Route Flap Damping with Assured Reachability

Pei-chun Cheng [*]
pccheng@cs.ucla.edu

Jong Han Park [*]
jpark@cs.ucla.edu

Keyur Patel [†]
keyupate@cisco.com

Lixia Zhang [*]
lixia@cs.ucla.edu

ABSTRACT

It is well known that a relatively small percentage of unstable routes exists in the global routing system which contributes an out of proportion number of routing updates. The route flap damping (RFD) was once considered an effective means to curtail such instability. However, both measurement studies and operational observations show that BGP path exploration can trigger false route damping which leads to prolonged periods of lost network reachability. As a result, many networks turned off RFD. In this paper we propose a simple solution, RFD+RG, dubbed *RFD with Reachability Guard*, to address the reachability problem in the RFD deployment. RFD+RG performs route flap damping without losing reachability, and the +RG enhancement component works independently from specific damping algorithms and can be integrated into any existing RFD scheme. We use collected BGP data to evaluate RFD+RG performance and our results show that RFD+RG can suppress up to 27% of route instabilities while faithfully preserving reachability.

Categories and Subject Descriptors

C.2.2 [**Computer Communication Networks**]: Routing Protocols

General Terms

Measurement, Algorithms, Performance

Keywords

BGP, Route Flap Damping, Reachability, Stability

[*] Computer Science Department, UCLA.
[†] Cisco Systems, Inc.

AINTEC'10, November 15–17, 2010, Bangkok, Thailand.

1. INTRODUCTION

Border Gateway Protocol (BGP) [20] tied together the Internet's global routing infrastructure. BGP routers exchange routing updates to adapt to topological connectivity changes caused by either intentional routing policy changes or more commonly unexpected software and hardware failures. Because BGP runs in a flat routing space, a single unstable route can cause a ripple effect which results in thousands of update messages propagating throughout the entire network. It is well known that a relatively small percentage of unstable routes exists in the global routing system and contributes an out of proportion number of routing updates [13, 18].

Two major mechanisms have been designed to mitigate the impact of unstable routes. The Minimal Route Advertisement Interval (MRAI) is used to pace out BGP updates by introducing a minimal gap between two consecutive update messages for the same prefix, with a default value of 30 seconds. Note that MRAI enforces a nondiscriminatory rate limit on *all* prefixes, regardless their status of (in)stability. The second mechanism, Route Flap Damping (RFD) [23], is designed to detect and suppress perpetual route instabilities. It was once considered an effective means to maintain the overall Internet routing stability [11].

However measurement studies and operation experience revealed that RFD deployment can lead to prolonged periods of route convergence and loss of reachability [17, 7, 25, 26], which are resulted from RFD's pathological interplay with the BGP path exploration. As a result, even though persistent routing instabilities are observed repeatedly [13, 8], the operational community is deeply concerned with RFD and suggested to turn it off [21]. Compounded with the fact that the MRAI timer is also being turned off at various places, the global routing system today faces a potential danger of melting down by excessive amounts of routing updates.

To revive the deployment of RFD in the global routing system, in this paper we propose a simple enhancement to RFD, RFD with Reachability Guard (RFD+RG), to achieve route flap damping *without loss of reachability*. Our solution is based on the observation that many unstable prefixes are covered by relatively stable prefixes, and that a router often can reach a prefix via multiple neighbors, among them only a subset experiences route instability. We evaluated RFD+RG using BGP feeds collected by RIPE [2], and our results show that it can reduce the total number of BGP updates by 5% to 27% without loss of reachability. We emphasize that our proposed enhancement, the +RG component in RFD+RG, is not a new damping algorithm itself. Rather, it is a generic and compatible *addition* to any existing RFD algorithm to prevent reachability losses.

2. ROUTE FLAP DAMPING

In this section, we briefly describe the operation of RFD, its known issues, and the existing solutions. Interested readers are directed to [23, 17, 7, 25, 26] for more detailed descriptions.

2.1 A Brief History

The original RFD algorithm was developed in the mid 1990s and standardized in RFC 2439 [23]. Its goal is to prevent sustained routing oscillations without sacrificing route convergence time for generally well behaved routes. RFD associates a penalty value with each route. The penalty increases when a new update message is received for the corresponding route. When the penalty value exceeds a predefined suppression threshold, the route is *suppressed* (or *damped*) and excluded from the BGP best path selection. The penalty value decays exponentially over time, and the route is *reused* (becoming eligible for the best path selection) when the penalty value decreases below a predefined reuse threshold.

Unfortunately, as the operational community put forth avid efforts in adopting flap damping [4, 5], RFD was shown to have some undesirable negative effects. In [17], Mao *et al.* showed that, as a path vector protocol, BGP can amplify a single route flap into many updates during *path exploration* which can falsely trigger route suppressions. As a result, RFD exacerbates convergence time for fairly stable routes and, even worse, hurts reachability when all existing routes to a given prefix are suppressed [7, 17].

