ROV-MI: Large-Scale, Accurate and Efficient Measurement of ROV Deployment

WENQI CHEN ET AL., TSINGHUA UNIVERSITY

PRESENTER: ZHAN JIN

DATE: MAR. 11, 2025

The Problem with BGP

- **B**order **G**ateway **P**rotocol (BGP) connects **A**utonomous **S**ystems (ASes).
- BGP ***doesn't*** authenticate routes: ASes can announce ***any*** prefix.
- **Prefix Hijacking!**

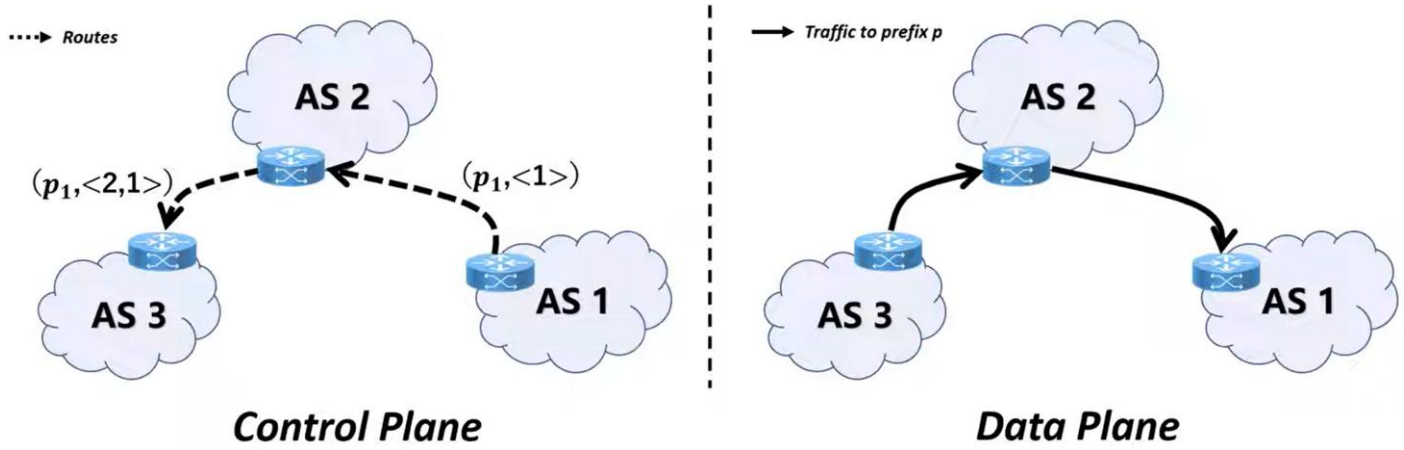

BGP Prefix Hijacking

- Announce a route containing an ***invalid prefix p***
- ***Hijacking the traffic to p*** in data plane

