Network Security and Measurement

- Network Tomography -

Prof. Dr. Thomas Schmidt

http://inet.haw-hamburg.de | t.schmidt@haw-hamburg.de

Agenda

Introduction to Network Tomography

Binary Network Tomography

Pinpointing Router Behavior in the Wild

Looking into the inner Core

TOMOGRAPHY

Problem

Most measurements are performed from endpoints at the network edge

- Which observations are caused by the network core?
- What are underlying characteristics of the network core?

In many cases, it is impossible to measure the Internet core directly

Problem

Can we study the internal characteristics of a network using only information visible at its edge?

Most measurements are performed from endpoints at the network edge

- Which observations are caused by the network core?
- What are underlying characteristics of the network core?

In many cases, it is impossible to measure the Internet core directly

The Concept of Network Tomography

Coined by Vardi in 1996

- Model the properties of your network
- Take appropriate measurements at many endpoints
- Correlate the measurements
- Invert the problem with the help of your model
- Infer properties of the internal network

Y. Vardi: *Network Tomography: Estimating Source-Destination Traffic Intensities from Link Data*
Journal of the American Statistical Association, Vol. 91, No. 433 (March, 1996), pp. 365-377

Internet Tomography

Many paths share links, why measurements correlate

Tomography problem:

$$y = \boldsymbol{A} \cdot \theta + \epsilon\text{, with}$$

$\boldsymbol{A}$ the routing matrix, θ network parameters, and ϵ random noise.

y are the measured observables.

This linear stochastic model needs inversion to infer inner link properties.

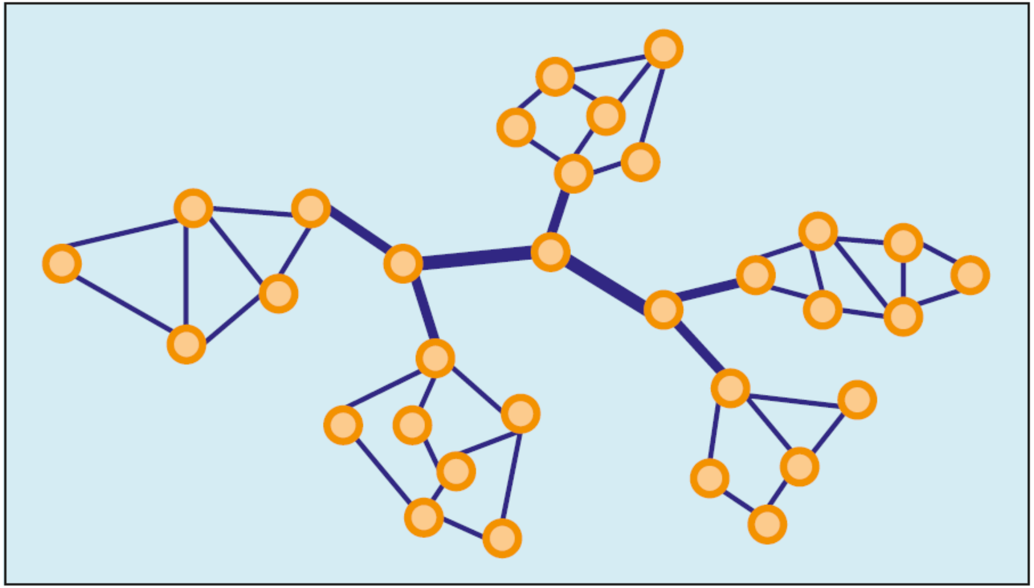

M. Coates, et al: *Internet Tomography, IEEE Signal Processing Magazine.* ***19*** *(3): 47–65, 2002*

Examples

Congested Link – is reduced available bandwidth due to congested link?
→ θ is the vector of traffic intensity per link

Lossy Link – is packet loss due to a link?
→ θ is the vector of link success probabilities

Packet delays – which links produce large delays?
→ θ is the vector of link delays

Reducing the Problem Space

BINARY NETWORK TOMOGRAPHY

Limitations of General Network Tomography

Modeling space can be very large.

Network conditions vary – models often need time-dependence.

Interpretation and Inversion often require complex mathematical models, e.g., about queuing.

Measurements often too sparse or too inconsistent to grant full insights.

A simplified approach promises more success.

The basics of Binary Network Tomography

If a path does not have property A, then **none** of the nodes have property A.

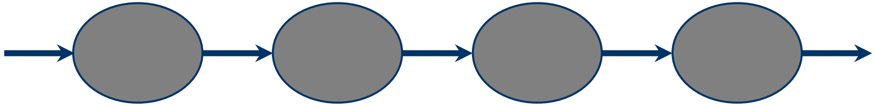

The basics of Binary Network Tomography

If a path does not have property A, then **none** of the nodes has property A.

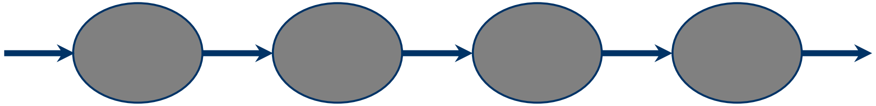

If a path has property A, then **at least one** of the nodes has property A.

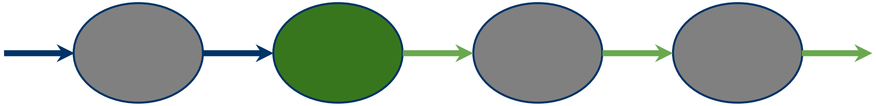

The Benefit of Binary

Observation of property A on a path is binary.

- We see only two types of paths

For each transition node *i*, we can model this with a variable

$$x_i = \begin{cases} 0 \text{ if node } i \text{ has property A,} \\ 1 \text{ if node } i \textbf{ does not} \text{ have property A.} \end{cases}$$

Then the property of a path *j* of nodes N_J

can be expressed as

$$y_j = \prod_{i \in N_j} x_i$$

Resolving Binary Network Tomography

With just one path of property A, we cannot decide which node is responsible for the observation.

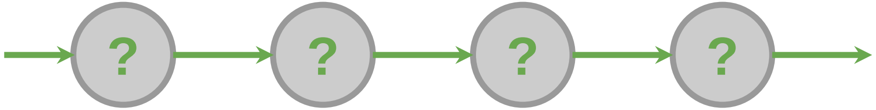

Resolving Binary Network Tomography

With just one path of property A, we cannot decide which node is responsible for the observation.

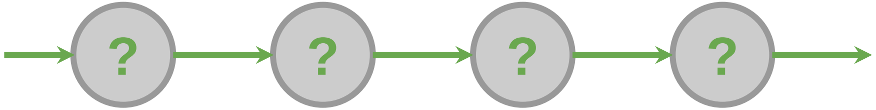

But with multiple paths we can solve a simple question.

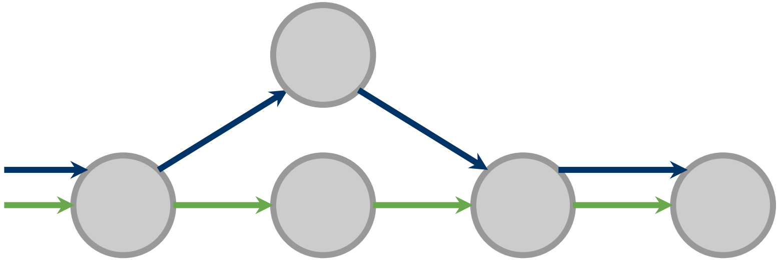

Which node is only on the green path?

Resolving Binary Network Tomography

With just one path of property A, we cannot decide which node is responsible for the observation.

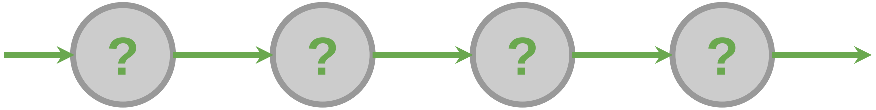

But with multiple paths we can solve a simple question.

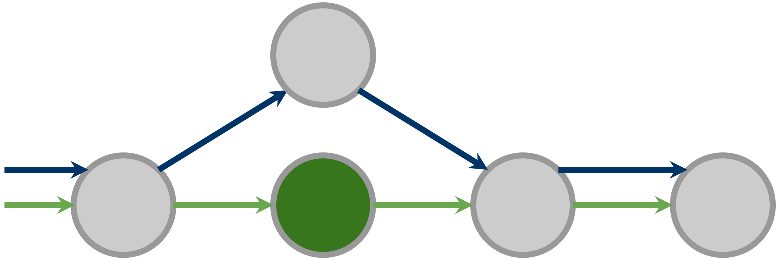

Which node is only on the green path?

Measurements in the wild. More challenges.

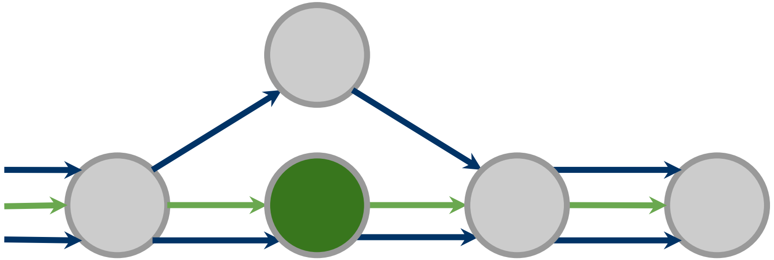

Measurements introduce noise.

Measurements in the wild. More challenges.

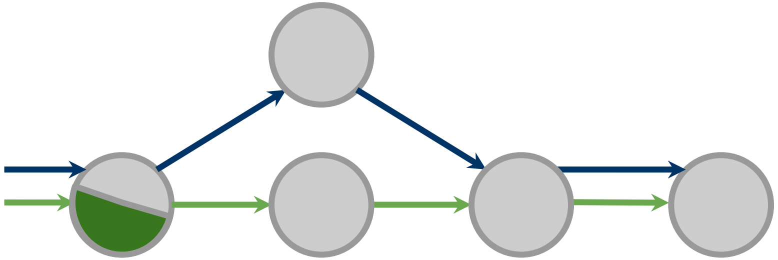

Measurements introduce noise.

Nodes behave inconsistently.

BeCAUSe –

Bayesian Computation for AUtonomous System

Pinpoint ASes based on path information.

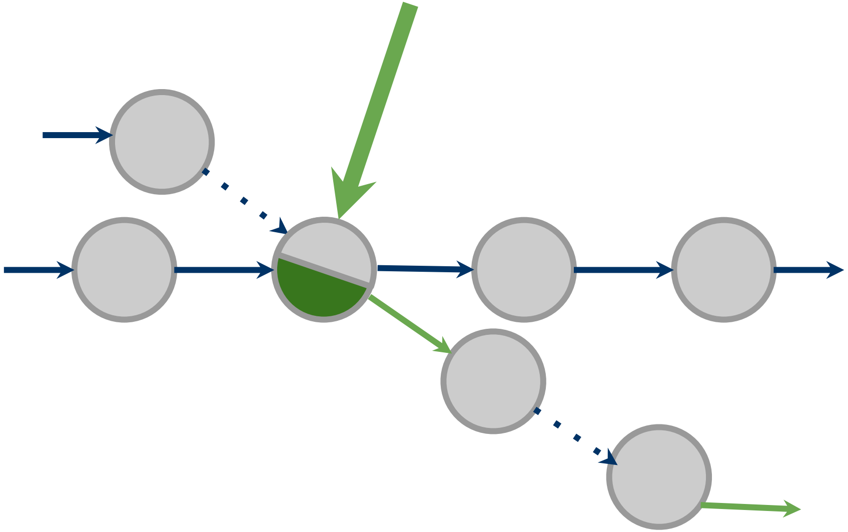

Approach: BeCAUSe

Bayesian Computation for AUtonomous System

Instead of framing pinpointing as a binary network tomography problem we consider a **probabilistic setting**.