A number of research efforts have been devoted to improve the accuracy of route flap detection to solve the false suppression problem. A common approach is to extract the *signatures* of path exploration and persistent route flaps. In [17] and [7], the authors show that during a path exploration, a BGP router selects and advertises the best route in a non-increasing order of route preference. Based on this observation, they suggest that the penalty value should be increased only when there is a change of direction in route preference. In [25], the authors propose to stamp BGP updates with a unique event identifier of the source of instability. By doing so, multiple updates of the same route due to a single flap event can be clustered together and the route is penalized only once. The authors of [26] observed that path exploration generally lasts for a shorter time period than persistent rout flaps, and proposed to penalize only the first update received within a time window.

However, none of above proposals is widely deployed. One possible explanation can be that these proposals usually require new BGP attributes to be added to routing updates that expose sensitive information such as route preference or failure locations. The new attributes could also impact Internet operations in an unexpected way [3]. Furthermore, these proposals only reduce false route flap detections by a certain degree and provide no guarantee on preserving reachability. As of today, RFD is turned off at many networks. [21] states that *"... the application of flap damping in ISP networks is NOT recommended. ... flap damping is harmful to the reachability of prefixes across the Internet."*

2.2 Why Revisit RFD?

Recent rises of real-time applications bring new requirements to the routing system. Compared to conventional non-real-time data communications, real-time voice and video applications are much less tolerant to delay jitters that are caused by frequent route changes. Measurement studies show that BGP events are highly correlated with 50% of Skype quality degradation and 90% of call drops [15]. At the same time, network operators gradually lose control over routing instability: not only the flap damping is largely turned off, but also the use of MRAI has been decreasing due to the desire for faster routing convergence [14]. We believe that if we can fix RFD's reachability loss problem, it could again play an important role in stabilizing the global routing system.

3. MANY ROADS LEAD TO ROME

Previous works in RFD [16, 7, 25, 26] all considered each prefix as an independent unit of reachability in the routing system. However in reality, a given destination network N can be reached through multiple *paths* in general, and N's address space is often covered by more than one *prefix* in the routing table. In this section, we first measure the existence of such alternative reachability, and then show that noisy prefixes can often be reaachable via alternative paths that are significantly stabler.

3.1 Prevalence of Alternative Routes

We use the BGP table snapshots collected in LINX[1] by RIPE [2] on December 1st, 2009 to measure the number of nexthop neighbors for a given destination from the collector's view. Given that we do not know the routing policies of each router connected at LINX, the results presented below do not necessarily represent the perspective of each individual router. Rather, we are interested in quantifying the potential alternative reachability at a popular peering site such as LINX. We also performed the same alternative route measurements using routing data collected from other Internet exchange points and the results are similar to that of LINX. Due to space limitation we only present the LINX results in this paper.

[1] London Internet Exchange Point

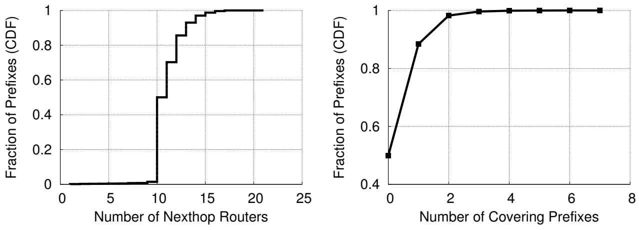

(a) Nexthop Routers (b) Covering Prefixes

Figure 1: Prevalence of Alternative Routes

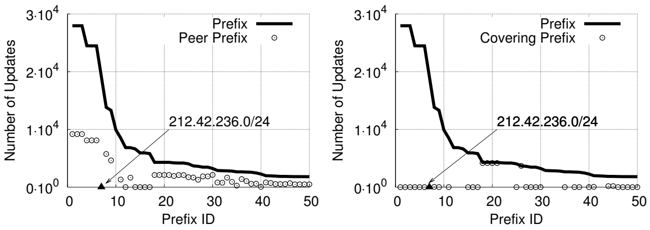

(a) Using Multiple Nexthop Routers (b) Using Covering Prefixes

Figure 2: Relative Dynamics of Alternative Routes

Figure 1(a) depicts the observed number of nexthops for each prefix in the global routing table. Except for a few prefixes that are originated by the local routers and have a low nexthop count, most prefixes can be reached via 10 or more different nexthops. This is because ten of the routers at LINX advertise the full routing table to the collector. Other routers advertise a partial routing table (*i.e.,* they treat the collect as a peer) and account for the nexthop counts greater than 10. Note that this result only represents a perspective from one particular BGP data collector, and different measurement settings may yield different results. However generally speaking, the number of nexthops to reach a given destination is approximately the same with the number of BGP neighbors announcing the full routing table, *i.e.,* the neighbor routers of one's providers.

