Effect of Minimal Route Advertisement Interval Timers on Border Gateway Protocol Convergence

Rajvir Gill
rajvirg@sfu.ca

Communication Networks Laboratory
http://www.ensc.sfu.ca/research/cnl
School of Engineering Science
Simon Fraser University

Roadmap

- Introduction
- BGP convergence
- Flexible Load Dispersing (FLD)-MRAI algorithm:
  - CPU utilization and modified DoP
  - modified reusable timers
  - duration of MRAI
  - space and time complexity
- Implementation of FLD-MRAI
- Performance evaluation
- Conclusions, future work, references

Introduction

- BGP is de-facto inter AS routing protocol
- Operates successfully in a network of the Internet's size
- Supports CIDR
- BGP 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

Minimal Route Advertisement Interval: MRAI

- MRAI is the interval limitation that affects BGP convergence
- Default value: 30 s
- MRAI timers control the MRAI value:
  - per-destination
  - per-peer
- Optimal MRAI value depends on:
  - network size
  - topology
  - traffic volume
  - network conditions

- Y. Rekhter and T. Li, "A border gateway protocol 4 (BGP-4)," *IETF* RFC 1771, Mar. 1995.

RFC : Request for Comments

MRAI timers

- Per-destination:
  - associated with each network destination
  - may not be used because of the Internet size
- Per-peer:
  - associated with each peer in the network
  - starts ticking when the source router sends a route advertisement to peers
  - adversely affect advertisements to each destination

Processing delay

- Total time of an update waiting in the queue and the time required for a BGP router to process
- Uniform processing delay:
  - BGP router processes update messages sequentially
  - delay in processing updates affects the processing time of update messages that follow
- Measurements:
  - update messages are processed within 200 ms
  - average processing time is 101 ms with the upper bound of 400 ms

- 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.
- A. Feldmann, H. Kong, O. Maennel, and A. Tudor, "Measuring BGP pass-through times," in *Proc. PAM*, Antibes Juan-les-Pins, France, Apr. 2004, pp. 267–277.

Motivation

- One of the major problems of BGP: a longer convergence time and a large number of update messages due to unreachable destinations
  - path to the destination failure
  - router failure
- Possible solution:
  - an algorithm that decreases the BGP convergence time and the number of update messages exchanged in the network

Contributions

- MRAI with Flexible Load Dispersing (FLD-MRAI) algorithm:
  - modifies reusable timers
  - employs MRAI durations based on BGP advertisement events
  - applies to heterogeneous and large networks
  - performs well in networks with unspecified traffic load

Roadmap

- Introduction
- BGP convergence
- Flexible Load Dispersing (FLD)-MRAI algorithm:
  - CPU utilization and modified DoP
  - modified reusable timers
  - duration of MRAI
  - space and time complexity
- Implementation of FLD-MRAI
- Performance evaluation
- Conclusions, future work, references

BGP convergence

- BGP convergence time: the time interval between first update message sent, until all update messages that are a consequence of the original update are received
- Adaptive MRAI timers decrease the BGP convergence time and guarantee network stability
- PED algorithm proposed timer: 35 s
- Processing efficiency of router's CPU affects the 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.

PED: Path Exploration Damping

BGP convergence

- SSLD and optimal values for MRAI reduce the BGP convergence time
- Delay due to router or link failure increases BGP convergence time
- Shortest path to destination decreases BGP convergence time
- Router's CPU load depends on the number of BGP messages
- High CPU utilization of a BGP router causes delay
- Mismatch in policy configurations between two ASes may also cause network instabilities

  - 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.
  - 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.

SSLD: Sender Side Loop Detection
CPU : Central Processing Unit

Roadmap

- Introduction
- BGP convergence
- Flexible Load Dispersing (FLD)-MRAI algorithm:
  - CPU utilization and modified DoP
  - modified reusable timers
  - duration of MRAI
  - space and time complexity
- Implementation of FLD-MRAI
- Performance evaluation
- Conclusions, future work, references

FLD-MRAI algorithm: scenarios

- Processing of updates:
  - normal load: DoP prefers the shortest path
  - high load: DoP prefers a longer path in the presence of the shortest path
  - empirical value for processing delay: 200 ms
  - MRAI consists of two states: idle and processing