Approach: BeCAUSe

Bayesian Computation for AUtonomous System

Instead of framing pinpointing as a binary network tomography problem we consider a **probabilistic setting**.

Each AS *i* has a **probability value** (**p_i**) **of** implementing A.

Approach: BeCAUSe

Bayesian Computation for AUtonomous System

Instead of framing pinpointing as a binary network tomography problem we consider a **probabilistic setting**.

Each AS *i* has a **probability value (**p_i**)** of implementing A .

Likelihood of a single path:

$$\mathbf{P}(\mathrm{path}|\mathrm{p}) = \begin{cases} \prod_{i \in \mathrm{path}} (1 - p_i), & \text{if path does not show A,} \\ 1 - \prod_{i \in \mathrm{path}} (1 - p_i), & \text{if path shows A.} \end{cases}$$

Bayesian inference of all paths

Posterior distribution of A given the observed data set D

$$\mathbf{P}(\mathbf{p}|D) \propto \underbrace{\mathbf{P}(D|\mathbf{p})}_{\text{Likelihood}} \cdot \underbrace{\mathbf{P}(\mathbf{p})}_{\text{Prior}}$$

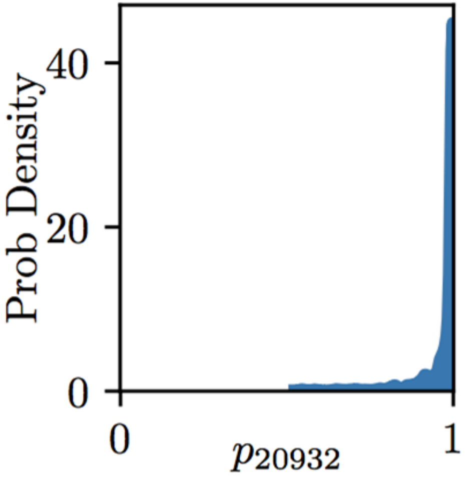

Bayesian inference of all paths

Posterior distribution of A given the observed data set D

$$\underbrace{\mathbf{P}(\mathbf{p}|D)}\propto\underbrace{\mathbf{P}(D|\mathbf{p})}_{\text{Likelihood}}\cdot\underbrace{\mathbf{P}(\mathbf{p})}_{\text{Prior}}$$

We take samples of the LHS using **Metropolis-Hastings** and **Hamiltonian Monte Carlo** to get a **probability distribution** for each AS.

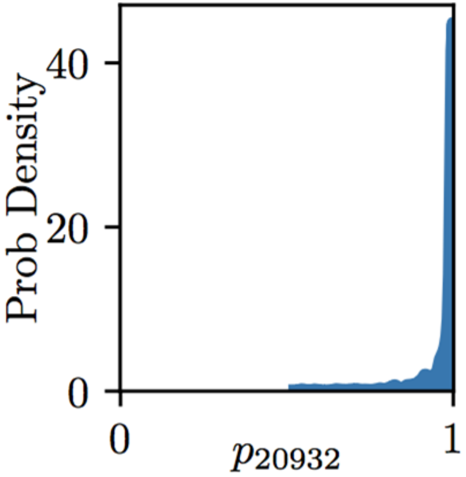

Applying BeCAUSe

PINPOINTING ROUTER BEHAVIOR IN THE WILD

Application Examples

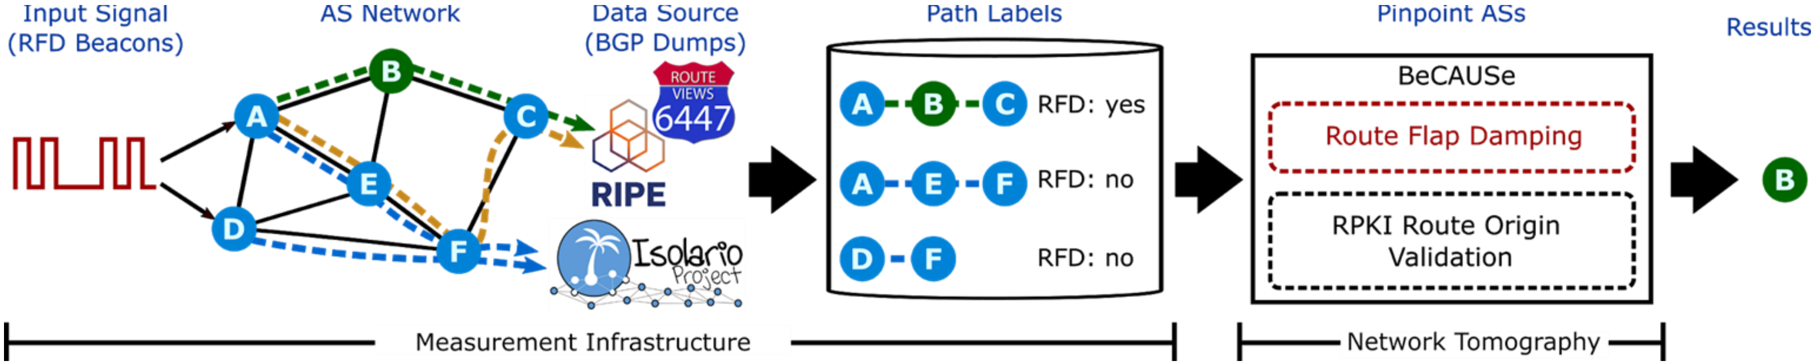

What is BGP Route Flap Damping?

Route

Flap

Damping

10.20.30.0/24

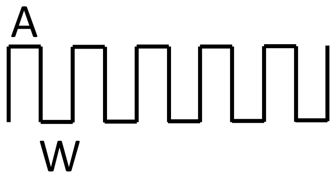

Why should you care about RFD deployment?

RFD **impacts** passive and active BGP **measurements**.
Current results on BGP Update noise may be an underestimation.

Deprecated default **parameters affect Internet reachability.**
Especially in today's rich topology.

Many **different recommendations** in the past two decades.
Different configurations may lead to conflicting goals.

Why should you care about RFD deployment?

RFD **impacts** passive and active BGP **measurements**.

Current results on BGP Update noise may be an

No measurements of BGP Route Flap Damping deployment.

y.

Many **different recommendations** in the past two decades.

Different configurations may lead to conflicting goals.

Let us measure deployment of RFD.

First we need path data.

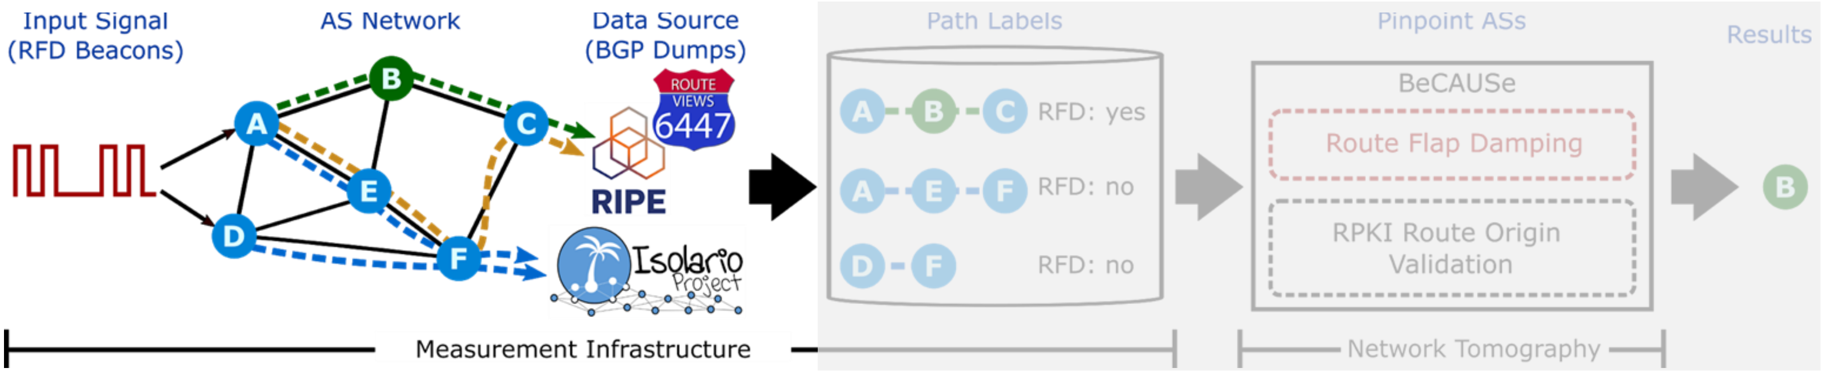

How does Route Flap Damping work?

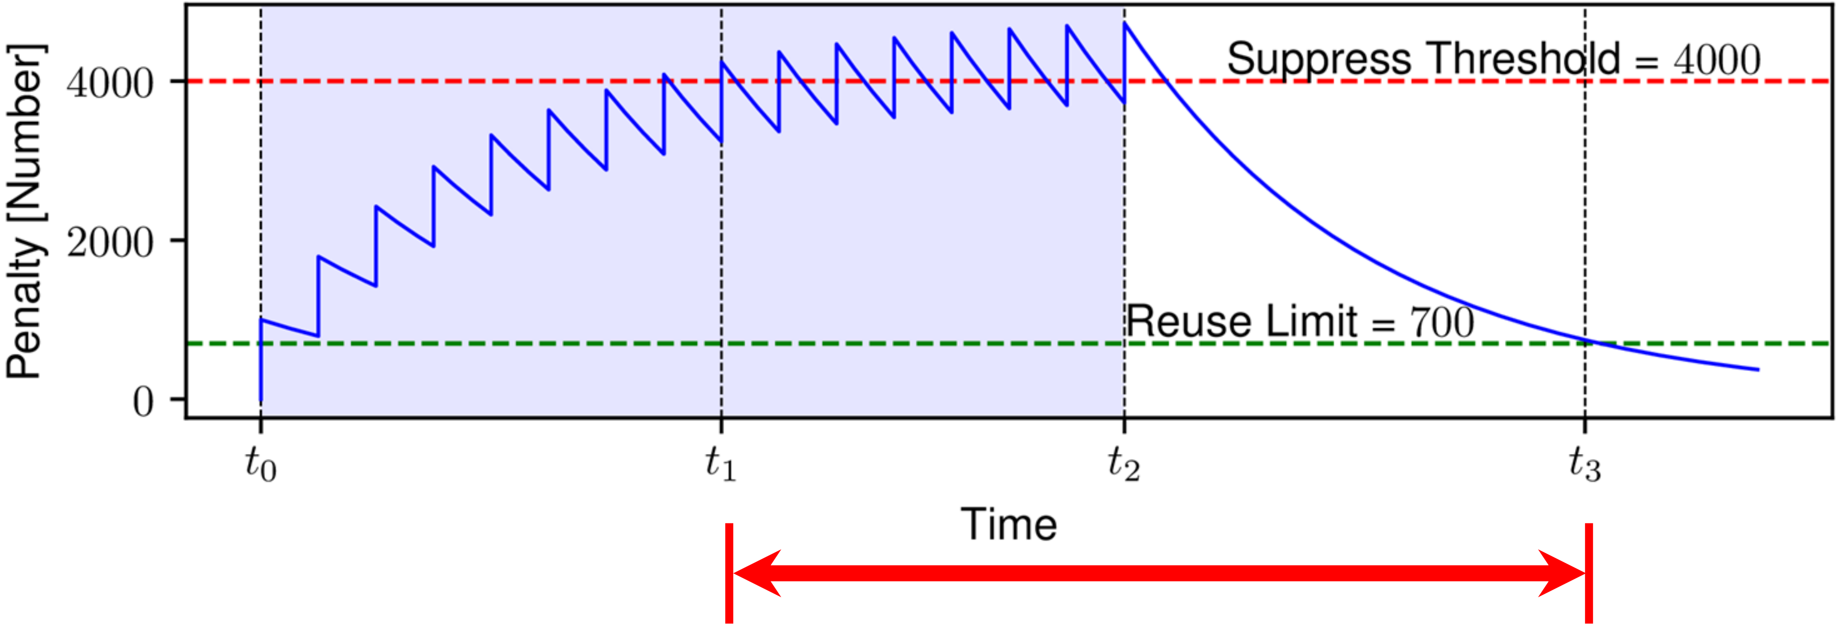

RFD active

How does Route Flap Damping work? AIMD principle.