In addition, a prefix can also be reached through any of its covering prefixes (*e.g.,* destination 10.1.1.0/24 can be reached through a route to 10.1.0.0/16). Figure 1(b) shows the distribution of the number of covering prefixes for a given prefix. Note that more than 50% of all prefixes in the global routing table have covering prefixes. This result echos our earlier observations in 2003 [24]. The majority of the covered prefixes (little less than 40% of all prefixes) have one covering prefix; about 10% of all prefixes have two, and the rest of the covered prefixes have 3 to 7 covering prefixes. We further checked the historical global routing tables to see whether the statistics for covered prefixes has changed over time. From 2005 to 2009, the fraction of covered prefixes remained steady between 45% and 55%. A similar observation is also made in [12] that approximately half of all prefixes are covered, and this percentage has not changed significantly over time.

3.2 The Stability of Alternative Routes

We have shown that alternative routes to a destination exist in general. The next question is whether alternative and stable routes exist for unstable prefixes. From the BGP updates collected from LINX during December 1st to 7th in 2009, we choose an example router and identify the top 50 most noisy prefixes (*i.e.,* prefixes with the largest number of updates) from the updates it sent out. We then check whether these prefixes have any alternative routes, and if they exist, whether these alternative routes have fewer updates.

For the top 50 most noisy prefixes, Figure 2(a) and Figure 2(b) show the number of updates received on the prefix itself from the router, together with the number of updates on their alternative routes, *i.e.,* through other routers or covering prefixes[2] respectively. When calculating the number of updates, we also check the reachability of the noisy prefixes and their alternative routes for a fair comparison.

We observed that a few stable alternative routes often exist for these noisy prefixes. For example, prefix $212.42.236.0/24$ is found noisy as a router belong to AS286 forwarded $17{,}635$ updates for this prefix during the week of December 1st. However, this prefix can also be reached using other routes either via AS8468, or via a covering prefix $212.42.224.0/19$, and these two alternative paths had only 1 and 0 updates respectively for the whole week! Overall, the number of updates received on the noisy prefixes is always greater than that of their alternative paths when they exist, and the difference is significantly large in most cases. Similar observation has been made by Huston *et al.* [10], and they conjectured that these noisy prefixes with stable covering prefixes could be due to poorly tuned (or lack of tuning) of automated traffic engineering processes.

The existence of both stable and unstable routes for a prefix sheds light on an opportunity to suppress routing instabilities while preserving the reachability. In the following sections, we derive a practical RFD enhancement and evaluate its performance.

4. RFD+RG: NOT ANOTHER ROUTE FLAP DAMPING ALGORITHM

In this section, we describe a simple addition to route flap damping; we call the combined scheme *Route Flap Damping with Reachability Guard* (RFD+RG). The basic idea is to suppress a flapping prefix p only when one or more alternative routes to p exist, *i.e.,* when the reachability to p can be preserved. We emphasize that this work is not another damping algorithm itself, but a complementary addition to any existing rout flap damping scheme.

In all the previous works, one common issue is that the prefix flap *detection* is bundled with route *suppression*. That is, a prefix is *suppressed* immediately once it is *detected* unstable. Consequently the performance of all these previous proposed solutions critically depends on the accuracy of the detection algorithm. Either the detection algorithm can cor-

[2] For clarity, if there exist multiple alternative routes, we only show the number of updates for the most stable one.

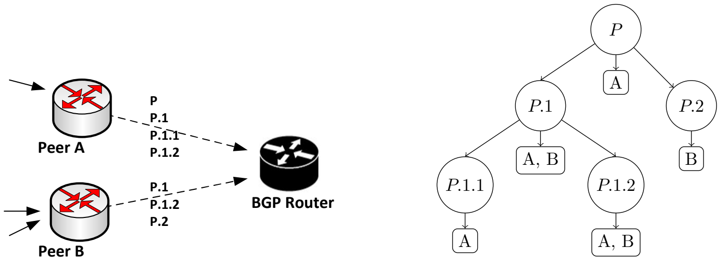

(a) Alternative Routes (b) Local Routing Table

Figure 3: Protector and Protectee Prefixes & Routes.

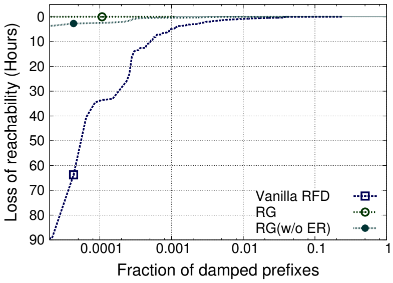

Figure 4: Reachability Loss

rectly distinguish true route flapping from path explorations all the time, which is difficult if not impossible, or otherwise any false detection would lead to false suppression and potential reachability losses.

