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 portion number of routing updates, and that the route flap damping (RFD) was once considered a major contributor 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 period of lost reachability. As a result, many networks have 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 instabilities while faithfully preserving reachability.

1. INTRODUCTION

Border Gateway Protocol (BGP) [17] tied together the Internet's global routing infrastructure. BGP routers versatilely exchange routing updates to adapt to topology changes, including intentional 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 portion number of routing updates [11, 16].

Two major mechanisms are employed to mitigate the impact of unstable routes. The Minimal Route Advertisement Interval (MRAI) is used to clock 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) [20], is designed to detect and suppress perpetual route instabilities. It was once considered a main contributor to the overall Internet routing stability [9].

Unfortunately, measurement studies and operation experience reveal RFD's pathological interplay with the BGP path exploration, which leads to prolonged periods of route convergence and loss of reachability [15, 6, 21, 22]. As a result, even though persistent instabilities are observed over time [11, 7], the operational community expressed concerns with RFD and suggested to disable it [18]. Compounded with the fact that the MRAI timer is also being turned off at various places, the global routing system may face a potential danger of melting down by excessive amounts of routing updates.

In this paper, we propose a simple additional component to RFD 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, while only a subset of them observes route instability. We evaluated our solution using BGP feeds collected by RIPE [2], and our results show that it can reduce the total BGP updates by 5% to 27% with no reachability losses. We emphasize that this work is not another new damping algorithm. Rather, we provide a compatible *addition* to the existing RFD algorithms to avoids reachability losses. By doing so, we hope to revive the deployment of RFD in the global routing system.

2. ROUTE FLAP DAMPING

In this section, we briefly describe the operation of RFD, its issues discovered so far, and the succeeding enhancements. Interested readers are directed to [20, 15, 6, 21, 22] for more detailed descriptions.

2.1 A Brief History

The original RFD algorithm was first designed in the mid 1990s and standardized in RFC 2439 [20]. Its goal is to *prevent sustained routing oscillations without sacrificing route convergence time for generally well behaved routes*. For each route, RFD assigns a penalty value to each update. The penalty increases when a new update message is received for the corresponding route. When the penalty exceeds a pre-

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

defined 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* when the penalty value decreases below a predefined reuse threshold.

Unfortunately, while the operational community put forth avid efforts in adopting flap damping [3, 4], RFD has been shown to have some undesirable negative effects. In [15] Mao *et al.* showed that 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 [6, 15].

Since then, extensive research efforts have been made to improve the accuracy of route flap detection. A common approach is to extract the *signatures* of path exploration and persistent route flaps. In [15] and [6], 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 only be increased when there is a change of direction in route preference. In [21], the authors propose to stamp BGP updates with a unique event identifier of the source of instability. By doing so, multiple updates rooted in a single flap event can be easily clustered and penalized only once. Observing that path exploration generally lasts for a shorter time period than persistent rout flaps, the authors of [22] propose 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 introduce additional complexities or require exposures of sensitive information such as route preference or failure locations. Furthermore, these proposals only reduce false route flap detections, but provide no guarantee on preserving reachability. As a result, various networks start turning off RFD. [18] 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?

In addition to the known fact that a relatively small percentage of unstable routes exists in the global routing system which contribute a relatively large number of updates (the top 50 prefixes contribute about 10% of total BGP updates) [11, 16], the recent rise of real-time applications raises new requirements on the routing system. Real-time (voice or video) applications are less tolerant to frequent route changes compared to conventional non-real-time data communications. Measurement studies show that BGP events are highly correlated with 50% of Skype quality degradation and 90% of call drops [13]. 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 [12]. 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 [14, 6, 21, 22] in RFD consider 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. We then show that the alternative reachability can often be more stable in the case of noisy prefixes.

3.1 Prevalence of Alternative Routes

First, based on recent BGP table snapshots collected in LINX[1] by RIPE [2] on December 1st, 2009, we measure the number of nexthop neighbors for a given destination per the collector's view. This result approximates a LINX member[2] router's view assuming that it peers with all other members.

Figure 1(a) depicts the observed number of nexthops for each prefix. Except a few prefixes that are locally originated by the member routers, the majority of prefixes can be reached via more than 10 different nexthops. This is mostly because there exist 10 routers in LINX which advertise the full routing table. Other routers advertise only a partial routing table (*i.e.* peers) and account for the nexthop counts greater than 10. Note that this result only represents a perspective from one particular exchange point, and different measurement settings shall yield different results. However in general, the number of nexthops to reach a given destination is approximately the same with the number of BGP peers announcing the full routing table, *i.e.* neighboring providers.