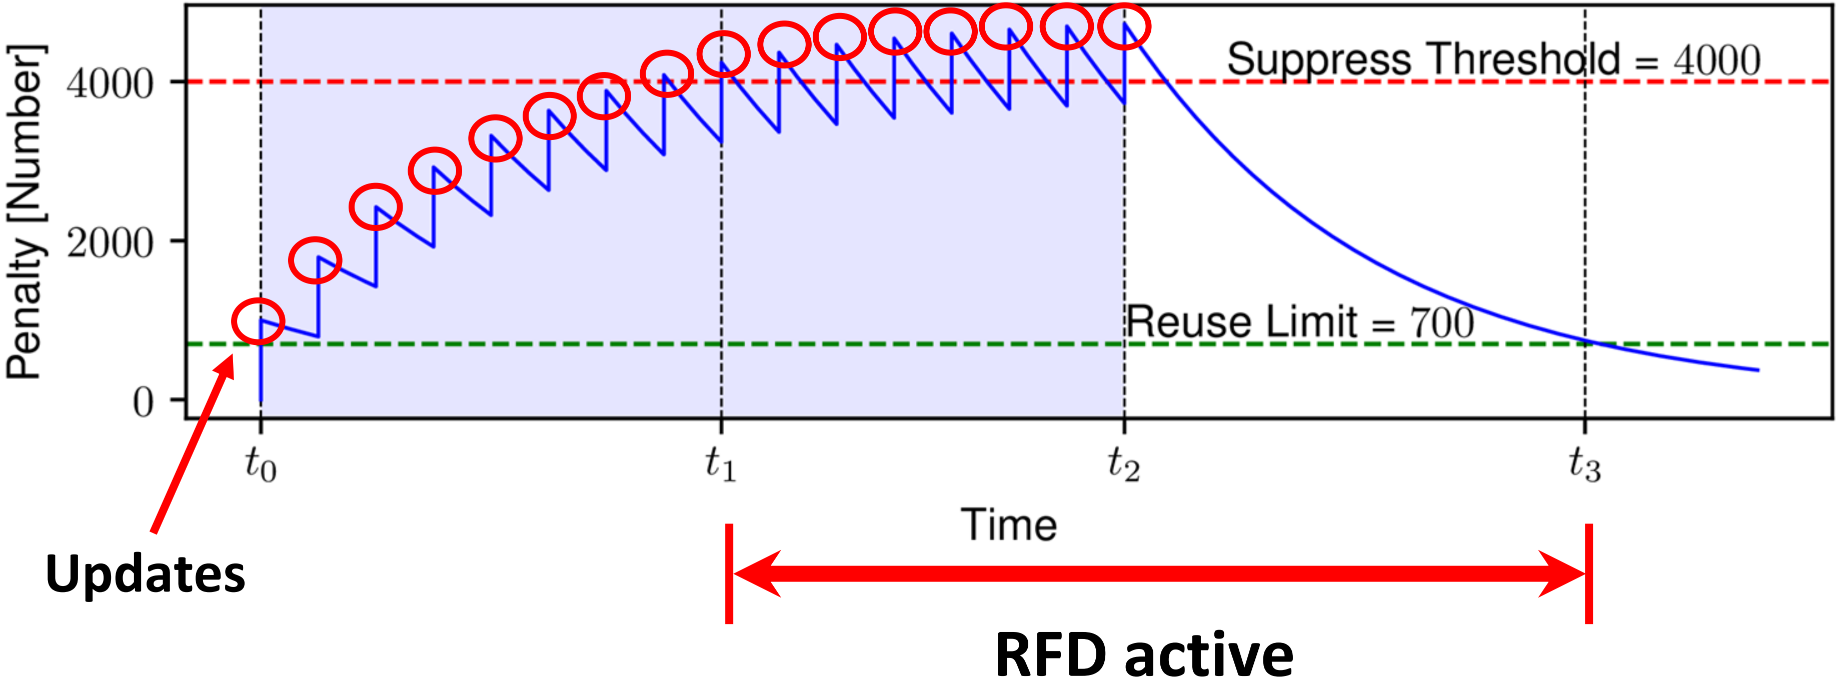

How does Route Flap Damping work? Start.

Damping starts

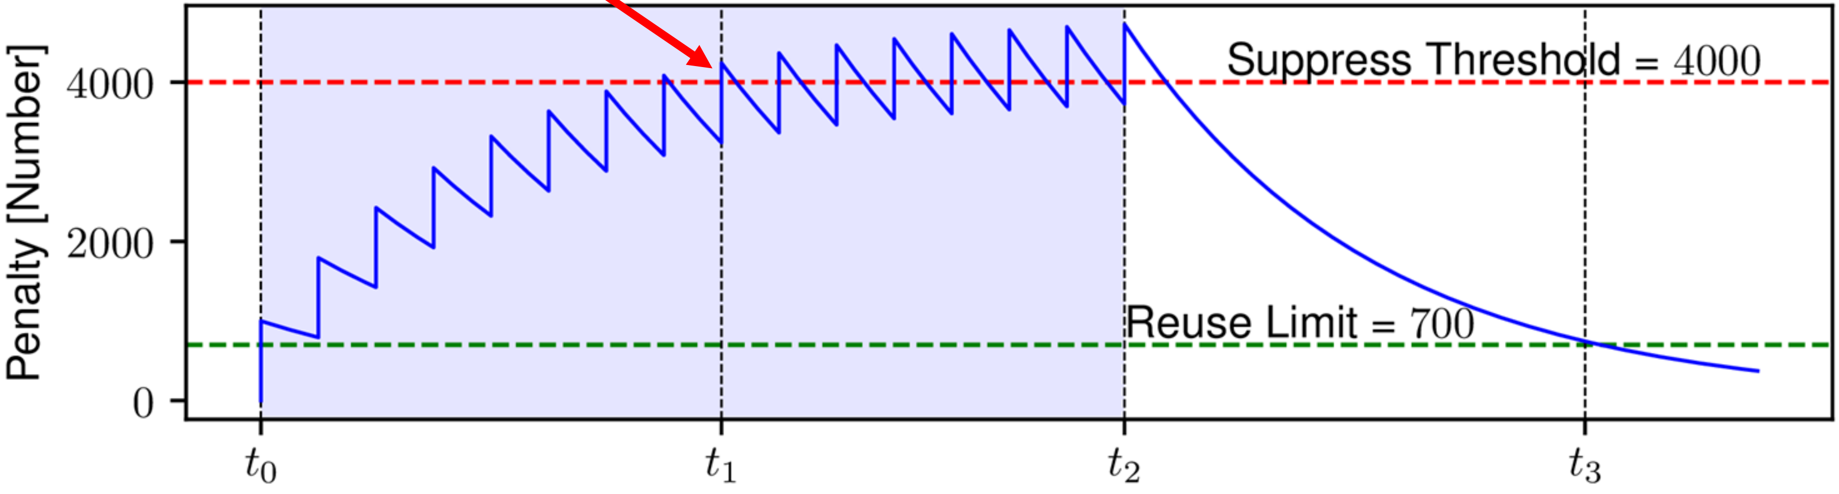

RFD active

How does Route Flap Damping work? Wait.

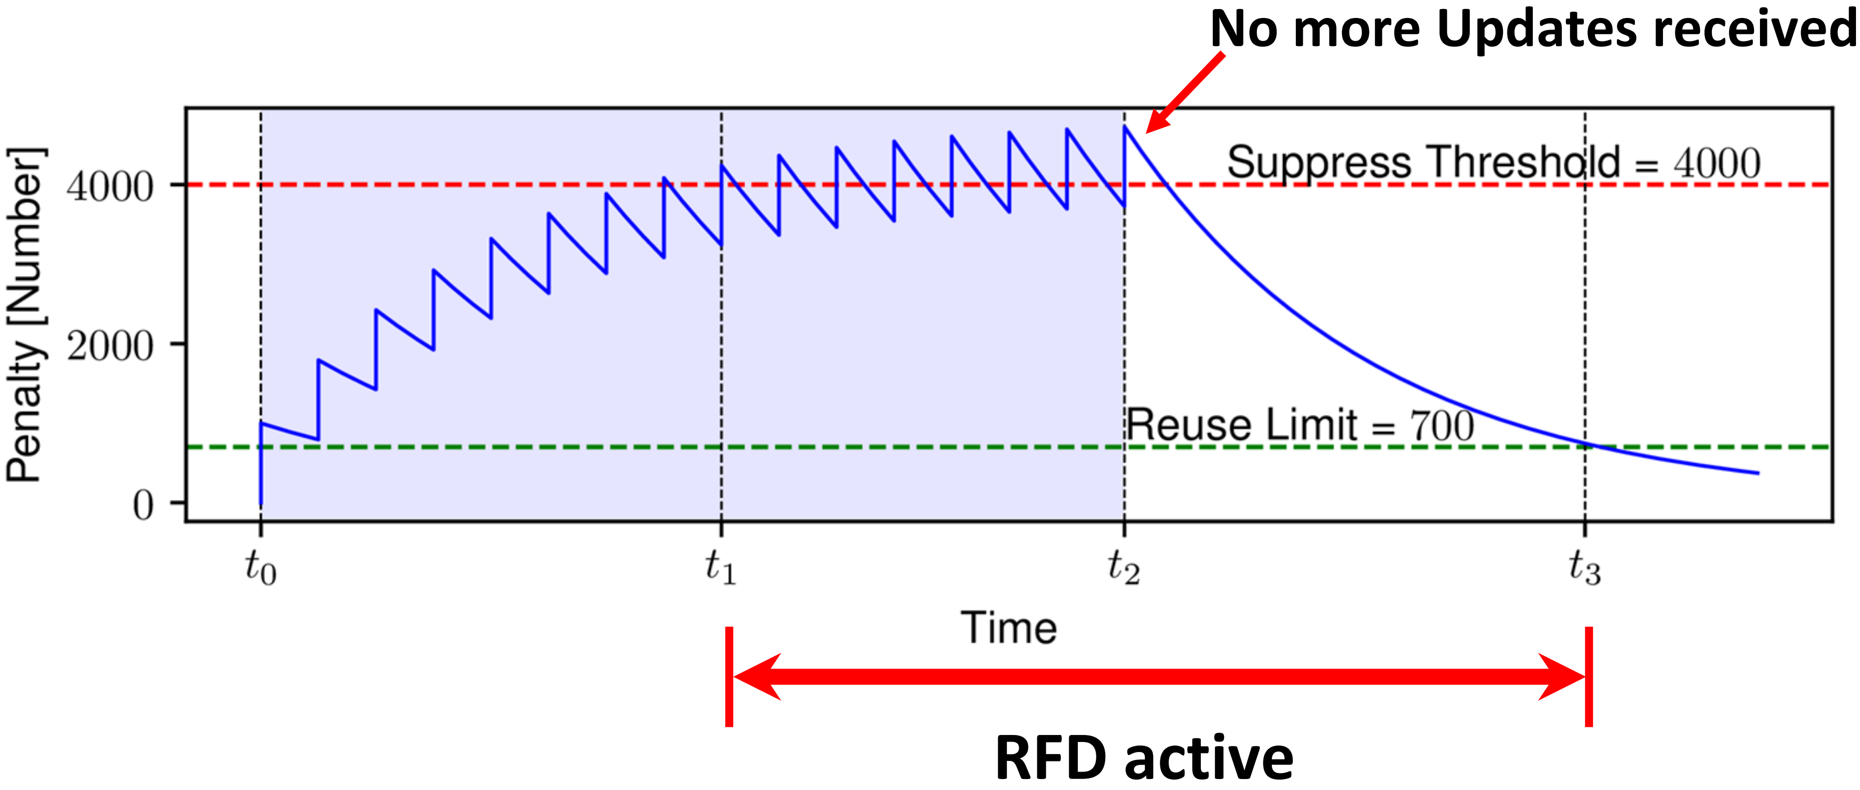

No more Updates received

RFD active

How does Route Flap Damping work? Release.

