cryptographically sound security proofs for basic and public-key Kerberos," in *ESORICS* (Gollmann, D., Meier, J., and Sabelfeld, A., eds.), vol. 4189 of *Lecture Notes in Computer Science*, pp. 362–383, Springer, 2006.

[14] Ballani, H., Francis, P., and Zhang, X., "A study of prefix hijacking and interception in the Internet," in *SIGCOMM'07*, 2007.

[15] Barak, B., Goldberg, S., and Xiao, D., "Protocols and lower bounds for failure localization in the Internet," in *EUROCRYPT 2008*, Apr. 2008.

[16] Barbir, A., Murphy, S., and Yang, Y., "Generic threats to routing protocols," *Network Working Group. IETF Request for Comments: 3962. Available at* `http://www.ietf.org/rfc/rfc4593.txt`, 2004.

[17] Bellare, M., Kohno, T., and Namprempre, C., "Authenticated encryption in SSH: provably fixing the SSH binary packet protocol," in *CCS '02*, ACM Press, 2002.

[18] Bellovin, S. M. and Gansner, E. R., "Using link cuts to attack internet routing," in *Tech. Rep., ATT Research, 2004, Work in Progress 2003 USENIX*, 2003.

[19] blog, M., "Hijacking of public DNS servers in Turkey, through routing."

[20] Boldyreva, A., Fischlin, M., Palacio, A., and Warinschi, B., "A closer look at PKI: Security and efficiency," in *Public Key Cryptography* (Okamoto, T. and Wang, X., eds.), vol. 4450 of *Lecture Notes in Computer Science*, pp. 458–475, Springer, 2007.

[21] Boldyreva, A. and Kumar, V., "Extended abstract: Provable-security analysis of authenticated encryption in Kerberos," in *IEEE Symposium on Security and Privacy*, pp. 92–100, IEEE Computer Society, 2007.

[22] Boldyreva, A. and Lychev, R., "Provable Security of (S-BGP) and other Path Vector Protocols: Model, Analysis, and Extenstions," in *ACM CCS 2012*.

[23] Brogle, K., Goldberg, S., and Reyzin, L., "Sequential Aggregate Signatures with Lazy Verification from Trapdoor Permutations," in *ASIACRYPT 2012*.

[24] Brown, M. A., "Renesys blog. Pakistan hijacks YouTube." `http://www.renesys.com/blog/2008/02/pakistan_hijacks_youtube_1.shtml`.

[25] Butler, K., Farley, T., McDaniel, P., and Rexford, J., "A survey of BGP security issues and solutions." Technical Report TD-5UGJ33, AT&T Labs, 2004.

[26] Butler, K., McDaniel, P., and Aiello, W., "Optimizing BGP security by exploiting path stability," in *CCS '06: Proceedings of the 13th ACM conference on Computer and communications security*, (New York, NY, USA), pp. 298–310, ACM Press, 2006.

[27] Chan, H., Dash, D., Perrig, A., and Zhang, H., "Modeling adoptability of secure BGP protocols," in *SIGMETRICS '06/Performance '06: Proceedings of the joint international conference on Measurement and modeling of computer systems*, (New York, NY, USA), pp. 389–390, ACM Press, 2006.

[28] Chang, H., Dash, D., Perrig, A., and Zhang, H., "Modeling adoptability of secure BGP protocol," in *SIGCOMM'06*, 2006.

[29] Chen, C., Jia, L., Loo, B. T., and Zhou, W., "Reduction-based security analysis of internet routing protocols," in *WRiPE*, 2012.

[30] Chi, Y.-J., Oliveira, R., and Zhang, L., "Cyclops: The internet as-level observatory," *ACM SIGCOMM Computer Communication Review*, 2008.

[31] Chi, Y. J., Oliveira, R., and Zhang, L., "Cyclops: The Internet AS-level observatory," *SIGCOMM CCR*, 2008.

[32] Chiesa, M., Battista, G. D., Erlebach, T., and Patrignani, M., "Computational complexity of traffic hijacking under bgp and s-bgp," in *ICALP 2012*, July 2012.

[33] Cisco, "Bgp best path selection algorithm: How the best path algorithm works." Document ID: 13753, 2012. `http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a0080094431.shtml#bestpath`.