In this work, we propose to decouple route flap detection from route suppression, and we develop a set of rules for making route suppression decisions based on engineering design tradeoffs. More specifically, our solution can take any existing algorithm for route flap detection. Once a flapping route is identified, we perform a reachability check to determine its eligibility for suppression. This approach eliminates reachability losses due to route suppression, which can be triggered by either correct or false flapping detections.

4.1 Protector and Protectee

We say that a prefix $p1$ is a *protector* of another prefix $p2$, if $p1$ fully covers $p2$'s address space. Furthermore, we call the routes to protector and protectee prefixes as *protector routes* and *protectee routes*, respectively.

Assuming a BGP router R connects to two neighbor routers A and B as shown in Figure 3(a), Figure 3(b) depicts R's routing table. We assume that R's prefixes are organized using a binary trie[3] and the route to one prefix via a particular neighbor is denoted as *(prefix, neighbor)*. In this example, the prefix $P.1$ is a protector prefix of $P.1$[4], $P.1.1$ and $P.1.2$, and the route $(P.1, A)$ is a protector route of routes $(P.1, B)$, $(P.1.1, A)$, $(P.1.2, A)$, and $(P.1.2, B)$. Note that even if $P.1.1$, $P.1.2$ could collectively cover the whole address range of $P.1$, each of them is not a protector of $P.1$.

For ease of presentation, we use $protector(p)$ to represent prefix p's protector prefixes, and $protector(p, r)$ to represent the set of route (p, r)'s protector routes. Similar notations are used for the protectee prefix and route. In addition, we say that a route is *valid* if it is not withdrawn nor damped. Without specified otherwise, *reachability loss* in the following sections means any reachability loss caused by route flap damping.

4.2 Reachability Guard

In order to preserve reachability, RFD+RG introduces two new steps in performing route flap damping: *reachability check* and *early release*.

4.2.1 Reachability Check

Assuming that persistent flapping is identified for a route (p, r), before suppressing this route, RFD+RG checks the set of $protector(p, r)$. If a valid protector route is found, the flapping route (p, r) can be safely suppressed without losing reachability; otherwise route (p, r) is left intact. If route (p, r) flaps persistently, its protector route set, $protector(p, r)$, will be evaluated continuously; as soon as a valid protector route appears, the flapping route will be suppressed. For example in Figure 3(b), supposing that route $(P.1.2, B)$ is found unstable, its protector routes, $(P.1.2, A)$, $(P.1.A)$, $(P.1.B)$ and $(P.A)$ will be be evaluated, and if any of them is valid, $(P.1.2, B)$ is suppressed.

For simplicity in implementation, we check protector routes bottom-up along the tree structure. However, based on our own observations and that from [12], shorter prefixes tend to be stable in general, thus given an unstable prefix, searching in the top-down manner could be a more efficient way to find a valid route.

Due to the dynamic nature of network routing, it is possible that a protector route itself may later become unstable or withdrawn. We address these cases next.

4.2.2 Early Release

The second new step introduced by RFD+RG, *early release*, is triggered whenever a route (p, r) is $withdrawn$, and p does not have any valid protector route. As a result, p's suppressed protectees (if any) should be examined to make sure that they are not losing reachability. As an example, let's assume a scenario that the both route $(P.1.2, B)$ and $(P.2, B)$ are unstable and suppressed, while all the other routes in Figure 3(b) are stable. Now supposing that route (P, A) is withdrawn, then we need to release route $(P.2, B)$ to preserve the reachability to $P.2$. We can safely keep suppressing route $(P.1.2, B)$ since it still has protector routes.

PROPOSITION 1. *Consider a BGP router that enables route flap damping with reachability guard. The router should not lose reachability due to damping.*

PROOF. Due to the limited space, we only sketch a proof by contradiction. Consider the router triggered an ordered

[3] A trie is a prefix tree, which is widely used in BGP implementations to organize routing tables internally.

[4] Based on our reflective definition, a prefix is also its protector or protectee prefix.

sequence of events, $\xi = e_1 \cdots e_{t-1}e_t$, upon update arrivals or damping timer expirations. In route flap damping, an event can be an announcement[5], a withdrawal, a suppression, or a reuse of a route. Now assume that after an event e related to a route (p, r), the router started to lose reachability to p. Since announcement and reuse events would not hurt reachability, the loss of reachability can only be due to one of the following four cases: (1) the event suppresses (p, r) while there are no protector routes, (2) (p, r) is already suppressed and the event further suppresses (p, r)'s last valid protector route, (3) the event withdraws (p, r) while there are no protector routes, (4) (p, r) is already suppressed and the event withdraws (p, r)'s last valid protector route. By case studies, one can easily show that, if the two checks of reachability guard are correctly implemented, case 1, 2 and 4 should not happen. Moreover, case 3 does not lead to reachability loss (*i.e.,* suppress a withdrawn route would not hurt reachability). □

4.3 Applicability and Limitation

Clearly the proposed +RG enhancement is not a standalone solution. The enhancement operates on top of an existing damping scheme. The +RG enhancement can be easily added to the recent works on RFD and lead to SRFD+RG [17], RFD-RCN+RG [25], FRFD+RG [26], respectively. In evaluating +RG's performance, we choose to use the original RFD as defined in [23], because it has been widely implemented by router vendors. This fact leads to a few important implications.