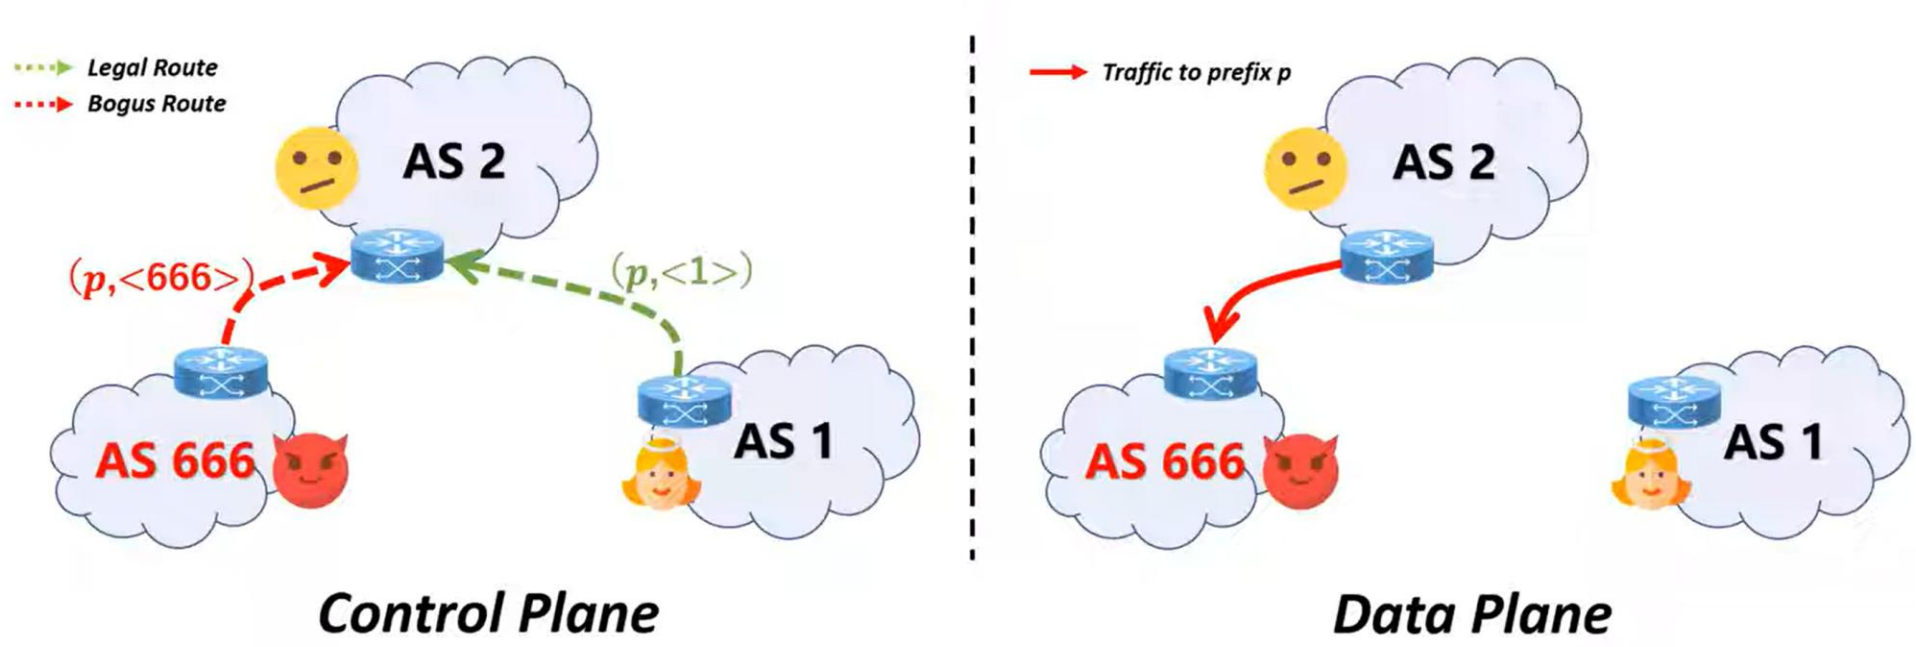

BGP Prefix Hijacking

- No authentication means **problems**:
  - Accidental misconfigurations (route leaks).
  - Malicious hijacking (e.g., Pakistan Telecom/YouTube, 2008).
  - Man-in-the-middle attacks.

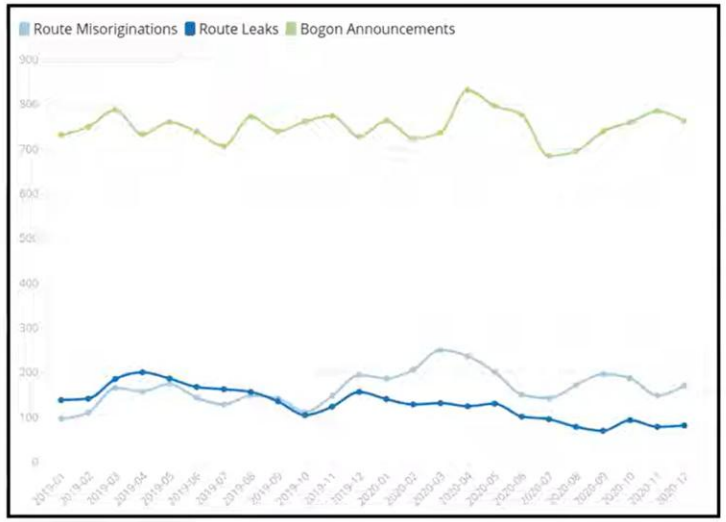

Statistics of BGP Hijacking Incidents in 2020

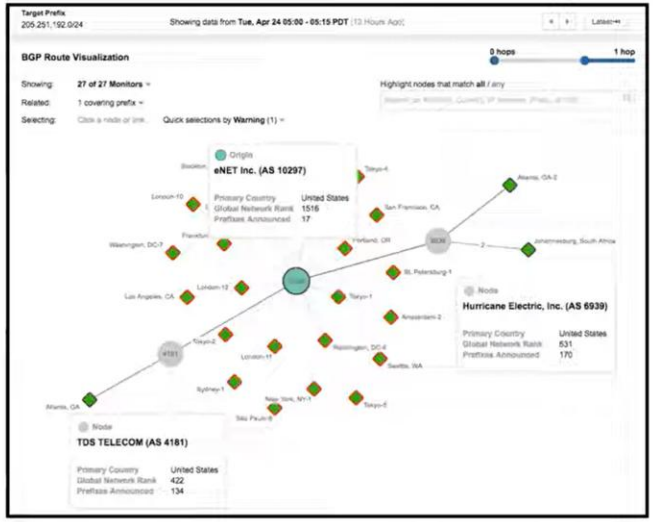

Hijacking towards MyEtherWallet in 2018

BGP Prefix Hijacking

- The **first** documented case of a BGP-based man-in-the-middle attack

Image source: https://www.kentik.com/blog/a-brief-history-of-the-internets-biggest-bgp-incidents/

BGP Prefix Hijacking

- No authentication means **problems**:
  - Accidental misconfigurations (route leaks).
  - Malicious hijacking (e.g., Pakistan Telecom/YouTube, 2008).
  - Man-in-the-middle attacks.
- **Impact:**
  - Reachability failures
  - data interception
  - network disruption
  - Undermining other Internet infrastructures (e.g., DNS)

RPKI: Securing BGP with Cryptography

- **R**esource **P**ublic **K**ey **I**nfrastructure (RPKI) aims to secure BGP.
- Providing ***authentication mechanism*** with **P**ublic **K**ey **I**nfrastructure (PKI).

