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.

CCS CONCEPTS

• **Mathematics of computing → Bayesian computation**; • **Networks → Public Internet**; **Routing protocols**.

KEYWORDS

Metropolis-Hasting, Hamiltonian Monte Carlo, RFD, RPKI

ACM Reference Format:
Caitlin Gray, Clemens Mosig, Randy Bush, Cristel Pelsser, Matthew Roughan, Thomas C. Schmidt, and Matthias Wahlisch. 2020. BGP Beacons, Network Tomography, and Bayesian Computation to Locate Route Flap Damping. In *ACM Internet Measurement Conference (IMC '20), October 27–29, 2020, Virtual Event, USA.* ACM, New York, NY, USA, 13 pages. https://doi.org/10.1145/3419394.3423624

IMC '20, October 27–29, 2020, Virtual Event, USA

ACM ISBN 978-1-4503-8138-3/20/10…$15.00
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 of the RFD parameters were safe and helpful, and consequently it was believed that operators would re-enable RFD. But at no time in all this history was the actual deployment of RFD measured.

Understanding RFD deployment and parameters is important because RFD can cause problems reaching Internet destinations [9, 24] and can obscure active and passive control plane measurements for researchers. More importantly, RFD is the poster child for a range of problems: those where we localise routing properties in the Internet ecosystem from external measurements. Heuristics [32, 41] have been used to tackle specific cases of such problems but we seek here to create a general approach.

Such problems fall under the heading of network tomography: inference about internal network behavior from end-to-end measurements. We present a new tomographic approach here—BeCAUSe—which is adapted to this class of problems. BeCAUSe uses computational Bayesian techniques, which have large advantages but have sometimes been discarded in favour of heuristics because naive approaches (*e.g.,* Gibb's sampling) to computational Bayes are computationally costly. In this paper, we show how more advanced approaches (Metropolis-Hastings [22, 25] and Hamiltonian Monte Carlo [13]) can find RFD ASs from path measurements. From this we can learn much about the existing deployment of RFD.

The paper makes the following key contributions:

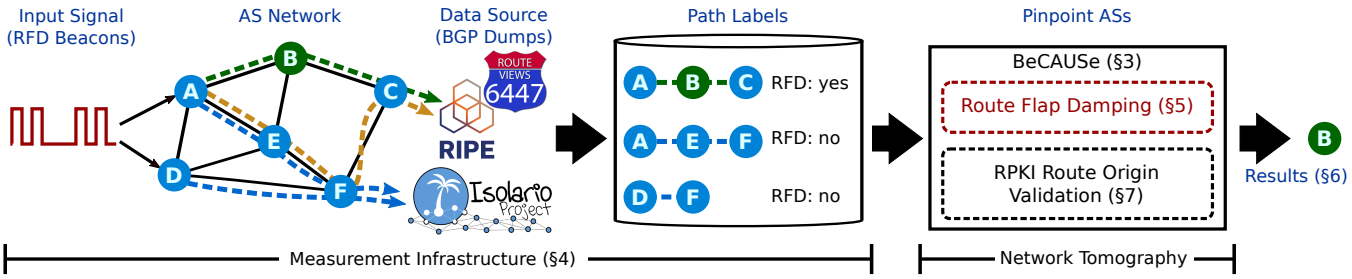

Figure 1: Overview of Measurement Infrastructure and BeCAUSe.

(1) We present a new group of network tomography algorithms adapted to the problem of large-scale routing inference christened BeCAUSe (BayEsian Computation for AUtonomous SystEms)—see § 3 and § 5.
(2) We develop a measurement infrastructure—RFD Beacons—to provide the inputs to the tomography problem—see Figure 1 and § 4.
(3) We perform the first large-scale study of the deployment of RFD in the wild, based on more than 2 months of data. The results suggest that RFD is used more widely and less carefully than one would hope or expect—see § 6.

We test BeCAUSe on the RFD and Route Origin Validation inference problems. We have limited ground truth data, but on that data BeCAUSe has a precision of 100% and a recall of 87%. More importantly, BeCAUSe reports not just a number, but also provides a degree of certainty in its estimates.

2 BACKGROUND

This section briefly introduces Route Flap Damping (RFD), which is our major measurement object, and binary network tomography. We explain how to apply binary network tomography inference of damping ASs from practical Internet measurement.

2.1 Route Flap Damping in BGP

A router configured to use RFD maintains a penalty value per prefix per BGP session that defines when a prefix should be suppressed or released. This value is additively increased with each announcement or withdrawal for that prefix, and decreases exponentially in between. When the penalty exceeds a threshold the prefix is suppressed until the penalty decays below a second threshold.

We illustrate the RFD mechanics and the interplay of the key configuration parameters in Figure 2. At t_0, the penalty is initialised to 0, and increases by a constant (1000) with each received announcement (green) or withdrawal (orange). Between each update, the penalty decreases based on the *half-life* parameter. When the penalty surpasses the *suppress-threshold* at t_1 the prefix is withdrawn. At t_2, the prefix stops oscillating and therefore the penalty reaches the *reuse-threshold* at t_3, leading to the release of the previously damped prefix.

RFD was introduced 25 years ago. Its patchy history is illustrated in Figure 3. Much has been learned about RFD's mechanics [5, 17, 24, 30], but little is known about its deployment.

RFD may not be uniformly deployed throughout an AS. A network operator can limit RFD to specific peers, *e.g.,* only customers. Reasons may be that some neighbors have proven to be particularly noisy while others provide the only transit to some part of the Internet and damping its routes may have major implications. RFD can also be configured differently depending on the prefix length. We encountered configurations where shorter prefixes were damped more aggressively in one network and less aggressively in a different AS.

We develop a method to perform controlled experiments (see § 4) to identify paths that contain RFD ASs. Having identified RFD for certain paths, only leaves us with the problem of tracing back the originating AS(s) on each path. The ideal measurement setup is a direct peering between measurement probe and vantage point, which is completely impractical given more than 70k ASs managed by almost the same number of different organisations across the globe. Luckily, pinpointing ASs that deploy RFD using path data can be formulated as a binary network tomography problem.

2.2 AS Inference Problems in General

RFD is one of a wider class of problems where we seek to localise particular routing policies or techniques in the inter-AS network.

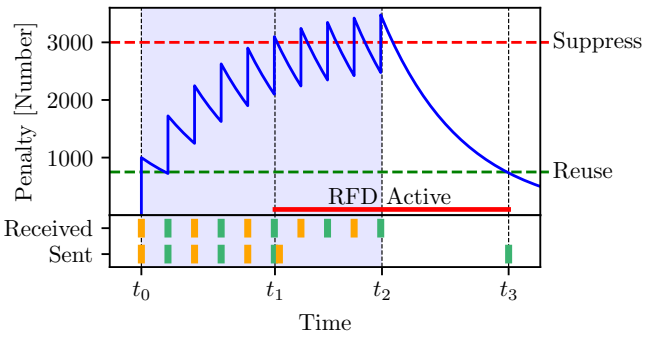

Figure 2: RFD router perspective: The penalty for a prefix that oscillates between announcement (green) and withdrawal (orange). The dashed, horizontal lines represent suppress and reuse-threshold. While RFD is active, the prefix is not advertised to neighboring routers, *i.e.,* a withdrawal is sent just after t_1, and only rescinded at t_3.

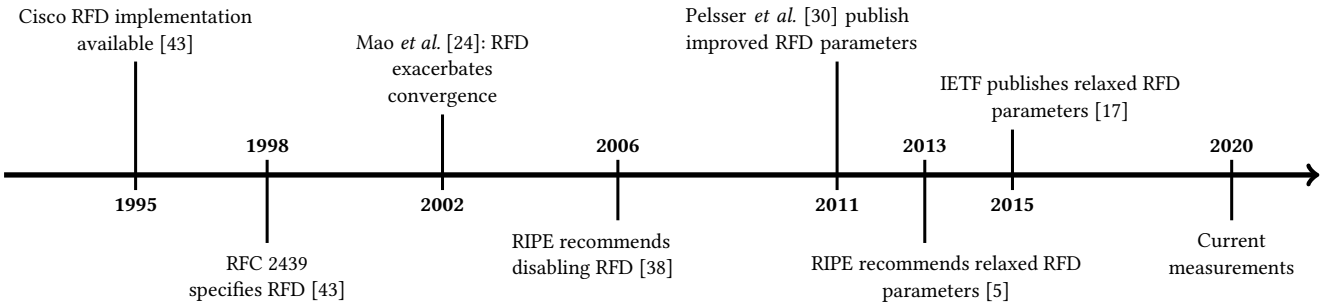

Figure 3: Timeline of Route Flap Damping.

Other problems in this space include finding ASs that use a particular community (some communities are transitive) or that black hole a particular type of traffic for instance traffic whose origin (signed through the RPKI) does not validate correctly. Apart from understanding RFD, we seek to be able to solve a range of such problems. To that end, we will also test the approach proposed here on RPKI origin validation.

2.3 Binary Network Tomography

In network measurement it is often impractical to interrogate network artefacts directly, either because of expensive overhead or (as in our case) because the artefacts have diverse owners who in many cases are competitors, and who have little interest in sharing such information. Network tomography can come to the rescue in these situations. However, our setting is unlike typical network tomography. In our problem the node properties we seek to find are not "problems" *per se*, so our approach is different though the differences may appear subtle at first.