First, the combined solution would inherit pros and cons of the underlying damping scheme. For example, RFD is known to reduce router load and help routing stability, at the potential cost of delayed convergence and reachability loss. This also applies to RFD+RG, however, with a critical difference that +RG would guarantee the fundamental objective of the Internet: reachability. Potentially +RG enhancement might reduce convergence time given that it never completely shuts down a prefix by damping. Olivier *et al.* showed that convergence is relatively fast once reachability is preserved [19].

Second, based on the existence of alternative reachability, RFD+RG shows an amphibious behavior. At one end, in a densely connected network where there always exist alternative routes, RFD+RG acts exactly the same as the original RFD. At the other end, in a single-homed edge network without any alternative route, no flapping routes can be damped, and RFD+RG simply degenerates to no flap damping! However, in practice, most networks fall in between, and +RG acts as a tuning knob which dynamically morphs the damping behavior to protect reachability.

Admittedly, the overall benefits and impact of RFD itself remain an open question. For example, is route flap damping absolutely necessary for the stability of the Internet routing system? Is it necessary to deploy flap damping throughout the network, or it is only needed at specific locations? How much overhead does RFD introduce to routers? Unfortunately, the research efforts on RFD have dwelt in the past few years due to the undesired reachability loss caused by RFD. Our work is the first to directly address the reachability loss problem, as an initial attempt to revive communities' interest in understanding route flap damping.

5. EVALUATION

As the first step toward evaluating the performance of the +RG component, we focus on understanding how a *single router* reacts to updates received from its neighbors. We implement an event driven BGP simulator [6] with different flap damping schemes, including (1) the vanilla RFD, *i.e.,* original RFD, (2) RFD+RG, and (3) a stripped version of RFD+RG which does not perform *early-release* and has less computation overhead. When selecting the best route, we used a simple shortest AS path selection algorithm.

The simulator is fed with the BGP update data collected from operational routers by RIPE's BGP collector at LINX exchange point. We randomly picked the BGP data from the week of December 1st – 7th, 2009 to emulate a semi-realistic scenario, assuming the collector is a BGP router R which connects to other operational routers at LINX, and receives routing updates from the neighbors. We picked different number of routers with full routing tables and no BGP session resets during that week. Due to the space limit, we only discuss the results for four neighbor routers. We emphasize again that this setting does not intend to emulate an actual operational router, as routers at an exchange point do not necessarily exchange full routing tables. Rather, our setting emulates a customer router with different number of providers.

We compare the performance by enabling or disabling route flap damping on R. Unless otherwise specified, the RFD implementation in the simulator uses Cisco default damping parameters [27]. To further validate our results, we also performed additional simulations on BGP feeds from BGP collectors at other exchange points. The results for different exchange points will also be presented later in this section.

5.1 Preserve Reachability

During our 1-week evaluation period (168 hours), the emulated router observes total 335,372 prefixes, of which 41,086 prefixes are damped at least once. For each damped prefix, we calculate its unreachable time duration due to damping. Figure 4 shows the reachability losses in time. Similar to previous observations, our results show that the vanilla RFD can significantly impact prefix reachability: 2% (822) of the damped prefixes lose reachability for more than 30 minutes; the top 50 unstable prefixes are suppressed for more than 5 hours. For the worst few prefixes, RFD blocks their reachability for longer than half a week.

[5] Announcement or withdrawal that does not trigger damping

[6] We only implement BGP functionalities related to our performance evaluation, namely *flap damping* and *best path selection*.

