BGP Route Flap Damping Algorithms

Wei Shen and Ljiljana Trajković[*]
Simon Fraser University
Vancouver, British Columbia, Canada
{wshen, ljilja}@cs.sfu.ca

Keywords: Routing protocols, BGP, route flap damping, network stability, ns-2 simulations.

Abstract

Route flap damping (RFD) plays an important role in maintaining the stability of the Internet routing system. It functions by suppressing routes that persistently flap. Several existing algorithms address the issue of identifying and penalizing route flaps. In this paper, we compare three such algorithms: *original RFD*, *selective RFD,* and *RFD+*. We implement these algorithms in ns-2 and evaluate their performance. We also propose possible improvements to the RFD+ algorithm.

1. INTRODUCTION

A route flaps when it exhibits routing oscillations. Route flap damping (RFD) mechanisms [1] are employed by the Border Gateway Protocol (BGP) to prevent persistent routing oscillations caused by network instabilities such as router configuration errors, transient data link failures, and software defects. Their goal is to reduce the number of BGP update messages sent within the network and to decrease the processing load imposed on BGP speakers. Well-designed route flap damping algorithms should not significantly delay the convergence of relatively stable routes.

A common approach in route flap damping is to assign a penalty to a route and increment the penalty value when the route flaps. When the penalty of a route exceeds the threshold *suppress limit*, the route is suppressed and not advertised further. The penalty of a route decays exponentially according to the parameter *half life,* which specifies the time for the penalty to be reduced by half. If the penalty decreases below the threshold *reuse limit*, the route is reused and may be advertised again.

The *original RFD* algorithm, defined in RFC 2439 [1], considers each route withdrawal or route attribute change as a flap and penalizes it accordingly. The pseudo code of the original RFD algorithm is shown in Algorithm 1. It was shown [2] that this approach could significantly delay the convergence of relatively well-behaved routes (routes that flap only occasionally). When a route is withdrawn, BGP searches for feasible alternatives leading to the desired destination. In the case when a particular destination becomes unreachable due to a link failure, a BGP speaker tries other feasible routes to the destination until it finds no alternatives. This path exploration due to a single route withdrawal could cause route suppressions elsewhere, resulting in a significantly delayed convergence.

```
when receiving a route r with prefix d from peer j
if (W(r) and !W(p))
        // W(x) returns true only if x is a withdrawn route
        // p is the previous route with prefix d from peer j
          a flap is identified: route withdrawal
else if (!W(r) and !W(p) and r ≠ p)
          a flap is identified: route attribute change
p = r
```

Algorithm 1. Pseudo code of the original RFD algorithm.

A new algorithm called *selective RFD* was proposed [2] to distinguish path explorations from genuine route flaps. It was observed that selection of routes during path exploration was based on the local preference in a non-increasing order. The selective RFD algorithm specifies that the sender attaches its local preference to each route advertisement. A flap is identified and the penalty value is incremented accordingly if the receiver detects a change of direction in route preference. An example is an increase in the route preference following a decrease. The pseudo code of the selective RFD algorithm is shown in Algorithm 2. Simulations of small networks indicated that selective RFD identifies genuine flaps better than the original RFD [2] algorithm. However, selective RFD does not always identify flaps correctly. Better paths may become available afterwards during path exploration due to topological dependencies and delays in message processing and propagation. This results in non-monotonicity of the route preference during path exploration [3].

```
when receiving a route r with prefix d from peer j
if (W(r) and !W(p))
       // W(x) returns true only if x is a withdrawn route
       // p is the previous route with prefix d from peer j
         tmp = 1
             //indicate a potential flap is temporarily ignored
         remember this potential flap: route withdrawal
else if (!W(r) and !W(p) and dop(r) > dop(p))
```

[*] This research was supported by the NSERC Grant 216844–03 and Canada Foundation for Innovation.

```
    // dop(x) returns the degree of preference of route x
      curBit = 1            // store comparison results
      if (preBit == −1)
          a flap is identified: route attribute change
          if (tmp == 1)
              add the temporarily ignored flap
  else if (!W(r) and !W(p) and dop(r) < dop(p))
      curBit = −1
      if (preBit == 1)
          a flap is identified: route attribute change
          if (tmp == 1)
              add the temporarily ignored flap
  p = r;  preBit = curBit;  tmp = 0
```

Algorithm 2. Pseudo code of the selective RFD algorithm.