DoP: Degree of Preference

FLD-MRAI algorithm: CPU utilization

- Calculation of available CPU:
  - CPU utilization is high, then the router responds slowly to subsequent requests in the queue
  - based on the priority of updates:

$$CPU_{available} = 100 - CPU_{active}$$
$$CPU_{active} = 100 * (CPU_{current}/CPU_{max})$$

$CPU_{available}$: percentage of available CPU of the neighboring router
CPU_{active} : percentage of active CPU utilization of the neighboring router
$CPU_{current}$: current CPU utilization
CPU_{max} : maximum CPU utilization

  - calculated every time a router receives the updates of a new or withdrawn route

FLD-MRAI algorithm: modified DoP

- Modified DoP (DoP_{mod}):
  - function of $Route_{info}$ and $CPU_{available}$, where $Route_{info}$ is the route having the shortest path
  - new, replaced, and/or withdrawn route
  - path with the highest value of DoP_{mod} is given the highest priority

FLD-MRAI algorithm: example with five routers

- R_0 is the source router and it advertises to destination R_2
- Two possible paths: R_0-R_1-R_2 and R_0-R_4-R_3-R_2
- Original BGP router: R_0-R_1-R_2
- FLD-MRAI: R_0-R_1-R_2 (normal load) and R_0-R_4-R_3-R_2 (high load)

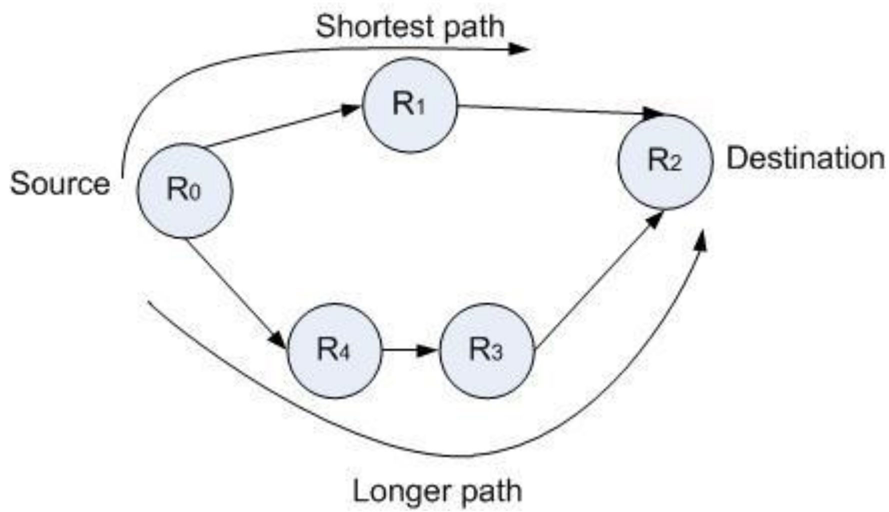

RFC : Request for Comments

Roadmap

- Introduction
- BGP convergence
- Flexible Load Dispersing (FLD)-MRAI algorithm:
  - CPU utilization and modified DoP
  - modified reusable timers
  - duration of MRAI
  - space and time complexity
- Implementation of FLD-MRAI
- Performance evaluation
- Conclusions, future work, references

FLD-MRAI algorithm: advertisement events

- MRAI timer depends on network conditions and advertisement events:
  - Tdown: link failure
  - Tup : link failure recovery
  - Tlong : router failure
  - Tshort: router failure recovery
- Occurrence of advertisement events:

Events	Number of events occurring during BGP convergence period
Tdown	43.4%
Tup	39.9%
Tlong	7.3%
Tshort	7.4%
Unidentified	2.0%

- D. Pei and J. V. Merwe, “BGP convergence in virtual private networks,” in *Proc. IMC*, Rio de Janeiro, Brazil, Oct. 2006, pp. 283–288.

FLD-MRAI algorithm: modified reusable timers

- After the processing time, three events may occur:
  - no new update received
  - new update message received
  - MRAI reusable timer expired
- Calculation of idle time:

$$T_{idle}(D) = MRAI_{total} - M_{last}$$

$T_{idle}(D)$: is the idle time of the destination

$MRAI_{total}$: total MRAI

M_{last} : time instance of the last message received

FLD-MRAI algorithm: modified reusable timers

- One reusable timer is required for all paths advertised during a short time interval
- Number of rounds per reusable MRAI timer controls the duration of MRAI round:

$$MRAI_{duration} = R_n * (t_n * g)$$

$MRAI_{duration}$: duration of MRAI round

R_n : number of rounds per reusable MRAI timer

t_n : number of reusable MRAI timers

g : granularity

Roadmap

- Introduction
- BGP convergence
- Flexible Load Dispersing (FLD)-MRAI algorithm:
  - CPU utilization and modified DoP
  - modified reusable timers
  - duration of MRAI
  - space and time complexity
- Implementation of FLD-MRAI
- Performance evaluation
- Conclusions, future work, references