The essential nature of network tomography is to reveal internal characteristics of a network from external observations. For example, we observe path properties, and wish to infer link or node properties.

Typical tomography is quantitative, *e.g.,* we observe volume of traffic, or size of delays, in which case the relationship between node and path properties are represented as linear equations, however, there is a strand of work on binary (or Boolean) tomography, which appears applicable here. Each node (each AS) is considered to either have property A or not. Paths have property A if at least one node on the path has property A, and if no nodes have the property, then the path will not.

We express this mathematically by defining variable x_i where,

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

Then, path j consisting of a set of nodes N_j ideally satisfies

$$y_j = \prod_{i \in N_j} x_i, \quad (2)$$

where

$$y_j = \begin{cases} 0 \text{ if path j has property A,} \\ 1 \text{ if path j } \textbf{does not} \text{ have property A.} \end{cases} \quad (3)$$

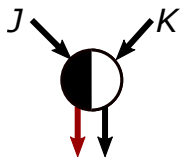

Figure 4: An AS may apply different policies at different ingress points or towards different peers resulting in measurement where one path (J) is subject to RFD and another (K) is not.

The result is a set of equations—one for each measurement[1]. Ideally, solving these equations would solve the localisation problem and reveal the damping ASs.

While simple in theory, most tomography problems have a number of challenges:

(1) There are rarely enough path measurements to obtain a unique solution and so some side information is needed to help refine solutions. Steering towards the sparsest solution is a common strategy when looking for network problems (which are hoped to be rare) but is not applicable here where one of the core goals is to identify the ASs that exhibit a property, not just localise a problem.
(2) The noise inherent in (any set of) measurements may mean the equations are inconsistent and hence have no solution. Robust approaches will look for approximate matches, but once again they often do so by introducing a model (*e.g.,* Gaussian noise) that is not appropriate in routing policy inference.

Moreover, an AS is not an atomic entity; it is a network in its own right. It is common for an AS to implement different policies at different ingress point or towards different peers [36] resulting in a situation such as shown in Figure 4 where measurement of paths J and K present contradictory results. However, this is not the same as in typical tomography problems where varying results are interpreted as a stochastic process (*e.g.,* a loss process) because the impact on paths J and K is (approximately) constant per path, just different between the paths.

[1]It is worth noting that path changes during a study may result in more than one measurement for each probe setting as in [10].

3 NETWORK TOMOGRAPHY WITH BeCAUSe

In this section, we present our algorithmic framework – BeCAUSe (BayEsian Computation for AUtonomous SystEms) – for inferring the cause of path observations. We reframe the binary tomography problem into this setting where pure binary choices are not possible, and use computational Bayesian inference. We apply this specifically to BGP in Section 5.1 using the data from the RFD measurement infrastructure. We demonstrate the applicability of this algorithm to more general tomography problems by inferring Route Origin Validation (ROV) in § 7.

3.1 Overview

We consider the binary property A of Autonomous Systems. Let each AS have some proportion p_i of routes to which it applies property A. For example, a network operator may apply RFD only to a particularly flappy set of customers, or have legacy configurations in old equipment. We also define and often use the complementary proportion $q_i = 1 - p_i$. We can think of these as probabilities, but they are subtly different from those used, for instance, in loss inference where it would indicate a probability of loss for an arbitrary packet.

Presuming that each AS contributes independently to the likelihood of RFD on a single path[2] the probability that a single path J *does not* have property A is

$$\mathbf{P}(J \text{ does not show A}) = \prod_{i \in J} q_i. \tag{4}$$

On the other hand, if **any** AS displays property A then the path will show A. As several ASs could potentially display A, the probability of any AS displaying the property is 1 minus the probability that no ASs display A. These two cases define the likelihood model for the observed dataset D given the probability vector $\mathbf{q}$ of the probabilities q_i for each AS in D.

First consider the probability of a single path J,

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

Then

$$\mathbf{P}(D|\mathbf{q}) = \prod_{J \in D} \mathbf{P}(J|\mathbf{q}), \tag{5}$$

describes the likelihood that we observe our data D given the set of values q_i for the ASs under consideration.

A Maximum Likelihood Estimator (MLE) would seek to find $\hat{\mathbf{q}}$ or $\hat{\mathbf{p}}$ that maximises (5). However, as we are in a probabilistic setting there is scope to instead infer the distribution of q_i or p_i. That is, gather many possible values of p_i, and the associated likelihood of producing the data we observe, to determine not only the most likely value, but also some measure of certainty about our results. The output distribution of all the p_i's together given the data D, denoted $\mathbf{P}(\mathbf{p}|D)$, is called the *posterior distribution*. The individual distributions $\mathbf{P}(p_i|D)$ are called marginal distributions and will provide information about which ASs are displaying our property of interest.

The basic probabilistic model introduced in (5) mirrors the binary tomography problem in § 2.3. The procedure can be generalised, though, to other likelihood models. It could, for example, incorporate specific types of errors in the measurements.

The likelihood in (5) is a variant of the Poisson binomial distribution and a closed form of $\mathbf{P}(\mathbf{p}|D)$ does not exist. Computational Bayesian methods are required to sample from the posterior distribution of interest.

3.2 Markov-Chain Monte Carlo

Markov-Chain Monte Carlo (MCMC) methods are a suite of computational Bayesian methods that use a stochastic simulation to approach complex inference problems. They are ideal when it is easy to simulate potential solutions and calculate their likelihood, but difficult to find the optimal solution, for example in the tomography problems described above. As highlighted above, instead of a single solution, we obtain many, in the form of a distribution, that can be used to calculate quantities of interest. In the case of the RFD problem we are estimating the RFD proportion of each node and how confident we are in these estimates.

In general, MCMC methods are used to infer the distribution of some set of parameters $\mathbf{p}$, given some dataset D—in tomography this will be the set of all measured paths (*i.e.,* infer $\mathbf{P}(\mathbf{p}|D)$). Using Bayes rule converts this into the form required for inference.

$$\mathbf{P}(\mathbf{p}|D) \propto \mathbf{P}(D|\mathbf{p}) \cdot \mathbf{P}(\mathbf{p}), \tag{6}$$

where $\mathbf{P}(D|\mathbf{p})$ is a likelihood model associated with the data as described in (5) and $\mathbf{P}(\mathbf{p})$ is the prior distribution, incorporating our knowledge of the parameters.

MCMC methods are designed to move around the space of possible solutions to the problem and take *samples* (*i.e.,* possible solutions) based on the associated likelihoods. The different methods to move around the space give rise to an active research area. Here we use two well known MCMC methods, Metropolis-Hastings [22, 25] and Hamiltonian Monte Carlo [4, 13]. These MCMC methods begin with some underlying knowledge of the parameters of interest, the prior distribution $\mathbf{P}(\mathbf{p})$. By using the information from the data, D, we update the likely values of each p_i according to the model likelihood, resulting in the posterior distribution $\mathbf{P}(\mathbf{p}|D)$. We can also use a uniform distribution for the prior if there is no underlying knowledge of the parameters.

Metropolis-Hastings (MH). The MH algorithm creates a Markov chain to explore the space of interest $\mathbf{P}(\mathbf{p}|D)$. At each step a new candidate for the probability vector $\mathbf{p}'$ is randomly proposed using the proposal distribution $\mathbf{Q}(\mathbf{p}'|\mathbf{p})$ (satisfying technical conditions [35]) conditioned on the current state $\mathbf{p}$. The proposal $\mathbf{p}'$ is accepted or rejected in a Metropolis update step with probability given by the acceptance probability

$$\alpha = \min\left(1, \frac{\mathbf{P}(\mathbf{p}'|D) \cdot \mathbf{Q}(\mathbf{p}|\mathbf{p}')}{\mathbf{P}(\mathbf{p}|D) \cdot \mathbf{Q}(\mathbf{p}'|\mathbf{p})}\right). \tag{7}$$

[2] We are not presuming that ASs are independent. We are well aware that, for instance, sibling ASs may have correlated policies. All we assume is that if a policy is *not ubiquitous* within their AS, then their decision about which routes to filter is not based on knowledge of the choices of their neighbours. This assumption is likely false in some places, but much less so than the alternative interpretation simply for the reason our technique is needed: in general, ASs share little of their internal business decision making with competitors. Moreover the assumption is not needed at all for ASs that have a uniform policy.

The acceptance probability α is calculated by substituting (6) and the choice of $\mathbf{Q}(\mathbf{p}'|\mathbf{p})$.

Hamiltonian Monte Carlo (HMC). HMC is related to Metropolis Hastings but uses Hamiltonian dynamics to explore the space by translating the density function of interest into a potential energy function and including a momentum variable [4]. The method uses a Markov Chain as in MH but new candidates are proposed by propagating the current state along a Hamiltonian trajectory using a Gaussian distributed momentum parameter. To obtain the samples of interest the auxiliary momentum parameters are ignored (marginalised over). This allows for multidimensional updates and allows the sampler to escape from local optima. HMC also uses a Metropolis update, and the acceptance probability uses the ratio of the auxiliary distribution of both the parameters of interest and the momentum.

