NDSS Symposium 2026

The Network and Distributed System Security (NDSS) Symposium 2026 will take place from 23 to 27 February 2026 in San Diego, California.

PathProb: Probabilistic Inference and Path Scoring for Enhanced and Flexible BGP Route Leak Detection

Yingqian Hao, Hui Zou, Lu Zhou, Yuxuan Chen, Yanbiao Li*

中国科学院
计算机网络信息中心
Computer Network Information Center, Chinese Academy of Sciences

中国科学院大学
University of Chinese Academy of Sciences

BGP is important yet vulnerable

- Border Gateway Protocol (BGP) is one of the key building blocks of the global Internet.
- However, BGP lacks built-in security protection and thus is vulnerable to route leaks.

GLOBAL BGP ROUTE LEAKS	Q3, 2025	GLOBAL BGP HIJACKS
1	July	1
1	August	0
2	September	0

CLOUDFLARE

Route leak incident on January 22, 2026

2026-01-23

On January 22, 2026, an automated routing policy config leak some Border Gateway Protocol (BGP) prefixes unint data center in Miami, Florida. While the route leak cause customers, multiple external parties were also affected b accidentally funnelled through our Miami data center location.

25 Minites
Congestion
Higer Latency
Discarded Traffic

The route leak lasted 25 minutes, causing congestion on some of our backbone infrastructure in Miami, elevated loss for some Cloudflare customer traffic, and higher latency for traffic across these links. Additionally, some traffic was discarded by firewall filters on our routers that are designed to only accept traffic for Cloudflare services and our customers.

[1]. Blog — Q3 2025 DDoS, bad bots, and BGP incidents statistics and overview
[2]. https://blog.cloudflare.com/route-leak-incident-january-22-2026/

Route Leaks: The propagation of route(s) beyond their intended scope

- A leak happens when an Autonomous System (AS) propagates a route in violation of standard export policies.
- **AS relationships** determine export policies.
  - Provider–Customer (***p2c***): provider sells transit service; customer pays.
  - Peer–Peer (***p2p***): settlement-free exchange

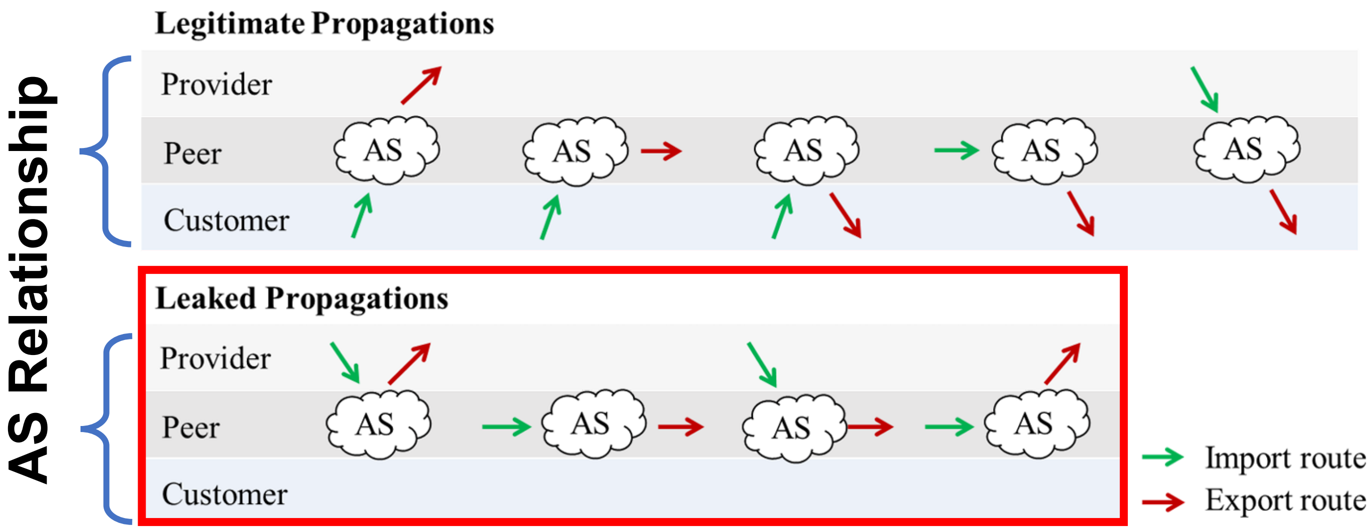

Route Leaks: Abnormal AS-path patterns

- A leak happens when an AS propagates a route in violation of standard export policies.
- This is observable in AS-path patterns

"ascend" through c2p links,
optionally include a single horizontal p2p link,
and then "descend" through p2c links.