FLD-MRAI algorithm: duration of MRAI

- Duration of MRAI timers:
  - longer duration: Tdown and Tlong
  - shorter duration: Tup and Tshort
- MRAI value:
  - same MRAI value: Tdown and Tlong
  - same MRAI value: Tup and Tshort
- Minimum duration: 15 s

FLD-MRAI algorithm: duration of MRAI

- Idle time is longer than 1 s:
  - Tshort or Tup
  - the shortest path becomes available
  - duration of MRAI round: 15 s
  - use one MRAI round: 15 s
- Idle time is shorter than 1 s:
  - Tlong or Tdown
  - router or link failure
  - duration of MRAI round: 30 s
  - use two MRAI rounds: 15 s

FLD-MRAI algorithm: example of reusable timers

- 15 reusable MRAI timers with granularity 1 s
- $Timer_0$ lasts one round of 15 s (Tshort or Tup)
- $Timer_2$ lasts one round of 30 s (Tdown or Tlong)
- After expiration, the duration of timers depends on the idle time

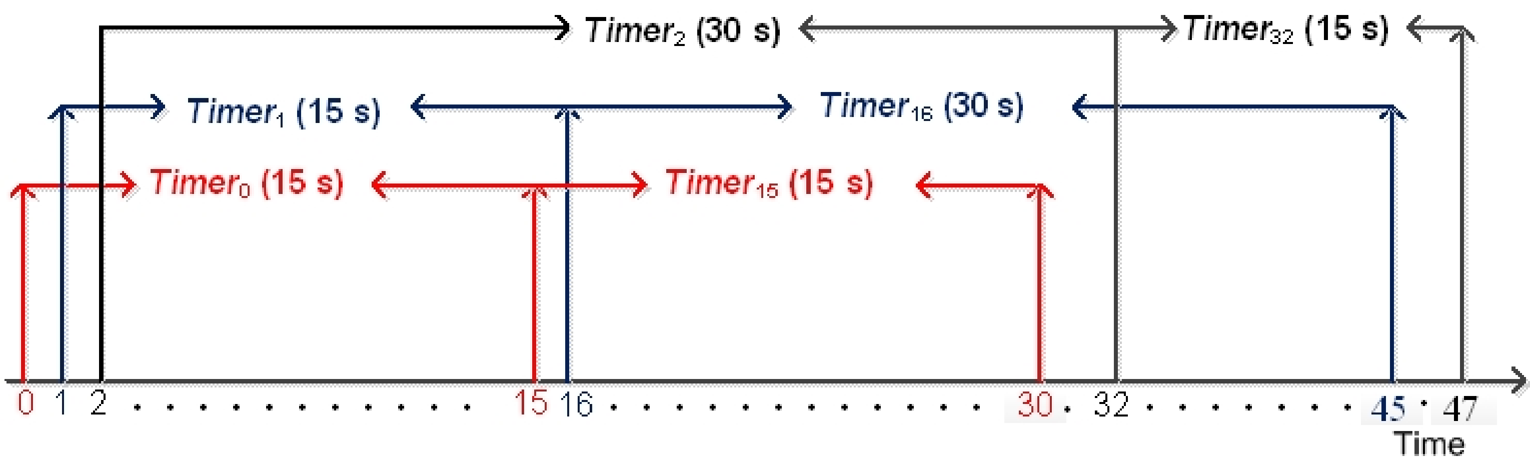

FLD-MRAI algorithm: pseudocode

when sending advertisement of the destination *D* to peers at t_o

set (S_i) // priority numbers on received updates according to the shortest path

if ($C_1 < C_2$) // calculate and compare the available CPU of the first and second priority neighboring router

if W(t) < T(t) // calculate and compare the waiting and transmission times

else (wait in queue of the first priority path)

if $dop_2 < dop_1$ // calculate and compare the degree of preference

choose the second priority path

MRAI = 30 s

goto processing state

else (wait in queue of the first priority path)

else if ($C_1 > C_2$)

wait in queue of the first priority path // duration of MRAI is based on the idle time

goto processing state

FLD-MRAI algorithm: pseudocode

when initiation of the new round

if (Idle(D) > 1 s) // *Tshort* or *Tup* may occur

set modified_reusable timer = 15 s

else if (e ∈ network failure)

