Performance Evaluation of Border Gateway Protocol with Route Flap Damping and Routing Policies

Ravinder Paul
rpa28@sfu.ca
Communication Networks Laboratory
http://www.sfu.ca/~ljilja/cnl
School of Engineering Science
Simon Fraser University

Roadmap

- Introduction
- Routing polices
- Route Flap Damping algorithms
- ns-2 implementation
- Performance evaluation
- Conclusions
- References

Introduction

- BGP-4 is the current de-facto inter AS routing protocol
- Exchanges network reachability information among BGP routers
- Supports CIDR
- Peer routers exchange four types of messages:
  - open
  - update
  - notification
  - keepalive
- BGP utilizes a path vector algorithm called the best path selection algorithm to select the best path

BGP : Border Gateway Protocol
AS : Autonomous System
CIDR: Classless Inter-Domain Routing

Important BGP terms

- Prefix: the 32-bit IP address block
- Route: a path to a particular destination
- Route preference: a metric indicating the degree of preference of a route

IP: Internet Protocol

Routing policies

- BGP allows a user to configure peer-to-peer policies
- May cause persistent route oscillations that affect the BGP stability
- Routing policies affect:
  - BGP convergence time
  - number of updates
  - number of flaps
- RFD mechanisms prevent BGP route oscillations caused by network instabilities

RFD: Route Flap Damping

Route Flap Damping (RFD)

- A route flaps when it oscillates from an available route state to an unavailable route state
- Routing oscillations may be caused by:
  - router configuration errors
  - transient data link failures
  - software defects
  - routing policies
- BGP employs RFD mechanisms to:
  - prevent persistent routing oscillations
  - reduce the number of BGP update messages
  - decrease the processing load

BGP convergence

- BGP convergence time is the time interval between the first update message sent, until all update messages that are a consequence of the original update received
- Mismatch in policy configurations between two ASes may also cause network instabilities and may increase the BGP convergence time
- Router or link failure may increases BGP convergence time
- The shortest path to destination decreases BGP convergence time

• N. Laskovic and Lj. Trajkovic, "BGP with an adaptive minimal route advertisement interval," in *Proc. IPCCC*, Phoenix, AZ, USA, Apr. 2006, pp. 142–151.

• G. Huston, M. Rossi, and G. Armitage, "A technique for reducing BGP update announcements through path exploration damping," *IEEE Journal on Selected Areas in Communications*, vol. 28, no. 8, pp. 1271–1286, Oct. 2010.

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

Motivation

- BGP problems:
  - route oscillations
  - longer convergence time due to BGP routing policies
  - large number of update messages
- Possible solutions employ:
  - RFD mechanism with modifications
  - routing policies suitable to network topology

Contributions

- Upgrade the RFD module from ns-2.27 to ns-2.34 version
- Implement BGP routing policies (ns-BGP-RP):
  - AS-path list
  - Community-path list
- RFD comparison of random graph with AS-level graph
- RFD modifications:
  - overcome unnecessary route suppressions
  - ensure better performance with large number of nodes

Roadmap

- Introduction
- Routing polices
- Route Flap Damping algorithms
- ns-2 implementation
- Performance evaluation
- Conclusions
- References

Routing polices

- BGP routing policies may be classified into:
  - customer-provider
  - peer-to-peer
  - sibling-to-sibling

Peer Type / Router / Origin	Customer	Peer	Provider	Sibling
Self	yes	yes	yes	yes
Customer	yes	yes	yes	yes
Peer	yes	no	no	yes
Provider	yes	no	no	yes
Sibling	yes	no	no	yes

• B. Premore, An analysis of convergence properties of the border gateway protocol using discrete event simulation, Ph. D. Thesis, Dartmouth College, 2003.

• G. Huston, "Interconnection, peering, and settlements," in *Proc. INET,* San Jose, CA, USA, June 1999, pp. 2-29.

Routing policies filters

- Commonly used routing policy filters are:
  - AS-path list
  - Community-path list
  - Prefix list
- Set up between two local ASes on the basis of:
  - traffic exchange
  - financial benefit
- ASes usually prefer to send traffic to the peer ASes having the same policy

Roadmap

- Introduction
- Routing polices
- Route Flap Damping algorithms
- ns-2 implementation
- Performance evaluation
- Conclusions
- References

RFD algorithms

- Assign a penalty to a route
- Route is suppressed when penalty exceeds the maximum suppression value
- Penalty of a route decays exponentially based on the half life parameter
- When penalty decreases below reuse limit, the route may be re-advertised

RFD algorithms

- RFD algorithms that identify flaps and penalize route flaps are:
  - Original RFD
  - Selective RFD
  - RFD+
  - Modified RFD+

• Original RFD: C. Villamizar, R. Chandra, and R. Govindan, "BGP route flap damping," *IETF RFC 2439*, Nov. 1998.
• Selective RFD: 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.
• RFD+: 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.
• Modified RFD+ algorithm: W. Shen and Lj. Trajkovic, "BGP route flap damping algorithms," in *Proc. SPECTS'05*, Philadelphia, PA, July 2005, pp. 488–495.