A new algorithm, named *RFD+* [3], correctly distinguishes between route flaps and path explorations in the case of an occasional flap. In RFD+, a flap is detected when the current route has a higher degree of preference than the previous one and the BGP speaker has received the current route more than once since its previous flap. The pseudo code of the RFD+ algorithm is shown in Algorithm 3. Simulations of RFD+ in small networks showed that RFD+ could correctly identify route flaps when a single flap occurred [3].

```
  when receiving a route r with prefix d from peer j
  if ($r \notin S(d, j)$)      // S(d, j) is the set of all routes with
                               // prefix d announced from peer j
      insert r into S(d, j)
  else if ($r \in S(d, j)$ and dop(r) > dop(p))
                               // degree of preference of route r is
                               // higher than for the previous route p
          a flap is identified and S(d, j) is cleared
```

Algorithm 3. Pseudo code of the RFD+ algorithm.

In Section 2, we describe the ns-2 implementation of route flap damping algorithms. Simulation results are presented in Section 3. In Section 4, we discuss the performance of selective RFD and RFD+. We also propose possible improvements to the route flap damping algorithms. We conclude with Section 5.

2. ns-2 IMPLEMENTATION OF THE RFD ALGORITHMS

The implementation of the three RFD algorithms is based on ns-BGP [4], a BGP module built for the network simulator ns-2 [5]. The original RFD and selective RFD algorithms were already implemented in the SSFNet BGP-4 v1.5.0 [6]. We ported relevant code from the SSFNet and made necessary modifications. The RFD algorithms were implemented by adding three C++ classes to ns-BGP: *DampInfo*, *ReuseTimer*, and *VecRoutes*. The *DampInfo* class stores the damping structure for a prefix advertised from a peer of a BGP speaker. This class also implements the three RFD algorithms. The *ReuseTimer* class keeps track of the reuse timer associated with a flapping route. The *VecRoutes* class maintains an array of the interim routes in the RFD+ algorithm. We modified several existing ns-BGP C++ classes to implement route flap damping mechanisms when update messages were received and routing decisions were made. Figure 1 shows the routing structure of the modified ns-BGP [4].

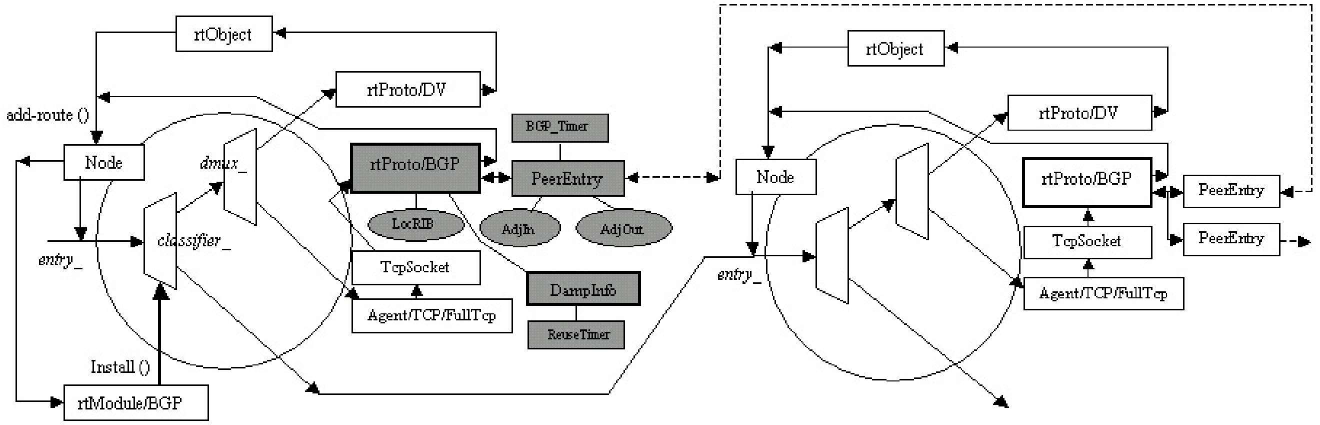

Figure 1. Routing structure of the ns-BGP with route flap damping.

3. PERFORMANCE ANALYSIS OF THE RFD ALGORITHMS

We consider four factors when designing simulation scenarios: network topology, inter-arrival time between updates, total simulation time, and the nature of flaps to simulate.

Network topology: We use the BRITE network topology generator [7] to create more realistic network topologies that also include link delays. We adopt BRITE's Generalized Linear Preference (GLP) model [8] to generate AS-level networks ranging from 100 to 500 nodes. We also use network topologies with 29 and 110 nodes, built from genuine routing tables [9].