RPKI: Securing BGP with Cryptography

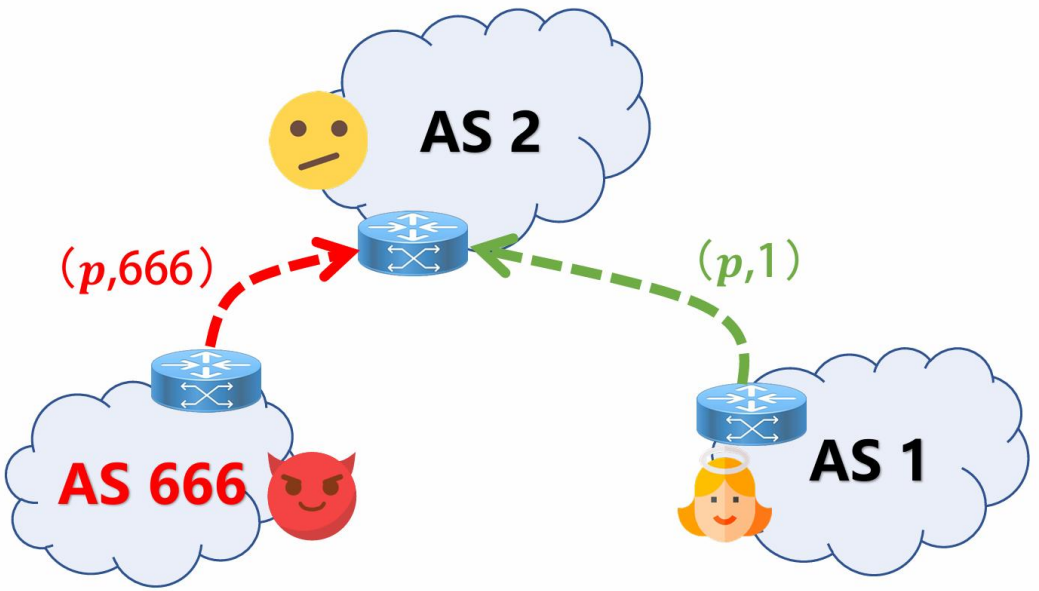

- Two Key Components**:**
  - **ROA (Route Origin Authorization):** A digitally signed record stating which AS is ***authorized*** to originate a prefix. Stored in a centralized, trusted repository (**R**egional **I**nternet **R**egistries, RIRs).
  - **ROV (Route Origin Validation):** ASes ***check*** incoming BGP updates against ROA records and ***filter*** invalid ones.

RPKI: Securing BGP with Cryptography

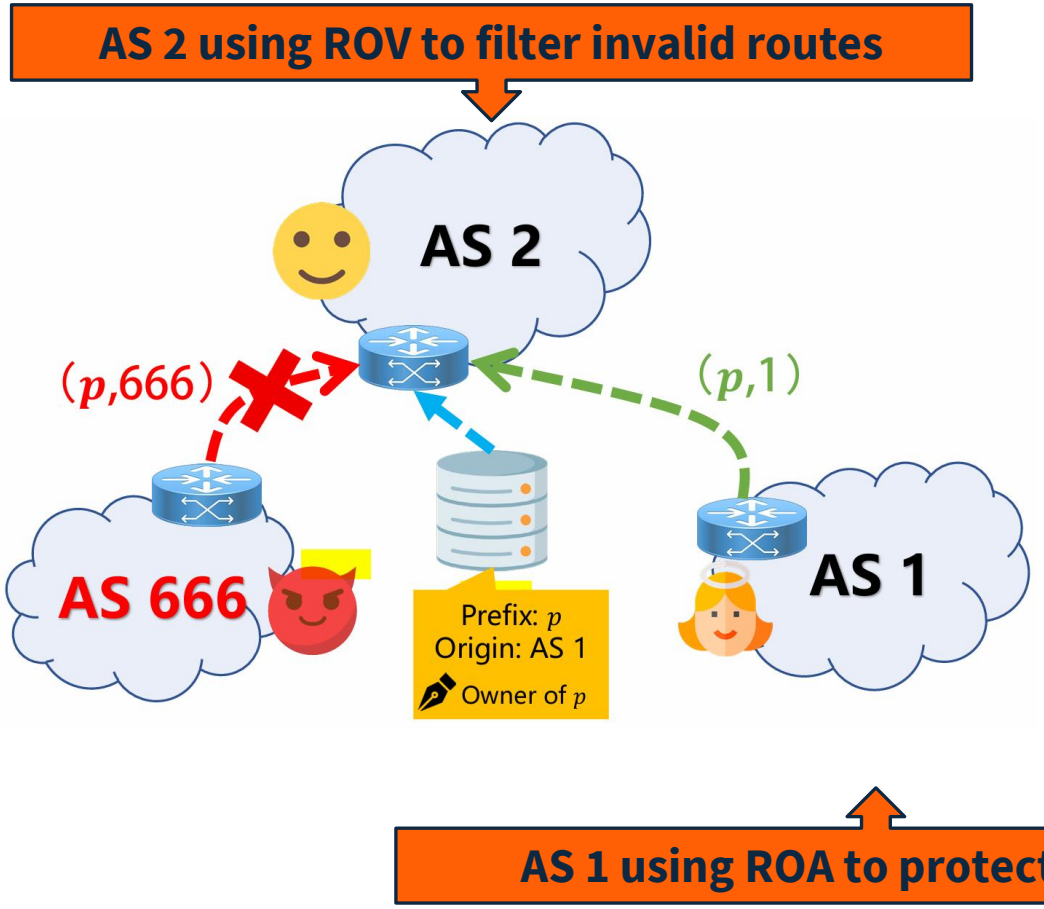

- Two Key Components**:**
  - **ROA (Route Origin Authorization):** A digitally signed record stating which AS is ***authorized*** to originate a prefix. Stored in a centralized, trusted repository (**R**egional **I**nternet **R**egistries, RIRs).
  - **ROV (Route Origin Validation):** ASes ***check*** incoming BGP updates against ROA records and ***filter*** invalid ones.

Deployment of RPKI

RPKI is getting ***more widely deployed***.