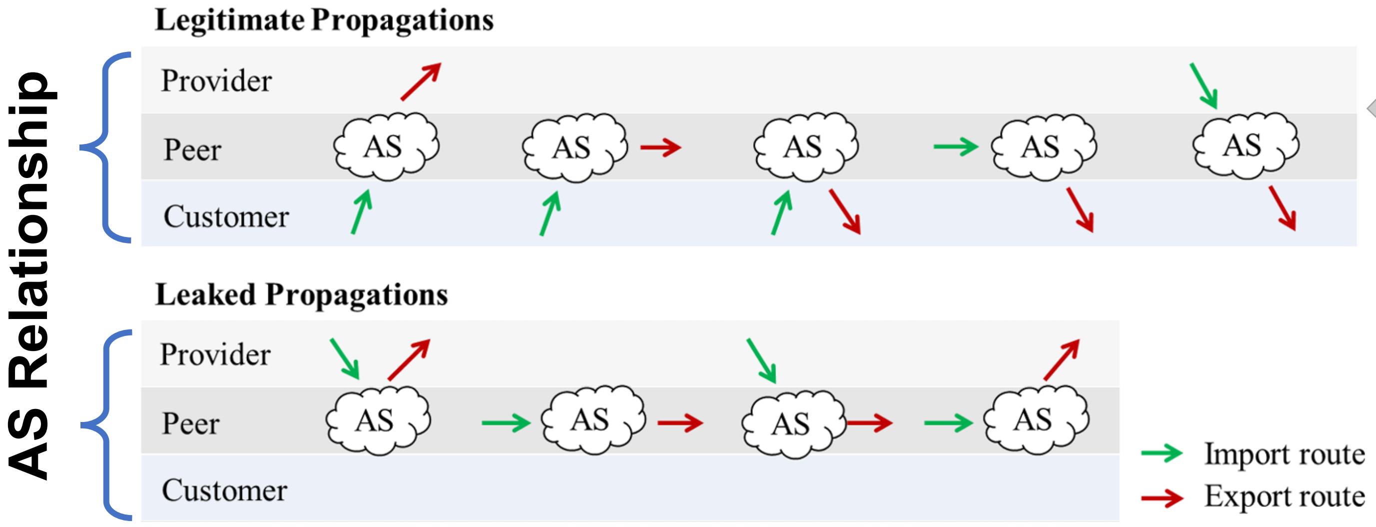

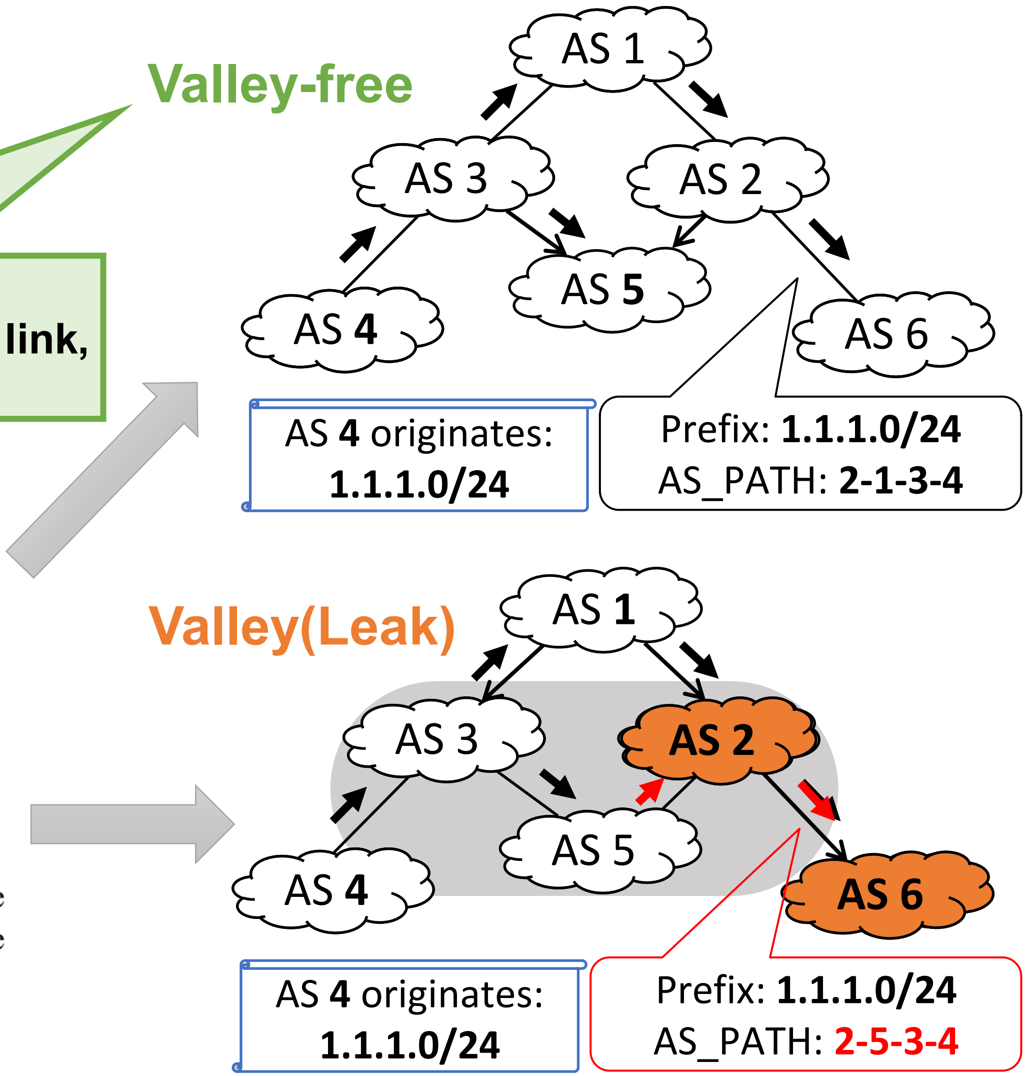

Motivation: Why existing solutions don't deploy

Existing Route Leak Solutions

Rule-based

Authoritative Declared AS Relationships (e.g. RPKI-based ASPA data)

→ High Confidence, **Limited Coverage**

→ Poor Defensive Capabilities

Deterministic Inferred AS Relationships (inferred from history BGP data)

→ **Low Confidence** Broad Coverage

→ High False Positive Rate

Machine Learning-based

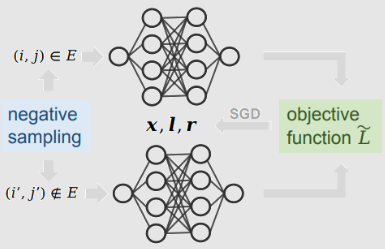

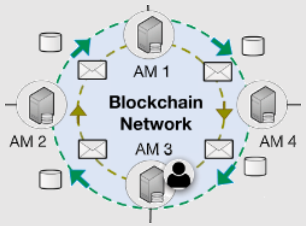

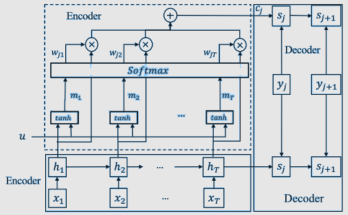

→ Captures rich features and highly adaptive **Low Interpretability & Sensitive to Data Quality**