In addition, as we mentioned earlier in this section, a prefix can be reached through any of its covering prefixes. 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. 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 [10] that approximately half of all prefixes are covered, and this percentage has not changed significantly over time.

[1] London Internet Exchange Point

[2] A member is a router in the exchange point

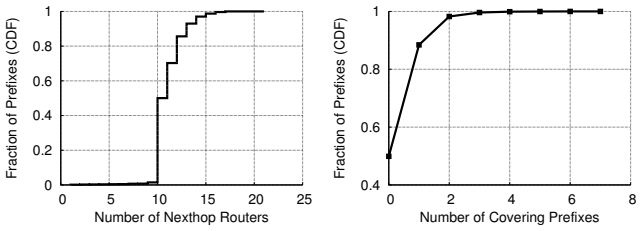

(a) Nexthop Routers (b) Covering Prefixes

Figure 1: Prevalence of Alternative Routes

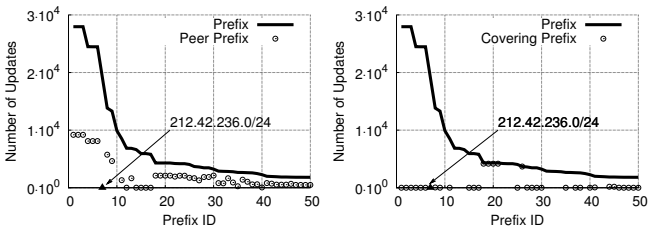

(a) Using Multiple Nexthop Routers (b) Using Covering Prefixes

Figure 2: Relative Dynamics of Alternative Routes

3.2 Alternative Routes Could Be More Stable

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 during December 1st to 7th in 2009 from LINX, we choose an example router and identify its top 50 noisy prefixes (with the largest number of updates). We then check whether these prefixes have alternative routes, and if alternative routes exist, we further calculate the number of updates received on those routes.

For the top 50 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[3] respectively. Also, when calculating the number of updates, we make sure that the noisy prefixes and their alternative routes are both reachable 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 via AS286[4] with 17,635 updates during the week of December 1st. However, this prefix can also be reached using alternative routes either via AS8468, or via a covering prefix $212.42.224.0/19$, with which only 1 and 0 updates observed respectively for the whole week! Overall, the number of updates received on the noisy prefixes is always greater than that of their stable alternatives when they exist, and the difference is significantly large in most cases. Huston *et al.* [8] made similar observation, and conjecture that these noisy prefixes with stable covering prefixes could be the outcome of poor tuning (or not tuning) of the automated traffic engineering process.

In this paper, we focus on the results of LINX for clarity. We also performed the same measurements using routing data collected from other exchange points. In general, we made the similar observation across different exchange points. This observation sheds light on an opportunity to significantly suppress instabilities while preserving the reachability. In the following sections, we derive a practical RFD enhancement and evaluate its performance.

[3]For clarity, we only show the number for the most stable covering prefix and neighbor.

[4]A LINX member router from AS286

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 can be preserved. We emphasize that our work is a complementary addition to all existing rout flap damping schemes, rather than another damping algorithm.

In the previous work, various route flap damping schemes all suppress prefixes immediately once they are found unstable. As a result, it is critical that the algorithms can correctly tell route flapping from path explorations, because false detection leads to false suppression and potential reachability losses. In contrast, our work decouples route *flapping detection* from route *suppression*. One can apply different algorithms to route flapping detection. Once a flapping route is identified, it *must* further pass a reachability test to be eligible for suppression. This approach eliminates reachability losses due to false flapping detections, which may be unavoidable in practice.

4.1 Protector and Protectee

As described in Section 3, it is often the case that a given destination network can be reached through multiple different routes. 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 peers with two neighbor routers 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 [5] and the route to one prefix via a particular peer is denoted as *(prefix, peer)*. In this example, the prefix $P.1$ is a protector prefix of $P.1$[6],$P.1.1$ 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 $(P.1.2, B)$. For ease of discussions, we use $protector(p)$ to represent prefix p's protector prefixes, and $protector(p, r)$ to represent the set of route (p, r)'s protec-

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

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

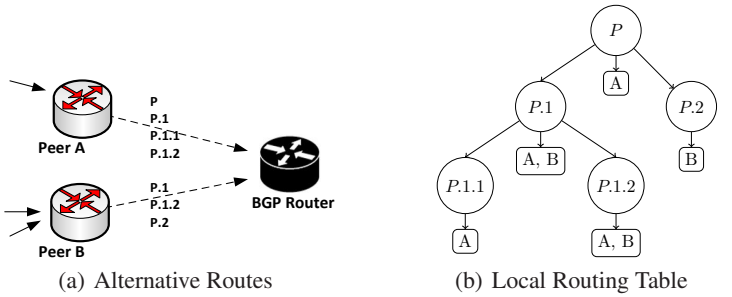

(a) Alternative Routes (b) Local Routing Table

Figure 3: Protector and Protectee Prefixes & Routes.