*"Recent reports show a promising trend in RPKI deployment, with over **30%** of prefixes all over the Internet registered in PRKI by April, 2021."*

-- Wenqi Chen et al.,

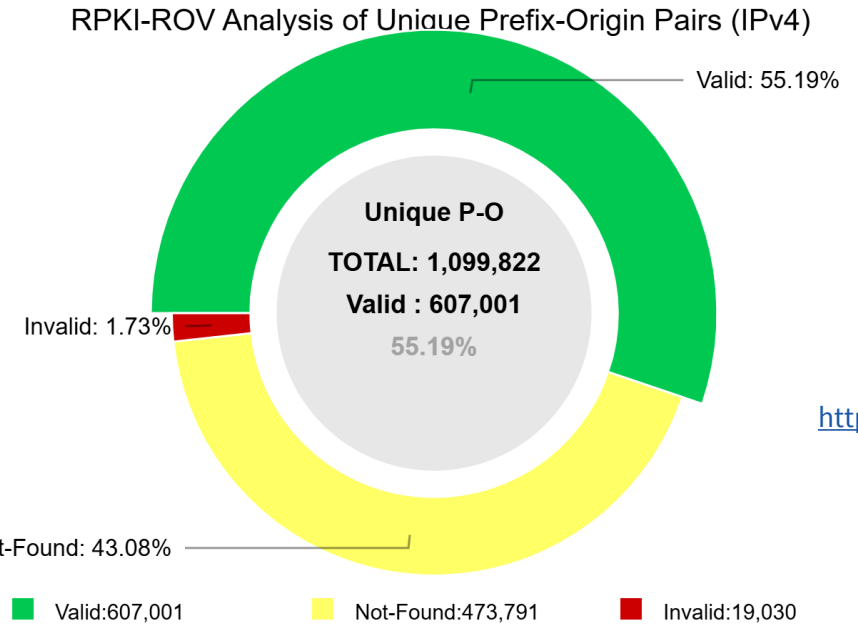

Image source: https://rpki-monitor.antd.nist.gov/

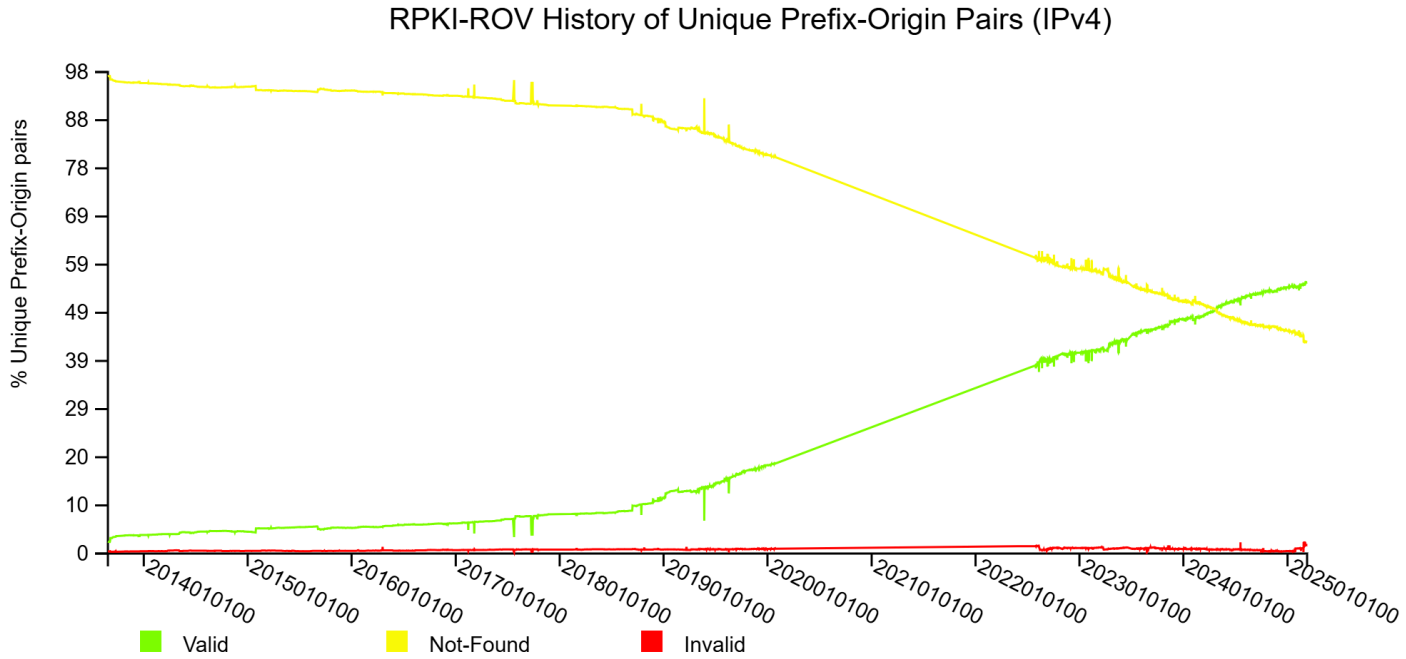

NIST RPKI Monitor: RPKI-ROV Analysis **Protocol:** IPv4 **RIR:** All **URL:** https://rpki-monitor.antd.nist.gov/ROV#div2

Deployment of ROV remains unclear

ROA Deployment	ROV Deployment
• **Public database** ○ Directly available data in the centralized PRKI repository • **Easy to analyze** ○ Can direct know which Ases/prefixes are protected by ROA	• **Private Configuration** ○ No central database. Need to observe the propagation of invalid routes along AS paths • **Hard to infer** ○ Hard to pinpoint the Ases that adopts ROV

Why is Measuring ROV Deployment Difficult?

- **Challenge 1: Large-scale measurement**
  - Limited origin of invalid routes (PEERING testbed, only direct peers)
  - Control plane observation: highly dependent on the vantage points
- **Challenge 2: Accurate and efficient inference**
  - Heuristic methods: Low accuracy (easily confused by other filtering).
  - PEERING + MCMC: Still limited scope; MCMC is slow and scales poorly.

Goal

Need a large-scale, accurate, and efficient way to measure ROV.