// events change due to the network failure

choose the second priority path // after expiration of the timer

set modified_reusable timer = 30 s

goto processing state

else if (e ∉ network failure)

goto processing state

else if (Idle(D) < 1 s) // *Tlong* or *Tdown* may occur

set modified_reusable timer = 30 s

else if ($P_t \in P_s$) // if the shortest path becomes available

choose the shortest path // after expiration of the timer

set modified_reusable timer = 15 s

goto processing state

else ($P_t \notin P_s$) // if the shortest path is not available

goto processing state

Roadmap

- Introduction
- BGP convergence
- Flexible Load Dispersing (FLD)-MRAI algorithm:
  - CPU utilization and modified DoP
  - modified reusable timers
  - duration of MRAI
  - space and time complexity
- Implementation of FLD-MRAI
- Performance evaluation
- Conclusions, future work, references

FLD-MRAI algorithm: space complexity

- Space complexity:
  - depends on the number of non-converged routes (n)
  - router keeps three variables of each non-converged route: $CPU_{current}$, CPU_{max}, and M_{last}
  - variables are integer counters that a router may easily store
  - space complexity is O(n)

FLD-MRAI algorithm: time complexity

- Time complexity:
  - division, multiplication, and subtraction operations: $CPU_{available}$, T_{idle}, and $MRAI_{duration}$
  - division and multiplication depend on input size n
  - subtraction is constant
  - approximate these variables with constants equal to their maximum values
  - calculation of variables do not depend on input size n
  - FLD-MRAI requires two subtractions, three multiplications, and one division (equations: pages 14, 19, and 20)

FLD-MRAI algorithm: time complexity

- Time complexity:
  - number of neighbors and non-converged routes during one MRAI round affect the maximum number of update messages
  - BGP router may send only one advertisement and one withdrawal during a single MRAI round
  - time complexity of the computation of variables is $O(n)$ if the number of neighbors is constant

Roadmap

- Introduction
- BGP convergence
- Flexible Load Dispersing (FLD)-MRAI algorithm:
  - CPU utilization and modified DoP
  - modified reusable timers
  - duration of MRAI
  - space and time complexity
- Implementation of FLD-MRAI
- Performance evaluation
- Conclusions, future work, references

Implementation: FLD-MRAI

- ns-2.34 network simulator
- ns-BGP 2.0 developed module (ported from the SSFNET simulator)
- 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: FLD-MRAI