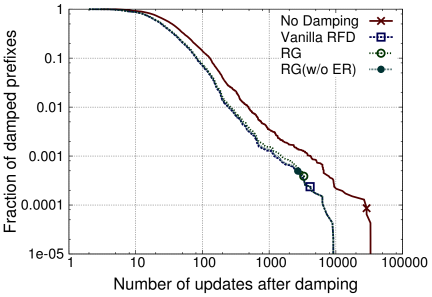

Figure 5: Update Count

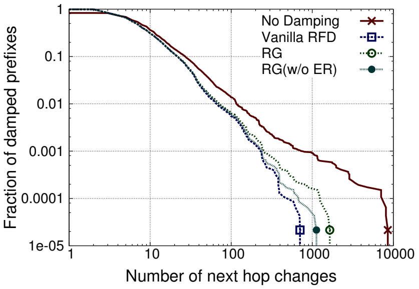

Figure 6: Nexthop Change Count

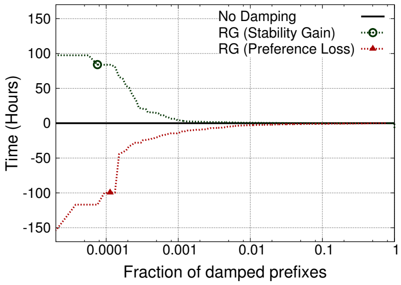

Figure 7: Stability and Preference

On the other hand, with RFD+RG, no prefix lost reachability due to damping, as shown in Table 1. With the stripped version of RFD+RG (without early release), some prefixes lose up to 4 hours of reachability as the worst case; this happens when the protector routes themselves become unstable or withdrawn. Nevertheless, compared with the vanilla RFD, the reduction of reachability loss is still more than an order of magnitude.

An interesting observation is that RFD behaves better than one may have expected: 80% of the damped prefixes did not lose reachability due to the help of its protector routes. However, such result is unreliable and depends on the opportunistic existence of protector routes. In contrast, RFD+RG is designed to assertively eliminates reachability losses caused by route damping.

5.2 Reduce Router Load

For the emulated router, Figures 5 and 6 show the complementary CDF for the number of BGP updates and nexthop changes, the two primary contributors of BGP processing load [23]. In Figure 5, the long tail distribution shows that a small number of prefixes contribute to a relatively large number of updates. This result also conforms to the observations made in [13, 18].

In our evaluation case, RFD+RG reduces 36% of BGP updates contributed by damped prefixes, or 24% of the total updates for all prefixes. Figure 5 also shows results for the vanilla RFD and the stripped RG for comparison. We observe that RFD+RG reduces nearly the same amount of updates compared to vanilla RFD without any prefix losing reachability.

Figure 6 shows the complementary CDF over the number of BGP nexthop changes. In BGP, it is important to minimize nexthop router changes in order to keep forwarding plane stability. Nexthop changes trigger changes to the forwarding table (FIB), and frequent FIB changes can degrade data plane performance [26]. Figure 6 shows that, compared to no RFD, RFD+RG reduces the number of nexthop changes by 28% when counting all nexthop changes by damped prefixes, or 21% when counting nexthop changes for all prefixes. For a few most unstable prefixes, the reduction of nexthop changes is up to an order of magnitude. Moreover, compared to the vanilla RFD, RFD+RG only allows a small number of additional nexthop changes to preserve reachability.

Table 1: Summary of Evaluation Results (LINX), 4 peers

	Reach. loss (hours)	Reduced updates (%)	Reduced NH changes (%)
RG	0	24.21	21.70
RG(w/o ER)	91	25.39	22.46
RFD	2018	26.00	23.55

Table 1 summarizes the overall performance amortized for all prefixes. Note that our aim here is to compare the simulated performance rather than the actual workload reduction on real routers. Although RFD+RG significantly reduces the number of updates and nexthop changes, we do not know whether this amount of workload reduction makes a significant difference to the latest generation of routers. But a recent work by Houidi *et al.* showed that still the recent commercial routers are not as effective in processing and exchanging instantaneous bursts of updates as one might expect, due to some limitations in the implementations [9]. Besides, updates and nexthop reduction is only one benefit of RFD+RG. In the next section, we discuss another major merit of RFD+RG in improving routing stability.

5.3 Make A Safer Trade-off

As a penalty-based system, the vanilla RFD improves the routing stability with an undesirable tradeoff of reachability losses [17, 21, 26]. In contrast, RFD+RG makes a conservative trade-off between route *preference* and *stability*: the router can suppress a preferred but flapping route and use a more stable route that may be less preferred.

Figure 7 illustrates this tradeoff. For each prefix damped by RFD+RG, we measure its *stability gain* and *preference loss*. We first calculate the continuous usage time for the router to use the same nexthop to reach a prefix, before switching to another nexthop. Considering that 10 minutes is the median duration of Skype calls with 95% confidence [6], we define any nexthop change in less than 10 minutes as unstable. Then we compare the difference of stable usage time between *No RFD* and *RFD+RG*. The *preference loss* is measured as the time period that the router uses a longer[7] route to reach a prefix. Figure 7 shows a clear tradeoff between stability and preference. For prefixes that have unstable shorter routes, RFD+RG would suggest the router to use longer alternative routes, which yields better stability at the cost of

[7] Recall that we used a simple shortest AS path selection algorithm.

Table 2: Results of Different Number of Peers

Num. Peers	Reach. loss (hours)	Reduced updates (%)	Reduced NH changes (%)
1	0	11.87	25.25
2	0	21.86	25.53
3	0	21.00	22.08
4	0	24.21	21.70
5	0	23.45	21.69
6	0	22.16	20.42

using a less preferred path.