Original RFD algorithm

- Original RFD algorithm was proposed in RFC 2439
- Each route withdrawal or route attribute change (route replacement) is considered to be a flap and is penalized accordingly

Message type	AS path
Advertisement	1 3 5
Advertisement	1 3 5 7
Withdrawal	

Flap

- May suppress legitimate route due to single announcement
- Slows down the convergence
- Causes service providers to turn off the RFD feature

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

Selective RFD algorithm

- Sender side attaches route preference to each route advertisement
- Receiver side compares the current route with previous route in terms of route preference
- Flap is identified if a change of direction in route preference is detected (a decrease followed by an increase)

Message type	AS path
Advertisement	1 3 5
Advertisement	1 3 5 7
Withdrawal	

⟸ Flap

- Better path may be identified with decreased route preference as a flap

• 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.

RFD+ algorithm

- Overcomes the problem of the Selective RFD algorithm
- Flap is identified:
  - current route preference is compared with the previously announced route preference
  - BGP speaker has received the current route more than once since its previous announcement

Message type	AS path
Advertisement	1 3 5
Withdrawal	
Advertisement	1 3 5

← Flap

- RFD+ better identifies the path exploration of BGP and route flaps

• 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.

Modified RFD+ algorithm

- A series of advertisement, withdrawal, re-advertisement, withdrawal, and re-advertisement, the modified RFD+ identify this event as two flaps

Message type	AS path
Advertisement	1 3 5
Withdrawal	
Advertisement	1 3 5 7
Withdrawal	
Advertisement	1 3 5 7 9

⟸ 2 Flaps

- Does not significantly increase the BGP convergence
- Modified RFD+ may be used in the network to identify genuine route flaps

• W. Shen and Lj. Trajkovic, "BGP route flap damping algorithms," in *Proc. SPECTS '05*, Philadelphia, PA, July 2005, pp. 488–495.

Roadmap

- Introduction
- Routing polices
- Route Flap Damping algorithms
- ns-2 implementation
- Performance evaluation
- Conclusions
- References

Ns-2 implementation

- Based on a BGP model developed for network simulator ns-BGP 2.0
- ns-BGP 2.0 was ported from SSFNET
- RFD-AMRAI was added in ns-BGP 2.0 and updated to ns-2.27 version
- BGP routing policies are added in ns-BGP 2.0 and updated to ns-2.34
- AS-path list and Community-path list are implemented in ns-BGP-RP
- Improvements to RFD algorithms are made in suppressing routes

• SSFNet: http://www.ssfnet.org/
• ns-BGP 2.0: http://www.ensc.sfu.ca/~ljilja/cnl/projects/BGP/
• RFD-AMRAI BGP: http://www.ensc.sfu.ca/~ljilja/cnl/projects/RFD-AMRAI/

Implementation: ns-BGP-RP

- Routing structure of a modified ns-2 node:
  - forwarding plane: categorizes the received packets whether to be processed or forwarded to neighboring nodes
  - control plane: controls computation, maintenance, and implementation of routes in routing tables