→ Poor Usability : still require experts with in-depth BGP knowledge to interpret BGP events

PathProb: Probabilistic Relationships for AS Links + Legitimacy Score for AS Paths

- Each AS link $(u, v) \rightarrow$ **[P(c2p), P(p2p), P(p2c)]**, sum=1
- Output: **continuous legitimacy score**, then threshold to decide leak

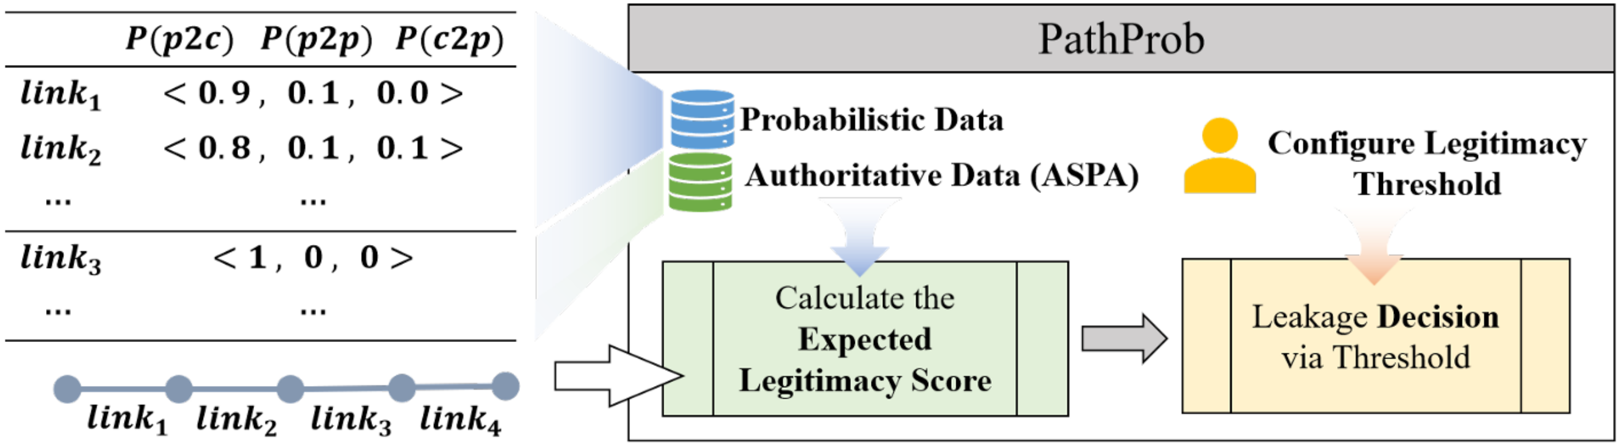

Path		202365 - 50673 - 6939 - 199524 - 58212 - 13627					Result
		$link_1$	$link_2$	$link_3$	$link_4$	$link_5$	
CAIDA		*c2p*	*p2p*	*p2p*	*p2c*	*p2c*	*leaked*
PathProb	*c2p*	**0.944**	0.451	0.001	0.004	0.0	$E_{leg} = 0.697$
	p2p	0.056	**0.532**	**0.551**	0.166	0.0	
	p2c	0.0	0.0	0.448	**0.830**	**1.0**	

Accommodates the complex relationships

Trade off sensitivity against false alarms via threshold

Seamlessly integrate with ASPA

Legitimacy Score for AS Paths: How is the Legitimacy Score Calculaed?

- A path is legitimate ***iff*** it satisfies **valley-free**
- Compute expected legitimacy score by **enumerating all possible valley-free realizations of the path** and **summing their probabilities** to derive the path's legitimacy score

$$E_{\text{leg}}(\text{path}) = \sum_{k=1}^{n-1} \left(\prod_{i=1}^{k-1} P(c2p_i) \cdot P(p2p_k) \cdot \prod_{i=k+1}^{n-1} P(p2c_i) \right) + \sum_{k=1}^{n} \left(\prod_{i=1}^{k-1} P(c2p_i) \cdot \prod_{i=k}^{n-1} P(p2c_i) \right)$$

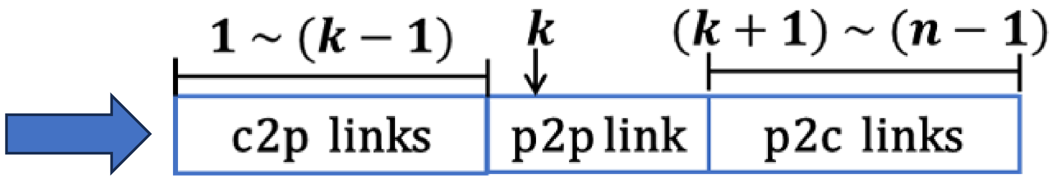

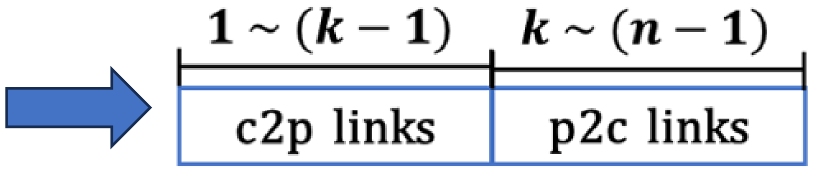

Legitimacy Score for AS Paths: From Full-Path to Triple-Minimum

- The full-path expectation method suffers from **length-induced uncertainty**
- Route leak is confined to a triple of ASes, the triple-minimum prevents the accumulation of uncertainty

$$E_{\text{leg}}(\text{path}) = \min_{i=1}^{n-2}(E_{\text{leg}}(\text{path}_{i:i+2}))$$
$$= \min_{i=1}^{n-2}(P(c2p_i) + P(p2c_{i+1}) - P(c2p_i) \times P(p2c_{i+1}))$$