In both MH and HMC the chain is generated from the proposed parameter $\mathbf{p}'$ as follows

$$\mathbf{p}^{(t+1)} = \begin{cases} \mathbf{p}', & \text{with probability } \alpha, \\ \mathbf{p}^{(t)}, & \text{otherwise,} \end{cases}$$

where $\mathbf{p}'$ is generated from either the MH or HMC proposal and α is the corresponding update probability. We invite the reader to pursue [4] for a thorough discussion of MCMC methods.

Prior Distributions. To finalise the algorithm we must also decide on a prior distribution $\mathbf{P}(\mathbf{p})$. This provides some flexibility of the method to incorporate our knowledge about the measurement system. In the RFD case, for example, we know that our Beacons do not dampen routes (see § 4). If there is no background knowledge a uniform (uninformative) distribution should be used.

The prior will dominate in cases where we do not have enough data about the parameter. We tested a variety of standard priors (*e.g.,* the uniform and β distributions) and found that there is sufficient data in the BGP setting for most ASs, so the choice of prior does not strongly influence the results. But a good choice of prior does make quantifying the uncertainty of inferences easier.

BeCAUSe generates samples from our distribution of interest using only a likelihood model, path measurements and the prior. There is no requirement of ground truth for 'training.' Techniques that require training are impractical here where there is little ground truth data.

The level of certainty about the inferences are implicit in the distributions and allows for informed decision making based on the desired application. The samples from $\mathbf{P}(\mathbf{p}|D)$ can be used in many different ways. In § 5, we highlight how they can be used to identify RFD-enabled ASs, and in § 7, we confirm its validity for RPKI Route Origin Validation (ROV).

4 RFD MEASUREMENT INFRASTRUCTURE

Controlled, active experiments have improved accuracy for network tomography problems [1, 19, 32, 37, 40] as they leverage well-defined input signals (*e.g.,* oscillating prefixes) to provoke observable events (*e.g.,* in BGP dumps). Using passive monitoring of uncontrolled BGP events from route announcement feeds to infer RFD deployment is not possible, because we need to know whether updates were sent or not, in order to know whether updates were damped.

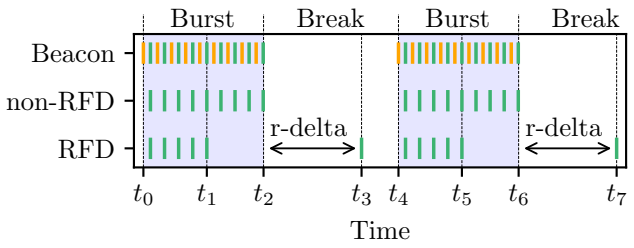

Figure 5: Beacon pattern and RFD signature for both RFD path and non-RFD path, and time until re-advertisement (r-delta).

We built an infrastructure that injects signals to stimulate route flap damping, which then becomes visible in common BGP data sets. We identify a clear signature for AS paths that include at least one autonomous system deploying RFD.

4.1 Generating Oscillating Prefixes

We use two-phase BGP Beacons, which oscillate at different frequencies on controlled schedules from geographically distributed peers. We measure and analyse the resulting signals in BGP dumps from common route collector projects while ignoring BGP update churn created by non-RFD causes, *e.g.,* MRAI. Adjusting the update intervals allows us to explore different RFD deployment configurations.

Two-phase BGP Beacons. These differ from previous BGP Beacons in order to cover the full RFD mechanics. Constant rate announcements and withdrawals would cause RFD ASs to constantly dampen Beacon prefixes, hiding further information on RFD parameters. Instead, we oscillate IP prefixes for specific periods of time, then allow RFD routers to reset the damping penalty, which will then cause re-advertisement of the Beacon prefixes. Therefore, the Beacons have two phases, the *Beacon pattern*:

Burst: a sequence of alternating announcements and withdrawals starting with a withdrawal and ending with an announcement.

Break: BGP announcements and withdrawals are paused.

When receiving Beacon announcements at vantage points, we need to identify which Beacon event caused this announcement. Similar to the RIPE Beacons [34], we encode the sending timestamps of our RFD Beacons in the transitive BGP *aggregator attribute*.

The RFD signature. Our Beacons are carefully designed to create a specific RFD penalty behavior resulting in a recognizable BGP update signature if RFD has occurred anywhere on the path between the Beacon router and the vantage point. Figure 5 shows the Beacon pattern and the observed signature; first the announcements are damped away ($t_1 - t_2$ and $t_5 - t_6$), thereafter a quite delayed re-announcement follows after the Reuse Threshold is reached (t_3 and t_7). The latter is released, because the last BGP update was an announcement during the Burst. In Section 4.2, we use this signature to decide whether an RFD-enabled AS exists on a path.

Preventing interference because of MRAI. Minimum Route Advertisement Interval (MRAI) [31] is another mechanism that

limits oscillating BGP updates. As MRAI induces a signal very different from RFD—delaying updates at most n seconds, where n is a configurable constant—it does not interfere with our recognition of the RFD signature.

4.2 Path Labeling

We search for our RFD signature in passively collected BGP update dumps of public route collector projects for each Beacon prefix and label each path individually. Paths are cleaned by removing AS path prepending and paths with loops were not present in our dataset. For our temporal analysis of the signals we need to consider that any BGP update will arrive at the vantage point only after a propagation delay of the BGP message. To identify the *re-advertisement* (*i.e.,* the delayed resending of the last announcement from the preceding the Burst phase), we argue that the time delta between the final update from the Burst and the re-advertisement during the Break, *r-delta*, must exceed the normal propagation time of the respective prefix at this vantage point.

To define the minimum propagation time for a re-advertisement (minimum r-delta), both the normal propagation delay for our Beacons and common MRAI configurations need to be considered. The propagation delay of our anchor prefixes is at most 1 minute (see Section 4.3). At the time of this writing, there are no studies measuring the values that are used to configure MRAI on the Internet, but there is at least one vendor defaulting MRAI to 30 seconds. Considering Cisco RFD default parameters, a prefix is suppressed for at least 21 minutes, for Juniper even longer. Given these distinct timescales, we find that setting the minimum propagation time for the re-advertisements to 5 minutes clearly separates the signals.

After analyzing all pairs of Burst-and-Break for each path, we arrive at a set of RFD paths and a set of non-RFD paths. To cope with unexpected infrastructure failures such as session resets, we label paths with RFD for which at least 90% of Burst-Break pairs match the above requirements.

4.3 Setup

Configuration. We deploy seven Beacon sets, in Europe, South and North America, Asia, and Africa and analyze all BGP dumps from RIPE RIS [33], RouteViews [42], and Isolario [23].[3] Beacons are a maximum of two AS hops away from a Tier 1 provider. We verified that our upstream networks do not use RFD and therefore do not influence our measurements.

At each of our seven Beacon sites we announced four different /24 IP prefixes (28 in total), one anchor prefix and three IP prefixes oscillating on different schedules. Beacon update intervals are configured identically across all locations. Anchor prefixes were announced and withdrawn every two hours, the same update interval as the RIPE Beacons, and are a control reference for propagation behavior. To prevent filtering of the prefixes in case of route origin validation deployment, we configured the corresponding route object entries in both the Internet routing registry and the RPKI for all prefixes.

We did not expect RFD configurations more strict than the vendor default values, which already suppress 14% of all prefixes [30]. A Juniper or Cisco router would start damping a prefix that flaps at least every 9 or 8 minutes respectively (see Appendix B for default values). To confirm this, we configured our Beacons with an update interval of 15, 30, and 60 minutes in August 2019. We observed measurable RFD for the fastest Beacon prefix (15 minutes).

[3]We will disclose details in the camera-ready version to allow for full reproducibility.

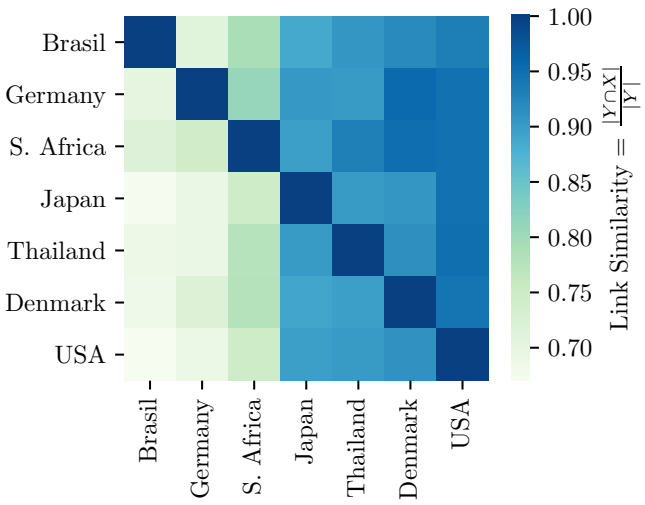

Figure 6: Similarity of links on AS paths compared between Beacon sites.