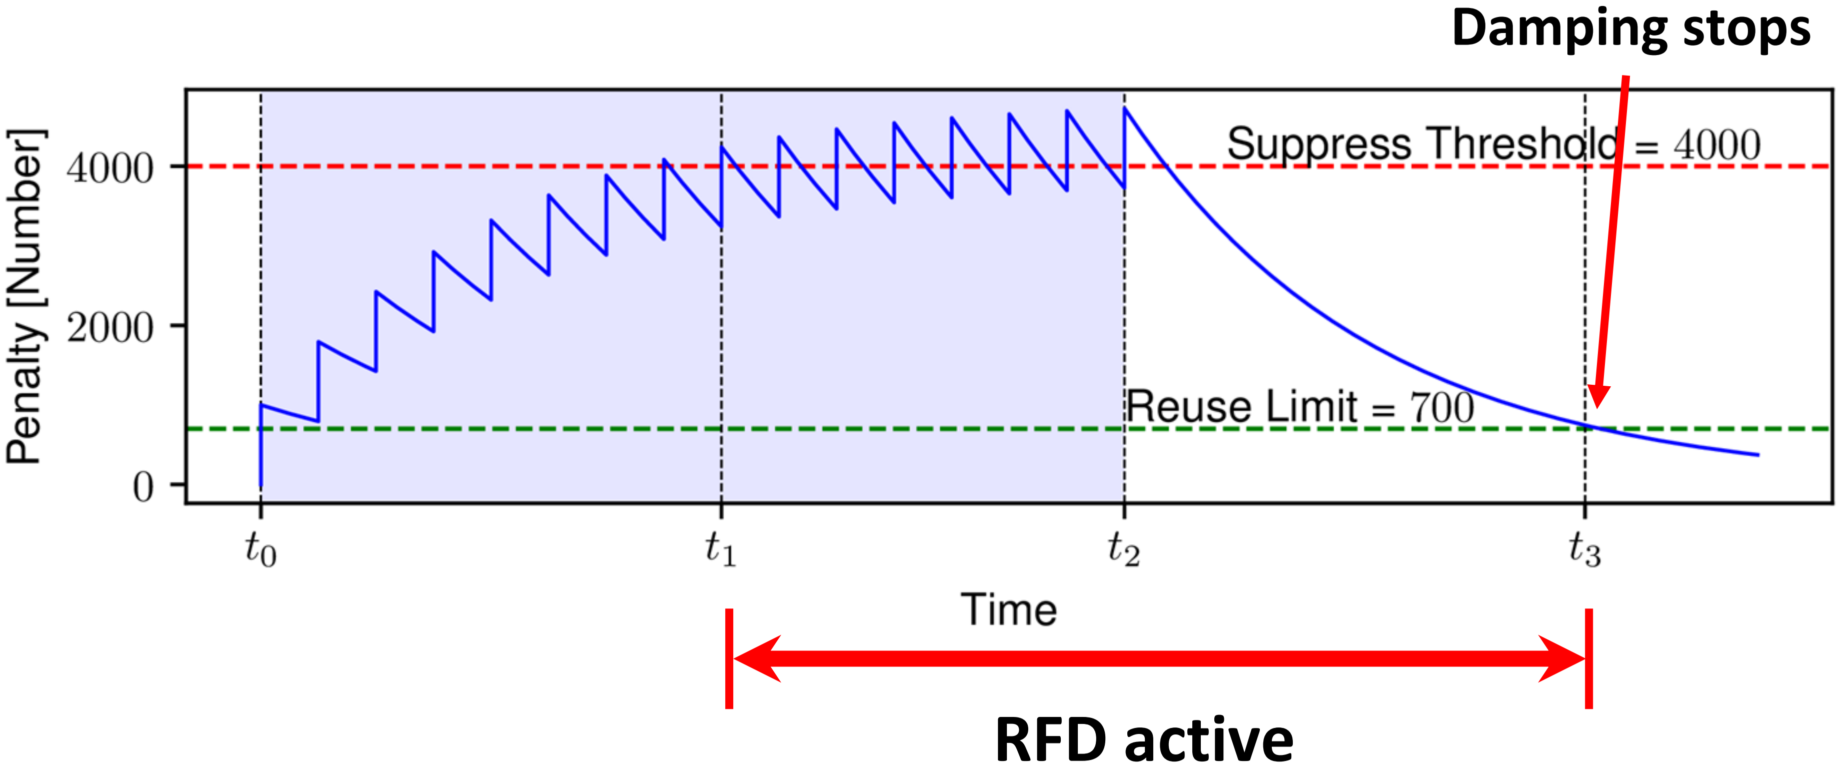

Generating the RFD update signature

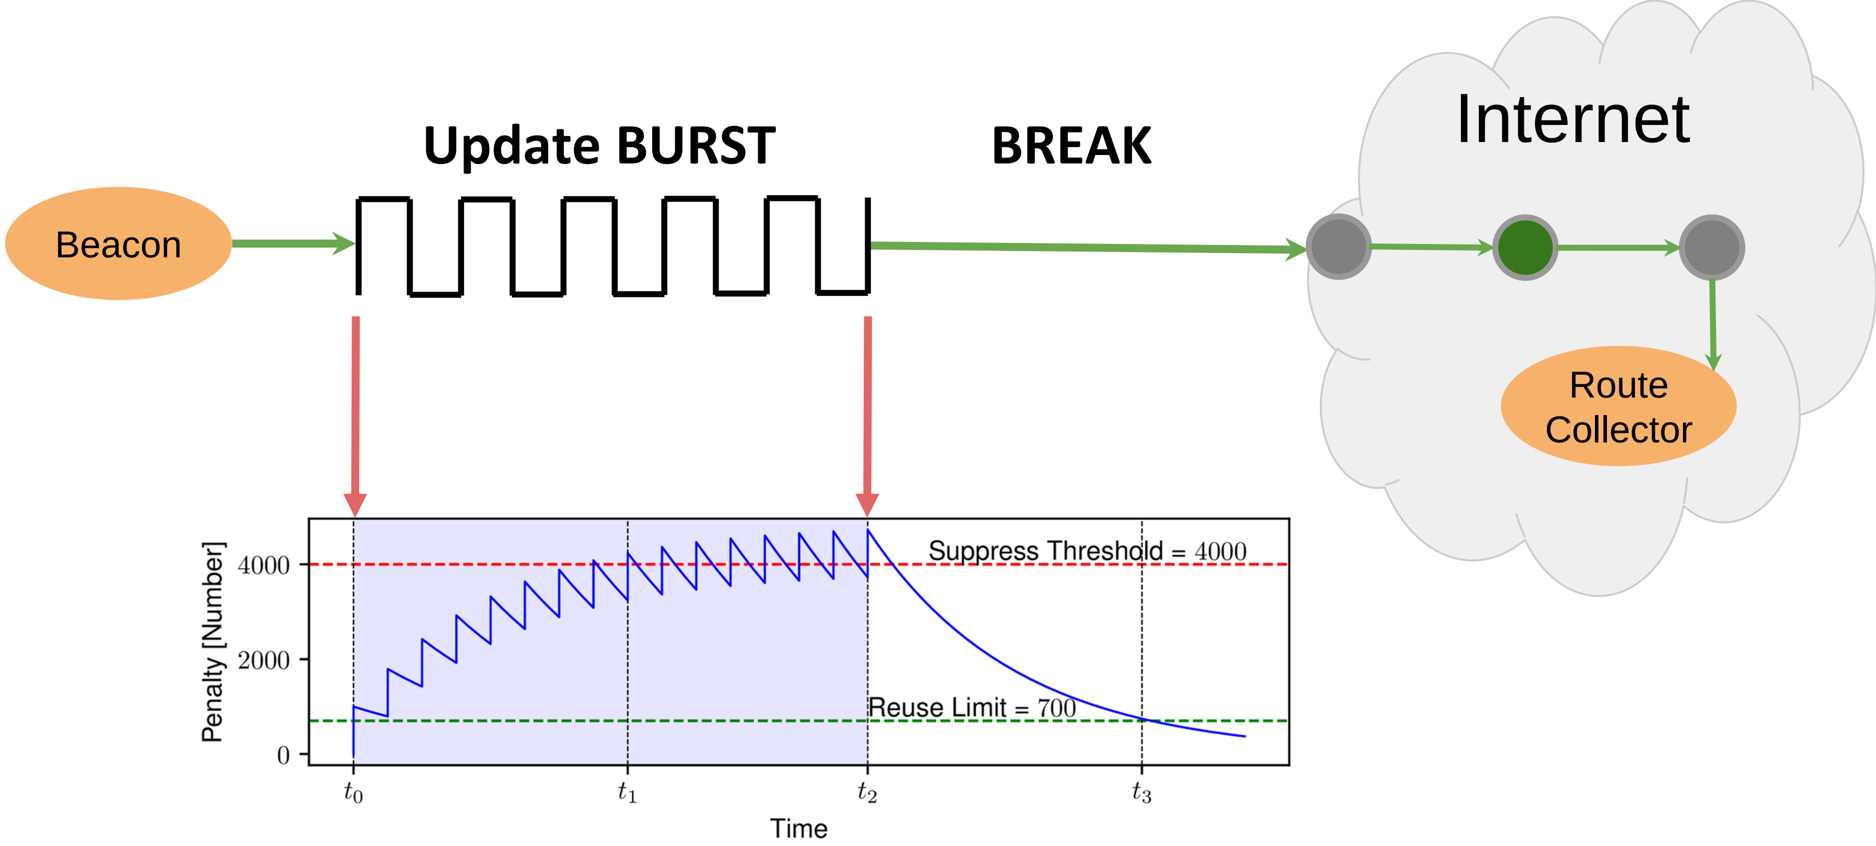

BGP Beacons with different frequencies

Beacon prefixes

Update patterns

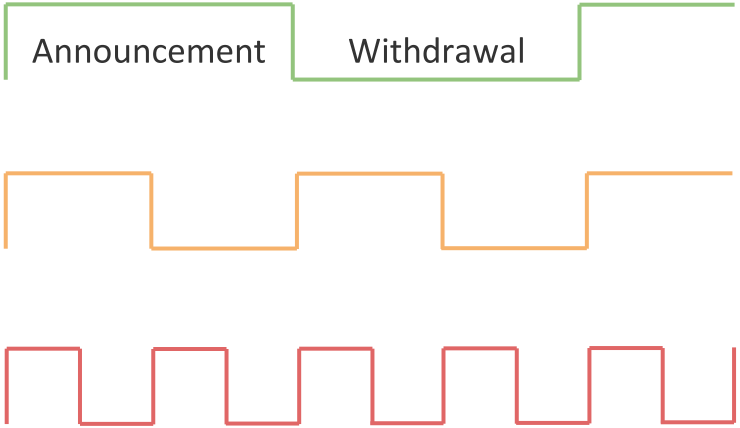

147.28.35.0/24 — long update interval

147.28.34.0/24

147.28.33.0/24 — short update interval

Locations of our BGP Beacons

Bangkok, TH

Johannesburg, ZA

København, DK

München, DE

São Paulo, BR

Seattle, US

Tokyo, JP

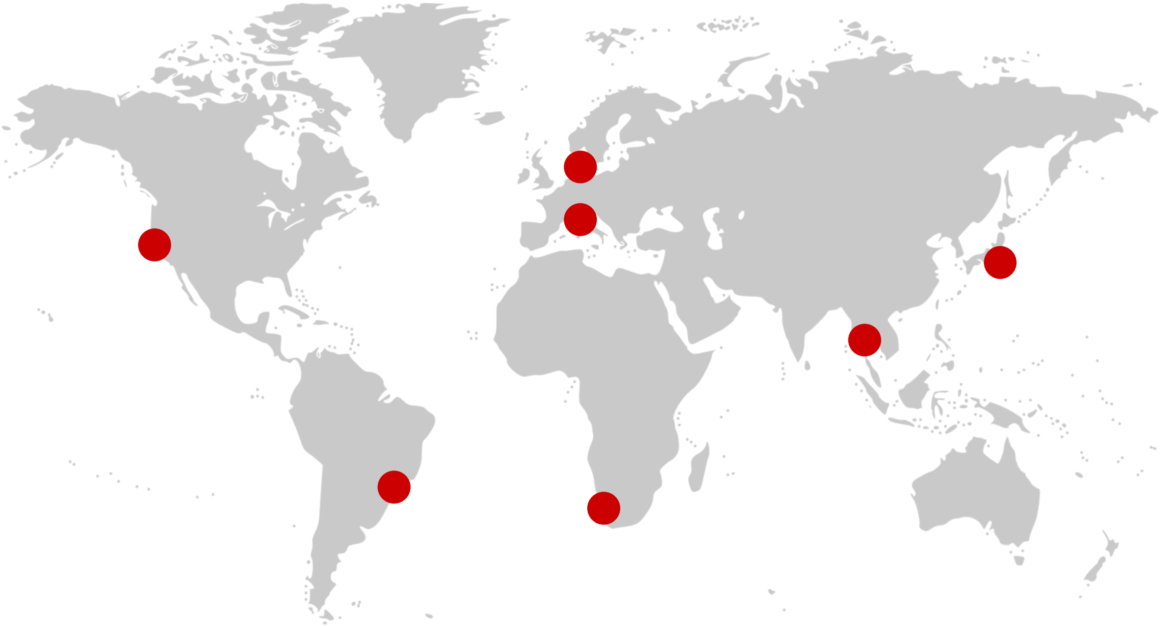

RFD causes a very recognizable pattern.

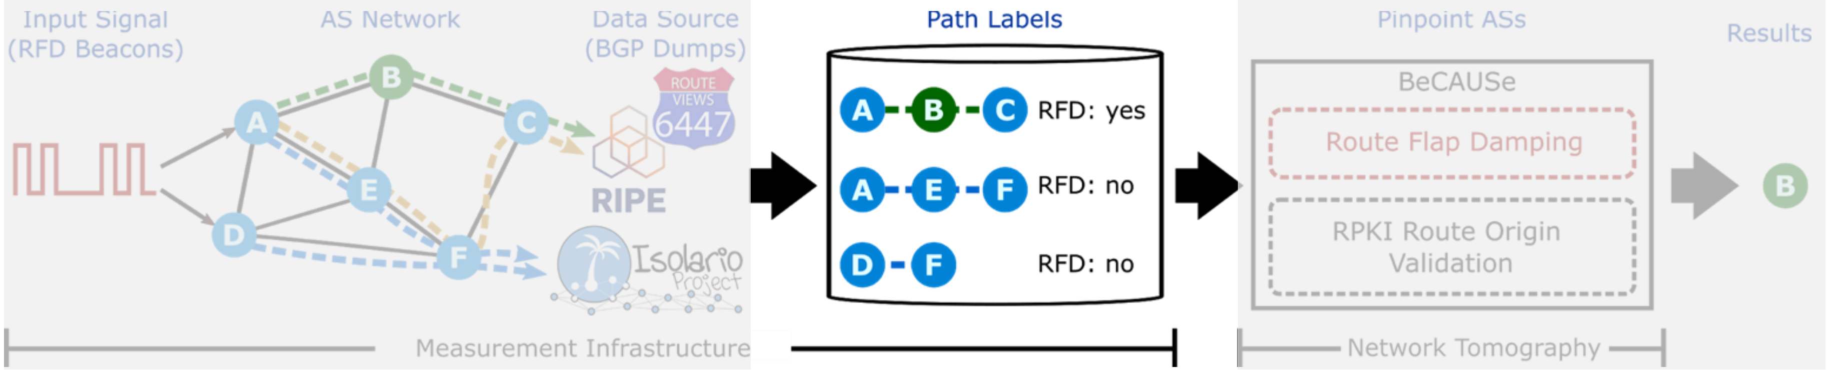

View from a Vantage Point

Our prefixes are damped during the Burst (Blue) and re-advertised during the Break (White).