Implementation: ns-BGP-RP

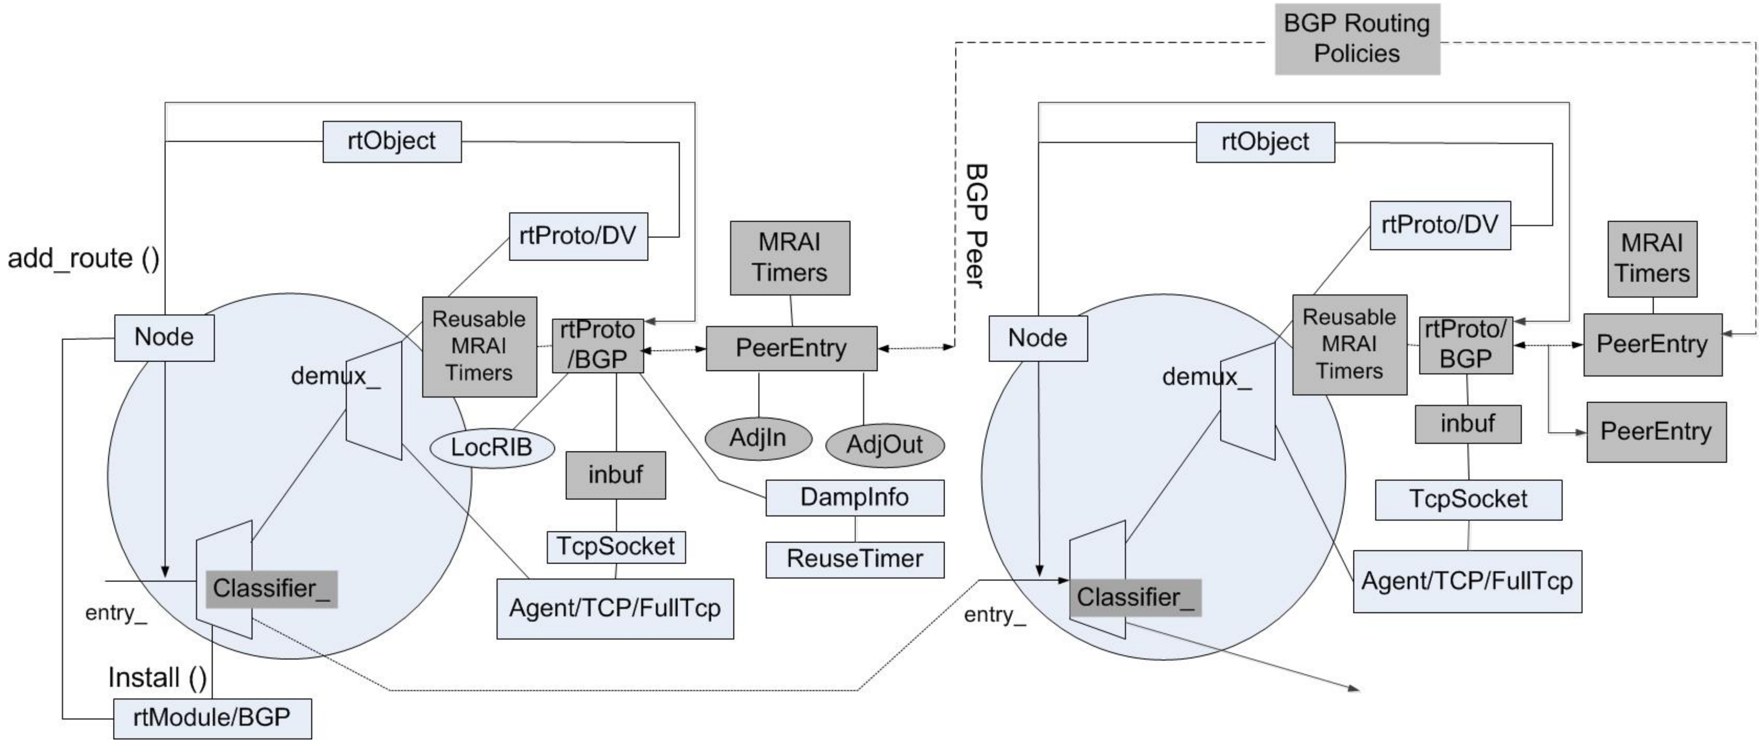

Implementation: ns-BGP-RP

- ns-BGP-RP configures AS-path list and Community-path list routing policies by using regular expressions
- Uses tre-0.8.0 library for regular expressions
- AS-path list:
  - filters the BGP AS path attributes that define the entire set of AS numbers
  - uses a regular expression string to identify the attribute pattern to deny or permit the list

Implementation: ns-BGP-RP

- Community-path list:
  - numbered or named
  - identifies and filters the routes according to common attributes between two networks

```
when receiving a route r with prefix d from neighbor
if (routingPolicy = 1)
        // routing policy of source matches with neighbor
update RIB
        // update the Routing Information Base
else if (routingPolicy = 0)
        // routing policy of source does not matches with neighbor
update RIB
        // update the Routing Information Base
```

Implementation: ns-BGP-RP

- Modifications to make RFD algorithms less aggressive with maximum suppression values of:
  - 2,000
  - 4,000
  - 6,000

• Route flap damping considered useable [Online]. Available: http://www.ripe.net/ripe/mail/archives/routing-wg/2012-July/002163.html

Simulation scenarios: performance comparisons

- Routing policies:
  - BGP without policies
  - BGP with AS-path list policy
  - BGP with Community-path list policy
- RFD comparison with various topologies created by two topology generators
- RFD simulation with different modified maximum suppression values

Topology from BCNET BGP Routing Information Base (RIB)

Origin AS 271 Local Preference 100 From 207.23.253.2 Prefix 1.0.4.0/22 to 223.255.254.0/24		Date 2-11-2011 To 216.6.50.9
Time	AS-path	Updates: A (advertise) and W (withdraw)
09:35:16	6327 7473 38040 9737 56120 I	A
03:18:40	6327 15412 18101	A
03:18:40		W
19:51:45	6327 9498 45528 45528 45528 45528 I	A
19:51:45	6327 9498 45528 I	A
06:16:06	6327 15412 18101 45528 I	A
03:55:53	6327 3549 55410 45528 I	A
09:33:16	6327 1273 37986 24186 45528 I	A
17:07:45	6327 1273 37986 24186 45528 I	A
09:33:16	6327 15412 18101 45528 I	A
02:51:47	6327 9498 9730 45528 I	A

RIB: Routing Information Base

GT-ITM topologies

- Number of nodes in a generated topology is calculated as:

$$N = T*N_t*[1+(K*N_s)]$$

N : number of nodes

T : fully connected transit domain

N_t: average number of nodes per transit AS

K : average number of stub ASes per transit AS

N_s: average number of nodes per stub AS

Symbols	100-node topology	200-node topology	300-node topology	500-node topology
N	100	200	300	500
T	1	1	1	1
N_t	4	8	12	20
K	3	4	4	3
N_s	8	6	6	8

GT-ITM: Georgia Tech Internetwork Topology Models

BRITE topologies

- BRITE generates different types of Internet topologies based on various models
- Generates AS-level topologies

Parameters	Values
Node placement	Random
Growth type (how nodes join in topology)	Incremental
Preferential connectivity	On
Bandwidth distribution	Constant
Alpha (GLP-specific exponent)	0.45
Beta (GLP-specific exponent)	0.65
Size of high level square	Incremental
N (number of nodes)	100, 200, 300, or 500

BRITE: Boston University Representative Internet Topology Generator

Roadmap

- Introduction
- Routing polices
- Route Flap Damping algorithms
- ns-2 implementation
- Performance evaluation
- Conclusions
- References

Performance evaluation: simulation of BGP routing policies

- Network topologies:

Topology	Number of nodes	Topology generator
Topology 1	67	Manually, from BCNET BGP traffic
Topology 2	300	GT-ITM
Topology 3	500	GT-ITM

GT-ITM: Georgia Tech Internetwork Topology Models

Performance evaluation: network Topology 1 (67 nodes)

- Comparison of BGP module with and without policies: convergence time, number of updates, number of flaps, and number of suppressed routes

Scenario	Convergence time (s)	Number of updates	Number of flaps	Number of suppressed routes
BGP without policies	129.29	745	302	0
BGP-RP with AS-path list	130.65	1,102	437	1
BGP-RP with Community-path list	147.13	970	430	1

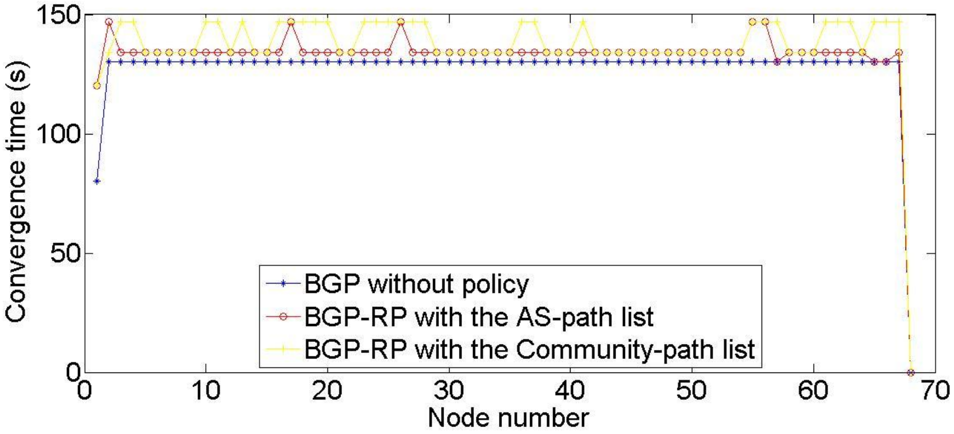

Performance evaluation: network Topology 2 (300 nodes)

- Comparison of BGP module with and without policies: convergence time, number of updates, number of flaps, and number of suppressed routes

Scenario	Convergence time (s)	Number of updates	Number of flaps	Number of suppressed routes
BGP without policies	1,109.59	27,113	16,062	977
BGP-RP with AS-path list	1,133.48	28,366	15,879	977
BGP-RP with Community-path list	1,157.25	28,575	20,257	979

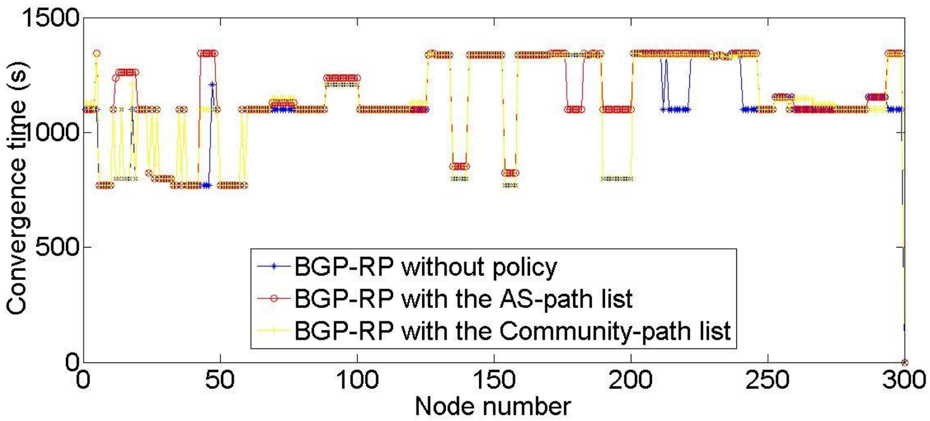

Performance evaluation: network Topology 3 (500 nodes)

- Comparison of BGP module with and without policies: convergence time, number of updates, number of flaps, and number of suppressed routes

Scenario	Convergence time (s)	Number of updates	Number of flaps	Number of suppressed routes
BGP without policies	498.92	24,822	19,225	1,308
BGP-RP with AS-path list	522.19	28,109	18,552	1,288
BGP-RP with Community-path list	562.55	28,446	20,258	1,314

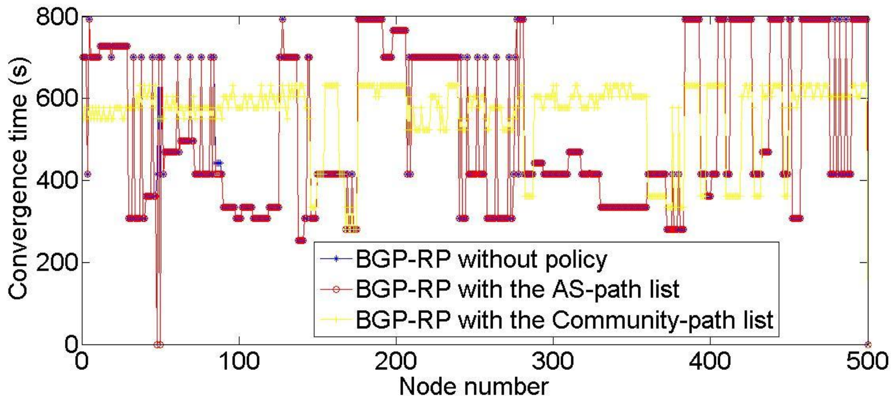

Performance evaluation: RFD analysis with GT-ITM and BRITE topologies

- Network topologies:

Topology	Number of nodes	Topology generator
Topology 1	100	GT-ITM and BRITE
Topology 2	200	GT-ITM and BRITE
Topology 3	300	GT-ITM and BRITE
Topology 4	500	GT-ITM and BRITE

Performance evaluation: network Topology 1 (100 nodes)

- Comparison of BGP convergence time, number of updates, number of flaps, and number of suppressed routes

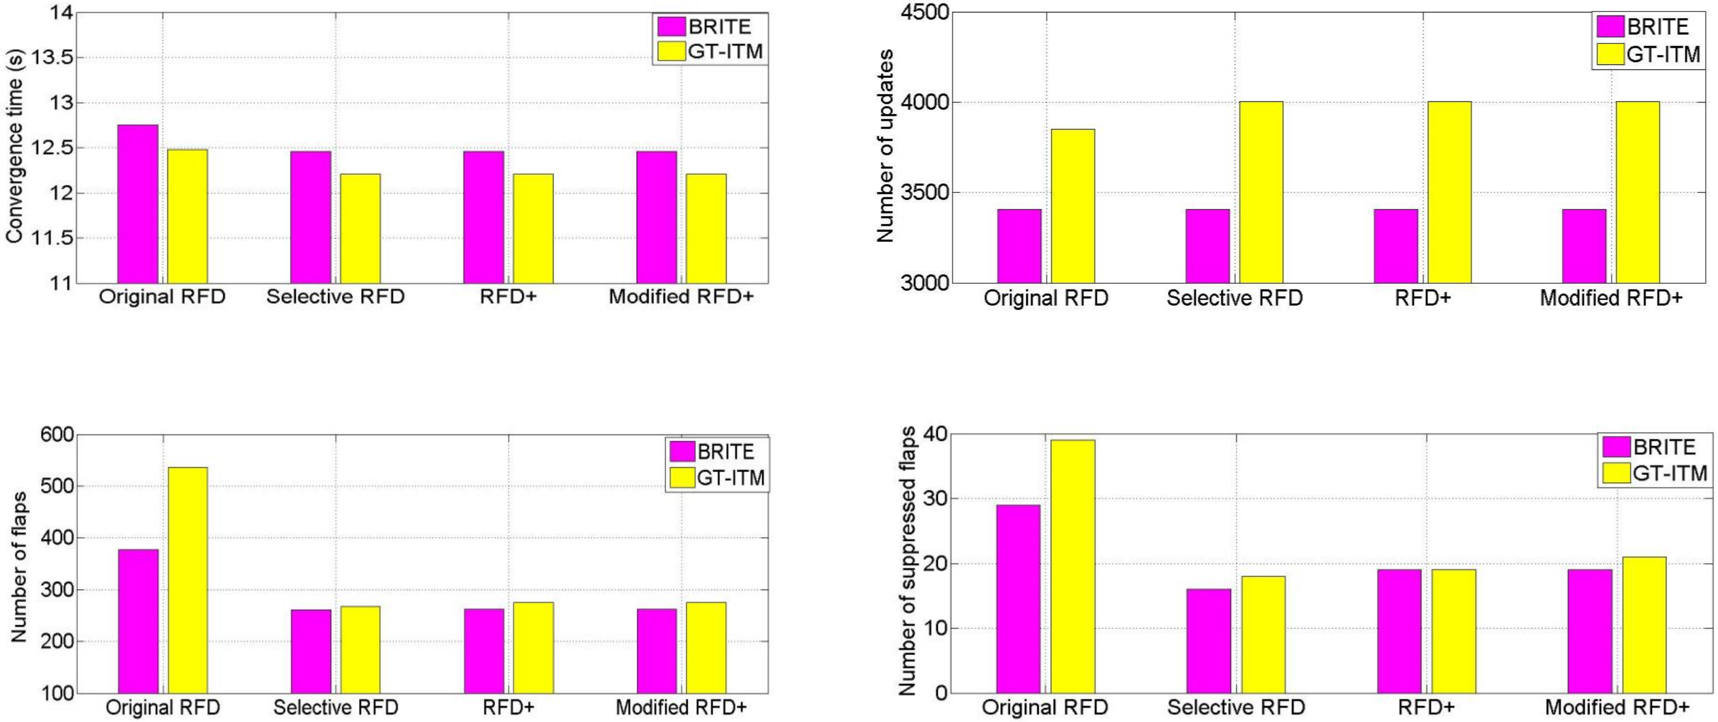

Performance evaluation: network Topology 2 (200 nodes)

- Comparison of BGP convergence time, number of updates, number of flaps, and number of suppressed routes

Algorithm	Convergence time (s)		No. of updates		No. of flaps		No. of suppressed flaps	
	BRITE	GT-ITM	BRITE	GT-ITM	BRITE	GT-ITM	BRITE	GT-ITM
Original RFD	1,351.42	1,347.85	8,054	16,702	871	2,286	142	304
Selective RFD	1,351.42	1,349.19	8,056	16,944	71	497	16	78
RFD+	1,351.42	1,363.37	8,056	16,944	71	499	21	81
Modified RFD+	1,351.42	1,363.37	8,056	16,944	101	499	27	81