After preliminary tests, we conducted two measurement campaigns. In March 2020, we chose 1, 2, and 3 minutes as update intervals during Bursts of two hours, because an update interval of 2 minutes would trigger RFD with the recommended parameters [5, 17]. We set the Break duration to 6 hours to account for very slowly decaying RFD penalties. If a router is configured such that the penalty does not decay during the Break, then the updates from next Burst will increase the penalty again, causing the router to suppress the prefix indefinitely. In April 2020, we chose 5, 10, and 15 minutes as update intervals to cope with RFD parameters that differ more significantly from recommended values—either because vendors ship deprecated default configurations, or manual adjustment by operators. We configure the Break to 2 hours, because the *max-suppress-time* is by default 1 hour and we did not observe suppress phases longer than 1 hour in the Break in March. The Burst length was still 2 hours.

In the following, we process each prefix per site separately, because they flap with different update intervals and thus belong to independent experiments.

Validation. Validating the baseline characteristics of the injected BGP announcements is crucial before analyzing the collected data further. During this validation period, we statically announced all prefixes.

Our Beacon prefixes were visible at 99% or more of all vantage points that deliver a full feed (*i.e.,* $\geq$ 700k IP prefixes) to the route collectors (416 full feed peers in RIPE RIS and RouteViews). Surprisingly, 1% of the announcements (270 million in total) included an empty, invalid aggregator IP field. We could not find any specific reason, though, we noticed that more than half of these announcements were sent by AS 32097, a peer in the Isolario route collector

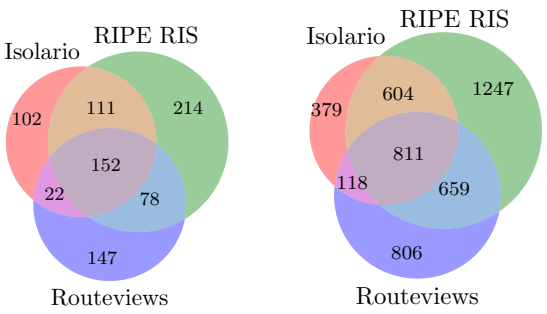

(a) # ASs on paths (b) # Ordered AS links on AS paths

Figure 7: Overlap of gathered data for RFD IP prefixes.

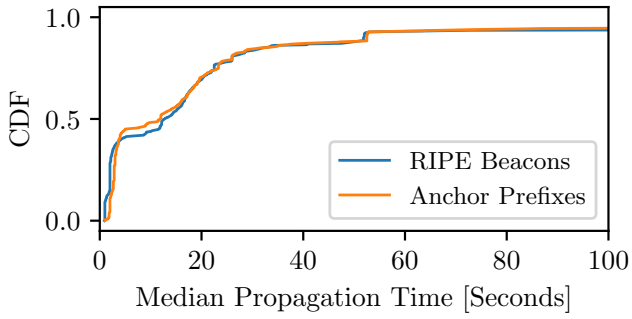

Figure 8: Comparison of propagation times between RIPE Beacons and RFD anchor prefixes across all vantage points.

project. This may be caused by misconfigured or malfunctioning routers. We discarded announcements with a missing or invalid aggregator IP, because our analysis would become less accurate without the encoded timestamps.

Our setup involves multiple sites with the aim to trigger RFD at multiple locations. Figure 6 depicts the relative amount of links (*i.e.,* adjacent ASs) each site shares with any other site to the vantage points. Between 70% and 95% of all AS links (4186 in total) in the public BGP feeds can be observed using a single of our Beacon sites. The median that a given link occurs on different paths, however, is 11 paths (not shown). This is a significant increase compared to using Beacons locations individually, which would lead to a median of 3 paths. Hence, observing AS links from multiple angles increases the confidence in our observations.

It is important to include diverse vantage points to enhance visibility of ASs and links. Figure 7 clearly shows that each route collection project contributes a substantial amount of additional data, which is the reason why we include all three data sources in our study.

These results underscore two advantages of our setup. First, we can observe the behavior of an AS from multiple Beacons and vantage points, which helps pinpointing the location of RFD and increases confidence in our observations, as will become evident in Section 6. Second, injecting updates from additional locations allows us to discover additional ASs and AS links.

To further validate our infrastructure we measured the propagation time, *i.e.,* the time it takes from sending the announcement from the Beacon routers until the first announcement of each router reaches the vantage points. We compare the results with the propagation behavior of RIPE BGP Beacons [34] in Figure 8. Both Beacon sets show the same characteristics. It is worth noting that the propagation delay depends on the collector project in use (not shown). Some vantage points in the RouteViews project export updates exactly 50 seconds after our Beacon routers sent the BGP updates. In contrast to this, vantage points in Isolario export updates for all but two Beacons within 30 seconds, whereas RIPE vantage points show a much more diverse behavior.

5 IDENTIFYING RFD-ENABLED ASs

We presented a method to trigger and determine RFD paths. Now, we pinpoint RFD deployment of specific ASs. Our Bayesian approach (§ 5.1) does not make any assumptions based on RFD mechanics, in contrast to heuristics (§ 5.2) that we use for comparison.

5.1 Bayesian Inferences

The output of BeCAUSe are many samples of $\mathbf{p}$: an N dimensional distribution across ASs. As we are interested in making decisions about each of the parameters individually, which tells us which AS may be showing RFD, we look at the marginal distributions of each p_i (*i.e.,* the distribution of each of the parameter separately). To identify RFD-enabled ASs using BeCAUSe we must establish distributions that are indicative of RFD.

5.1.1 Explanation of Algorithm Output. In contrast to many classification algorithms, the output of this method are diagnostic pictures (distributions) for each AS about its behavior. Here, we highlight the diagnostic ability of these distributions, and describe a basic summarisation and classification process to provide automatic insights.

Figure 9 depicts the marginal distributions of 4 ASs that are indicative of behaviors of interest.

(a) The distribution is heavily skewed with most mass at 1. There is very little spread suggesting there is strong evidence the AS is damping.
(b) The distribution is heavily skewed with most mass at 0. There is almost no spread suggesting there is strong evidence that the AS is not damping.
(c) Mass centred around mean 0.1 with comparatively higher spread suggests contradictory information about RFD, *i.e.,* some paths that damp and others that don't. In fact, AS 701 damps inconsistently.
(d) The distribution we see here is the β prior distribution. As it persisted, it is likely that we did not see any meaningful data about this AS. Interestingly, this AS is on damped paths; nevertheless, there is already another AS on these paths that is likely to damp, so we cannot extract any information about this AS.

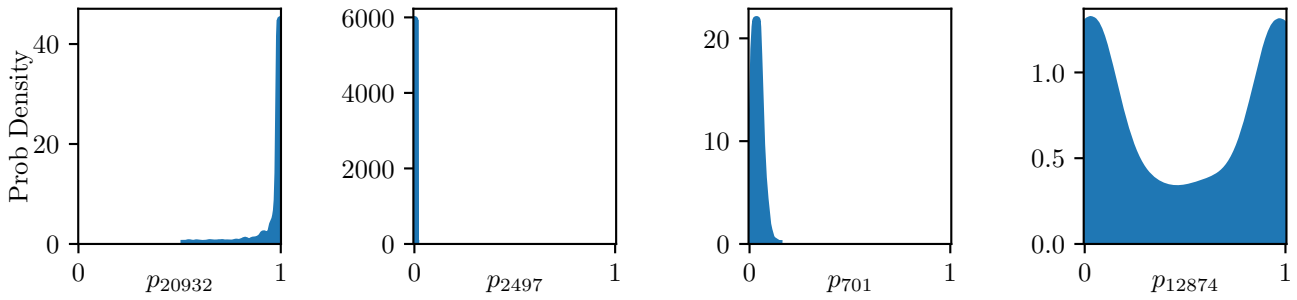

(a) AS 20932: the mass at 1 is strong evidence for RFD.

(b) AS 2497: the mass at 0 is strong evidence that there is no RFD.

(c) AS 701: the mean ~ 0.1 indicates data suggesting a mix of RFD on different paths.

(d) AS 12874: the distribution is similar to the β prior suggesting low evidence about this AS.

Figure 9: Example output probability distributions of RFD demonstrating their diagnostic ability to detect RFD and quantify the uncertainty in the inference.

5.1.2 Summarising the distributions. Inferring the distributions of each AS gives us flexibility in identifying ASs that are implementing RFD. There are several methods, ranging from simple to complex that we could use. The distributions could be summarised by taking the average (or some other point estimate) of the distribution as a metric. A threshold could then be used on these to determine RFD-enabled ASs. We can also preserve some of the information regarding the shape or spread of the distributions to use the implicit level of certainty in our decisions. We will focus here on one possible way of summarising the distributions with two metrics to measure the expected value and the certainty and use these to categorise the ASs and identify RFD.

Summary metrics. We generate 2 summaries of the distribution for AS from each method:

- The **mean** of the distribution; and
- The **Highest Posterior Density Interval** (HDPI)

The first is just the average $\bar{p}_i$ of the distribution, $P(p_i|D)$, and gives an estimate of the expected value for this AS. HDPI finds the smallest interval that contains $\gamma = 0.95$ of the mass. Otherwise know as the smallest Bayesian credible interval, it is the interval $[A, B]$ where γ of the mass falls between A and B such that $B - A$ is minimised. The width of the HDPI measures the (asymmetric) spread of the distribution and gives an idea of the uncertainty in our mean estimate.