5.4 Additional Evaluation Results

5.4.1 Impact of the Number of Neighbors

In order to gain an insight on the impact of the number of neighbors on the number of alternative paths, we performed the same emulation of the example router R using an increasing number of neighbors from 1 to 6. Table 2 summarizes the reduction in updates and nexthop changes when a router using RFD+RG peers with different number of neighbors. As expected, the results show that +RG is able to guarantee reachability independent from the number of neighbors, while overhead reduction depends on the specific behavior of each neighbor. That is, if the number of neighbors is greater than 2, more (or fewer) neighbors do not necessarily result in higher (or lower) reduction. Note that for the case of 1 neighbor, there is no alternative next hops. That is, the reachability guarantee is provided either by covering prefixes or otherwise unstable prefixes undamped, and all the overhead reduction is due to suppressing noisy prefixes due to the existence of covering prefixes.

5.4.2 Applying to Other Internet Exchange Points

The evaluation results presented so far are based on BGP feeds from one specific exchange point. In this section, we extend the evaluation by using BGP feeds from 7 other exchange points. Table 3 summarizes our evaluation results based on BGP data collected from 8 different Internet exchange points. We make two observations. First, even though different exchange points are in different geographic locations and have different operational environments, a RFD+RG enabled router can run RFD without reachability losses at all the locations. Second, different exchange points observe different magnitude of instability during the same measurement period. For some exchange points, such as London and Stockholm, RFD+RG helps reduce routing updates and nexthop changes by more than 20%, while for other places with relatively stabler routing, such as Geneva, the damping is triggered less often and the update reduction percentage is lower. We have further examined Geneva's raw BGP feeds and verified that many noisy prefixes observed by other exchange points somehow showed much stabler behavior at Geneva. Overall, RFD+RG is able to reduce the total BGP updates by 5% to 27% across different exchange points.

6. DISCUSSION AND FUTURE WORK

Table 3: Results for Different Exchange Points

	Location	Reach. loss	Damped prefixes	Reduced updates (%)	Reduced NH changes (%)
LINX	London	0	46,488	24.21	21.70
AMS-IX	Amsterdam	0	13,464	13.28	19.09
CIXP	Geneva	0	9,125	5.82	16.40
NETNOD	Stockholm	0	45,496	27.16	20.30
MIX	Milan	0	31,543	11.10	14.33
NYIIX	New York	0	15,907	8.99	15.20
DE-CIX	Frankfurt	0	26,708	17.89	27.07
MSK-IX	Moscow	0	29,314	12.97	19.77

In this paper, our goal is to demonstrate that a simple addition to RFD can effectively eliminate the loss of reachability due to route damping. Admittedly, this gain does not come for free. Compared to vanilla RFD, RFD+RG may require additional data structures and/or computations. First, RFD+RG needs to keep track of the protector-protectee relations among prefixes. Fortunately, BGP implementations usually organize routing tables using the aforementioned trie structures which already embed such links between covering and covered prefixes, and the +RG component does not consume any extra memory space [22, 1]. Second, a router must check multiple routes before making one damping decision. In the worst case, one may need to traverse all protector or protectee routes. However, as the results in Section 3 show, the majority of prefixes have a fairly number of protectors or protectees (mostly less than 3), thus we expect that +RG only introduces a reasonable amount of computation load in general. More evaluation studies are needed to further quantify the incurred overhead.

One remaining issue is how to handle unstable *orphan* routes, *i.e.,* routes that do not have protectors. In this work we follow a simple principle of *reachability first*, and therefore we deliberately let go the updates of such unstable routes. It is our belief that to (re)enable RFD we must provide an effective means to address operators' concerns in reachability losses [21]. In this paper we showed that a simple solution can both eliminate reachability losses and remove a significant portion of updates generated by unstable prefixes. For operators who are most concerned with routing stability and willing to trade some small loss of reachability for better routing stability protection, a router may be configured to damp these orphan routes, perhaps with some remedial actions such as shortening the suppression period.

Another open issue concerns routing convergence delay. One limitation of this work is that we only focused on evaluating the damping behavior and reachability losses at specific routers, our emulation setting does not allow us to measure the impact of +RG on the route convergence time in the global routing system. Quantifying such system wide impacts would require a large scale synthetic simulation with realistic Internet scale topologies.

7. SUMMARY

We deem it necessary to maintain in the global routing system some basic defensive measures against potential excessive update flooding. To address the reachability loss problem that has been observed with the existing route flap damping schemes, this work presents a simple addition to RFD to prevent undesirable reachability losses caused by route flap damping. The solution is built upon the densely connected Internet AS level topology and diversified nature of BGP routing, which enables one to reach a given network destination via multiple paths and different prefixes.