Performance evaluation: network Topology 3 (300 nodes)

- Comparison of BGP convergence time, number of updates, number of flaps, and number of suppressed routes

Algorithm	Convergence time (s)		No. of updates		No. of flaps		No. of suppressed flaps	
	BRITE	GT-ITM	BRITE	GT-ITM	BRITE	GT-ITM	BRITE	GT-ITM
Original RFD	961.92	956.52	14,126	21,848	1,286	2,791	117	224
Selective RFD	961.92	966.90	14,126	22,852	113	813	13	89
RFD+	961.92	966.90	14,126	22,852	113	843	17	86
Modified RFD+	961.92	966.90	14,126	22,948	137	843	21	88

Performance evaluation: network Topology 4 (500 nodes)

- Comparison of BGP convergence time, number of updates, number of flaps, and number of suppressed routes

Algorithm	Convergence time (s)		No. of updates		No. of flaps		No. of suppressed flaps	
	BRITE	GT-ITM	BRITE	GT-ITM	BRITE	GT-ITM	BRITE	GT-ITM
Original RFD	578.14	580.29	22,283	36,494	1,431	1,459	58	76
Selective RFD	578.14	578.73	22,283	36,332	158	927	24	34
RFD+	578.14	578.73	22,283	36,332	158	927	29	38
Modified RFD+	578.14	578.73	22,283	36,332	170	957	33	40

Performance evaluation: RFD analysis with various suppress values

- Network topologies :

Topology	Number of nodes	Topology generator
Topology 1	67	Generated from the BCNET traffic routes
Topology 2	200	GT-ITM
Topology 3	300	GT-ITM
Topology 4	500	GT-ITM

- Comparison of number of flaps with Original RFD, RFD+ , and Modified RFD+ algorithms

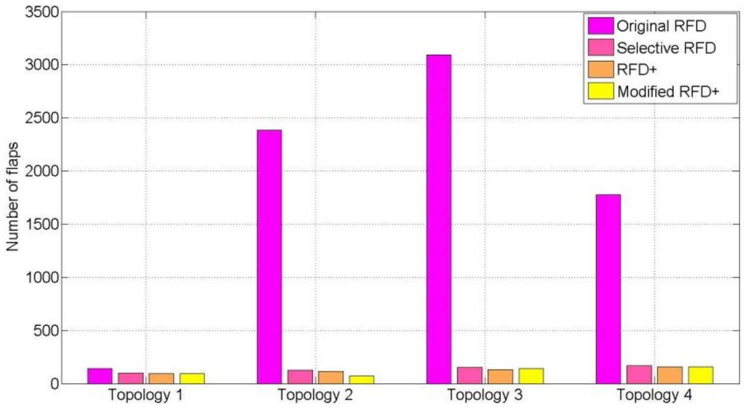

Performance evaluation: network Topology 1 (67 nodes)

- Number of flaps suppressed with maximum suppression value 4,000

Algorithm	Convergence Time (s)	No. of updates	No. of flaps suppressed
Original RFD	132.660	745	2
Selective RFD	131.431	745	1
RFD+	130.256	745	1
Modified RFD+	130.256	745	1

- Number of flaps suppressed with maximum suppression value 6,000

Algorithm	Convergence time (s)	No. of updates	No. of flaps suppressed
Original RFD	132.660	745	0
Selective RFD	131.431	745	0
RFD+	130.256	745	0
Modified RFD+	130.256	745	0

Performance evaluation: network Topology 2 (100 nodes)

- Number of flaps suppressed with maximum suppression value 4,000

Algorithm	Convergence Time (s)	No. of updates	No. of flaps suppressed
Original RFD	1,402.02	16,957	170
Selective RFD	1,401.50	16,957	28
RFD+	1,400.02	16,957	4
Modified RFD+	1,400.02	16,957	7

- Number of flaps suppressed with maximum suppression value 6,000

Algorithm	Convergence time (s)	No. of updates	No. of flaps suppressed
Original RFD	1,402.02	16,957	0
Selective RFD	1,401.50	16,597	0
RFD+	1,400.02	16,957	0
Modified RFD+	1,400.02	16,957	0

Performance evaluation: network Topology 3 (300 nodes)

- Number of flaps suppressed with *maximum suppression value* 4,000

Algorithm	Convergence Time (s)	No. of updates	No. of flaps suppressed
Original RFD	956.92	22,857	209
Selective RFD	966.47	22,857	82
RFD+	966.91	22,857	8
Modified RFD+	966.91	22,857	10

- Number of flaps suppressed with *maximum suppression value* 6,000

Algorithm	Convergence time (s)	No. of updates	No. of flaps suppressed
Original RFD	956.92	22,857	0
Selective RFD	966.47	22,857	0
RFD+	966.91	22,857	0
Modified RFD+	966.91	22,857	0

Performance evaluation: network Topology 4 (500 nodes)

- Number of flaps suppressed with *maximum suppression value* 4,000

Algorithm	Convergence Time (s)	No. of updates	No. of flaps suppressed
Original RFD	580.29	36,332	39
Selective RFD	579.34	36,332	26
RFD+	578.73	36,332	12
Modified RFD+	578.73	36,332	13