1) Categorising. The objective of the algorithm is to allow a user to gain a specific level of certainty depending on the application. The metrics provide us with valuable information, but for the purpose of this work we must translate these metrics into a 'decision'. We maintain some of the information about certainty by mapping results to a category from 1 to 5, where 1 and 2 are highly likely and likely not damping and 4 and 5 are likely and highly likely damping. Category 3 is uncertain, either because of contradictory data, or, most often, due to lack of specific data about this AS. Note that not enough data does not necessarily mean the AS is not on many paths. Nodes that are regularly on paths with other damping ASs do not display any specific information. Categorisation based on the summaries are given in Table 1.

Table 1: Categories based on distribution summaries. 'Else' indicates the flag if no other category is assigned. The highest category is chosen for each AS.

	Average: $\bar{p}_i$	HDPI: $[A_i, B_i]$
Category 1	$\bar{p}_i \in [0, 0.15)$	$A_i \in [0, 0.15)$
Category 2	$\bar{p}_i \in [0.15, 0.3)$	$A_i \in [0.15, 0.3)$
Category 3	$\bar{p}_i \in [0.3, 0.7)$	else
Category 4	$\bar{p}_i \in [0.7, 0.85)$	$B_i \in [0.7, 0.85)$
Category 5	$\bar{p}_i \in [0.85, 1]$	$B_i \in [0.85, 1]$

The cut-off values are chosen to automatically implement the insights from the output distributions in Figure 9, and we provide data driven justification in § 6.1. After summarising and categorising both the MH and HMC distributions by mean and HDPI, we use the highest flag.

2) Identifying ASs that use RFD inconsistently. It is evident from the data and marginal posterior distributions that some AS use RFD inconsistently. For example, AS 701 damps all neighbours except AS 2497. The distribution of p_{701} in 9(c) highlighted contradictory data. After categorising based on the thresholds above, we utilise the marginal distributions to determine ASs that damp inconsistently.

Recall, that if the path displays RFD then there is at least one AS on the path that damps. Therefore, there should be at least one path on each RFD path that is labelled in Category 4 or 5. If the path does not contain an inferred RFD AS, we use the posterior marginal distributions to determine the AS that is most likely causing RFD. Specifically, for each AS X on the path we determine the posterior probability that AS X is the most likely to be causing RFD on the path. If

$$\mathbf{P}(\min(p_i \text{ for } i \in J) = X) > 0.8, \tag{8}$$

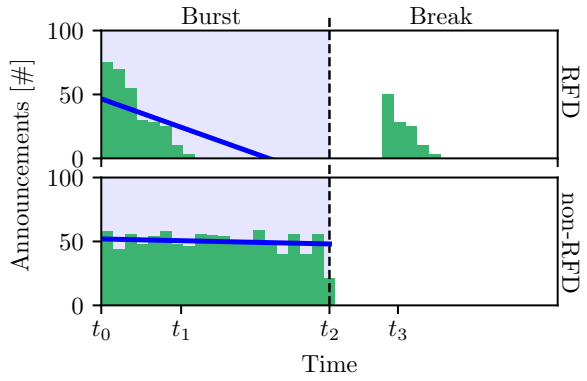

Figure 10: Typical distribution of announcements during a Burst-Break pair for an RFD AS (top) and a non-RFD AS (bottom), with the linear regression function (blue) of histogram heights.

then there is sufficient evidence that X is the damping AS on path J. For each damping path J in the data, if there is an AS that is most likely damping we label this AS as a Category 4.

In general, the summarisation, cut-off, and flagging methods can be tailored depending on the desired confidence level of the outcomes. We accept Category 4 and 5 to be an RFD-enabled AS; however, if higher certainty is required, we could use only Category 5 ASs or change the thresholds appropriately.

5.2 Passive Measurement Heuristics

An alternative to Bayesian Inference for identifying ASs that deploy RFD is to rely on heuristics. We now present 3 metrics toward that aim. For each AS we take the average of the metrics as the final output. These will be used for comparison purposes to highlight the power and simplicity of BeCAUSe. We will see that heuristics are less precise and need tuning that is absent from the Bayesian approach. Additionally, the heuristics would become very inaccurate if RFD was deployed in the majority of networks.

5.2.1 RFD Path Ratio. This heuristic quantifies the relative occurrences of an AS on a path showing the RFD signal compared to the total number of paths this AS appears on. We calculate for each AS:

$$M_1(\text{AS}) = \frac{\#\text{RFD paths (AS)}}{\#\text{RFD paths}(AS) + \#\text{non-RFD paths}(AS)}$$

This metric is robust for richly connected ASs, *i.e.,* Tier 1 provider and transit networks. Stub ASs tend to be biased towards the RFD configuration of their upstream provider(s). False positives will occur for ASs, which only have one upstream with RFD enabled.

5.2.2 Inferring RFD ASs Based on Alternative Paths. This metric is motivated by two observations. (*i*) damped prefixes will reveal alternative paths between a Beacon prefix and a vantage point because of path hunting. (*ii*) An AS that actively damps prefixes will not be part of an alternative path.

For each damped path, we determine a set of alternative paths between the Beacon and the vantage point. We expect that alternative paths are used more frequently after the original path has been damped. Then, for each AS we determine the average share of alternative paths without the AS across all damped paths.

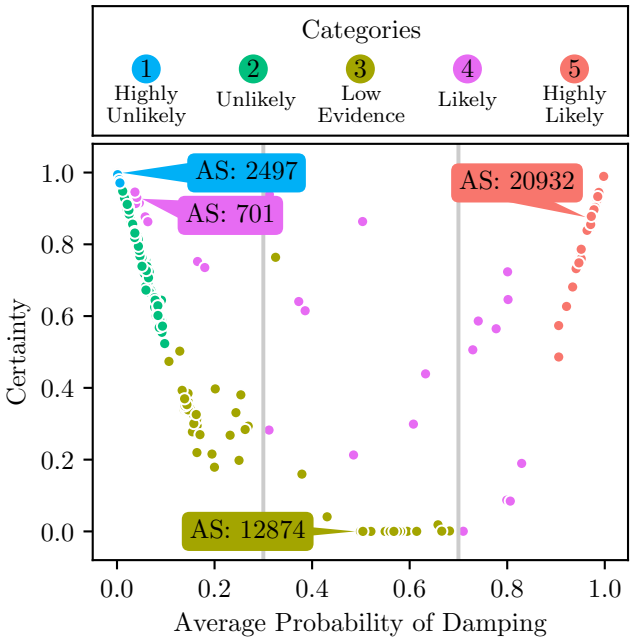

Figure 11: Scatter plot of the mean of the marginal posterior distribution for each AS (x-axis) with cut-offs (Table 1) in grey, and a measure of certainty about the estimate (y-axis) for the 1 min update interval. ASs are colored based on the assigned category.

Table 2: Total and share of assigned categories for the 1 minute update interval.

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

5.2.3 Announcement Distribution across Bursts. This metric is motivated by the observation that a damping AS sends fewer BGP updates near the end of a Burst compared to non-damping ASs. Figure 10 visualizes an average case of an RFD AS compared to a non-RFD AS. The blue and white areas, separated by a vertical red line, indicate the Burst and Break phases respectively. Both plots show a histogram of received announcements grouped in 40 time intervals. The blue line displays the linear regression function of the histogram heights during the Burst. Based on the slope and relative change of this linear regression function, we map this behavior to a score between 0.0 and 1.0.

6 RESULTS

6.1 Pinpointing RFD ASs with BeCAUSe

Figure 11 highlights the output of BeCAUSe and detailed category shares are shown in **??**. The scatterplot depicts two summary metrics of the output distributions to plot the average probability of damping against the certainty of this estimate. We use one minus the length of the HDPI to quantify certainty, so more confident estimates are closer to 1. The x-axis is a measure of how likely an AS uses RFD, and the y-axis is a measure of the spread—how sure we are of our x-axis estimate. The ASs are colored by category determined by the process introduced in Section 5.1. There is a characteristic U shape. On the right we have ASs that are likely using RFD as shown by the high average $\bar{p}_i$. ASs in the top right are on a large number of RFD paths and so have a high certainty, *e.g.,* AS 20932 in Figure 9(a). For ASs where there is less data (but not contradictory), we see our confidence decrease but the average $\bar{p}_i$ still suggests RFD. Conversely, on the top left we have ASs that are on many non-RFD paths, *e.g.,* AS 2497 in Figure 9(b). For ASs that are on less paths we have less information, but the average remains high as we see in the Category 1 and 2 ASs (blue and green). The base with low evidence suggests ASs for which we have little information, and we recover the prior with high spread, *e.g.,* AS 12874 in Figure 9(d).

The ASs in Category 4 that are spread across the plot are quite interesting. Such cases suggest we have contradictory information about these ASs—probably due to inconsistent damping. For example, recall the distribution of AS 701 in Figure 9(c) that has a low probability of damping, because on the majority of labeled paths it receives updates from AS 2497, which is a neighbor that is not being damped. Despite the low mean probability, our pinpointing method has identified this AS, and others, as RFD as they are the most likely ASs to be causing RFD on some damped paths as described in § 5.1. Figure 11 shows grey vertical lines at $\bar{p} = 0.3$ and $\bar{p} = 0.7$. These are the category cut-offs from § 5.1, chosen to segment the region into the three distinct different regions over the different update intervals.