Inter-arrival time: For each network topology, we select at least three values for the inter-arrival time between update messages: a value smaller than the default Minimum Route Advertisement Interval (MRAI) of 30 s (10 s or 20 s), an intermediate value (50 s or 100 s), and a value large enough for BGP to converge (500 s or 1000 s).

Scenario simulation time: The duration of a simulation depends on whether or not the scenario contains the route suppression period when BGP nodes wait for routes to become reused and advertised again. By comparing scenarios with and without the suppression period, we can evaluate the impact of route suppression on individual BGP speakers and on the network.

Nature of flaps to simulate: To test the effectiveness of the damping algorithms, we mimic occasional and persistent flaps by using one and five flaps within a certain period of time, respectively. The process is shown in Figure 2.

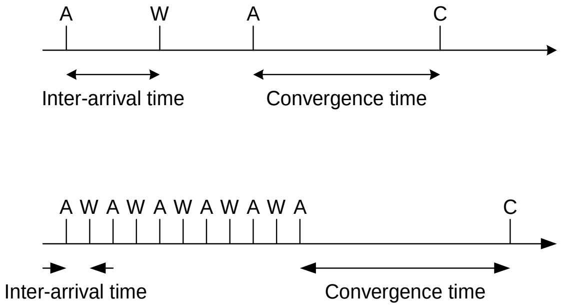

Figure 2. Timeline for occasional flaps (top) and persistent flaps (bottom): A (advertise), W (withdraw), and C (converge).

We use the default MRAI value of 30 s and apply jitter to it. The default Cisco settings for the RFD parameters are adopted, as shown in Table 1. We examine the overall number of updates, overall number of reported flaps, number of flaps reported by each BGP speaker, maximum number of flaps associated with a peer of each BGP speaker, total number of suppressions caused by all the flaps, and convergence time. Unless otherwise specified, convergence time is the time gap between the instance when the origin router announces a route and the instance when the network sees the last update message [10]. The three damping algorithms are compared in cases of occasional and persistent flaps, in networks generated by BRITE and built from genuine routing tables. We also examine advertisement and withdrawal phases and the effects of the inter-arrival time and the location (core or edge) of the origin router.

Table 1. Default CISCO setting for route flap damping.

Suppress limit	2000
Reuse limit	750
Half life (s)	900
Withdrawal penalty	1000
Attribute change penalty	500
Maximum suppression time (s)	3600

3.1 Advertisement and Withdrawal Phases

Due to the lengthy path exploration process after a withdrawal, a withdrawal message causes BGP to converge significantly slower than in the case of an advertisement message. This holds for all damping algorithms.

Table 2 shows the convergence time of the advertisement and withdrawal phases for networks of various sizes. In the case of the 500-node network, the withdrawal phase takes ~20 times longer than the advertisement phase. During the withdrawal phase, original RFD has the fastest convergence because of the most aggressive route suppression. In these simulation scenarios, RFD disabled, selective RFD, and RFD+ algorithms behave identically because for a single flap selective RFD and RFD+ do not cause a sufficient number of route suppressions in specific nodes to affect the convergence time. Simulation results also indicate that the convergence time of the advertisement phase may not vary with the increase of network size. Damping algorithms have little effect on the advertisement phase because advertisement triggered suppressions are rare. They play an important role in the withdrawal phase because withdrawal triggered suppressions are common. The withdrawal phase depends on the network size and the network topology (dense or sparse).

3.2 Effect of the Inter-arrival Time

Table 3 indicates that there is no visible relationship between the inter-arrival time and the BGP convergence time. The increase of inter-arrival time beyond a certain threshold has no effect on the BGP convergence time. If the

Table 2. Occasional flaps: convergence times (in seconds) for the advertisement and withdrawal phases for various networks.

RFD algorithm	Phase	Network size (no. of nodes)				
		100	200	300	400	500
RFD disabled	Advertisement	27.017	27.017	27.017	27.017	27.018
	Withdrawal	216.21	297.31	405.3	594.21	675.3
Original RFD	Advertisement	27.017	27.017	27.017	27.017	27.018
	Withdrawal	189.21	297.31	270.31	486.21	567.21
Selective RFD	Advertisement	27.017	27.017	27.017	27.017	27.018
	Withdrawal	216.21	297.31	405.3	594.21	675.3
RFD+	Advertisement	27.017	27.017	27.017	27.017	27.018
	Withdrawal	216.21	297.31	405.3	594.21	675.3

Table 3. Occasional flaps: effect of inter-arrival time on network performance for a 200-node network (without considering route suppression period).