AS path‘s score = minumun score among all its triples

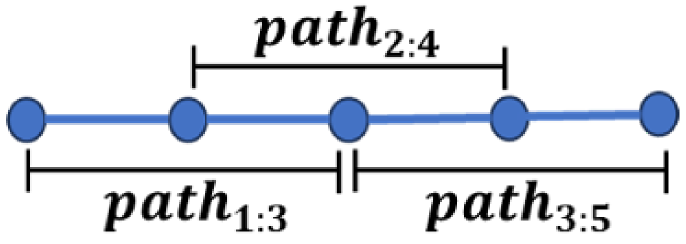

Path		58057 - 34549 - 2914 - 6762 - 34984 - 60051 - 208293 - 203214						Full-path	Triple-minimum
		$link_1$	$link_2$	$link_3$	$link_4$	$link_5$	$link_6$		
PathProb	*c2p*	**0.893**	**0.631**	0.222	0.018	0.011	0.0	**0.286**	**0.679**
	p2p	0.107	0.352	**0.649**	0.39	0.166	0.0		
	p2c	0.0	0.017	0.129	**0.592**	**0.823**	**1.0**		

Probabilistic Relationships for AS Links: Core and Edge Links

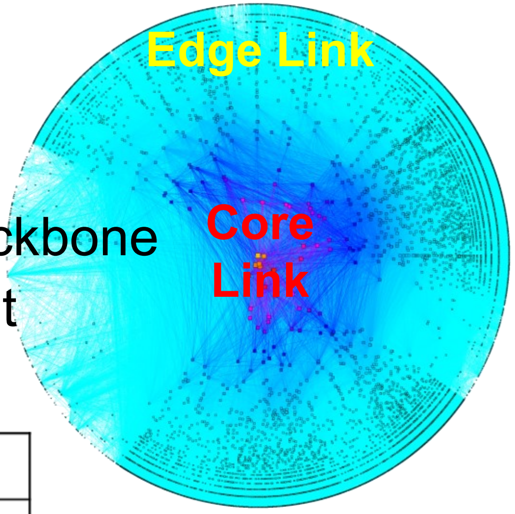

- Hierarchical inference:
  - **Core links**: densely interconnected and complex, form the Interne's backbone
  - **Edge links**: sparse and simple, located at the boundaries of the Internet

[1]. https://asrank.caida.org/

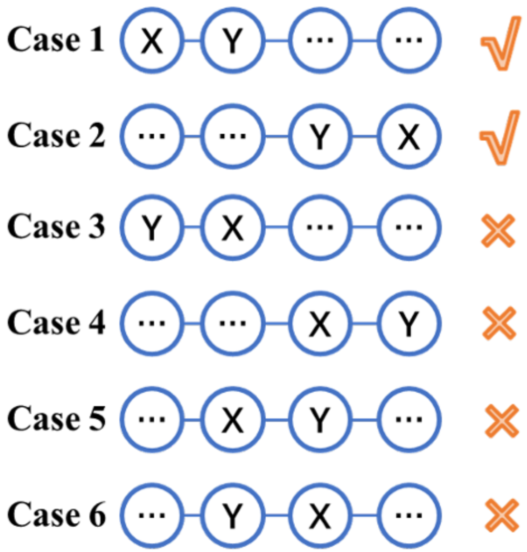

X-Y is an edge link: Exist Case 1 or Case 2 and no other Cases

AS Paths
8 4 3 7 1 2
8 5 7 1 2
5 8 3 2 1 7
5 6 7 3 1

Round 1	Round 2	Round 3	Round 4
$PathSet_1$	$PathSet_2$	$PathSet_3$	$PathSet_4$
8 4 3 7 1 2	4 3 7 1 2	3 7 1 2	7 1 2
8 5 7 1 2	8 5 7 1 2	8 5 7 1 2	8 5 7 1 2
5 8 3 2 1 7	5 8 3 2 1 7	5 8 3 2 1 7	5 8 3 2 1 7
5 6 7 3 1	6 7 3	7 3	
$EdgeLink$	$EdgeLink$	$EdgeLink$	$EdgeLink$
(8,4) (5,6) (3,1)	(4,3) (6,7)	(3,7)	null

$$EdgeLinkSet = \{(8,4),(5,6),(3,1),(4,3),(6,7),(3,7)\}$$

$$CoreLinkSet = \{link \mid link\ in\ PathSet_4\}$$

$$CorePath = PathSet_4$$

Probability Inference——ILP Model

Maximize/Minimize $\mathbf{c}^T \mathbf{x}$
Subject to: $\mathbf{A}\mathbf{x} \leq \mathbf{b}, \mathbf{x} \in \mathbb{Z}^n$

- ILP formulations for modeling valley-free constraints
  - **Strict Model**: enforces the valley-free property as a set of hard constraints, strictly forbidding any violation in the AS path structure.
  - **Loose Model**: allows a limited number of valley-free violations (e.g., link skips) to maintain feasibility

Model / **Valley-free constraints**	**Strict Model** $\begin{array}{ccc} & x_i = 0 & x_i = 1 \\ \hline y_i = 0 & c2p & - \\ y_i = 1 & p2c & p2p \end{array}$	**Loose Model** $\begin{array}{ccc} & x_i = 0 & x_i = 1 \\ \hline y_i = 0 & z_i = 0,\ c2p & z_i = 1,\ \text{skipped} \\ y_i = 1 & z_i = 0,\ p2c & z_i = 0,\ p2p \end{array}$
The *p2p* or *p2c* links appear after all *c2p* links	$y_{n,m} \leq y_{n,m+1}$, $\forall n = 1, ..., n\ max$ and $\forall m = 1, ..., m\ max$	$y_{n,m_1} \leq y_{n,m_2} + z_{n,m_2}$, $\forall n = 1, ..., n\ max$ and $\forall m_2 = 2, ..., m_2\ max$ and $\forall m_1 = 1, ..., m_2 - 1$
All *p2c* links appear after all *c2p* or *p2p* links	$y_{n,m} - x_{n,m} \leq y_{n,m+1} - x_{n,m+1}$, $\forall n = 1, ..., n\ max$ and $\forall m = 1, ..., m\ max$	$y_{n,m_1} - x_{n,m_1} \leq y_{n,m_2} - x_{n,m_2} + 2z_{n,m_2}$, $\forall n = 1, ..., n\ max$ and $\forall m_2 = 2, ..., m_2\ max$ and $\forall m_1 = 1, ..., m_2 - 1$
The path contains at most one *p2p* link	$\sum_m x_{n,m} \leq 1,\ \forall n = 1, ..., n\ max$	$\sum_m (x_{n,m} - z_{n,m}) \leq 1,\ \forall n = 1, ..., n\ max$