- Number of flaps suppressed with *maximum suppression value* 6,000

Algorithm	Convergence time (s)	No. of updates	No. of flaps suppressed
Original RFD	580.29	36,332	4
Selective RFD	579.34	36,332	2
RFD+	578.73	36,332	1
Modified RFD+	578.73	36,332	0

Performance evaluation: BCNET

- We examined the routing table of AS 271 (BCNET)
- If RFD was enabled on 2-11-2011, then the number of flaps identified by RFD algorithms would be:
  - Original RFD algorithm:19 flaps
  - Selective RFD algorithm: 4 flaps
  - RFD+ algorithm: 2 flaps
  - Modified RFD+ algorithm: 5 flaps
- AS path lengths of BGP routes between 2 and 4

BCNET relationship graph

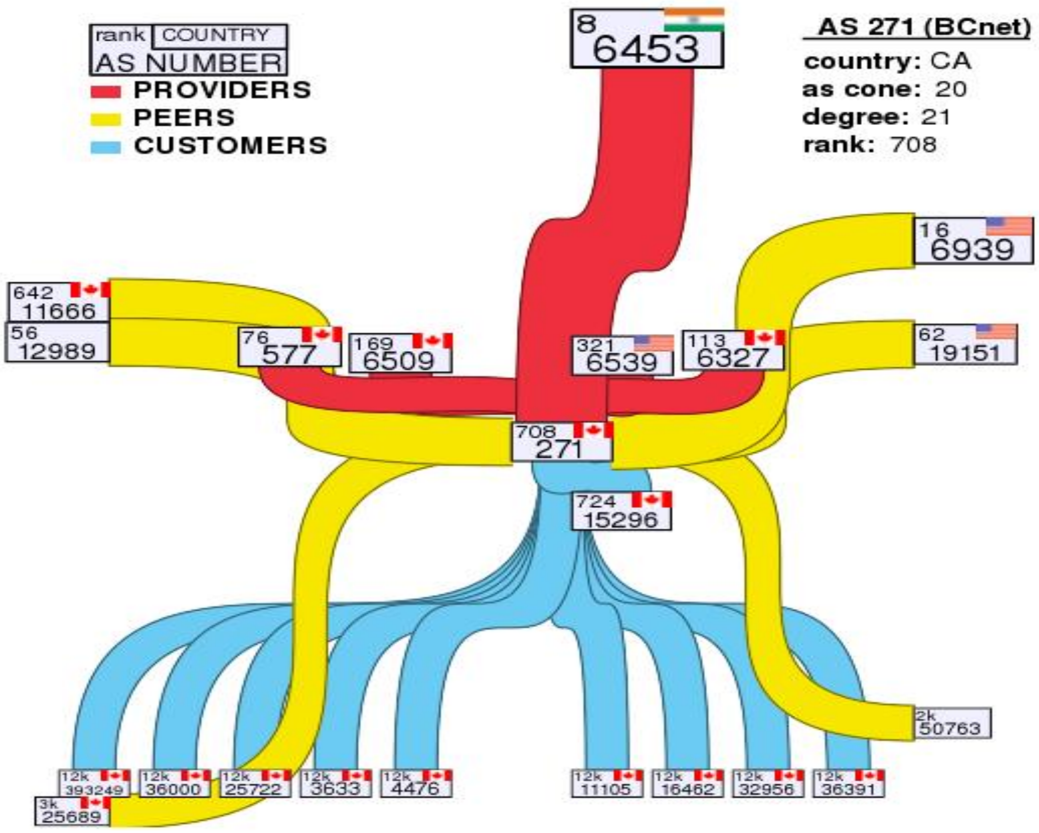

• AS Rank: Information for a single AS: AS Relationship Graph (AS 271) [Online]. Available: http://as-rank.caida.org/?mode0=as-info&mode1=as-graph&as=271.

Roadmap

- Introduction
- Routing polices
- RFD algorithms
- ns-2 implementation
- Performance evaluation
- Conclusions
- References

Conclusions

- Routing policies in networks increase:
  - convergence time
  - number of updates
  - number of flaps
- BRITE topologies have:
  - smaller number of updates
  - smaller number of flaps
- Convergence time of network topologies generated by GT-ITM is smaller
- RFD with modified suppressed values may help to counter route fluctuations

Conclusions

- Simulation results also show that the AS-path list policy performs similarly to BGP without any policy setup and performs better than the Community-path list policy
- ns-BGP-RP with the AS-path list policy exhibit the best performance in networks with large diameter and may help improve performance of today's Internet
- An adaptive approach to the route flap damping may help achieve an agreement between the BGP policy configuration without adversely affecting the network stability and network security

References: http://www.sfu.ca/~ljilja/cnl

- R. Gill, R. Paul, and Lj. Trajković, "Effect of MRAI timers and routing policies on BGP convergence times," in *Proc. IPCCC*, Austin, TX, USA, Dec. 2012, pp. 314–323.
- T. Farah, S. Lally, R. Gill, N. Al-Rousan, R. Paul, D. Xu, and Lj. Trajković, “Collection of BCNET BGP traffic,” in *Proc. ITC*, San Francisco, CA, USA, Sept. 2011, pp. 322–323.
- S. Lally, T. Farah, R. Gill, R. Paul, N. Al-Rousan, and Lj. Trajković, “Collection and characterization of BCNET BGP traffic,” in *Proc. IEEE PACRIM*, Victoria, BC, Canada, Aug. 2011, pp. 830–835.
- N. Laskovic and Lj. Trajković, “BGP with an adaptive minimal route advertisement interval,” in *Proc. IPCCC*, Phoenix, AZ, USA, Apr. 2006, pp. 142–151.
- W. Shen and Lj. Trajković, "BGP route flap damping algorithms," in *Proc. SPECTS* 2005, Philadelphia, PA, July 2005, pp. 488-495.
- T. D. Feng, R. Ballantyne, and Lj. Trajković, “Implementation of BGP in a network simulator,” in *Proc. ATS*, Arlington, VA, USA, Apr. 2004, pp. 149–154.