[34] Cowie, J., "Renesys blog. China's 18-minute mystery." `http://www.renesys.com/blog/2010/11/chinas-18-minute-mystery.shtml`.

[35] Degarbiele, J. P. and Paterson, K. G., "On the (In)security of IPsec in MAC-then-encrypt configurations," in *ACM CCS 2010*.

[36] Dhamdhere, A. and Dovrolis, C., "Twelve years in the evolution of the internet ecosystem," *Trans. Netw.*, 2011.

[37] Dolev, S. and Tzachar, N., "Empire of colonies: Self-stabilizing and self-organizing distributed algorithms," in *OPODIS*, pp. 230–243, 2006.

[38] Gao, L. and Rexford, J., "Stable internet routing without global coordination," *SIGMETRICS Perform. Eval. Rev.*, vol. 28, pp. 307–317, June 2000.

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

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

[41] Gill, P., Schapira, M., and Goldberg, S., "Let the market drive deployment: A strategy for transitioning to BGP security," in *ACM SIGCOMM 2011*, Aug. 2011.

[42] Gill, P., Goldberg, S., and Schapira, M., "Nanog'56. A survey of interdomain routing policies," 2012.

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

[44] Goldberg, S., Halevi, S., Jaggard, A., Ramachandran, V., and Wright, R., "Rationality and traffic attraction: Incentives for honestly announcing paths in BGP," in *ACM SIGCOMM 2008*, Aug. 2008.

[45] Goldberg, S., Halevi, S., Jaggard, A. D., Ramachandran, V., and Wright, R. N., "Rationality and traffic attraction: Incentives for honest path announcements in BGP," in *SIGCOMM'08*, 2008.

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

[47] Goldberg, S., Schapira, M., Hummon, P., and Rexford, J., "How secure are secure interdomain routing protocols?," in *ACM SIGCOMM 2010*, Aug. 2010.

[48] Goldberg, S., Xiao, D., Barak, B., Rexford, J., and Tromer, E., "Path-quality monitoring in the presence of adversaries," in *ACM SIGMETRICS 2008*, June 2008.

[49] Goodell, G., Aiello, W., Griffin, T., Ioannidis, J., McDaniel, P., and Rubin, A., "Working around bgp: An incremental approach to improving security and accuracy in interdomain routing," 2003.

[50] Griffin, T. and Huston, G., "BGP wedgies." RFC 4264, 2005.

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

[52] Griffin, T. and Wilfong, G., "A safe path vector protocol," in *IEEE INFOCOM 2000*.

[53] Griffin, T. and Wilfong, G., "An analysis of BGP convergence properties," in *ACM SIGCOMM*, 1999.

[54] Hu, Y., Perrig, A., and Johnson, D., "Efficient security mechanisms for routing protocols," 2003.

[55] Hu, Y.-C., Perrig, A., and Sirbu, M., "SPV: secure path vector routing for securing BGP," in *SIGCOMM '04: Proceedings of the 2004 conference on Applications, technologies, architectures, and protocols for computer communications*, (New York, NY, USA), pp. 179–192, ACM Press, 2004.

[56] Huston, G., "Peering and settlements - Part I," *The Internet Protocol Journal (Cisco)*, 1999.

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

[58] Kent, S. and Chi, A., "Threat model for bgp path security." Internet draft: draft-ietf-sidr-bgpsec-threats-04, 2013.

[59] Kent, S., Lynn, C., and Seo, K., "Secure border gateway protocol (S-BGP)," *JSAC*, 2000.

[60] Kent, S., Lynn, C., and Seo, K., "Secure border gateway protocol (S-BGP)," *IEEE Journal on Selected Areas in Communications*, vol. 18, no. 4, pp. 582–592, 2000.

[61] Kent, S. and Seo, K., "Security architecture for the internet protocol." IETF RFC 4301, 2005. Available at `http://tools.ietf.org/html/rfc4301#page-4`.

[62] Kent, S. T., Lynn, C., Mikkelson, J., and Seo, K., "Secure Border Gateway Protocol (S-BGP) - real world performance and deployment issues," in *NDSS*, The Internet Society, 2000.