ROV-MI: A New Framework for ROV Measurement

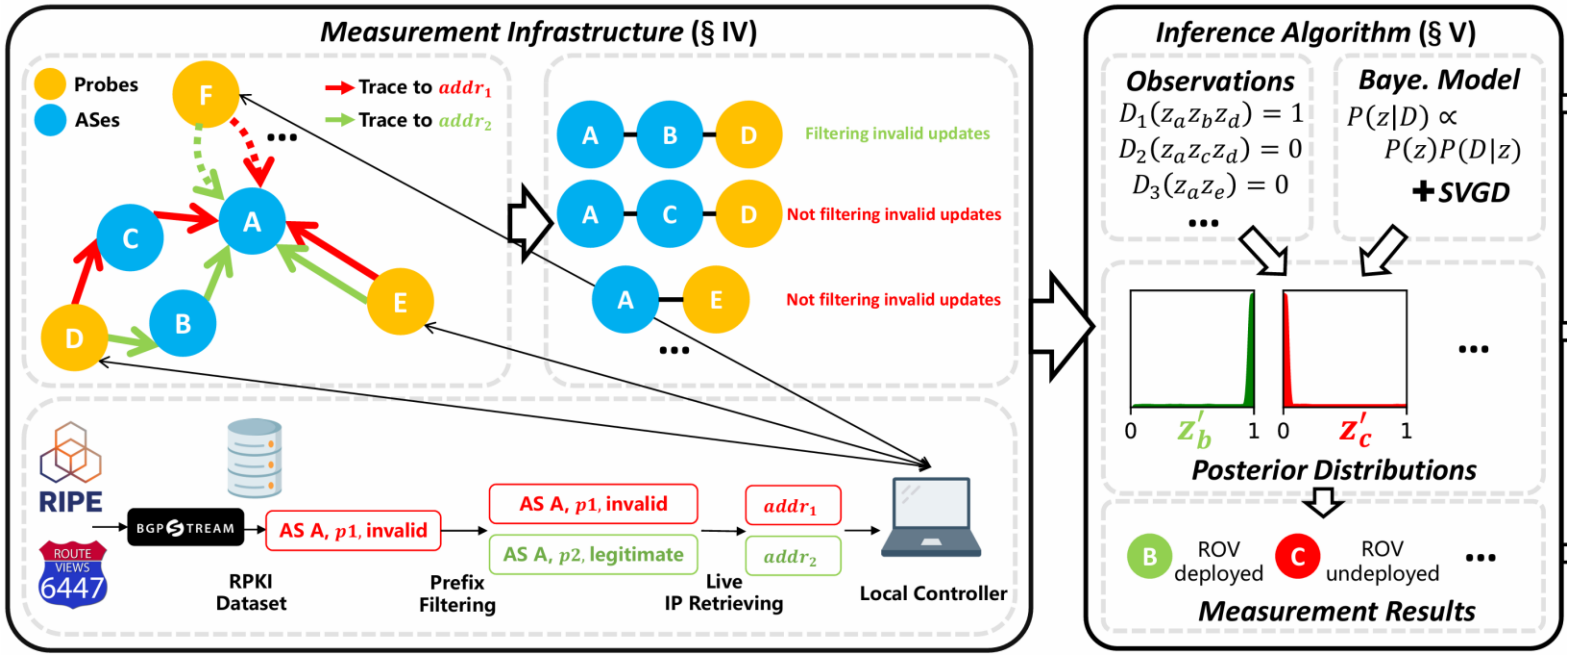

- **Large-Scale Measurement Infrastructure:**
  - Uses ***in-the-wild*** invalid prefixes (from BGPStream). **Key Innovation!**
  - Active probing (traceroute) to label paths: filtering or not filtering?
- **Accurate and Efficient Inference Algorithm:**
  - Bayesian inference model (probabilistic).
  - ***Stein Variational Gradient Descent (SVGD)***: Faster and more scalable than MCMC. **Key Innovation!**

Measurement Infrastructure: Finding Invalid Prefixes

- **Challenge:** PEERING testbed only has a few origins.
- **Solution:** Use real invalid prefixes observed in global BGP updates!
  - Source: BGPStream (collects updates from RIPE RIS and RouteViews).
  - Validate prefixes against ROA records (using Routinator).
- Benefit: ~10x more origins, leading to **~10x more paths**.
- Filtering Multi-homing prefixes and those covered by other legitimate prefixes are removed.