Core-link Inference: ILP + MRF + Gibbs Sampling

- **Markov Random Field (MRF):** models **local dependencies** among core links
- **Gibbs sampling:** calculates **marginal probabilities** for core links **efficiently**
- **Warm start**: uses Loose Model's output for **faster convergence**

ILP Loose Model

MRF Model

$$P(X) = \frac{1}{Z} \prod_i \phi_i(x_i, x_{N(i)})$$

$$Z = \sum_X \prod_i \phi_i(x_i, x_{N(i)})$$

$$\phi_i(x_i, x_{N(i)}) = \sum_k \alpha_{ik}(x_i, x_i^{\mathrm{prev}(k)}, x_i^{\mathrm{next}(k)})$$

$$P(x_i \mid X \setminus \{x_i\}) = P(x_i \mid x_{N(i)}) = \frac{\phi_i(x_i, x_{N(i)})}{\sum_{x_i'} \phi_i(x_i', x_{N(i)})}$$

Gibbs Sampling

1) **Initialization**: Assign initial values to all variables $X^{(1)} = (x_1^{(1)}, \ldots, x_n^{(1)})$, either randomly or via heuristic methods.
2) **Conditional Sampling**: For each sample $X^{(i)} = (x_1^{(i)}, \ldots, x_n^{(i)})$ (where $i = 1$ to K), compute each $x_j^{(i)}$ (where $j = 1$ to n) in sequence using its conditional probability:

$$P(x_j^{(i)} \mid x_1^{(i)}, \ldots, x_{j-1}^{(i)}, x_{j+1}^{(i-1)}, \ldots, x_n^{(i-1)})$$

   Once all variables x_j have been updated, they collectively form the i-th sample.
3) **Iteration**: The Gibbs sampling procedure iteratively repeats this process until K samples have been generated.

Edge-link Inference: Propagation + Context-aware Resolution

- Path1: If an adjacent core link is highly likely *p2c/p2p* ⇒ infer edge link as *p2c*
- Path2: Recursive propagation across overlapping paths

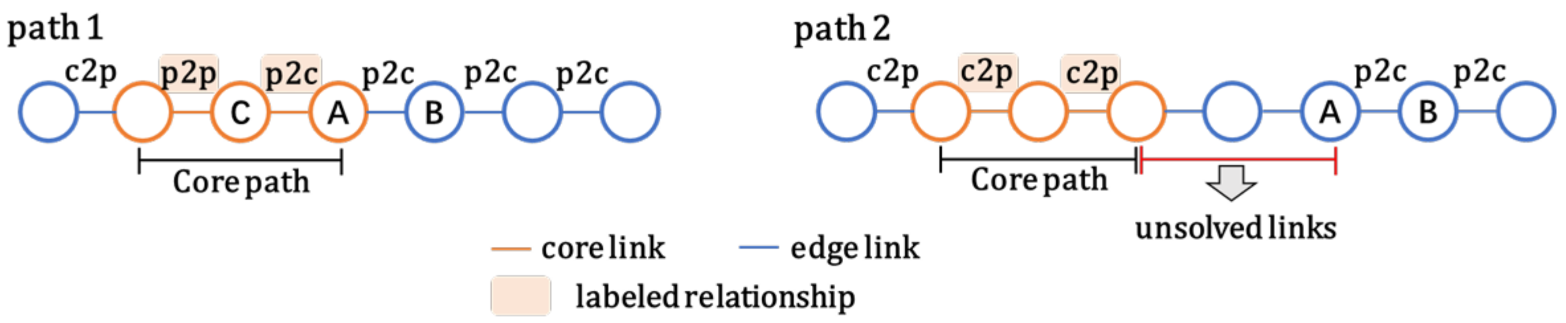

- Path3: isolated edge links use uniform prior distribution: (1/3, 1/3, 1/3)
- Path4 and Path5: contextual links solved with **Strict Model**

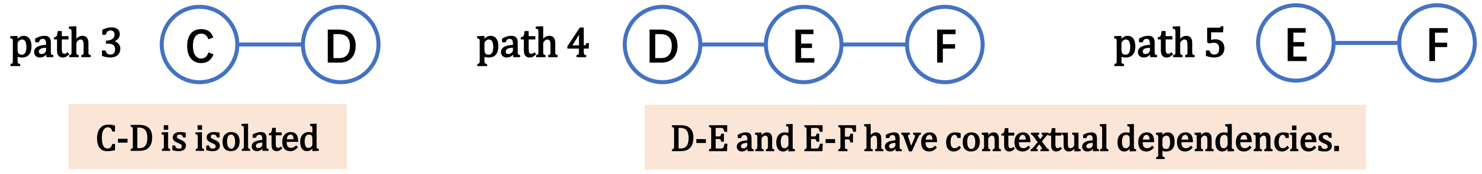

Evaluation Setup

Index	Steps	Metrics
1	Infer probabilistic AS relationships	Accuracy
2	Detect route leaks	Precision, Recall, FPR
3	Simulate the global Internet topology, route leak events, and deployment scenarios.	LIR (leakage infection rate) LCR (legitimate connection rate)