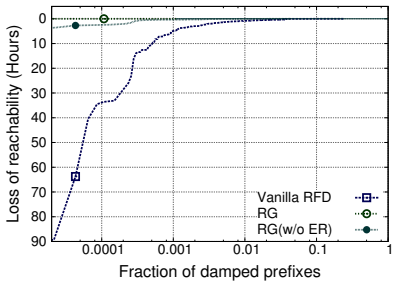

Figure 4: Reachability Loss

tor 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 otherwise specified, *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 check stops and the flapping route (p, r) can be safely suppressed; otherwise route (p, r) is left intact. If route (p, r) flaps persistently, its protector route set, $protector(p, r)$, will be evaluated continuously in the hope that the flapping can eventually be suppressed as soon as a valid protector route appears. 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 [10], 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 such cases in the following section.

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 $P.2$'s reachability. We can safely keep suppressing route $(P.1.2, B)$ since it still has protector routes.

PROPOSITION 1. *Consider a BGP router that enabled route flap damping with reachability guard. The router should not lose reachability caused by damping.*

PROOF. Due to the limited space, we only sketch a proof by contradiction. Consider the router triggered an ordered sequence of events, $\xi = e_1 \cdots e_{t-1} e_t$, upon update arrivals or damping timers expires. In flap damping, one event can be an announcement[7], withdrawal, suppression, or 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 an announcement and reuse events would not hurt reachability, there can only be 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 could not happen. Moreover, case 3 does not lead to reachability loss (*i.e.* suppress a withdrawn route doesn't hurt reachability). □

5. EVALUATION

In this section, we evaluate the effectiveness and impact of RFD+RG by using an event driven simulator, in which we implemented (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. We feed the simulator with BGP update data collected from operational routers by RIPE's BGP collector at the exchange point LINX. We randomly picked the BGP data from the week of December 1st – 7th, 2009 to emulate a realistic scenario, assuming the collector as a new BGP router R in LINX which peers with four other routers at LINX [8], and receive routing updates from the neighbors.

We compare the performance by enabling or disabling route flap damping on R. Unless otherwise specified, the

[7] Announcement or withdrawal that does not trigger damping

[8] We picked four routers with full routing tables and no session resets during that week.

RFD implementation in the simulator uses Cisco default damping parameters [23]. When selecting the best route, we used a simple shortest AS path selection algorithm. To further validate our results, we also performed similar simulation 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 observations made previously, our results show that the Vanilla RFD can significantly impact prefix reachability: 2% (822) of of the damped prefixes lose reachability for more than 30 minutes; the top 50 unstable prefixes are suppressed for more than 5 hours. And for a few worst prefixes, RFD blocks their reachability for even longer than half a week!

On the other hand, with RFD+RG, no prefix ever loses reachability due to damping. With the stripped version (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 existence of shorter and stable prefixes that cover the address space of the damped routers. Nevertheless, RFD+RG can eliminate reachability losses by route damping.

5.2 Reduce Router Load

For the emulated router, Figures 5 and 6 depict the complementary CDF for the number of BGP updates and nexthop changes, the two primary contributors of BGP processing load [20]. In Figure 5, the long tail distribution shows that a small number of prefixes contribute to relatively large number of updates, and thus BGP path changes. This result also conforms to the observations made in [11, 16].

In this 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 reference. We observe that RFD+RG reduces nearly the same amount of updates compared to vanilla RFD without any prefix losing reachability.

Figure 6 depicts the complementary CDF over the number of BGP nexthop changes. In BGP, it is important to minimize nexthop router changes to keep forwarding plane stability. Nexthop changes trigger the updates to interface FIB, and frequent FIB changes can further degrade forwarding performance [22]. Figure 6 shows that RFD+RG reduces the number of nexthop changes by 28% if counting all nexthop changes by damped prefixes, or 21% of nexthop changes for all prefixes. The reduction is an order of magnitude for a few most unstable prefixes. Moreover, RFD+RG only allows a small number of additional nexthop changes than the vanilla RFD to preserve reachability.

Table 1: Summary of Evaluation Results (LINX)

	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. Our simulation results serve to only compare the relative performance, but not necessarily the actual load reduction on real routers.