Measurement Infrastructure: Labeling Paths

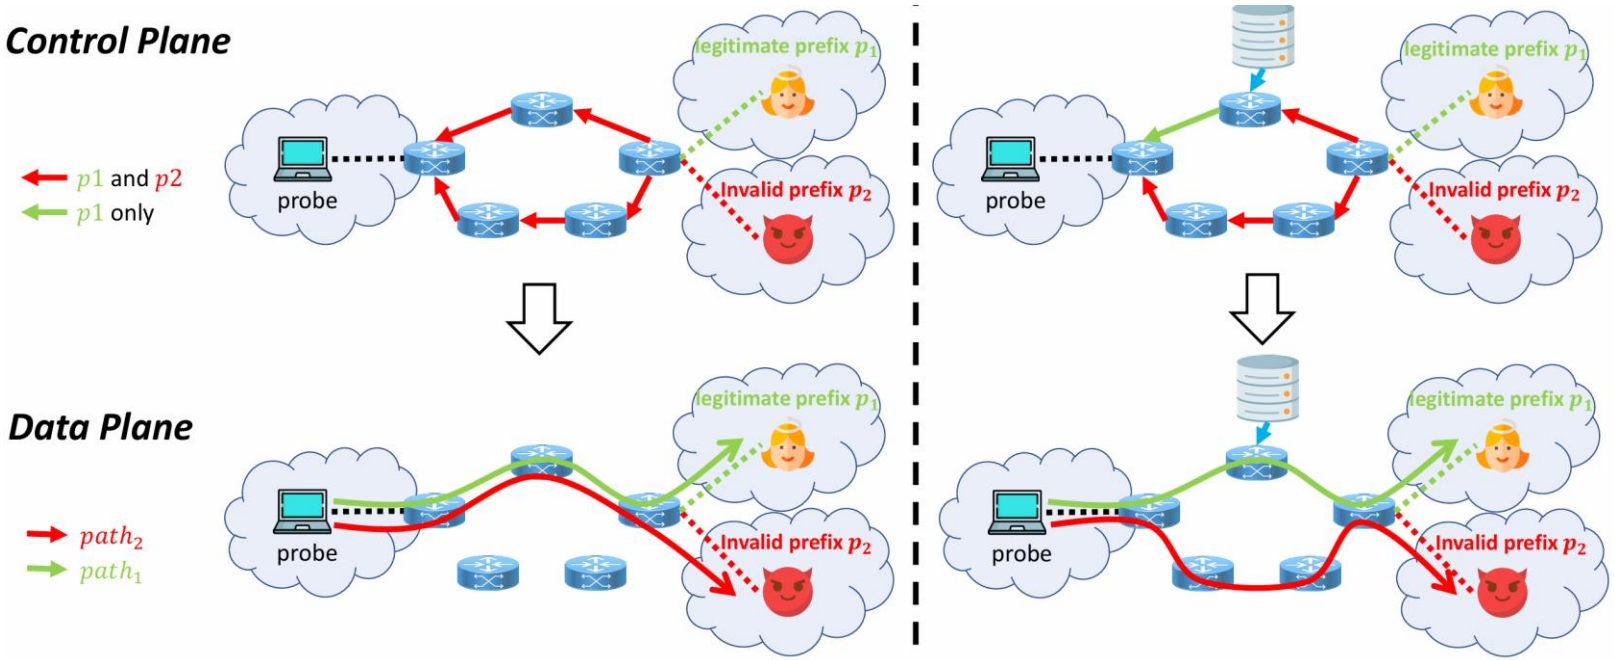

- For each invalid prefix (p_1) and origin AS (A):
  - Find a legitimate prefix (p_2) from the same origin AS (A).
  - Retrieve live IP addresses ($addr_1$, $addr_2$) from p_1 and p_2.
  - Use public probes (RIPE Atlas, perfSONAR) to traceroute to $addr_1$ and $addr_2$.
- Compare AS-level paths:
  - **Identical paths:** Path doesn't filter invalid updates.
  - **Different paths:** Path does filter invalid updates (and the other doesn't).

Inference Algorithm: Bayesian Approach

- **Goal:** Infer the probability that each AS doesn't filter invalid routes (z_i).
- **Model:** Probabilistic (Bayesian inference).
  - Treat ROV deployment as a random variable ($z_i \in [0, 1]$).
  - Labeled paths are observed data (D).
  - Calculate posterior distribution: $p(Z|D) \propto p(Z)p(D|Z)$
- **Challenge:** High-dimensional, correlated variables (one z_i for each AS).
- **Prior Work:** MCMC is slow, scales poorly, and may not converge.

Inference Algorithm: SVGD

- **S**tein **V**ariational **G**radient **D**escent (SVGD)
- **Key Idea:** Deterministic approximation of the posterior distribution.
- Starts with a set of initial "particles" (possible values of Z).
- Iteratively updates particles using a gradient-based approach.
  - Gradient direction calculated using Stein's method.
  - Kernel function (k) influences convergence.

$$\hat{\phi}^*(\theta) = \frac{1}{m}\sum_{k=1}^{m}\left[k(\theta_k^l, \theta)\nabla_{\theta_k^l}\log p(\theta_k^l|D) + \nabla_{\theta_k^l}k(\theta_k^l, \theta)\right]$$

- Advantages over MCMC:
  - Deterministic descent (faster, more efficient).
  - Doesn't require large sample sizes.
  - Better convergence.

Inference Algorithm: SVGD

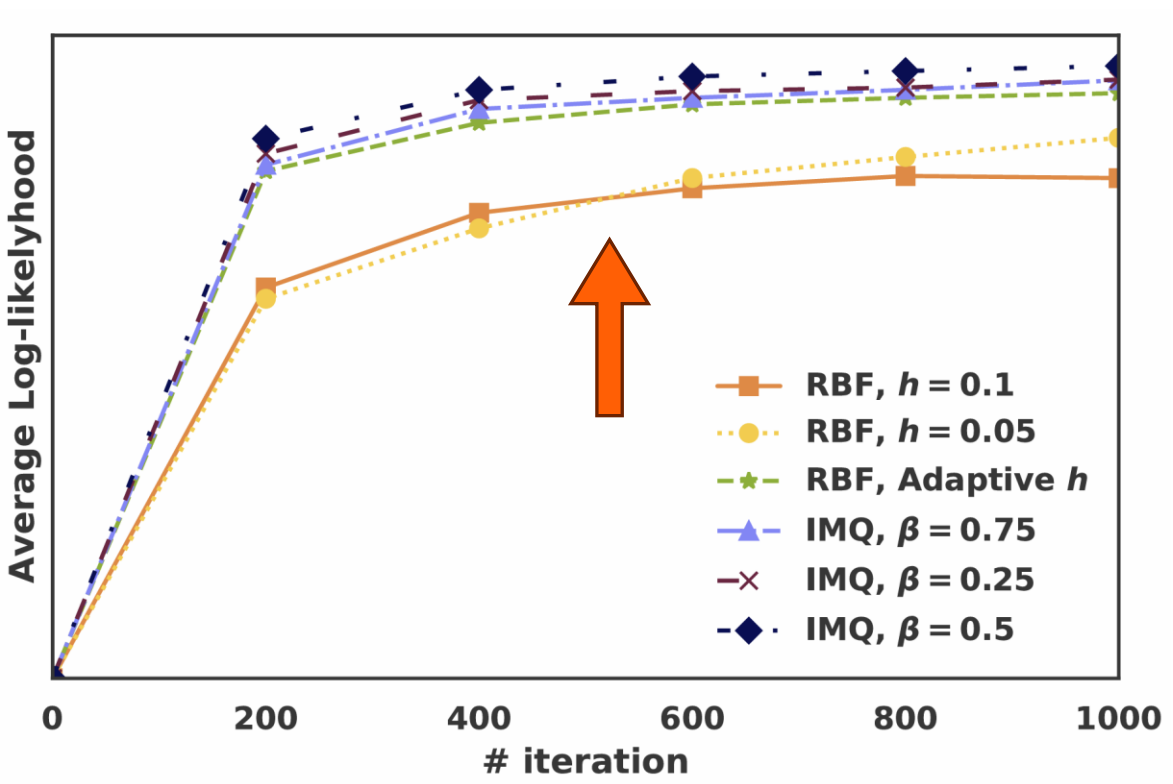

- **Kernel function (k) influences convergence.**
  - *Reviewer 17D: "comparing RBF with h=0.1 and h=0.05, I wonder what causes the 'swap' at the ~500th iteration."*