Our results suggest that 9.1% (sum of Category 4 and 5) is the lower bound of RFD deployment. There a three reasons why we may have labeled a damping AS as non-damping. First, an AS damping solely customers is not detectable with our setup because our Beacons are located in or close to Tier-1 providers and thus the Beacon signals travel only from a provider to a customer or between peers in the Internet topology. This is confirmed by the observation that less than 3% of links on the measured paths are customer links. Second, an update interval of 1 minute may not be small enough to trigger some configurations. Third, a damping ASs may be hiding behind another damping AS, so our updates are already being suppressed before they can reach this AS. This last issue would be much more significant if RFD deployment was larger. With the above challenges and visibility issues, it is, with our measurement setup, impossible to establish an upper bound for RFD deployment.

6.2 Deployed RFD Parameters

RFD is configurable in various ways. To target different configurations, we used 6 different update intervals. Figure 12 visualizes

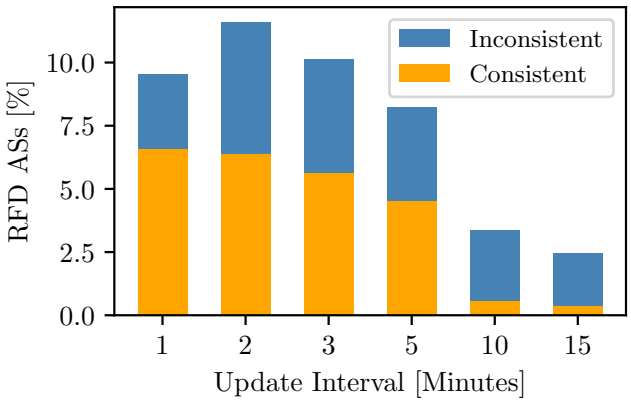

Figure 12: Share of damping ASs from total (534) for each update interval. Only ASs measured in all 6 experiments are counted.

the share of ASs using RFD for a given update interval. The orange bars indicate the share of RFD-enabled ASs for which we have non-contradicting data, *i.e.,* damping all neighbors consistently. These ASs have been labeled using only the probability of damping (step (1) in § 5.1). The blue bars include inconsistently damping ASs that were labeled with step (2). We observe an unexpected spike at 2 minutes as a single AS with a large customer cone damps inconsistently.

While Figure 12 illustrates how quickly a prefix needs to flap to get damped, we cannot infer the exact value of the *suppress-threshold* in use because one Beacon event may cause multiple updates distant in topology (*e.g.,* path hunting). We assume, however, that many operators use predefined configurations, and try to find confirmation in our data. Currently, there are two sources of parameter sets: (*i*) the recommendations by the IETF and RIPE [5, 17], and (*ii*) vendors that ignore these recommendations and pre-configure a deprecated *suppress-threshold* (see Appendix B). The largest fraction of ASs stop dampening for update intervals larger than 5 minutes (Figure 12). A router with deprecated default values would start damping at the 5 minutes update interval. We suspect the continuous increase of RFD ASs for the smaller update intervals is caused by some network operators following the current recommendations. The very few damping ASs at 10 or 15 minutes are likely induced by updates amplified by topology properties. Based on feedback from almost 50 network operators we were able to confirm that there is a significant tendency (≈60%) to use vendor default values.

To expose the announcement pattern of damped paths we analyzed the *max-suppress-time* values that are used in practise. Figure 13 visualizes the distribution of the mean time delta between the end of the Burst and the re-advertisement across all damped paths for the entire measurement period. First, we notice that the time until re-advertisement (r-delta) rarely surpasses 60 minutes, suggesting a large *max-suppress-time* is uncommon. For the smallest update interval (1 minute) we find three plateaus, starting at 10, 30, and 60 minutes, indicating that these are the most commonly configured values for the *max-suppress-time* parameter. The same

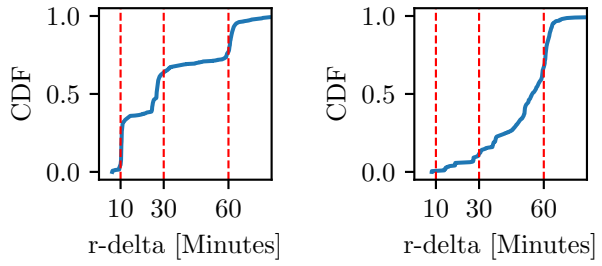

(a) 1 Minute update interval (b) 3 Minutes update interval

Figure 13: CDF of *re-advertisement delta* in the Break for each damped path. Only the 1 minute update interval is small enough to unveil the *max-suppress-times* (red lines).

pattern cannot be observed for the slightly larger 3 minute update interval and even less for 5, 10, and 15 minutes (not shown). For these larger update intervals, the penalty decreases naturally faster below the reuse-threshold and before the *max-suppress-time* expires.

6.3 Comparison to Operator Ground Truth

Pinpointing RFD ASs on paths that include only a vantage point and our Beacon AS is not challenging. Limiting ourselves to these scenarios (≈ 3%) would, however, significantly reduce the amount of ASs we can draw a conclusion about. BeCAUSe allows us to include paths of arbitrary length giving much wider coverage. To validate our findings we directly contacted network operators of every measured AS and received 75 replies in total. Unfortunately, we cannot map the provided configurations to exact update intervals from our experiment (1, 2, 3, 5, 10, 15 min) because effects such as path hunting increase the number of updates along the path, and make our results less sharp. Therefore, we chose to compare our results for the smallest update interval (1 minute) to ground truth. Overall, BeCAUSe performs very well on this small dataset with 100% precision compared to the heuristics, which have one false positive. The recall is 87% (80% for the heuristics) mainly due to visibility issues. We summarize the main reasons of divergence in Table 2. We removed two ASs, namely AS 8218 and AS 7575, from the ground truth dataset, as they are not detectable with our current measurement setup, hence it is not possible for the pinpointing methods to locate these ASs.

The MCMC methods perform better than the heuristics regarding precision, although the heuristics already yield 97% precession (see Table 3). In some challenging cases, the heuristics will incorrectly label ASs with RFD True, when they are on many RFD paths but are not causing them. This is the case for AS 5645, for example. BeCAUSe considers the entire path data in the likelihood and so accounts for this but only flags the upstream AS and identifies that there is no information about these downstream ASs. However, in these cases when evidence from the labelled paths is lacking BeCAUSe labels the AS as *unsure* (*e.g.,* AS 37474), while two of the metrics for the heuristics use additional data from the raw update dumps to identify these nodes. The MCMC algorithm flags inconsistently damping ASs, which cannot be found by the heuristics, by identifying ASs that are most likely to be causing the damping signal (recall § 5.1).

As network tomography problems assume nodes act consistently, the most challenging scenario is the deployment of heterogeneous RFD configurations, *e.g.,* an AS damps only customers. We could instead pinpoint individual AS links, but, unfortunately, when considering links, our data is too sparse to gain reasonable results from BeCAUSe or the heuristics.

Although the overall results of BeCAUSE appear to be on par with our heuristics, it is important to note that the heuristics are tailored to a specific pinpointing use case. In contrast to this, BeCAUSe is generic, which we will show in the next section.

7 APPLYING BeCAUSe BEYOND RFD

The BeCAUSe algorithm uses a simple likelihood model and does not require domain knowledge, *e.g.,* RFD. In this section, we present the effective application of the same algorithm to locate a different AS property, route origin filtering, and briefly discuss general usage.

7.1 Pinpointing ROV ASs with BeCAUSe

We are interested in locating ASs that drop invalid routes (*i.e.,* those IP prefixes that are announced from incorrect origin ASs) using RPKI route origin validation (ROV) [26]. In this section we first simulate the output of an ROV measurement using real-world AS paths and a set of ASs known to use ROV. Then, we benchmark BeCAUSe using this dataset. Therefore, this experiment does not uncover new ROV ASs, but simply benchmarks BeCAUSe in a different plausible use case.

Table 3: Overview of reasons of divergence between methods to pinpoint that an AS deploys (✓) or does not deploy (✘) RFD, compared to operator feedback.

			Pinpointing Method		
# Cases	Example AS	Ground Truth	BeCAUSe	Heuristics	Reason for Divergence
56	IIJ (AS 2497)	✘	✘	✘	-
10	Atom86 (AS 8455)	✓	✓	✓	-
3	Verizon (AS 701)	✓	✓	✘	Heterogeneous configuration
2	JINX (AS 37474)	✓	✘	✓	Upstream uses RFD
2	TekSavvy (AS 5645)	✘	✘	✓	Upstream uses RFD

Table 4: Summary of algorithm performance on ground truth. Overall, BeCAUSe and the heuristics perform well on locating RFD ground truth. BeCAUSe generalises to locate Route Origin Validation (§ 7).

	BeCAUSe		Heuristics	
	Precision	Recall	Precision	Recall
RFD	100%	87%	97%	80%
ROV	100%	64%	n/a	n/a

To label AS paths as *ROV* (or *non-ROV*) we make use of existing data sources that accurately pinpoint ROV ASs, either based on strictly controlled experiments [32] or ground truth [11]. From Isolario, RouteViews, and RIPE RIS we collect all AS paths of two RPKI Beacon prefixes (`147.28.241.0/24` and `147.28.249.0/24` [32]), and label the AS paths ROV (or non-ROV) if one of the ROV ASs is on path (or non-ROV otherwise).