RFD algorithm	Evaluation parameters	Inter-arrival time (s)						
		10	20	30	50	100	500	1000
RFD disabled	Convergence time (s)	61.017	41.017	48.017	8.017	35.018	27.02	27.02
	No. of updates	1832	1930	2434	2434	3725	8056	8056
	No. of flaps	/	/	/	/	/	/	/
Original RFD	Convergence time (s)	61.017	41.017	48.017	8.017	35.016	0.02	0.02
	No. of updates	1832	1930	2434	2434	3678	7202	7202
	No. of flaps	961	1001	1430	1430	2229	3829	3829
Selective RFD	Convergence time (s)	61.017	41.017	48.017	8.017	35.018	27.02	27.02
	No. of updates	1832	1930	2434	2434	3725	8055	8056
	No. of flaps	487	530	563	563	624	800	802
RFD+	Convergence time (s)	61.017	41.017	48.017	8.017	35.018	27.02	27.02
	No. of updates	1832	1930	2434	2434	3725	8056	8056
	No. of flaps	435	454	491	491	493	497	497

inter-arrival time is not sufficiently large for BGP to converge, the difference between the instances when an update is ready to be sent and when the MRAI timer expires may strongly affect the length of the convergence time. This is illustrated in Table 3 for short inter-arrival times (10 s or 20 s). In the case of a single flap, when the inter-arrival time is short, damping algorithms do not affect the convergence time and the number of updates, as also shown in Table 3. As the inter-arrival time increases, the number of updates and the number of reported flaps increases or remains constant. The number of flaps and route suppressions in RFD+ is the least sensitive to changes in the inter-arrival time. The original RFD algorithm is the most sensitive.

3.3 Location of the Origin Router

When the sender is at the edge of the network, it often takes BGP up to ~20% longer to converge than when the sender is located at the core of the network. This difference in convergence times may increase for certain network topologies. In the 300-node network with original RFD enabled, the convergence time during the withdrawal phase increases by ~50%, as shown in Table 4. The number of generated update messages also increases when the sender is located at the edge of the network. Positioning the origin router at the edge of the network often results in up to ~25% increase in the number of updates compared to placing the origin router at the core of the network. In the 500-node network, this difference in the number of updates during the advertisement phase increases to ~80%. An exception is the withdrawal phase in the 200-node network under original RFD: both the convergence time and the number of updates decrease when the origin router is located at the edge of the network. This suggests that the effect of the origin router's location on the convergence time and the number of updates depends on the network topology, the phase (advertisement or withdrawal), and the damping algorithm.

3.4 Occasional and Persistent Flaps

We simulated the performance of the three damping algorithms under both occasional and persistent flaps.

Original RFD does not perform well in the case of occasional flaps. One flap in the network may result in many network nodes suffering from a significant delay in convergence (ranging from ~20 min to over 1 h) due to the extensive route suppressions. Figure 3 shows that in a network with 200 nodes, ~35% of the nodes suffer from a long (up to ~2000 s) convergence delay because of the route suppressions. In the network with 500 nodes, ~20% of the nodes suffer from a convergence delay up to ~3500 s. Negative values imply that the nodes do not receive the route re-advertisement after withdrawal and will wait until other nodes become reused and start to advertise. Negative values are shown to indicate the number of such nodes. For the 200-node (500-node) network, ~12% (~10%) of the nodes do not receive the route re-advertisement after withdrawal because of the route suppression.

Table 4. Occasional flaps in original RFD: effect of the origin router's location on network performance for various networks during the withdrawal phase.

Location of origin router	Evaluation parameters	Network size (no. of nodes)				
		100	200	300	400	500
Connected to core	Convergence time (s)	189.21	297.31	270.31	486.21	567.21
	No. of updates	2450	6237	8695	16896	26892
Connected to edge	Convergence time (s)	189.21	270.3	405.2	486.21	594.31
	No. of updates	2450	5531	12321	16896	28488