Results

- **Measurement Scale:**
  - ~10x more ASes measured compared to prior work.
  - ~10x more paths labeled.

TABLE II. Data collected from data plane probing.

	#unique origins	#covered ASes	#unique paths
Gray *et al.* [14]	1	1,265	12,634
ROV-MI	678	11,074	115,427

TABLE VII. Validation on ground truth.

Method	Precision(%)	Recall(%)
Heuristic [9]	68	73
MCMC [14]	100	100
SVGD	100	100

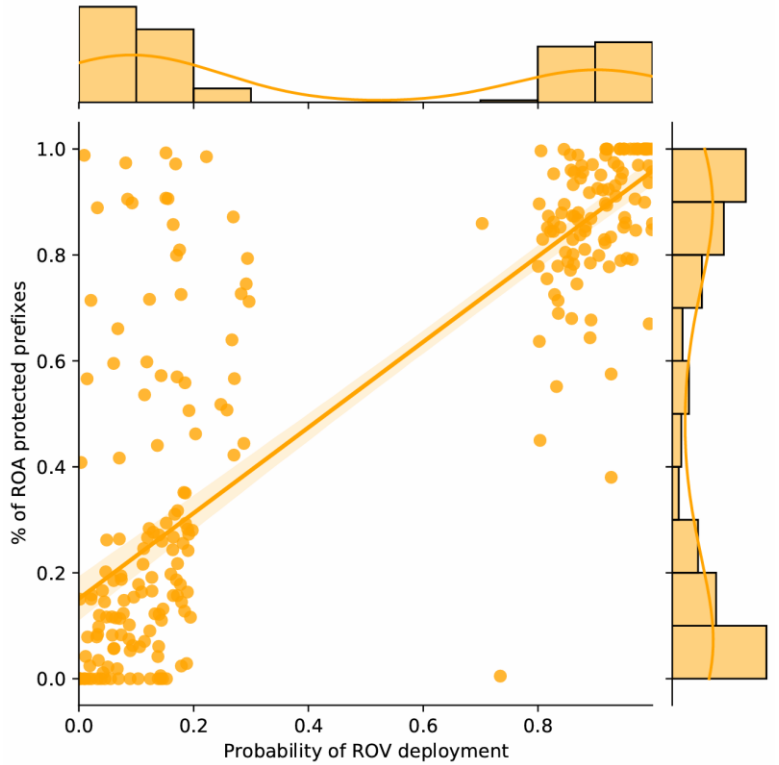

Results

- **Accuracy:**
  - Validated against "is-bgp-safe-yet" ground truth: ***Near-perfect precision and recall***.
  - ***High correlation*** between ROA and ROV deployment (as expected).

Results

- **Efficiency:**
  - SVGD converges ***~5x faster*** than MCMC.
  - Requires ***~500x fewer samples*** than MCMC.

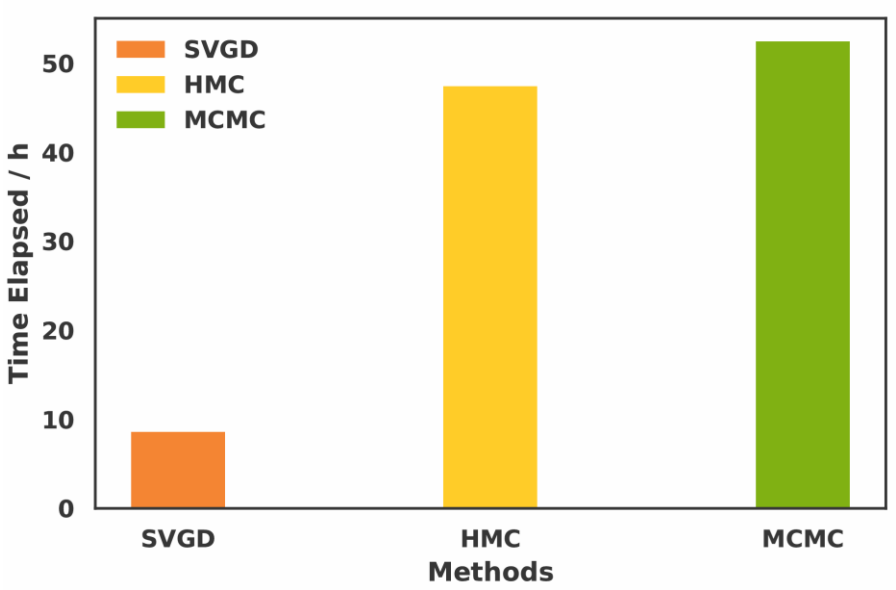

TABLE VIII. Number of samples when reaching convergence.

Method	SVGD	HMC	MCMC
#Particles	1,000	463,569	573,645

TABLE IV. The number/proportion of different types of ASes.