[63] Kushman, N., Kandula, S., and Katabi, D., "Can you hear me now?!: it must be BGP," *SIGCOMM CCR*, 2007.

[64] Labovitz, C., "Arbor blog: Battle of the hyper giants." `http://asert.arbornetworks.com/2010/04/the-battle-of-the-hyper-giants-part-i-2/`.

[65] Labovitz, C., Ahuja, A., and Jahanian, F., "Experimental study of internet stability and backbone failures," in *Fault-Tolerant Computing*, 1999.

[66] Labovitz, C., "Internet traffic 2007 - 2011.," 2011. Global Peering Forum. Santi Monica, CA. `http://www.monkey.org/~labovit/papers/gpf_2011.pdf`.

[67] Labovitz, C., Iekel-Johnson, S., McPhersony, D., Oberheiden, J., and Jahanian, F., "Internet inter-domain traffic," in *ACM SIGCOMM*, 2010.

[68] Lepinski, M., "Bgpsec protocol specification: draft-ietf-sidr-bgpsec-protocol-06." Internet-Draft, 2012.

[69] Lepinski, M., "BGPSEC Protocol Specification (v4)." IETF Internet Draft, 2012. Available at `http://tools.ietf.org/html/draft-ietf-sidr-bgpsec-protocol-04`.

[70] Lepinski, M., "An infrustructure to support secure internet routing." IETF RFC 6480, 2012. `http://www.tools.ietf.org/html/rfc6480`.

[71] Lychev, R., Goldberg, S., and Schapira, M., "Network Destabilizing Attacks," in *ACM PODC 2012*, July 2012.

[72] Lychev, R., Goldberg, S., and Schapira, M., "Network destabilizing attacks." Arxiv Report 1203.1281, march 2012.

[73] Lychev, R., Goldberg, S., and Schapira, M., "BGP Security in Partial Deployment: Is the Juice Worth the Squeeze?," in *ACM SIGCOMM 2013*, Aug. 2013.

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

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

[76] McPherson, D., Amante, S., Osterweil, E., and Mitchell, D., "Route-Leaks and MITM Attacks Against BGPSEC." IETF Internet Draft, Aug. 2013. Available at `http://tools.ietf.org/html/draft-ietf-grow-simple-leak-attack-bgpsec-no-help-02`.

[77] Misel, S. A., "Wow, as7007!," 1997. `http://www.merit.edu/mail.archives/nanog/1997-04/msg00340.html`.

[78] Mityagin, A., Panjwani, S., and Raghavan, B., "Analysis of the SPV secure routing protocol." Cryptology ePrint Archive, Report 2006/087, 2006. `http://eprint.iacr.org/`.

[79] Murphy, S., "BGP security vulnerabilities analysis," *Network Working Group. IETF Request for Comments: 3962. Available at* `http://www.ietf.org/rfc/rfc4272.txt`, 2006.

[80] Network, M., "Internet Routing Registry." `http://www.irr.net`.

[81] Osterweil, E., Amante, S., and McPherson, D., "TASRS: Towards a Secure Routing System Through Internet Number Resource Certification." Verisign Labs Technical Report 1130009, 2013.

[82] Palse, P., "Serving ROAs as RPSL route[6] Objects from the RIPE Database." RIPE Labs, June 2010. `https://labs.ripe.net/Members/Paul_P_/content-serving-roas-rpsl-route-objects`.

[83] Paseka, T., "Cloudflare blog: Why google went offline today," Nov. 2012. `http://blog.cloudflare.com/why-google-went-offline-today-and-a-bit-about`.

[84] Paterson, K. G. and Watson, G. J., "Plaintext-dependent decryption: A formal security treatment of SSH-CTR," in *EUROCRYPT* (H.Gilbert, ed.), vol. 6110 of *Lecture Notes in Computer Science*, pp. 345–361, Springer, 2010.

[85] Pilosov, A. and Kapela, T., "Stealing the Internet: An Internet-scale man in the middle attack," 2008. DEFCON'16.

[86] Renesys Blog, "Con-Ed steals the 'Net." http://www.renesys.com/blog/2006/01/coned_steals_the_net.shtml.