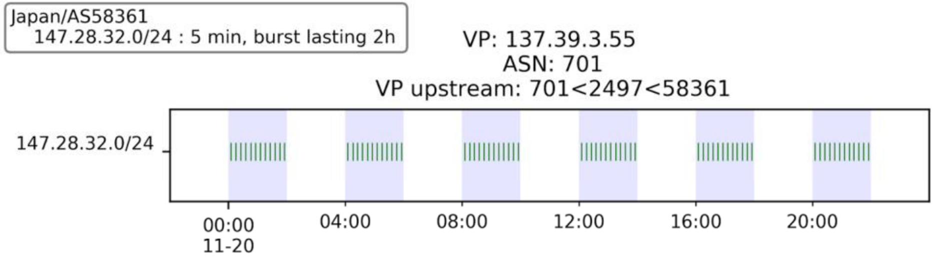

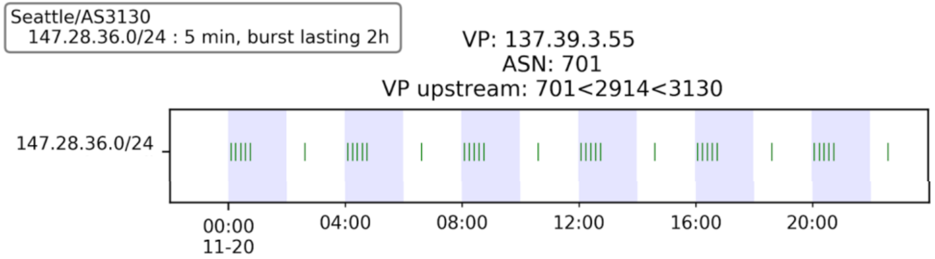

View from a Vantage Point

Our prefixes are damped during the Burst (Blue) and re-advertised during the Break (White).

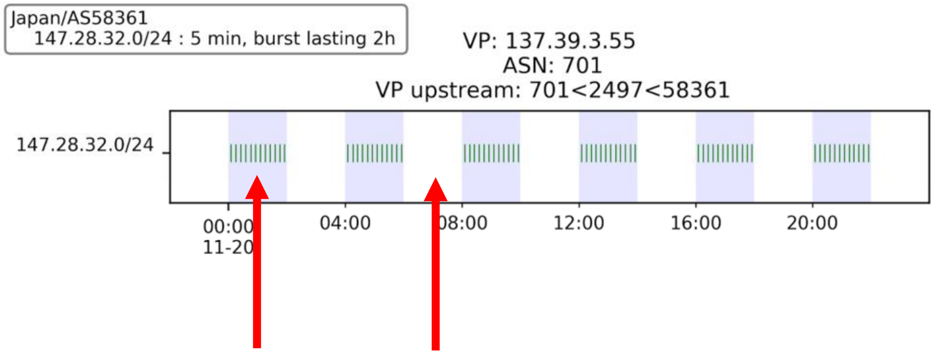

Burst

Break

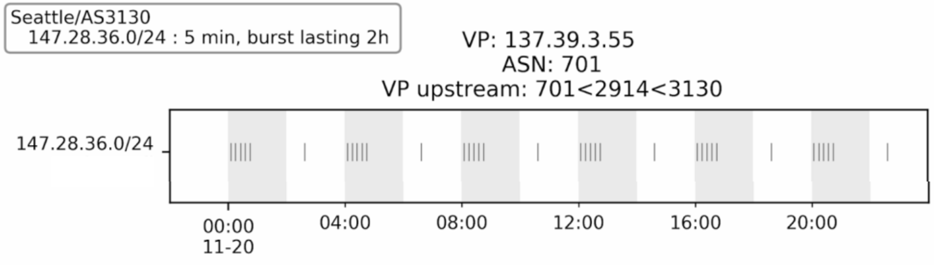

View from a Vantage Point

Our prefixes are damped during the Burst (Blue) and re-advertised during the Break (White).

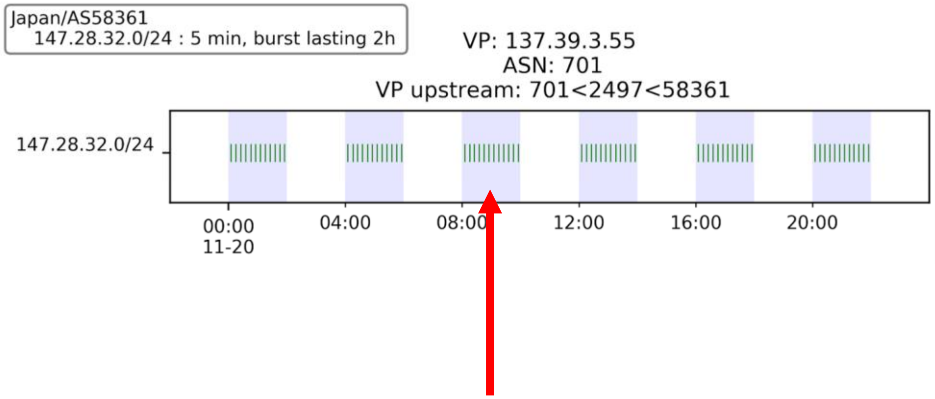

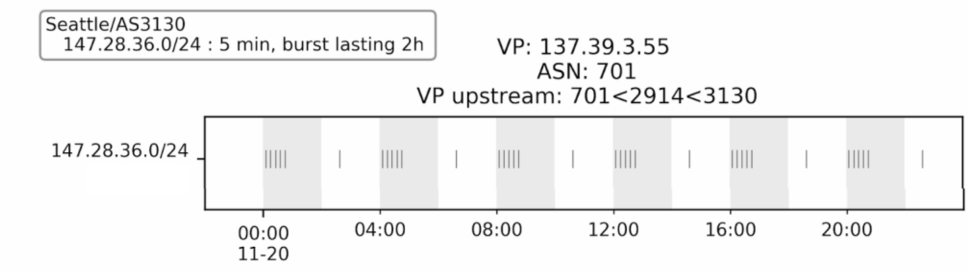

NO Route Flap Damping

View from a Vantage Point

Our prefixes are damped during the Burst (Blue) and re-advertised during the Break (White).