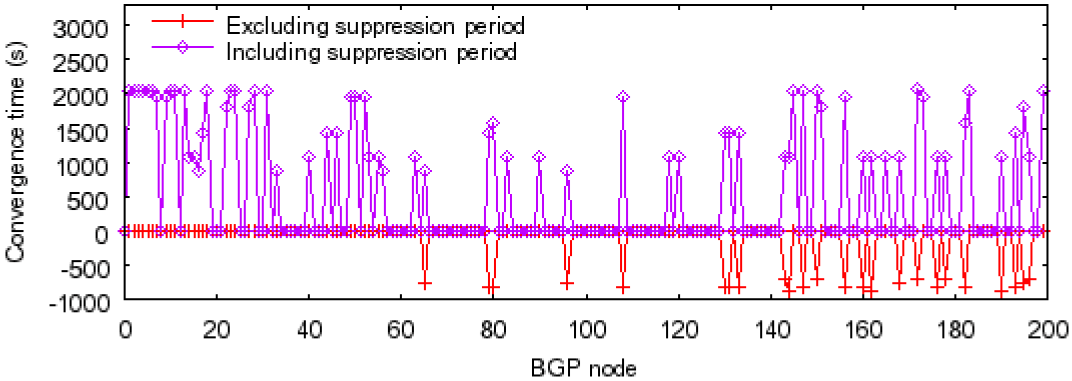

Figure 3. Occasional flaps: impact of route suppression on the convergence time of BGP nodes in a 200-node network.

Under occasional flaps, selective RFD performs better than original RFD in terms of the number of flaps and suppressions. RFD+ achieves the best result and it does not mistake path explorations for route flaps. Table 5 shows the maximum number of flaps reported by a BGP node for one of its peers for various damping algorithms. Original RFD and selective RFD have different degrees of route suppression. There is no route suppression in RFD+. For example, a single flap in a 200-node network under original RFD, selective RFD, and RFD+, causes a maximum of 16, 6, and 1 flap reported on a certain BGP node, respectively. The maximum number of flaps identified by each node for one of its peers is shown in Figure 4.

In the case of persistent flaps, original RFD prevents the spread of routing oscillations in the network as early as possible. A series of 4 flaps within a relatively short time often leads to the suppression of the route. Selective RFD may require the occurrence of additional flaps in order to suppress a flapping route. This number may increase significantly as the inter-arrival time increases. The number of flaps, calculated for the original and selective RFD algorithms, is shown in Table 6. Selective RFD may cause a delay in the route suppression. RFD+ may also underestimate the number of genuine flaps, causing a delay in the route suppression. For example, a peer of the origin router reports only 3 flaps when the origin router experiences a failure (“down” phase) followed by a recovery (“up” phase) for five consecutive times. Figure 5 shows the effects of the three RFD algorithms in the case of persistent flaps in a network with 300 nodes. Route suppression affects all the nodes in original RFD and a small percentage of the nodes in selective RFD. There is no route suppression in RFD+. Selective RFD and RFD+ are more lenient than original RFD in the suppression of persistently flapping routes, which is not a desirable property.

An interesting observation is that the advertisement of a route after reuse may cause new route suppressions in the network, causing a cascading effect.

Table 5. Occasional flaps: maximum number of flaps reported by a BGP node for one of its peers. Network topologies are generated by BRITE or built from genuine routing tables.

RFD algorithm	BRITE (no. of nodes)					Routing tables (no. of nodes)	
	100	200	300	400	500	29	110
Original RFD	10	16	15	22	25	9	23
Selective RFD	3	6	4	5	4	4	7
RFD+	1	1	1	1	1	1	1

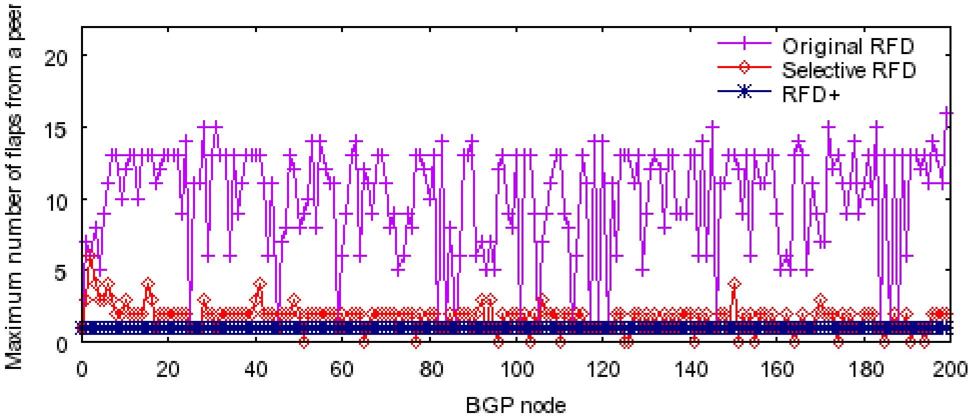

Figure 4. Occasional flaps: maximum number of flaps reported by each BGP node in a network with 200 nodes.