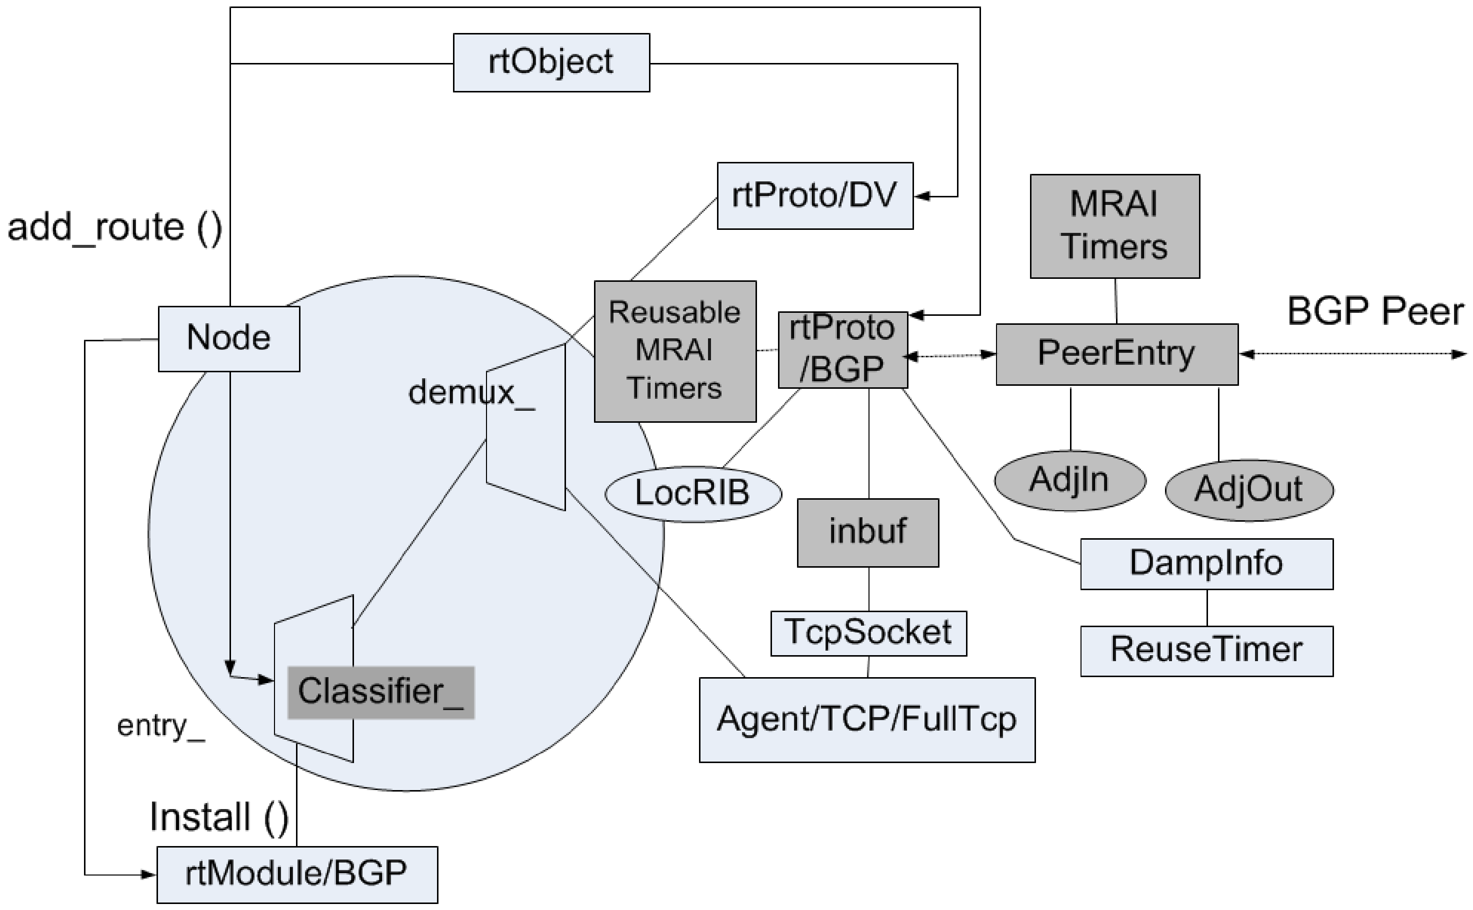

Simulation scenarios: performance comparisons

- FLD-MRAI algorithm:
  - FLD-MRAI-30
  - FLD-MRAI-15
  - default-MRAI-30
  - default-MRAI-15
  - adaptive MRAI

Simulated topologies

Topology	Number of nodes	Topology generator
Topology 1	67	Manually from BCNET BGP traffic
Topology 2	100	GT-ITM
Topology 3	200	GT-ITM
Topology 4	300	BRITE
Topology 5	500	BRITE

GT-ITM: Georgia Tech Internetwork Topology Models
BRITE : Boston university Representative Internet Topology gEnerator

BGP routing table (RIB)

Time	Peer's IP	Peer's AS	Source IP	AS path
2011-10-24, 05:18	216.6.50.9	6327	207.23.253.2	6327-7575-56203-1221
2011-10-24, 05:18	207.23.253.34	6453	207.23.253.2	6453-2914-2519-1221
2011-10-24, 05:18	216.6.50.9	6327	207.23.253.2	6327-2516-2519-1221
2011-10-24, 05:18	207.23.253.34	6453	207.23.253.2	6453-4725-7670-18144-1221
2011-10-24, 05:18	216.6.50.9	6327	207.23.253.2	6327-2516-7670-18144-1221
2011-10-24, 05:18	207.23.253.34	6453	207.23.253.2	6453-4725-1221
2011-10-24, 05:18	216.6.50.9	6327	207.23.253.2	6327-4725-1221

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
T	1	1
N_t	4	8
K	3	4
N_s	8	6
N	100	200