5.3 Make A Safer Trade-off

As a penalty-based system, the vanilla RFD improves the routing stability with an undesirable tradeoff of of reachability losses [15, 18, 22]. In contrast, RFD+RG makes a conservative trade-off between route *preference* and *stability*: the router can suppress a preferred but flapping route and instead 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. Any nexthop change in less than 10 minutes[9] is considered 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 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 with a cost of longer path.

5.4 Additional Evaluation Results

The results presented so far are based on BGP feeds from one particular exchange point. In this section, we extend the evaluation by using BGP feeds from 7 other exchange points. Table 2 summarizes the evaluation results. We make two observations. First, even though different exchange points are in different topological locations and have different operational environments, RFD+RG enables route flap damping with no reachability losses. 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%, whilst for relatively stable ones, such as Geneva, the damping is triggered less often and the update reduction is lower. We have further examined Geneva's raw BGP feeds and verified that

[9] 10 minutes is the 95% confidence of median Skype call duration [5].

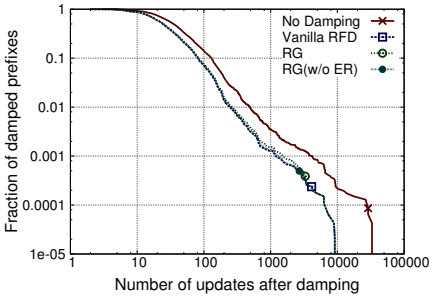

Figure 5: Update Count

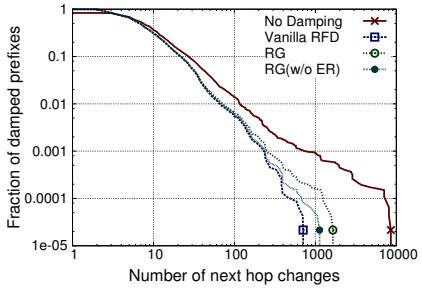

Figure 6: Nexthop Change Count

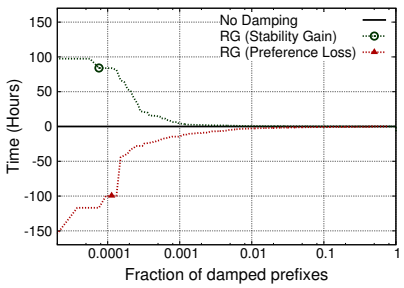

Figure 7: Stability and Preference

Table 2: 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

many noisy prefixes observed by other exchange points are rather stable in Geneva. Overall, RFD+RG is able to reduce the total BGP updates by 5% to 27% o across the exchange points.

6. DISCUSSION AND FUTURE WORK

In this paper, our goal is to demonstrate that a simple addition to RFD can effectively eliminate the loss of reachability. Admittedly, this gain does not come for free. Compared to vanilla RFD, RFD+RG requires additional data structures and computations. First, RFD+RG needs to keep track of the protector-protectee relationship among prefixes. Fortunately, BGP implementations usually organize routing tables using the aforementioned trie structures which already maintain such links between covering and covered prefixes [19, 1]. Second, a router must check multiple routes before making one damping decision. In the worst case, one may traverse all protector or protectee routes. However, as we observed in Section 3, the majority of prefixes have only a fair number of protectors or protectees (mostly less than 3), thus we expect that a reasonable amount of computation load are introduced. More evaluations are necessary to understand the incurred overhead.

A remaining issue is how to handle unstable *orphan* routes, *i.e.* do not have protectors. In this work we follow a simple principal: *never lose reachability*, thus 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 eliminate reachability losses. In this paper we showed that a simple solution can both eliminating reachability losses and removing a significant number of updates generated by unstable prefixes. In the future design, for aggressive users who are willing to sacrifice reachability for more update reduction, a router may damp these orphan routes with some remedial actions, such as shortening the suppression period.

Yet another open issue concerns routing convergence. One limitation in this work is that we only focused on evaluating the damping behavior and reachability losses at specific routers, and the emulation setting does not allow us to measure the impact of +RG on the route convergence time in the global routing system. To quantify such system wide impacts requires 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 flap damping schemes, this work presents a simple addition to RFD to prevent undesirable reachability losses due to route flap damping. The solution is built upon the diversified nature of BGP routing: one can reach a given network destination via multiple paths and different prefixes. The observation itself is not new. Our main contribution is a practical RFD enhancement RFD+RG derived from this observation, and a systematic evaluation of its effectiveness using real BGP data. Our 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 all damping schemes 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.

8. REFERENCES

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

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

[3] 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.

[4] 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.

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

[6] 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.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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