While it is possible to change parts of our pinpointing algorithm, we use the same implementation to locate ROV as in locating RFD ASs. There are two key differences in contrast to the RFD dataset: (*i*) 90% of paths are labeled ROV (versus 18% for RFD) and (*ii*) noise is absent.

BeCAUSe has good performance on this dataset and leads to 100% precision and 64% recall (see Table 3). The ASs that were missed are only seen on paths with another ROV filtering AS. For these ASs it is impossible to infer ROV usage, because they are 'hiding' behind another ROV AS. This is a common issue in network tomography, unrelated to BeCAUSe, where two nodes only ever appearing together on ROV paths are unable to untangle which (or both) are displaying our property of interest.

7.2 Towards Other Scenarios

There are many options to extend BeCAUSe, within binary tomography or even more generally. The algorithm itself remains the same; however, we can use different models and summarisation techniques should our application or research question change. One useful extension is to include error explicitly in the likelihood model $P(D|\mathbf{p})$. For example, using our measurement method for RFD, it is possible that paths containing an RFD AS do not get recorded as RFD paths. We can model this error in the likelihood. Using a new likelihood model can enable the application of this method to network tomography problems beyond the binary problem introduced here.

8 RELATED WORK

Network Tomography: A good survey of the early work on network tomography is provided in [8]. Early approaches aimed at inferring the origins of performance events using highly correlated multicast [6, 15] or striped unicast packets [12, 16]. The ideas were extended to summary statistics using alternative measurements, *e.g.,* passive measurement [29], however early work concentrated on *additive* metrics, where the relationships between internal network properties could be expressed as linear relationships leading to deterministic algorithms even where the underlying model was stochastic.

Very little work in this large literature considers MCMC approaches. The only approach that has been trialled is Gibb's sampling, *e.g.,* [14, 29] a special case of Metropolis Hastings. Both [14, 29] applied the additional condition of sparsity, *i.e.,* that the network measurements should be explained by larger loss at a few places, rather than small losses spread across the network. Both [14, 29] only considered trees in detail (on trees the sparsity has the nominally beneficial impact of forcing loss higher in the tree, aiding in the diagnosis problem). In our setting, sparsity is not always true (see AS origin validation) and we measure across general graphs not just trees. It is worth noting that piecing together trees would be a non-trivial component of the prior work, so a comparison would not actually be against the prior approaches, and thus is out of scope in this work.

Binary (or Boolean) Network Tomography aims to classify links as "good" or "bad," a simpler and more practical inference in some cases [2, 3, 14, 28, 29]. More recently binary tomography was applied to general networks to find censoring ASs by formulating the problem using logical constraints as SAT [10]. This approach has advantages: it can use the large body of work on solving SAT, but it should be noted that the SAT problem so created can (*i*) have many solutions, requiring some means to choose one, or (*ii*) can have zero solutions in the case of measurement noise. Our approach aims for a middle ground between earlier probabilistic approaches with the large body of assumptions and limitations they brought, and the more practical binary tomography. We aim to ascertain if an AS has a particular property (though unlike typical binary tomography having a property is not intrinsically bad or rare). However, concomitant with that inference, we desire to allow for the possibility of partial properties, and to provide a measure of uncertainty of the inference. Binary tomography, by itself does not typically allow such.

We did not compare to binary approaches as they cannot derive meaningful results in scenarios of inconsistent deployment. SAT would lead to zero valid solutions, based on our data. AS 701 is one concrete example in the current Internet, which damps routes on some paths and does not on others (see § 5.1).

Despite some limitations practical network tomography instantiations have been built [20, 39] showing that methods such as that proposed here can be scaled to create useful systems for networks as a whole.

Heuristic Approaches: There have also been multiple studies which attempted to pinpoint another behavior—the origin a routing changes—using heuristics [7, 18]. These are tenuously related to this work in the nature of the task being attempted. Closer is a prototype [9], that tries to find the reason for a missing route. However, their approach does not investigate the exact reasons for the unreachability (they state that misconfigured RFD was one of the possible explanations). Closer still is [27], which considers the general problem of inferring AS properties. Other papers also aimed at this, but most have presumed an atomic model of ASs, *i.e.,* that each AS has a deterministic behavior. A major difference between our work and these is that we perform controlled experiments using Beacons, rather than relying on natural churn.

Locating ASs that deploy route filtering based on RPKI attracted recently attention. Gilad *et al.* [21] locate route divergences between invalid and valid prefix announcements and infer RPKI deployment based on that. Testart *et al.* [41] presented another heuristic assuming that vantage points which provide a full feed export less invalid announcements if they deploy filtering. Both approaches implement uncontrolled experiments and thus are prone to false positives, *e.g.,* because of traffic engineering and incorrectly configured RPKI data. Reuter *et al.* [32] introduced controlled, active experiments to precisely pinpoint ASs that deploy RPKI filtering. This, however, requires a strict measurement setup, *i.e.,* direct peering between the experiment AS and vantage point.

9 CONCLUSION AND OUTLOOK

We presented BeCAUSe, an algorithmic framework to infer network properties. In contrast to heuristic methods, which are commonly applied to tackle this challenge, our proposal does neither make restrictions on the topology setup nor does it require specific participation of the network elements under investigation, except their usual packet forwarding. The underlying features are two Bayesian computation techniques, which allow for rigorous network tomography.

We demonstrated BeCAUSe to pinpoint autonomous systems that deploy route flap damping and route origin validation. To the best of our knowledge, route flap damping was not measured before on global scale. We uncovered that at least 9% measured ASs use RFD, of which $\approx$ 60% rely on deprecated, harmful vendor default configurations. We compared BeCAUSE with different heuristics, which exhibit less precission and recall. Most importantly, in contrast to heuristics, BeCAUSe is not designed for a specific pinpointing use case but a generic framework. BeCAUSe can help researchers also measuring deployment of RPKI-based filtering as we showed. Tracking down censorship or other similar topics might be application scenarios in the future.

We believe that a better understanding of operational practices can improve the Internet in the long term. We presented our results to the operators at RIPE 80 to raise awareness about the surprising usage of deprecated RFD parameters. Analyzing changes will be part of our future work.

With the data set that we provide, verified by ground truth, we hope to stimulate further discussions on the implications of suppressed routes on Internet measurements.

A Note on Reproducibility. We explicitly support reproducible research. All artifacts are available on https://rfd.rg.net.

ACKNOWLEDGMENTS

We would like to thank the anonymous reviewers and our shepherd Dave Choffnes for their valuable feedback. We also gratefully acknowledge the operators and RIPE that provide us infrastructure and Internet resources to conduct our experiments.

Caitlin Gray and Matthew Roughan were partially supported by the ARC Centre of Excellence for Mathematical and Statistical Frontiers (grant CE140100049). Caitlin Gray acknowledges the financial support of Data to Decisions CRC and CSIRO's Data61. Clemens Mosig, Thomas Schmidt, and Matthias Wählisch were partially supported by the German Ministry of Education and Research (grant X-Check).