GT-ITM: Georgia Tech Internetwork Topology Models

BRITE topologies

- BRITE generates different types of Internet topologies from different models
- Generate 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
M (number of links per new node)	1
N (number of nodes)	300 or 500

BRITE: Boston university Representative Internet Topology gEnerator

Assumptions

- Do not consider route flap damping when evaluating the performance of FLD-MRAI
- Each AS consists of a single BGP router
- BGP convergence procedure is complete if BGP router receives no update message from other BGP routers within 60 s

Roadmap

- Introduction
- BGP convergence
- Flexible Load Dispersing (FLD)-MRAI algorithm:
  - CPU utilization and modified DoP
  - modified reusable timers
  - duration of MRAI
  - space and time complexity
- Implementation of FLD-MRAI
- Performance evaluation
- Conclusions, future work, references

Validation test: network with 5 routers

- Topology with the minimum number of nodes is used to analyze the BGP convergence time
- FLD-MRAI algorithm is validated by using a simple network of five routers
- We tested the employed modifications in ns-BGP for both normal and high loads
- BGP convergence time of both scenarios with FLD-MRAI is compared to default-MRAI-30
- Node 0: source node
- Node 2: destination node

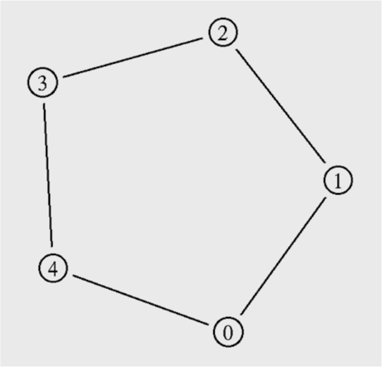

Validation test: network with 5 routers

- Normal load scenario events:
  - Tlong : n1 fails
  - Tshort : n1 recovers
  - Tdown: link between n0 and n1 fails
  - Tup : link between n0 and n1 recovers
- High load scenario: high traffic load to n1
- Simulation results indicate that FLD-MRAI performs as expected

Scenarios	default-MRAI-30 (s)	FLD-MRAI (s)
Tshort	88.70	52.70
Tlong	93.10	71.59
Tup	88.60	55.50
Tdown	93.05	60.90
High load	102.91	56.81

Validation test: completely connected graph

- Validate performance of the FLD-MRAI algorithm
- Choose the completely connected network with fifteen nodes
- We match simulation results of the convergence time and the number of update messages with results of the previous studies
- Simulate only Tdown event
- FLD-MRAI decreases the number of update messages from 3,200 to 1,500

- N. Laskovic and L. Trajkovic, “BGP with an adaptive minimal route advertisement interval,” in *Proc. IPCCC*, Phoenix, AZ, Apr. 2006, pp. 142–151.
- T. G. Griffin and B. J. Premore, “An experimental analysis of BGP convergence time,” in *Proc. ICNP*, Riverside, CA, Nov. 2001, pp. 53–61.
- J. Nykvist and L. Carr-Motyckova, “Simulating convergence properties of BGP,” in *Proc. ICCCN*, Miami, FL, Oct. 2002, pp. 124–129.

Validation test: completely connected graph

- Minor changes may exist due to the different simulation setups
- Results of the BGP convergence time and the number of update messages for default-MRAI-30 are similar to the results reported in the previous studies