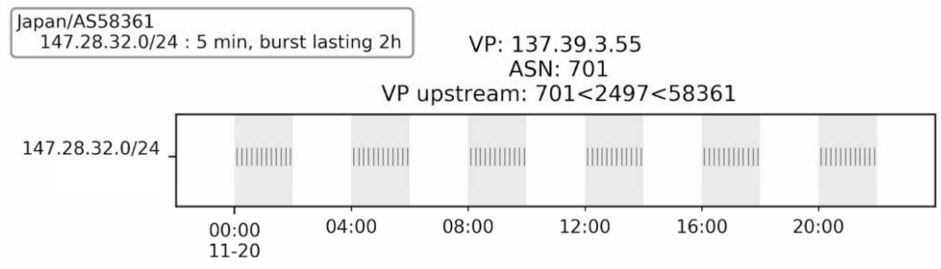

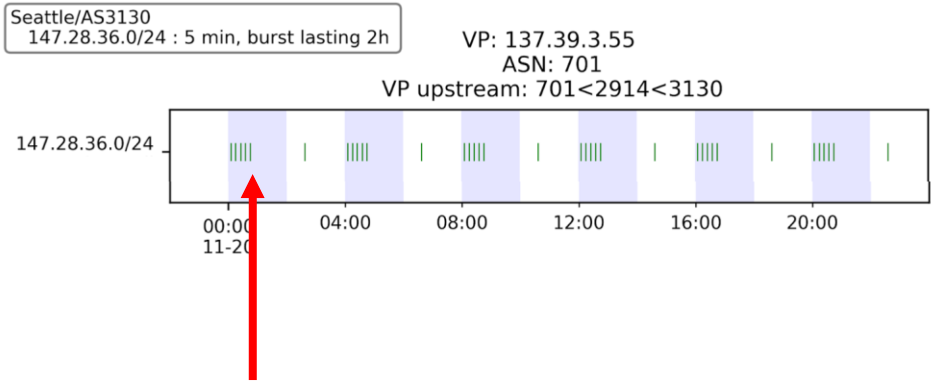

Signal Stops

View from a Vantage Point

Our prefixes are damped during the Burst (Blue) and re-advertised during the Break (White).

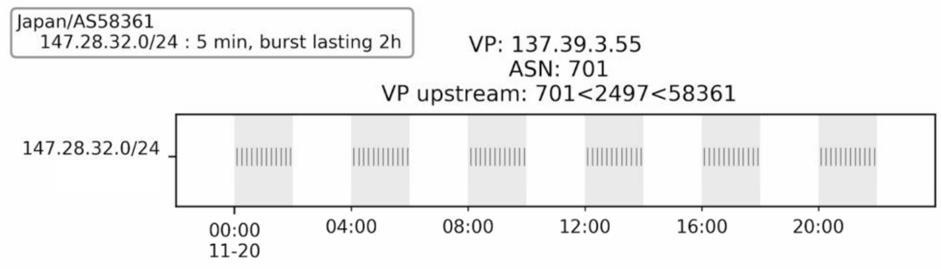

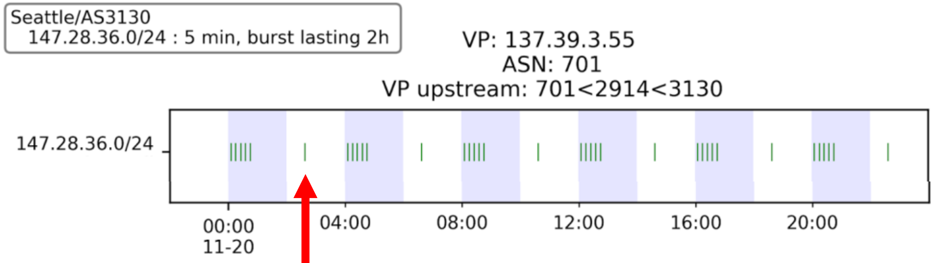

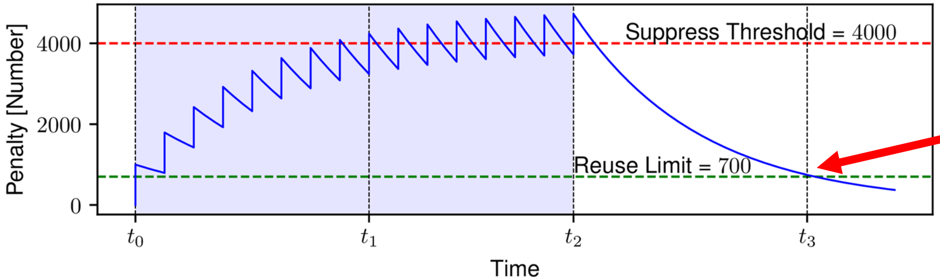

Re-advertisement

View from a Vantage Point

Our prefixes are damped during the Burst (Blue) and re-advertised during the Break (White).

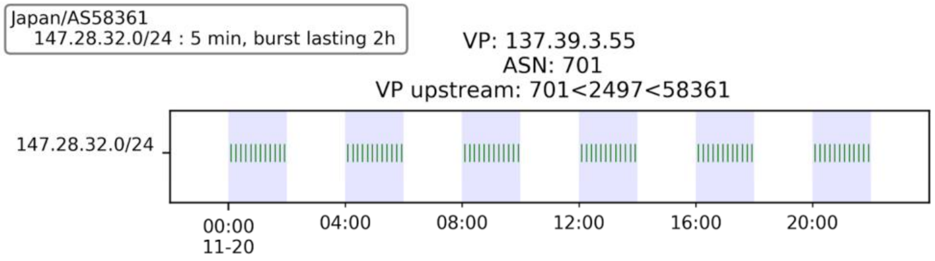

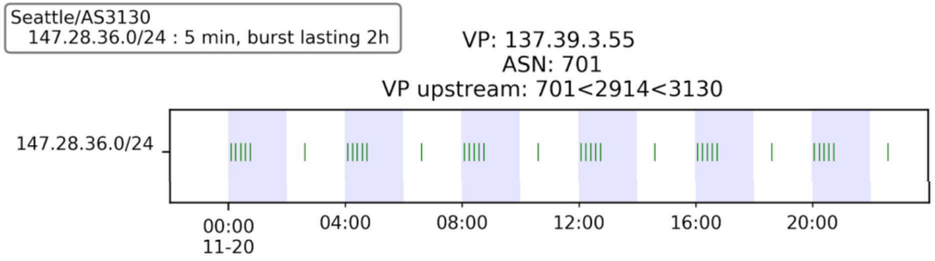

NO Route Flap Damping

Route Flap Damping

Pinpointing ASs based on path information.

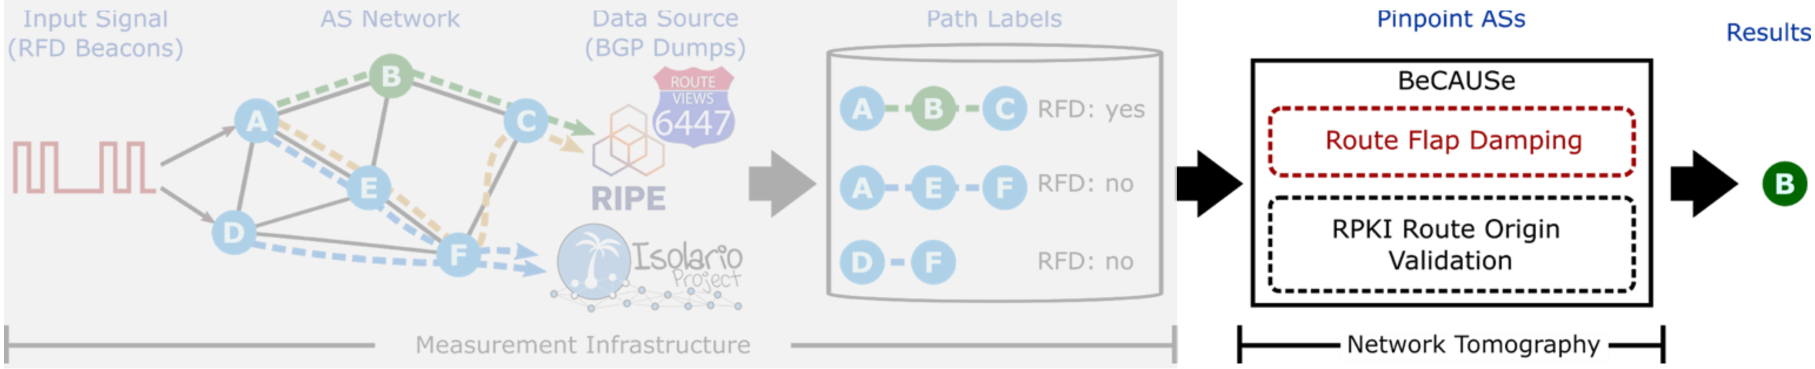

Challenges when pinpointing RFD ASs

1) If we find the Route Flap Damping pattern at some vantage points, **the damper could be anywhere on the AS path**.

2) Some ASs use **Route Flap Damping selectively** on a subset of neighbors.

3) **Noise** inherent in the measurement cannot result in a binary solution.

Applying BeCAUSe

Measure RFD on all path available via Vantage Points.

Compute RFD probabilities for paths.

Perform Monte Carlo sampling to get obtain a distribution for each AS.

Evaluate and classify results.