REFERENCES

[1] Ruwaifa Anwar, Haseeb Niaz, David Choffnes, Ítalo Cunha, Phillipa Gill, and Ethan Katz-Bassett. 2015. Investigating Interdomain Routing Policies in the Wild. In *Proc. of ACM IMC*. ACM, New York, NY, USA, 71–77.
[2] P. Barford, N. Duffield, A. Ron, and J. Sommers. 2009. Network Performance Anomaly Detection and Localization. In *Prof. of IEEE INFOCOM*. IEEE Press, Piscataway, NJ, USA, 1377–1385.
[3] A. Batsakis, T. Malik, and A. Terzis. 2005. Practical Passive Lossy Link Inference. In *Proc. of PAM (LNCS)*, Vol. 3431. Springer–Verlag, Berlin, Heidelberg, 362–367.
[4] Steve Brooks, Andrew Gelman, Galin Jones, and Xiao-Li Meng (Eds.). 2011. *Handbook of Markov Chain Monte Carlo*. CRC Press, Boca Raton, FL, USA.
[5] Randy Bush, Cristel Pelsser, Mirjam Kuhne, Olaf Maennel, Pradosh Mohapatra, Keyur Patel, and Rob Evans. 2013. *RIPE Routing Working Group Recommendations on Route Flap Damping*. RIPE Document ripe-580. RIPE.
[6] R. Cáceres, N.G. Duffield, J. Horowitz, and D. Towsley. 1999. Multicast-based inference of network-internal loss characteristics. *IEEE Trans. in Information Theory* 45, 7 (1999), 2462–2480.
[7] Matthew Caesar, Lakshminarayanan Subramanian, and Randy H. Katz. 2003. *Towards Localizing Root Causes of BGP Dynamics*. Technical Report UCB/CSD-03-1292. EECS Department, University of California, Berkeley. http://www2.eecs.berkeley.edu/Pubs/TechRpts/2003/6364.html
[8] Rui Castro, Mark Coates, Gang Liang, Robert Nowak, and Bin Yu. 2004. Network Tomography: Recent Developments. *Statist. Sci.* 19, 3 (2004), 499–517.
[9] Di-Fa Chang, Ramesh Govindan, and John Heidemann. 2004. Locating BGP Missing Routes Using Multiple Perspectives. In *Proc. of the ACM SIGCOMM Workshop on Network Troubleshooting: Research, Theory and Operations Practice Meet Malfunctioning Reality (NetT)*. ACM, New York, NY, USA, 301–306.
[10] Shinyoung Cho, Rishab Nithyanand, Abbas Razaghpanah, and Phillipa Gill. 2017. A Churn for the Better: Localizing Censorship using Network-level Path Churn and Network Tomography. In *Proc. of ACM CoNext*. ACM, New York, NY, USA, 81–87.
[11] Cloudflare. 2020. Is BGP safe yet? https://isbgpsafeyet.com/.
[12] M. Coates and R. Nowak. 2000. Network loss inference using unicast end-to-end measurements. In *Proc. of ITC Specialist Seminar on IP Traffic Measurement, Modeling and Managemen*. Monterey, CA, 28–1–28–9. Preprint https://hdl.handle.net/1911/19810.
[13] Simon Duane, A.D. Kennedy, Brian J. Pendleton, and Duncan Roweth. 1987. Hybrid Monte Carlo. *Physics Letters B* 195, 2 (1987), 216 – 222.
[14] N. Duffield. 2006. Network Tomography of Binary Network Performance Characteristics. *IEEE Transactions on Information Theory* 52, 12 (2006), 5373–5388.
[15] N.G. Duffield, J. Horowitz, F. Lo Presti, and D. Towsley. 2002. Multicast topology inference from measured end-to-end loss. *IEEE Transactions in Information Theory* 48, 1 (2002), 26–45.
[16] N.G. Duffield, F. Lo Presti, V. Paxson, and D. Towsley. 2001. Inferring link loss using striped unicast probes. In *Proc. of IEEE Infocom*. IEEE Press, Piscataway, NJ, USA, 22–26.
[17] J. Durand, I. Pepelnjak, and G. Doering. 2015. *BGP Operations and Security*. RFC 7454. IETF.
[18] Anja Feldmann, Olaf Maennel, Z. Morley Mao, Arthur Berger, and Bruce Maggs. 2004. Locating Internet Routing Instabilities. In *Proc. of ACM SIGCOMM*. ACM, New York, NY, USA, 205–218.
[19] Romain Fontugne, Esteban Bautista, Colin Petrie, Yutaro Nomura, Patrice Abry, Paulo Gonçalves, Kensuke Fukuda, and Emile Aben. 2019. BGP Zombies: An Analysis of Beacons Stuck Routes. In *Proc. of PAM Conf. (LNCS)*, Vol. 11419. Springer, Berlin Heidelberg, 197–209.
[20] D. Ghita, H. Nguyen, M. Kurant, K. Argyraki, and P. Thiran. 2010. Netscope: Practical Network Loss Tomography. In *Proc. of IEEE INFOCOM*. IEEE Press, Piscataway, NJ, USA, 1–9.
[21] Yossi Gilad, Avichai Cohen, Amir Herzberg, Michael Schapira, and Haya Shulman. 2017. Are We There Yet? On RPKI's Deployment and Security. In *Proc. of NDSS*. ISOC, Reston, USA, 15.
[22] W. K. Hastings. 1970. Monte Carlo Sampling Methods Using Markov Chains and Their Applications. *Biometrika* 57, 1 (1970), 97–109.
[23] IIT-CNR. 2020. Isolario Project. https://www.isolario.it/.
[24] Zhuoqing Morley Mao, Ramesh Govindan, George Varghese, and Randy H. Katz. 2002. Route Flap Damping Exacerbates Internet Routing Convergence. In *Proc. of ACM SIGCOMM*. ACM, New York, NY, USA, 221–233.
[25] Nicholas Metropolis and S. Ulam. 1949. The Monte Carlo Method. *J. Amer. Statist. Assoc.* 44, 247 (1949), 335–341. http://www.jstor.org/stable/2280232
[26] P. Mohapatra, J. Scudder, D. Ward, R. Bush, and R. Austein. 2013. *BGP Prefix Origin Validation*. RFC 6811. IETF.
[27] W. Mühlbauer, A. Feldmann, O. Maennel, M. Roughan, and S. Uhlig. 2006. Building an AS-topology model that captures route diversity. In *Proc. of ACM SIGCOMM*. ACM, New York, NY, USA, 195–206.

[28] H.X. Nguyen and P. Thiran. 2007. The Boolean Solution to the Congested IP Link Location Problem: Theory and Practice. In *Proc. of IEEE INFOCOM*. IEEE Press, Piscataway, NJ, USA, 2117–2125.
[29] Venkata N. Padmanabhan, Lili Qiu, and Helen J. Wang. 2002. Passive network tomography using Bayesian inference. In *Proc. of ACM Internet Measurement Workshop*. ACM, New York, NY, USA, 93–94.
[30] Cristel Pelsser, Olaf Maennel, Pradosh Mohapatra, Randy Bush, and Keyur Patel. 2011. Route Flap Damping Made Usable. In *Proc. of PAM Conf. (LNCS)*, Vol. 6579. Springer, Berlin Heidelberg, 143–152.
[31] Y. Rekhter, T. Li, and S. Hares. 2006. *A Border Gateway Protocol 4 (BGP-4)*. RFC 4271. IETF.
[32] Andreas Reuter, Randy Bush, Italo Cunha, Ethan Katz-Bassett, Thomas C. Schmidt, and Matthias Wählisch. 2018. Towards a Rigorous Methodology for Measuring Adoption of RPKI Route Validation and Filtering. *ACM Sigcomm Computer Communication Review* 48, 1 (January 2018), 19–27.
[33] RIPE 2020. Routing Information Service (RIS). http://www.ripe.net/projects/ris/rawdata.html
[34] RIPE NCC. 2020. Current RIS Routing Beacons. https://www.ripe.net/analyse/internet-measurements/routing-information-service-ris/current-ris-routing-beacons.
[35] Christian P. Robert and George Casella. 2005. *Monte Carlo Statistical Methods (Springer Texts in Statistics)*. Springer-Verlag New York, Inc., Secaucus, NJ, USA.
[36] Matthew Roughan, Walter Willinger, Olaf Maennel, Debbie Perouli, and Randy Bush. 2011. 10 Lessons from 10 Years of Measuring and Modeling the Internet's Autonomous Systems. *IEEE Journal on Selected Areas in Communications* 29, 9 (2011), 1810–1821.
[37] Brandon Schlinker, Todd Arnold, Italo Cunha, and Ethan Katz-Bassett. 2019. PEERING: Virtualizing BGP at the Edge for Research. In *Proc. of ACM CoNEXT*. ACM, New York, NY, USA, 51–67.
[38] Philip Smith and Christian Panigl. 2006. *RIPE Routing-WG Recommendation For Coordinated Route-flap Damping Parameters*. RIPE Document ripe-378. RIPE.
[39] Joel Sommers, Paul Barford, Nick Duffield, and Amos Ron. 2007. Accurate and Efficient SLA Compliance Monitoring. In *Proc. of ACM SIGCOMM*. Association for Computing Machinery, New York, NY, USA, 109–120. https://doi.org/10.1145/1282380.1282394
[40] Florian Streibelt, Franziska Lichtblau, Robert Beverly, Anja Feldmann, Cristel Pelsser, Georgios Smaragdakis, and Randy Bush. 2018. BGP Communities: Even More Worms in the Routing Can. In *Proceedings of the Internet Measurement Conference 2018*. ACM, New York, NY, USA, 279–292.
[41] Cecilia Testart, Philipp Richter, Alistair King, Alberto Dainotti, and David Clark. 2020. To Filter or Not to Filter: Measuring the Benefits of Registering in the RPKI Today. In *Proc. of PAM Conf. (LNCS)*, Vol. 12048. Springer, Berlin Heidelberg, 71–87.
[42] University of Oregon. 2017. Route Views Project. http://www.routeviews.org/.
[43] C. Villamizar, R. Chandra, and R. Govindan. 1998. *BGP Route Flap Damping*. RFC 2439. IETF.

A ETHICS

When performing active BGP measurements one needs to avoid impacting real-world operations. As we are sending many BGP Updates, especially when high update burst rates are active, we need to make sure that our Beacons do not overwhelm other routers. In the first measurement period, we caused 0.48% of all IPv4 control plane traffic seen in RIPE RIS, RouteViews, and Isolario data, whereas in the second period our Beacon caused 0.54% of all IPv4 BGP updates. Interestingly, the prefixes oscillating every minute were still causing a lot fewer updates than other prefixes on the Internet. As an example, we picked March 1th and measured how many announcements belonged to each prefix. We found $\approx$ 50 prefixes causing 3 times as many updates as one of our Beacon prefixes and 4 prefixes caused 17 times more updates individually than one of our Beacon prefixes.

B RFD DEFAULT PARAMETERS

RFD parameter	Cisco	Juniper	RFC 7454
Withdrawal penalty	1000	1000	1000
Readvertisement penalty	0	1000	0/1000
Attributes change penalty	500	500	500
Suppress-threshold	**2000**	**3000**	**6000**
Half-life (min)	15	15	15
Reuse-threshold	750	750	750
Max suppress time (min)	60	60	60