[87] Reuters, "Internet providers pledge anti-botnet effort," March 22 2012.

[88] Rikhter, Y., Li, T., and Hares, S., "A Border Gateway Protocol 4 (BGP-4)," *Network Working Group. IETF Request for Comments: 4271. Available at* http://www.ietf.org/rfc/rfc4271.txt, 2006.

[89] Ristenpart, T. and Yilek, S., "The power of proofs-of-possession: Securing multiparty signatures against rogue-key attacks," in *EUROCRYPT* (Naor, M., ed.), vol. 4515 of *Lecture Notes in Computer Science*, pp. 228–245, Springer, 2007.

[90] Rogaway, P., "Practice-oriented provable security and the social construction of cryptography." University of California Davis, 2009. Available at http://www.cs.ucdavis.edu/~rogaway/papers/cc.pdf.

[91] Roughan, M., Willinger, W., Maennel, O., Perouli, D., and Bush, R., "10 lessons from 10 years of measuring and modeling the internet's autonomous systems," *IEE JSAC*, vol. 29, no. 9, pp. 1810–1821, 2011.

[92] Sami, R. S., Schapira, M., and Zohar, A., "Searching for stability in interdomain routing," in *IEEE INFOCOM*, 2009.

[93] Sandvine, "Global internet phenomena report. 1h 2012.," 2012.

[94] Schuchard, M., Thompson, C., Hopper, N., and Kim, Y., "Taking routers off their meds: Why assumptions of bgp stability are dangerous," in *NDSS 2012*, Feb. 2012.

[95] Sobrinho, J. L., "Network routing with path vector protocols: Theory and applications," in *ACM SIGCOMM 2003*, pp. 49–60, Aug. 2003.

[96] Song, Y., Venkataramani, A., and Gao, L., "Identifying and addressing protocol manipulation attacks in "Secure" BGP"," in *ICDCS*, 2013.

[97] Sriram, K., "BGPSEC design choices and summary of supporting discussions." Internet-Draft: draft-sriram-bgpsec-design-choices-05, 2014.

[98] Subramanian, L., Roth, V., Stoica, I., Shenker, S., and Katz, R., "Listen and whisper: Security mechanisms for BGP," in *Proc. First Symposium on Networked Systems Design and Implementation (NSDI)*, (San Francisco, CA), Mar. 2004.

[99] Sundaresan, S., Lychev, R., and Valancius, V., "Preventing attacks on BGP policies: One bit is enough." Technical Report GT-CS-11-07, Georgia Institute of Technology, 2011.

[100] The BGP TTL Security Hack. `http://tools.ietf.org/html/draft-gill-btsh-02`.

[101] Varadhan, K., Govindan, R., and Estrin, D., "Persistent route oscillations in inter-domain routing," *Computer Networks*, vol. 32, no. 1, pp. 1–16, 2000.

[102] Wan, T., Kranakis, E., and van Oorschot, P. C., "Pretty secure BGP, psBGP.," in *NDSS*, The Internet Society, 2005.

[103] White, R., "Deployment considerations for secure origin BGP (soBGP)." draft-white-sobgp-bgp-deployment-01.txt, June 2003, expired.

[104] White, R., "Securing BGP through secure origin BGP," *The Internet Protocol Journal*, vol. 6, Sept. 2003. `http://www.cisco.com/web/about/ac123/ac147/archived_issues/ipj_6-3/ipj_6-3.pdf`.

[105] Zhang, X., Hsiao, H.-C., Hasker, G., Chan, H., Perrig, A., and Andersen, D. G., "Scion: Scalability, control, and isolation on next-generation networks," *In Proceedings of IEEE Symposium on Security and Privacy (Oakland)*.

[106] Zhao, M., Smith, S. W., and Nicol, D. M., "Aggregated path authentication for efficient BGP security," in *CCS '05: Proceedings of the 12th ACM conference on Computer and communications security*, (New York, NY, USA), pp. 128–138, ACM Press, 2005.

[107] Zhao, M., Zhou, W., Gurney, A. J. T., Haeberlen, A., Sherr, M., and Loo, B. T., "Private and verifiable interdomain routing decisions," in *ACM SIGCOMM*, 2012.