Table 6. Number of flaps required to suppress a route for original RFD and selective RFD.

RFD algorithm	Inter-arrival time (s)												
	100	200	250	270	290	300	310	312	314	316	318	320	321
Original RFD	3	3	3	3	3	3	3	3	4	4	4	4	4
Selective RFD	3	4	4	5	6	6	7	8	8	9	10	12	14

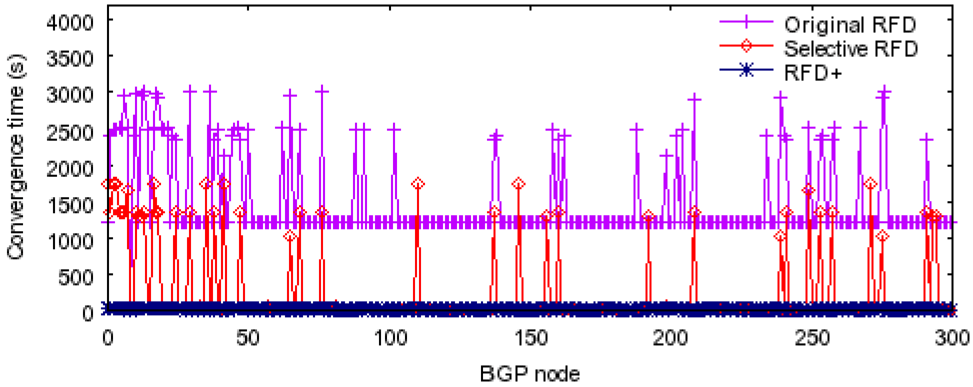

Figure 5. Persistent flaps: effects of various RFD algorithms on a 300-node network. Inter-arrival time is 300 s.

Table 7. Occasional and persistent flaps: total number of updates. Inter-arrival times are 1000 s and 300 s for occasional and persistent flaps, respectively. Network topologies are generated by BRITE or built from routing tables.

Nature of flaps	RFD algorithm	BRITE (no. of nodes)					Routing tables (no. of nodes)	
		100	200	300	400	500	29	110
Occasional	RFD disabled	3407	8056	14126	23344	36642	1209	17621
	Original RFD	2984	7380	10464	19179	29924	1149	11514
	Selective RFD	3407	8056	14126	23344	36642	1211	17616
	RFD+	3407	8056	14126	23344	36642	1209	17621
Persistent	RFD disabled	15935	38052	53789	75674	104236	5109	45607
	Original RFD	8016	11784	30822	53271	78519	1675	18308
	Selective RFD	13251	32535	52852	70924	100726	3291	44568
	RFD+	15850	37569	53452	75205	103892	5106	45694

3.5 Network Stability

Original RFD is optimal in achieving network stability: it always leads to the least number of update messages. The overall number of update messages for various damping algorithms is shown in Table 7. The simulation time is sufficient for BGP to converge after route reuse. For example, in a network with 200 nodes under persistent flaps, original RFD generates ~70% less updates than RFD+. In the 110-node network under persistent flaps, RFD+ generates more updates than the case when RFD is not enabled, as shown in Table 7. This is due to the additional updates caused by the reuse and re-advertisement of suppressed routes in RFD+.

4. IMPROVEMENTS OF RFD ALGORITHMS

Upon receiving a withdrawal message, selective RFD temporarily ignores the update. Nevertheless, it remembers the imposed withdrawal penalty and decays it exponentially. This decayed value is added to the penalty when the next received advertisement is identified as a flap. As a result, selective RFD may require additional flaps in order to suppress a flapping route. This situation worsens when the inter-arrival time between updates increases, as it was shown in Table 6. Based on the default Cisco RFD setting, selective RFD will not suppress any route if the inter-arrival time between route updates is larger than 322 s.