Result 1: Relationship Inference Accuracy

- Accuracy with CAIDA & ASPA validation dataset
  - CAIDA: broad coverage, albeit with best-effort precision
  - ASPA: high precision but with limited coverage

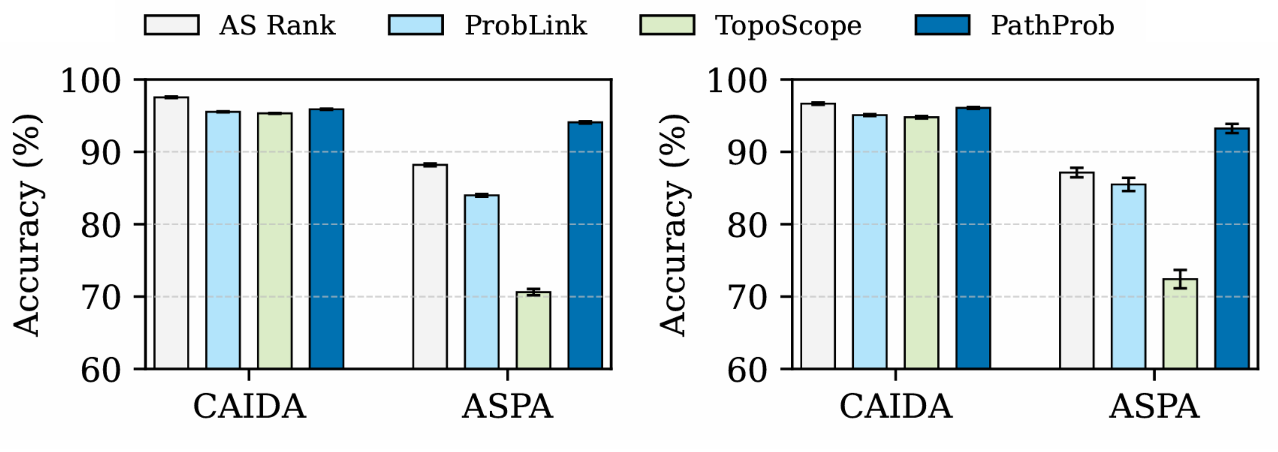

(a) RIB_Day path dataset (b) RIB_Year path dataset

PathProb achieves the **best accuracy**, outperforming other schemes by **5.87 ~ 23.46** and **6.08 ~ 20.81** percentage points on long-term (RIB_Year) and short-term (RIB_Day) path datasets with non-overlapping 95% confidence intervals with the highest-precision ASPA validation dataset

Result 2: Score Distribution + Threshold Choice

We use Cloudflare Radar's route-leak reports as ground truth, which provide timestamps and the leaked AS triples. For each event, we collect BGP updates from RIPE RIS and RouteViews around the reported time. Paths containing the reported triple are labeled as "**leaked**"; otherwise, they are labeled as "**legitimate**"

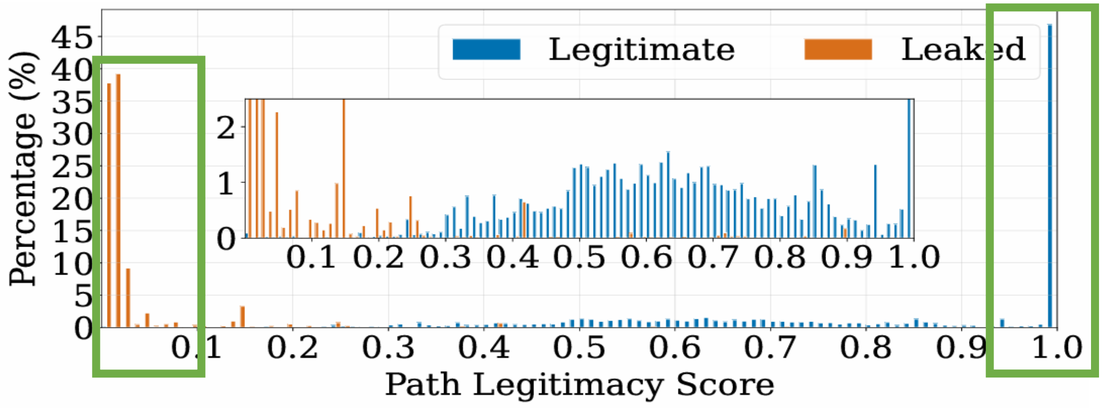

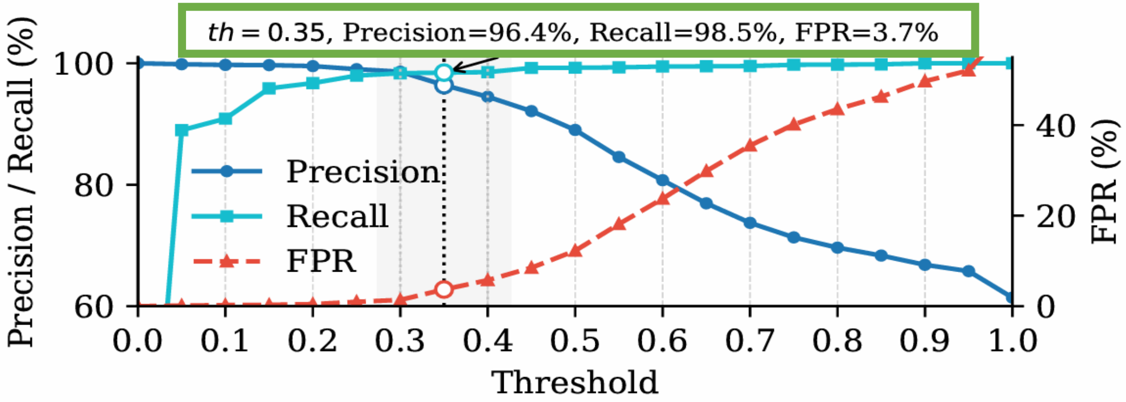

- Leaked path scores are concentrated in **[0, 0.1]**
  - 90.88% links: **[0, 0.1]**
- Legitimate paths exhibit higher scores
  - 46.83% links: **> 0.99**, 37.45% links: **[0.5, 0.9]**