Output distribution types

Low Variance / High Certainty

Route Flap Damping

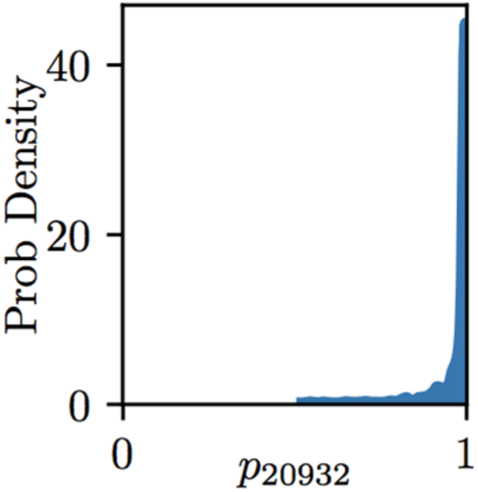

No Route Flap Damping

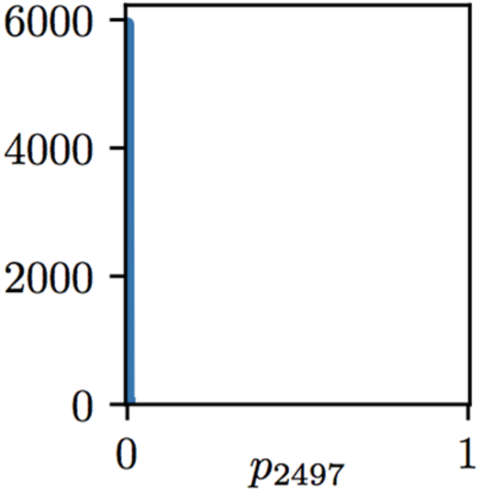

Output distribution types

Special cases

Inconsistent Damping

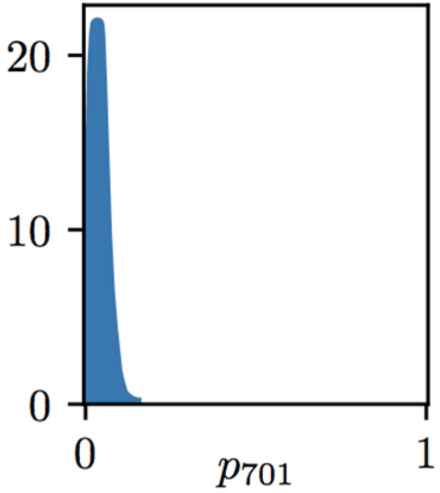

Lacking Data

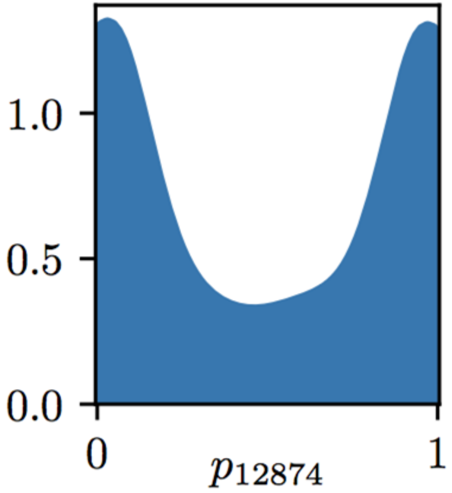

Distribution summary

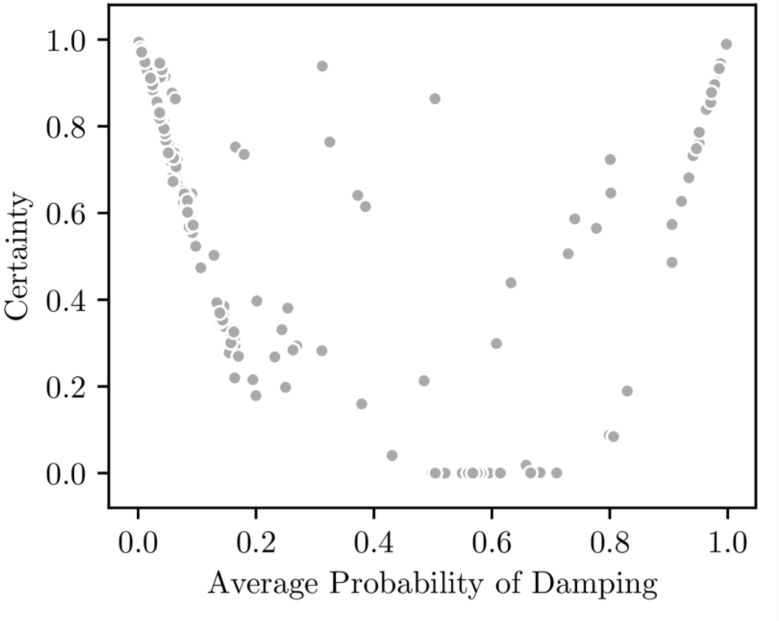

Goal: **Hard decisions** on RFD deployment.

Distribution summary - Damping ASs

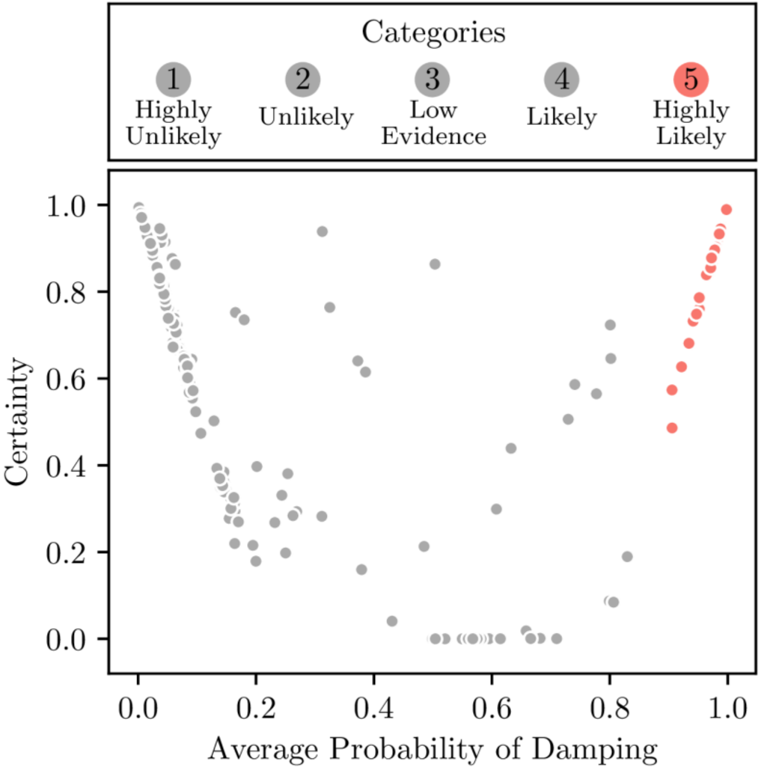

High damping probability combined with **high certainty** are labeled damping ASs.

Distribution summary - Inconsistently damping ASs

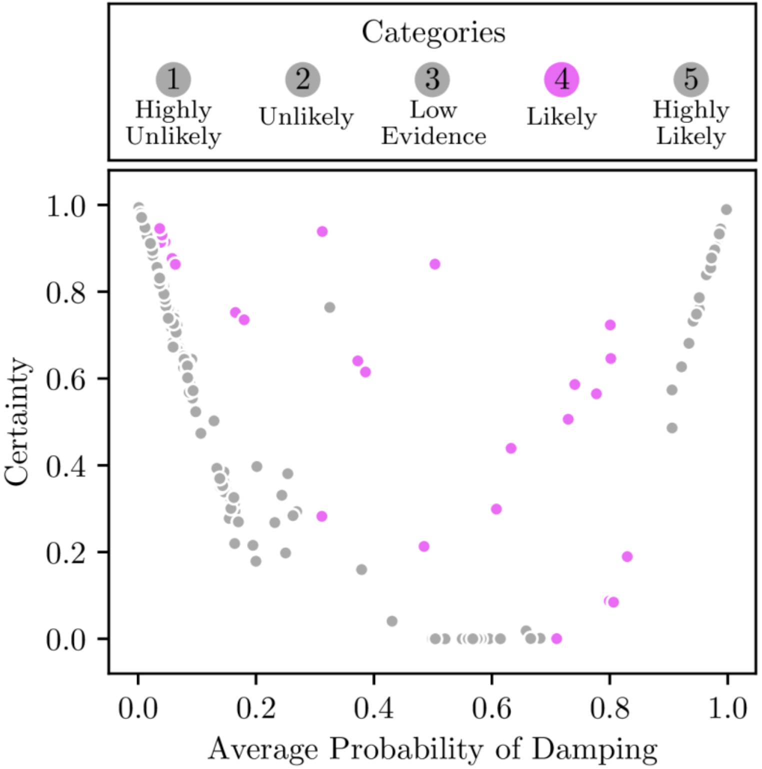

For **RFD labeled paths without** nodes labeled as damping, we find and label the node that is **most likely causing RFD**.

Distribution summary - Non-damping ASs

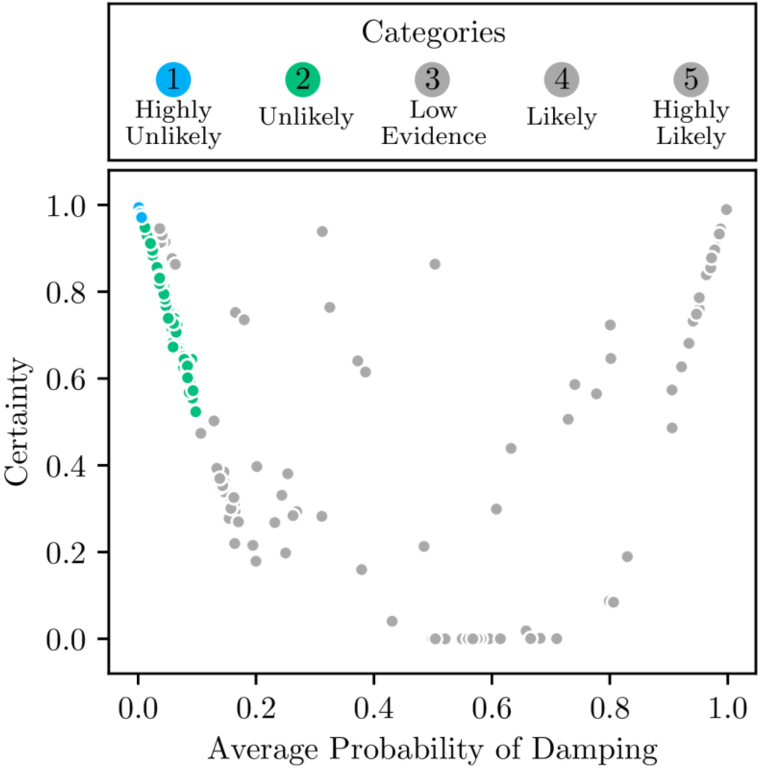

Low damping probability combined with **high certainty** are labeled non-damping ASs.

Distribution summary - Lacking evidence

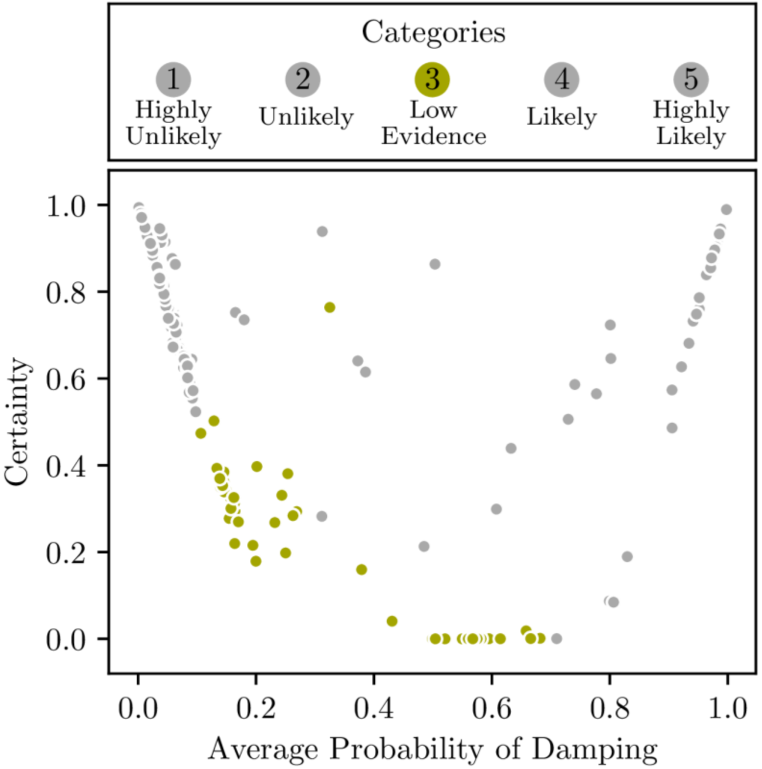

For the **remaining ASs** we cannot draw conclusions about RFD deployment.