RFD+ underestimates the number of real flaps in the case of persistent flaps. RFD+ treats a series of 5 updates (*advertisement, withdrawal, re-advertisement, withdrawal,* and *re-advertisement again*) as only one flap rather than two flaps. The last two updates are not sufficient to cause an additional flap. As a result, a peer of the origin router reports only *floor((N+1)/2)* number of flaps if the origin router experiences a failure ("down" phase) followed by a recovery ("up" phase) for *N* consecutive times. We propose a simple remedy to this problem by keeping track of the existence of "up-down-up" state of a route, which has also been implemented in selective RFD. A flap is identified either when detected by RFD+ or when a route is advertised, withdrawn, and advertised again. The pseudo code for the modified RFD+ is shown in Algorithm 4. With this simple modification, a peer of the origin router could identify all *N* flaps when the origin router fails and then recovers for *N* consecutive times. Table 8 shows the comparison of the RFD+ algorithm and its modified version in terms of convergence time and total numbers of updates, flaps, and suppressions. The inter-arrival time between updates is 300 s. *Modified RFD+* has a much longer convergence time because the flapping route is suppressed by the peer of the origin router in all cases. This is not the case for RFD+, where nodes do not suffer from suppression of routes except in the 200-node network. As a result of the route suppression, the total number of updates is reduced (by up to ~19%) in the modified RFD+ algorithm. The suppression of a persistently flapping route is the desired behavior. Simulation results also suggest that in rare cases the modified RFD+ algorithm may cause a BGP node to report additional flaps.

```
when receiving a route r with prefix d from peer j
if (W(r))
            // W(x) returns true only if x is a withdrawn route
    preUpdate = 0     // remember the update type
                      // 0: indicates withdrawal and
                      //1: indicates advertisement
else                // current route r is an advertisement
    if (preUpdate == 0 and dop(r) == preDop)
        // 'up-down-up' state is detected
        // dop(x) returns the degree of preference of route x
            a flap is identified
            clear R(d, j)          // R(d, j) is the set of all
    else if (r ∉ R(d, j))          // routes with prefix d
                                   // announced from peer j
            insert r into R(d, j)
    else if (r ∈ R(d, j))
            if (preUpdate == 0)
                a flap is identified
                clear R(d, j)
            else               // preUpdate == 1
                if (dop(r) > preDop)
                    a flap is identified
                    clear R(d, j)
    preDop = dop(r)      // remember the degree of
    preUpdate = 1        // preference and update type
                         // of route r
```

Algorithm 4. Pseudo code of the modified RFD+ algorithm.

Another pitfall of RFD+ is its potentially large memory consumption because a BGP speaker needs to store all the interim routes during path exploration for each prefix from each peer. This memory consumption may reach several gigabytes for a core Internet router. One way to reduce memory consumption is to hash each interim route into a simpler data type (e.g., integer) and store it rather than storing the complete route. Hashing may also reduce a router's processing time since comparing two integers is faster than comparing two routes. It is effective because RFD+ requires a large number of comparisons between routes.

While original RFD is designed to work well with persistent flaps, selective RFD and RFD+ perform better in the case of occasional flaps. There is a trade-off between the stability and availability of routes. Good availability of a route implies no route suppression and, hence, either no damping or a rather "lenient" damping algorithm. This causes more update messages to be generated and, hence, results in network instability. In terms of the convergence time, stability demands more "aggressive" damping algorithms for the reduction of generated update messages. This results in route suppression and delayed convergence. To achieve a balance between stability and availability, an adaptive RFD algorithm is desired.

A simple adaptive RFD algorithm named *combined RFD* integrates the original RFD and the modified RFD+ algorithms. Within a certain period of time, a BGP speaker uses the modified RFD+ algorithm for the first two identified flaps. It then switches to the original RFD algorithm starting with the third flap. The motivation is not to suppress a route anywhere if it flaps only once or twice. However, when a route flaps the third time within a certain period, damping starts to become "aggressive". Table 9 shows the simulation results when an existing route becomes unavailable and then available repeatedly for 8 consecutive times. The inter-arrival time between updates is set to 120 s. The total number of update messages in the case of combined RFD is reduced by up to ~16%, compared to the modified RFD+ algorithm. The combined RFD algorithm also tends to generate fewer updates than selective RFD and RFD+. In most cases, combined RFD generates less than 7% of additional updates compared to the original RFD algorithm. However, unlike original RFD, it does not suppress a route that flaps occasionally.

Table 8. Persistent flaps: comparison between the original RFD+ and the modified RFD+ algorithms in the case of 5 flaps.

Algorithm	Evaluation parameters	Network size (no. of nodes)				
		100	200	300	400	500
Modified RFD+	Convergence time (s)	1555.71	1555.71	1555.71	1555.71	1555.71
	No. of updates	12805	30541	45859	68038	98194
	No. of flaps	692	1335	2016	2684	3417
	No. of suppressions	5	22	31	22	23
RFD+	Convergence time (s)	27.02	1369.21	27.02	51.02	51.01
	No. of updates	15850	37569	53452	75205	103892
	No. of flaps	856	1642	2557	3391	4417
	No. of suppressions	44	91	100	121	201

Table 9. Comparison of various RFD algorithms when a route flaps 8 times. Inter-arrival time between updates is 120 s.