Categories	Number	Proportion
deployed	3,107	28%
undeployed	4,716	43%
partially-deployed	357	3%
unknown	2,894	26%
total	11,074	100%

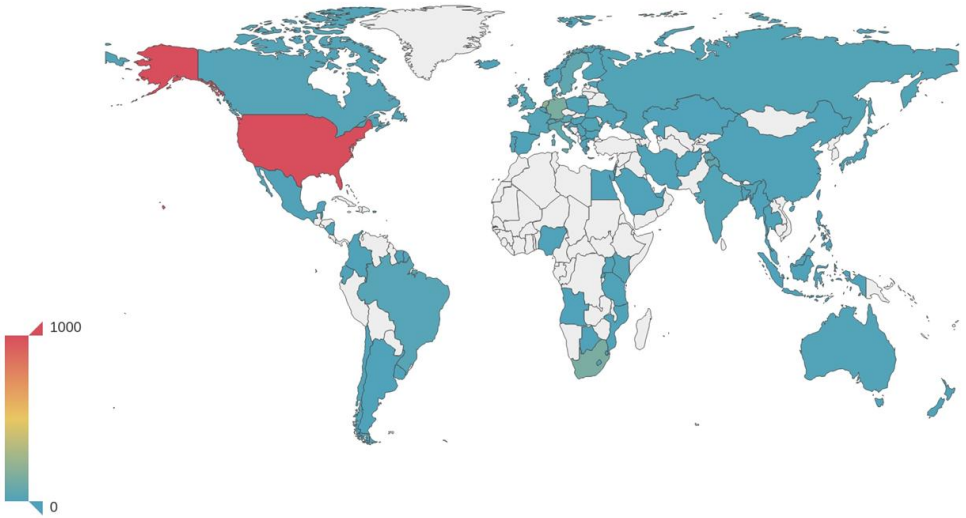

Country	#ASes with ROV
the United States	1,171
South Africa	219
Germany	193
Netherlands	186
Italy	157

What Did ROV-MI Find?

- ***~28%*** of measured ASes deploy ROV.
- Most deployed ASes are ***large transit ASes*** (significant impact).
- ROA and ROV deployment are highly correlated.
- ***Geographical disparity***: Mostly North America, Europe, and South Africa.

Ratings

1. Should not have been published
2. Marginal paper
3. Good paper, but could use some improvement
4. Excellent paper, well placed at a top conference
5. Award-quality paper

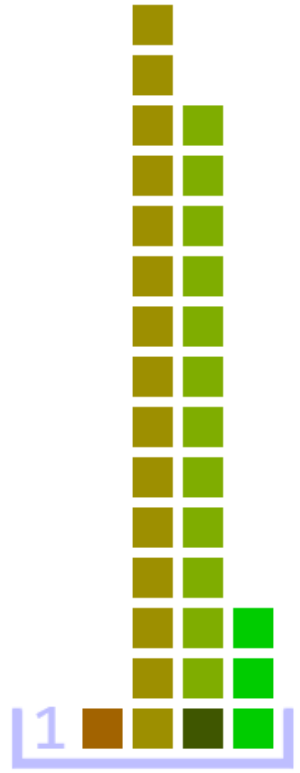

Paper Quality

Paper Interest

1. I want my time back
2. Not my favorite
3. Worthwhile read
4. Really enjoyed it, will recommend to friends / colleagues
5. New favorite paper!

Review & Discussion

- **IPv6:** Currently only IPv4; extending to IPv6 is important.
- **Continuous Measurement:** RPKI deployment is dynamic; need ongoing monitoring.
- **Geographical Bias:** Need more probes in underrepresented regions.
- **Why not ROA and ROV together?** Some ASs are using ROV but not ROA (Reviewer 17A, 17C, 17Y, 17R)
- **RPKI databases:** why not they also utilize RPKI databases to measure ROV usage as well? (Reviewer 17A)
- **Applying ROV-MI:** Applying ROV-MI to different regions or less-connected ASes could offer a more comprehensive view of global ROV adoption. (Reviewer 17C)
- **AS Topology:** The paper infers utilization of ROV by comparing two paths from both a legitimate and illegitimate IP address. However, there could be *numerous reasons for two paths to differ* in BGP pathing. *While these factors could be dealt with*, this paper does not mention how they deal with these confounding factors. (Reviewer 17H)
- **Ground Truth:** The 370 ASes which the paper utilizes as ground truth are collected *by communication with the network operators*. This collection method most likely will *not be a random sample* of all possible ASes which weakens the analysis when measuring the accuracy of the several thousand ASes not in the ground truth. (Reviewer 17H)

More Review & Discussion

- **Why is RPKI/ROV adoption slow?** (Reviewers 17A, 17C, 17H, 17T, 17X, 17Z)
- **Can ROV-MI be generalized to other problems?** (Reviewers 17B, 17D, 17F, 17Z)
- **How can we encourage ROV adoption?** (Reviewers 17B, 17O, 17AD)
- **Geographical disparities in ROV adoption.** (Reviewers 17D, 17G, 17L, 17J, 17AB, 17AC)
- **Ethical considerations of scanning.** (Reviewer 17N)
- **Trust in centralized authorities (RIRs).** (Reviewer 17AE)
- Why not have *an authoritative third-party to manually verify the deployment of ROV*, or any kind of security methods? Similar to the system currently used in CA and digital signature. In this case we can have *a single root of trust*. (Reviewer 17M)
- Deploying *ROV still isn't enough to avoid malicious behaviors*, it merely *increased some difficulty*. A 2010 paper "How Secure are Secure Interdomain Routing Protocols" (Goldberg) has shown that malicious ASes can still perform interception attacks under multiple mechanisms including ROV. So, instead of detecting for the deployment of certain security mechanism, we could consider *detecting anomalys in the routing path itself*, which is the ultimate check for malicious behavior. (Reviewer 17M)

Reference

- ROV-MI Paper link: https://www.ndss-symposium.org/wp-content/uploads/2022-214-paper.pdf
- ROV-MI author's presentation video: https://www.youtube.com/watch?v=4yKPDoZRujA