Summarised distributions

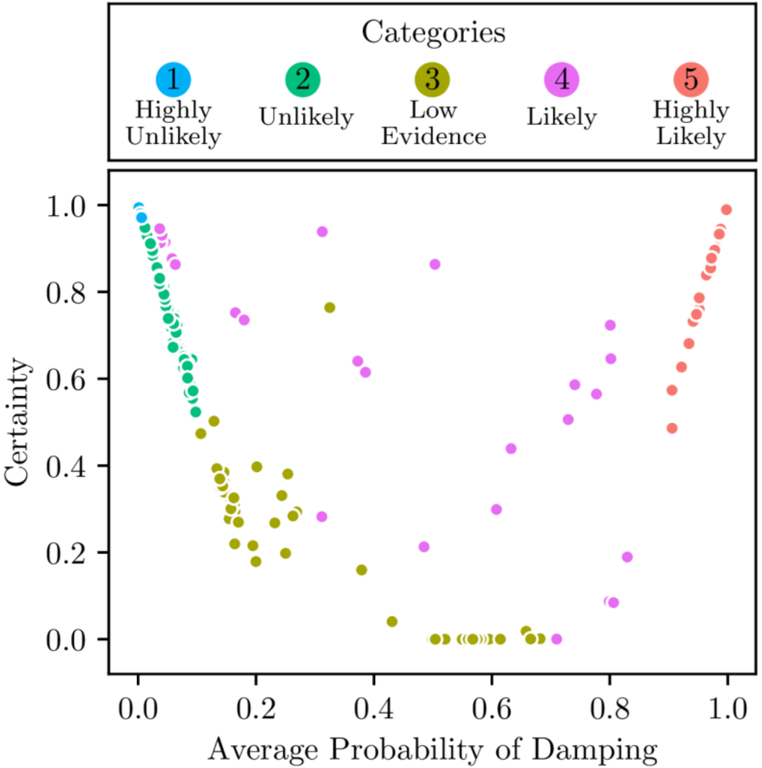

RFD

Category	1	2	3	4	5
Total	166	283	72	25	28
Share	28.9%	49.3%	12.5%	4.3%	4.8%

Route Flap Damping deployment

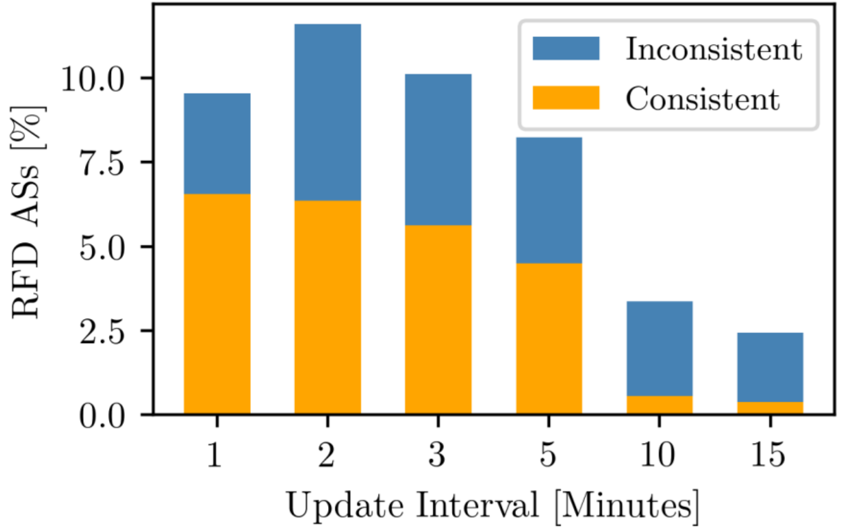

9% is the **lower bound** of RFD deployment due to lacking visibility.

Verified by 75 ASs BeCAUSe has **100% precision** and **87% recall.**

Which RFD parameters are deployed?

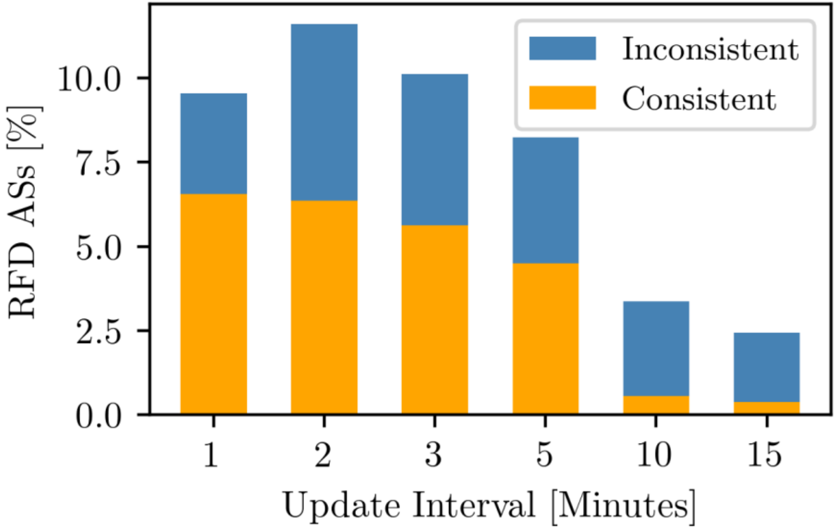

Routers with **RFD default params** start damping at the **5 minute update interval and lower.**

Most ASs use **deprecated vendor default configurations**.

Using BeCAUSe for ROV detection

Route Origin Validation: RPKI-based filtering of invalid prefix origins, as increasingly deployed on the Internet.

Using BeCAUSe for ROV detection

Route Origin Validation: RPKI-based filtering of invalid prefix origins, as increasingly deployed on the Internet.

ROV is a binary property that can be measured in similar controlled experiments. Ground truth exists from a previous operator study.

Using BeCAUSe for ROV detection

Specific to research question: summarisation techniques of probability and certainty results.

ROV is a binary property that can be measured in similar controlled experiments. Ground truth exists from a previous operator study.

In a separate simulation, we used **BeCAUSe to pinpoint ROV** deployment. We achieved identical precision, but lower recall due to high ROV usage.

Literature

C .Gray, C. Mosig, R. Bush, C. Pelsser, M. Roughan, T.C. Schmidt, M. Wählisch (2020).

BGP Beacons, Network Tomography, and Bayesian Computation to Locate Route Flap Damping.

ACM Internet Measurement Conference (IMC), p. 492–505, ACM : New York, 2020.

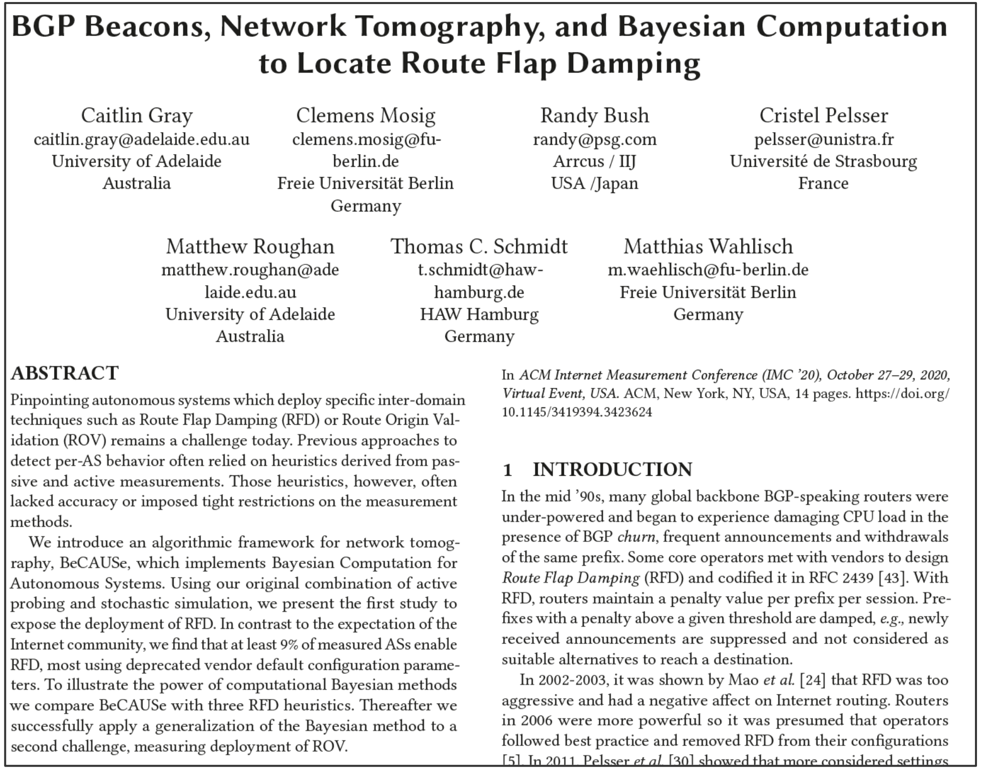

BGP Beacons, Network Tomography, and Bayesian Computation to Locate Route Flap Damping

Caitlin Gray
caitlin.gray@adelaide.edu.au
University of Adelaide
Australia

Clemens Mosig
clemens.mosig@fu-berlin.de
Freie Universität Berlin
Germany

Randy Bush
randy@psg.com
Arrcus / IIJ
USA /Japan

Cristel Pelsser
pelsser@unistra.fr
Université de Strasbourg
France

Matthew Roughan
matthew.roughan@adelaide.edu.au
University of Adelaide
Australia

Thomas C. Schmidt
t.schmidt@haw-hamburg.de
HAW Hamburg
Germany

Matthias Wahlisch
m.waehlisch@fu-berlin.de
Freie Universität Berlin
Germany

ABSTRACT

Pinpointing autonomous systems which deploy specific inter-domain techniques such as Route Flap Damping (RFD) or Route Origin Validation (ROV) remains a challenge today. Previous approaches to detect per-AS behavior often relied on heuristics derived from passive and active measurements. Those heuristics, however, often lacked accuracy or imposed tight restrictions on the measurement methods.

We introduce an algorithmic framework for network tomography, BeCAUSe, which implements Bayesian Computation for Autonomous Systems. Using our original combination of active probing and stochastic simulation, we present the first study to expose the deployment of RFD. In contrast to the expectation of the Internet community, we find that at least 9% of measured ASs enable RFD, most using deprecated vendor default configuration parameters. To illustrate the power of computational Bayesian methods we compare BeCAUSe with three RFD heuristics. Thereafter we successfully apply a generalization of the Bayesian method to a second challenge, measuring deployment of ROV.

In *ACM Internet Measurement Conference (IMC '20), October 27–29, 2020, Virtual Event, USA.* ACM, New York, NY, USA, 14 pages. https://doi.org/10.1145/3419394.3423624

1 INTRODUCTION

In the mid '90s, many global backbone BGP-speaking routers were under-powered and began to experience damaging CPU load in the presence of BGP *churn*, frequent announcements and withdrawals of the same prefix. Some core operators met with vendors to design *Route Flap Damping* (RFD) and codified it in RFC 2439 [43]. With RFD, routers maintain a penalty value per prefix per session. Prefixes with a penalty above a given threshold are damped, *e.g.*, newly received announcements are suppressed and not considered as suitable alternatives to reach a destination.

In 2002-2003, it was shown by Mao *et al.* [24] that RFD was too aggressive and had a negative affect on Internet routing. Routers in 2006 were more powerful so it was presumed that operators followed best practice and removed RFD from their configurations [5]. In 2011, Pelsser *et al.* [30] showed that more considered settings