- **th = 0.35,** achieves the opimal balance
  - Precision: 96.4%
  - Recall: 98.5%
  - FPR: 3.7%

Result 3: Route Leak Detection

- To evaluate PathProb's robustness in a fine-grained, multi-dimensional manner
  - Temporal partitioning: 6-30 June 2025
  - Collector-based grouping: RIPE NCC RIS 18 collectors and RouteViews 15 collectors
  - Network role filtering: Tier-1, stub, transit, content provider, sibling, partial transit, hybrid

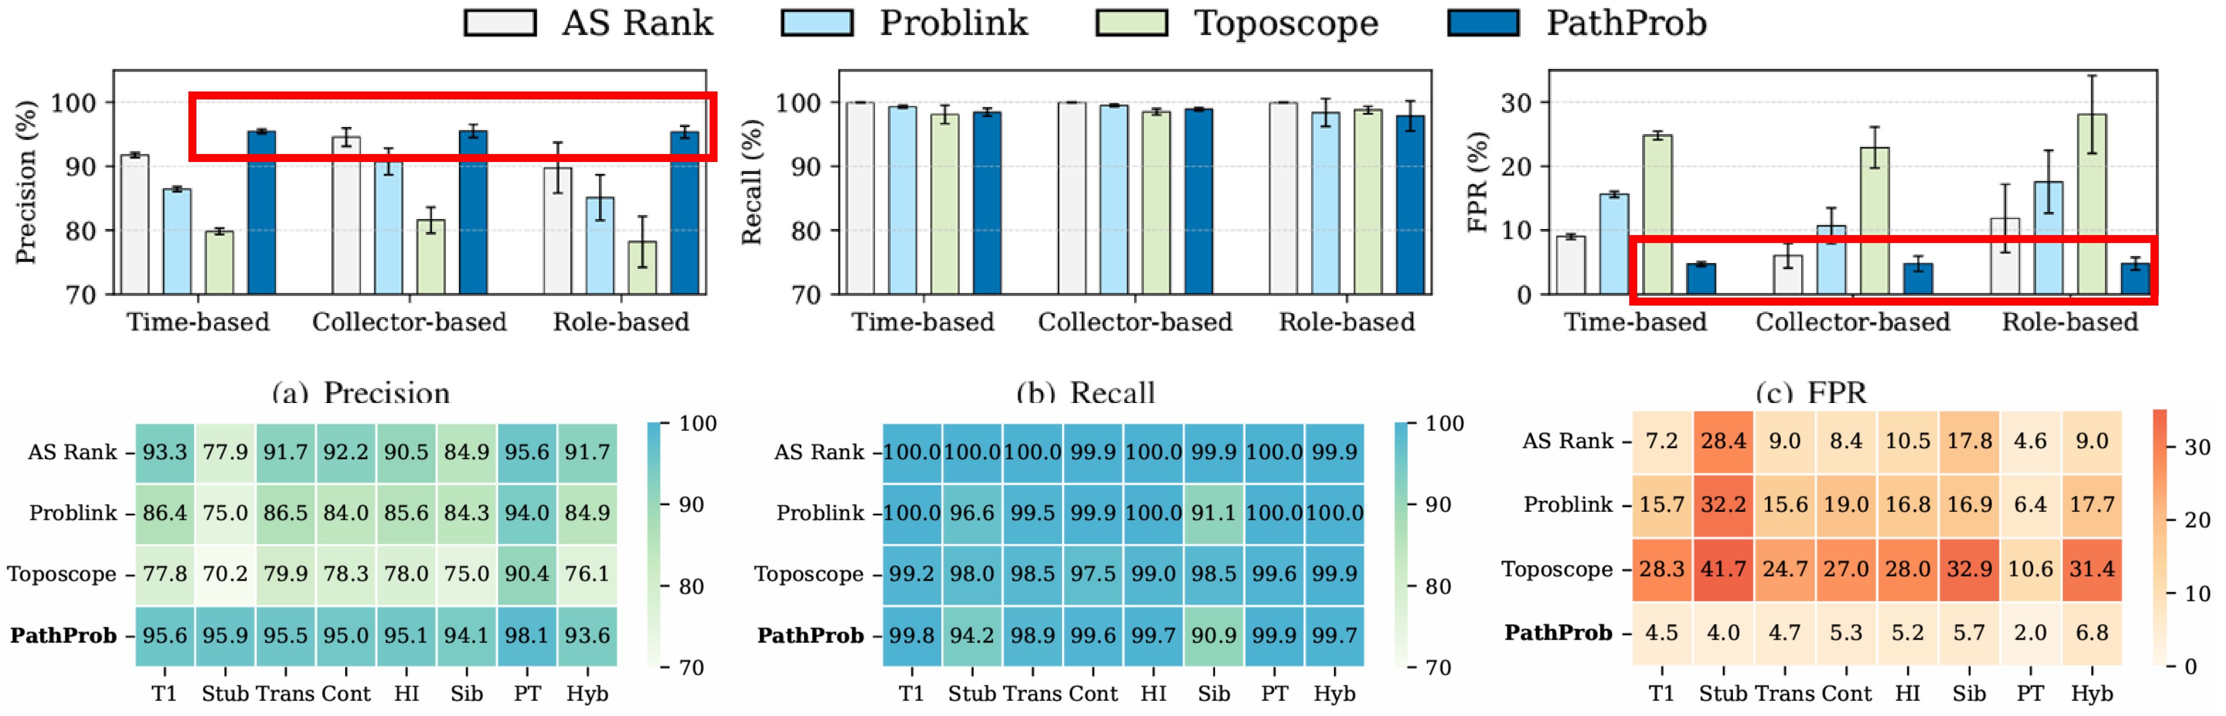

Highest precision: 95.4%, Lowest FPR: ≈ 4.7%, Recall: ≈ 98~100%

Result 4: Deployment with ASPA (Simulation)