Although this observation itself is not new, our main contribution is a practical RFD enhancement, +RG, derived from this observation, and a systematic evaluation of its effectiveness using real BGP data. The proposed enhancement is incrementally deployable without coordination between routers. Our preliminary results show that RFD+RG can reduce the total routing updates by up to 27% without inducing reachability loss. The +RG component can be integrated with any damping scheme being used and provide operators a safer tuning knob, one that trades off route preference, rather than loss of reachability, for route stability and overhead reduction.

Acknowledgment

We express our deepest gratitude to Ricardo Oliveira and Michael Meisel for their help to launch this project. We also thank anonymous reviewers for many helpful comments on this paper.

8. REFERENCES

[1] Quagga Software Routing Suite. http://www.quagga.net/.

[2] RIPE Routing Information Service. http://www.ripe.net/projects/ris/.

[3] RIPE NCC and Duke University BGP Experiment. http://labs.ripe.net/Members/erik/ripe-ncc-and-duke-university-bgp-experiment, 2010.

[4] T. Barber, S. Doran, D. Karrenberg, C. Panigl, and J. Schmitz. RIPE Routing-WG Recommendation for coordinated route-flap damping parameters. RIPE 178, May 1998.

[5] T. Barber, S. Doran, D. Karrenberg, C. Panigl, and J. Schmitz. RIPE Routing-WG Recommendation for coordinated route-flap damping parameters. RIPE 210, May 2000.

[6] K.-T. Chen, C.-Y. Huang, P. Huang, and C.-L. Lei. Quantifying skype user satisfaction. In *SIGCOMM ’06*.

[7] Z. Duan, J. Chandrashekar, J. Krasky, K. Xu, and Z.-L. Zhang. Damping BGP route flaps. In *Proc. IEEE International Conference on Performance, Computing, and Communications*, 2004.

[8] A. Elmokashfi, A. Kvalbein, and C. Dovrolis. BGP Churn Evolution: a Perspective from the Core. In *INFOCOM 2010*, 2010.

[9] Z. B. Houidi, M. Meulle, and R. Teixeira. Understanding Slow BGP Routing Table Transfers. In *IMC 2009*, Nov. 2009.

[10] G. Huston. Update Damping in BGP.

[11] G. Huston. Commentary on Inter-Domain Routing in the Internet. RFC 3221 (Informational), Dec. 2001.

[12] G. Huston. BGP Statistics. http://bgp.potaroo.net/index-bgp.html, 2010.

[13] G. Huston. The BGP Instability Report. http://bgpupdates.potaroo.net/instability/bgpupd.html, 2010.

[14] P. Jakma. Revisions to the BGP ’Minimum Route Advertisement Interval. http://tools.ietf.org/html/draft-ietf-idr-mrai-dep-02, 2010.

[15] N. Kushman, S. Kandula, and D. Katabi. Can you hear me now?!: it must be BGP. *SIGCOMM Comput. Commun. Rev.*, 37, 2007.

[16] Z. M. Mao, R. Bush, T. G. Griffin, and M. Roughan. BGP beacons. In *IMC ’03*, 2003.

[17] Z. M. Mao, R. Govindan, G. Varghese, and R. H. Katz. Route flap damping exacerbates internet routing convergence. In *SIGCOMM ’02*, 2002.

[18] R. Oliveira, R. Izhak-Ratzin, B. Zhang, and L. Zhang. Measurement of highly active prefixes in BGP. In *GLOBECOM’05*.

[19] R. Oliveira, B. Zhang, D. Pei, and L. Zhang. Quantifying path exploration in the internet. *IEEE/ACM Trans. Netw.*, 17(2):445–458, 2009.

[20] Y. Rekhter, T. Li, and S. Hares. A Border Gateway Protocol 4 (BGP-4). RFC 4271 (Draft Standard), Jan. 2006.

[21] P. Smith and C. Panigl. RIPE Routing Working Group Recommendations on Route-flap Damping. RIPE 378, May 2006.

[22] K. Solie and L. Lynch. CCIE Practical Studies. In *Practical Studies*, page 1032, Indianapolix, Indiana, USA, 2003. Cisco Press.

[23] C. Villamizar, R. Chandra, and R. Govindan. BGP Route Flap Damping. RFC 2439 (Proposed Standard), Nov. 1998.

[24] L. Z. Xiaoqiao Meng, Zhiguo Xu and S. Lu. An analysis of bgp routing table evolution. Technical Report 030046, Department of Computer Science, UCLA, 2003.

[25] B. Zhang, D. Pei, D. Massey, and L. Zhang. Timer Interaction in Route Flap Damping. In *ICDCS 2005*, 2005.

[26] K. Zhang and S. Wu. Filter-Based RFD: Can We Stabilize Network Without Sacrificing Reachability Too Much. In *Internet Federation for Information Processing*, 2007.

[27] R. Zhang and M. Bartell. *BGP Design and Implementation*, chapter Route Flap Dampening, pages 91–94. Cisco Press, 2005.