Algorithm	Evaluation parameters	Network size (no. of nodes)				
		100	200	300	400	500
RFD disabled	Convergence time (s)	42.02	42.02	15.21	42.1	42.02
	No. of updates	15675	30904	45823	59441	78906
	No. of flaps	/	/	/	/	/
	No. of suppressions	/	/	/	/	/
Original RFD	Convergence time (s)	2908.06	3067.35	3426.2	3705.34	3910.4
	No. of updates	7519	16062	28054	37699	53436
	No. of flaps	4553	10235	19103	25563	36896
	No. of suppressions	261	503	815	1048	1389
Selective RFD	Convergence time (s)	2254.92	2254.92	2254.92	2254.92	2254.92
	No. of updates	8468	16207	33859	40549	60068
	No. of flaps	1874	3803	5952	7610	10479
	No. of suppressions	178	390	663	857	1161
RFD+	Convergence time (s)	1349.38	1349.38	1349.38	1349.38	1502.64
	No. of updates	14344	28673	50167	66713	87916
	No. of flaps	1100	2154	3329	4326	5729
	No. of suppressions	66	120	148	152	238
Modified RFD+	Convergence time (s)	2374.92	2374.92	2374.92	2374.92	2374.92
	No. of updates	9190	19347	30958	45119	65630
	No. of flaps	676	1270	1975	2552	3300
	No. of suppressions	5	17	15	16	15
Combined RFD	Convergence time (s)	2271.41	2271.41	2271.41	2971.71	3583.96
	No. of updates	8202	17204	28608	37716	55477
	No. of flaps	2006	4622	9012	12137	19476
	No. of suppressions	254	499	788	1043	1339

In addition to switching between different RFD algorithms, it may also be useful to dynamically change certain RFD parameters, such as attribute change penalty and half life. For example, there are occasions when a suppressed route needs to be reused earlier rather than wait for a minimum of ~20 min.

5. CONCLUSIONS

In this paper, we compared three algorithms: original RFD, selective RFD, and RFD+. Simulation results suggested that no algorithm performs optimally under all circumstances. Original RFD is more efficient than selective RFD and RFD+ in suppressing persistently flapping routes and achieving network stability. However, it can cause significant convergence delay in the case of occasional flaps. RFD+ has the advantage of reporting no additional flaps. Nevertheless, it may underestimate the number of genuine flaps in the case of persistent flaps. This delays the suppression of flapping routes, which also exists in selective RFD. Selective RFD and RFD+ are more lenient in suppressing flapping routes. Future implementations need to consider a compromise between the original RFD and RFD+ algorithms, with balancing between network stability and availability of routes.

REFERENCES

[1] C. Villamizar, R. Chandra, and R. Govindan, "BGP route flap damping," *IETF RFC 2439*, Nov. 1998.

[2] Z. Mao, R. Govindan, G. Varghese, and R. Katz, "Route flap damping exacerbates Internet routing convergence," in *Proc. SIGCOMM 2002*, Pittsburgh, PA, Aug. 2002, pp. 221–233.

[3] Z. Duan, J. Chandrashekar, J. Krasky, K. Xu, and Z. Zhang, "Damping BGP route flaps," in *Proc. IPCCC 2004*, Phoenix, AZ, Apr. 2004, pp. 131–138.

[4] T. D. Feng, R. Ballantyne, and Lj. Trajković, "Implementation of BGP in a network simulator," in *Proc. ATS'04*, Arlington, VA, Apr. 2004, pp. 149–154.

[5] ns-2 [Online]. Available: http://www.isi.edu/nsnam/ns.

[6] SSFNet [Online]. Available: http://www.ssfnet.org/.

[7] BRITE [Online]. Available: http://www.cs.bu.edu/brite.

[8] T. Bu and D. Towsley, "On distinguishing between Internet power law topology generators," in *Proc. INFOCOM 2002*, New York, NY, June 2002, pp. 638–647.

[9] Multi-AS topologies from routing tables [Online]. Available: http://www.ssfnet.org/Exchange/gallery/asgraph.

[10] T. G. Griffin and B. J. Premore, "An experimental analysis of BGP convergence time," in *Proc. ICNP 2001*, Riverside, CA, Nov. 2001, pp. 53–61.