References: BGP

- S. Aggarwal and M. Aggarwal, “Dynamic load balancing based on CPU utilization and data locality in distributed database using priority policy,” in *Proc. ICSTE*, Phuket, Thailand, Oct. 2010, vol. 2, pp. 388–391.
- D. Pei and J. V. Merwe, “BGP convergence in virtual private networks,” in *Proc. IMC*, Rio de Janeiro, Brazil, Oct. 2006, pp. 283–288.
- R. Teixeira, A. Shaikh, T. Griffin, and J. Rexford, “Dynamics of hot-potato routing in IP networks,” in *Proc. ACM SIGMETRICS*, New York, NY, USA, June 2004, pp. 307–319.
- S. Agarwal, C. Chuah, S. Bhattacharyya, and C. Diot, “Impact of BGP dynamics on router CPU utilization,” in *Proc. PAM*, Antibes Juan-les-Pins, France, Apr. 2004, pp. 278–288.
- G. Siganos, M. Faloutsos, P. Faloutsos, and C. Faloutsos, “Power-laws and the AS-level Internet topology,” *IEEE/ACM Trans. Networking*, vol. 11, no. 4, pp. 514–524, Aug. 2003.
- C. Labovitz, A. Ahuja, A. Bose, and F. Jahanian, “Delayed Internet routing convergence,” in *Proc. ACM* SIGCOMM, Stockholm, Sweden, Aug. 2000, pp. 175–187.
- M. Faloutsos, P. Faloutsos, and C. Faloutsos, “On power-law relationships of the Internet topology,” in *Proc. SIGCOMM*, Cambridge, MA, USA, Sept. 1999, pp. 251–262.
- T. Griffin and G. Wilfong, “An analysis of BGP convergence properties,” in *Proc. SIGCOMM*, Cambridge, MA, USA, Aug. 1999, pp. 277– 288.
- C. Labovitz, G. Malan, and F. Jahanian “Origins of Internet routing instability,” in *Proc. INFOCOMM*, New York, NY, USA, Mar. 1999, pp. 218–226.
- Y. Rekhter and T. Li, “A border gateway protocol 4 (BGP-4),” *IETF* RFC 1771, Mar. 1995.
- Operational Experience with the BGP-4 protocol [Online]. Available: http://tools.ietf.org/html/draft-ietf-idr-bgp4-op-experience-01.

References: routing policies and RFD

- W. Sun, Z. M. Mao, and K. G. Shin, “Differentiated BGP update processing for improved routing convergence,” in *Proc. ICNP*, Santa Barbara, CA, USA, Nov. 2006, pp. 280–289.
- Z. Mao, R. Govindan, G. Varghese, and R. Katz, “Route flap damping exacerbates Internet routing convergence,” *Computer Communication Review Journal*, vol. 32, no. 4, pp. 221–233, Aug. 2002.
- L. Gao and J. Rexford, “Stable Internet routing without global coordination,” in *Proc. ACM SIGCOMM*, San Diego, CA, USA, Dec. 2001, pp. 681–692.
- C. Labovitz, R. Wattenhofer, S. Venkatachary, and A. Ahuja, “The impact of Internet policy and topology on delayed routing convergence,” in *Proc. INFOCOMM*, Anchorage, AL, USA, Apr. 2001, pp. 537–546.
- K. Varadhan, R. Govindan, and D. Estrin, “Persistent route oscillations in inter-domain routing,” *Computer Networks Journal*, vol. 32, no. 1, pp. 1–16, Mar. 2000.
- C. Villamizar, R. Chandra, and R. Govindan, “BGP route flap damping,” *IETF* RFC 2439, Nov. 1998.
- Route flap damping considered useable [Online]. Available: http:\\www.ripe.net/ripe/mail/archives/routing-wg/2012-July/002163.html.
- BGP++ [Online]. Available: http://www.ece.gatech.edu/research/labs/MANIACS/BGP++/.
- TRE [Online]. Available: http://laurikari.net/tre/.

References: topology generator

- BCNET [Online]. Available: https://wiki.bc.net/atl-conf/display/Content/Home.
- GT-ITM [Online]. Available: http://www.cc.gatech.edu/projects/gtitm/.
- BRITE [Online]. Available: http://www.cs.bu.edu/brite.

Acknowledgements

- Examiner:
  - Prof. Emeritus William A. Gruver
- Supervisor:
  - Prof. Emeritus R. H. Stephen Hardy
- Chair:
  - Prof. John D. Jones
- Senior supervisor:
  - Prof. Ljiljana Trajković

Thank you!