- Simulate the Internet topology using BGPy[1], and randomly select attacker and victim pairs
- Integrate with other AS relationship sources (e.g. ASPA) and alternative route-leak mitigation mechanisms (e.g. OTC) under different deployment scenarios

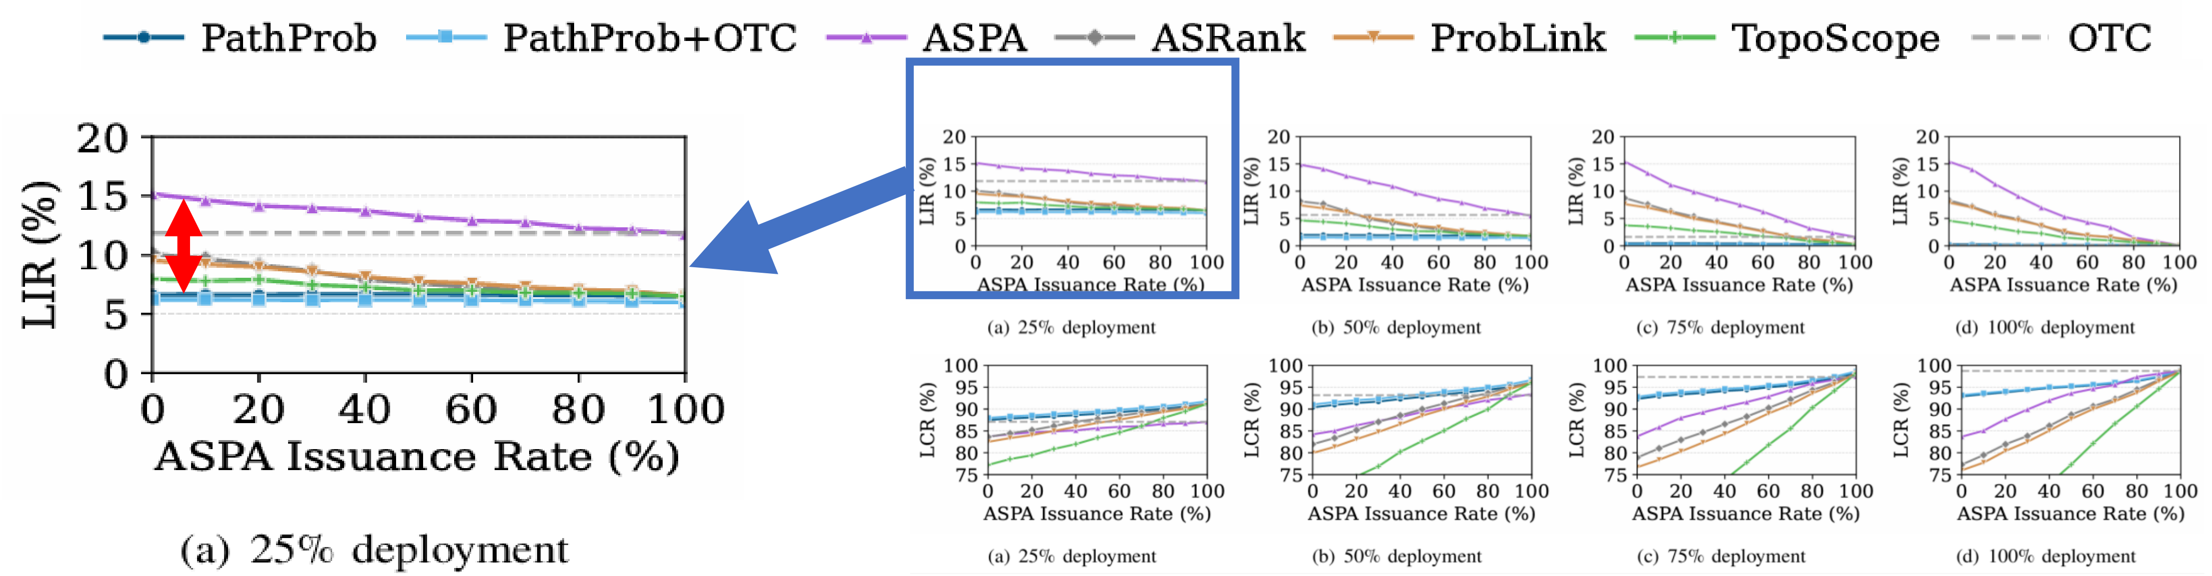

Early deployment: reduces LIR from 15.17% to 6.63% (~56%↓) and improves LCR from 83.70% to 87.49%

PathProb can compatible with other technical approaches (PathProb+OTC), achieving better results: LIR = 6.19% (~59.2%↓) and LCR = 87.99%.

[1]Furuness, J., Morris, C., Morillo, R., Herzberg, A., & Wang, B. (2023, August). Bgpy: The bgp python security simulator. In Proceedings of the 16th Cyber Security Experimentation and Test Workshop (pp. 41-56).

Conclusion

01

PathProb is **enhanced**

PathProb is a **enhanced** route-leak detector that is **interpretable** yet **robust** to complex AS relationships.

02

PathProb is **flexible**

It outputs a continuous **legitimacy score**, **reducing false positives** while allowing operators to **flexibly** tune the sensitivity–false-alarm trade-off via a threshold.

03

PathProb is **compatible**

It integrates **naturally with other AS relationship sources**, such as ASPA, and delivers **strong performance gains** even under **partial deployment**

Thank you!

- For more details, please read our paper: **PathProb: Probabilistic Inference and Path Scoring for Enhanced and Flexible BGP Route Leak Detection**

- Contact: lybmath@cnic.cn
- Code: https://github.com/hyq8868/PathProb and https://doi.org/10.5281/zenodo.17920055

中国科学院
计算机网络信息中心
Computer Network Information Center, Chinese Academy of Sciences

中国科学院大学
University of Chinese Academy of Sciences