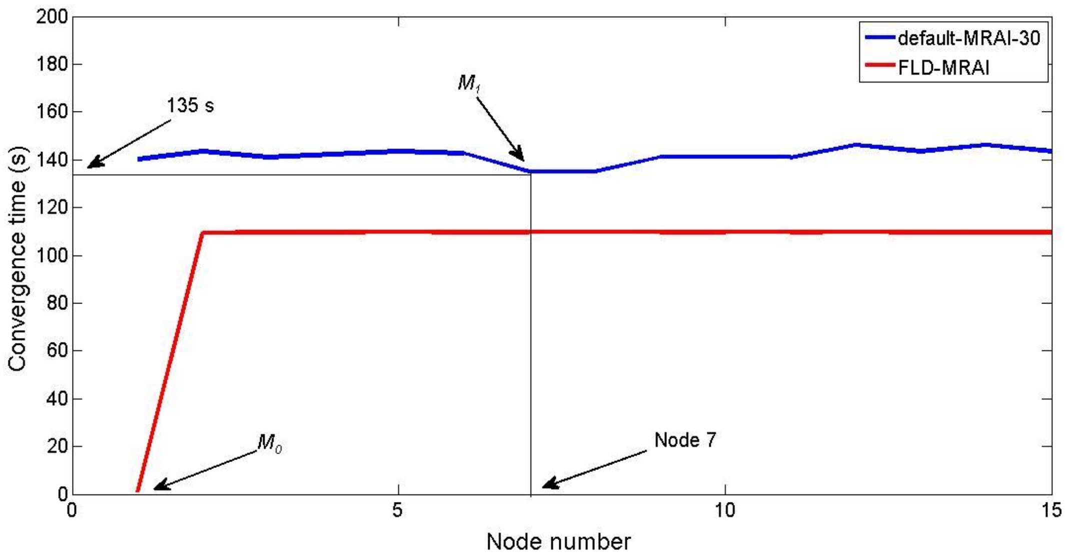

M_o : Optimal MRAI value for FLD-MRAI
M_1 : Optimal MRAI value for default-MRAI-30

Performance evaluation: network Topology 1

- FLD-MRAI algorithm: convergence time (s) of the Tshort event for network Topology 1 for various BGP options

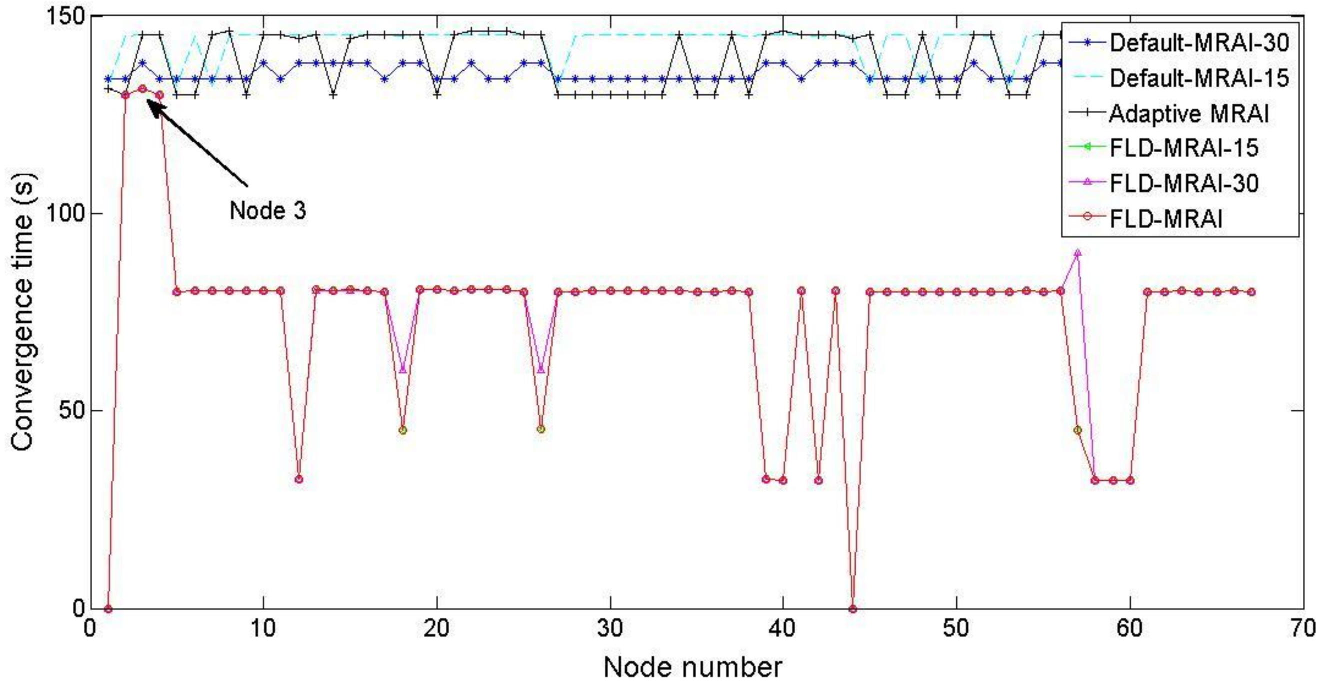

Performance evaluation: network Topology 1

- FLD-MRAI algorithm: convergence time (s) of the Tlong, Tup, and Tdown events for network Topology 1 for various BGP options

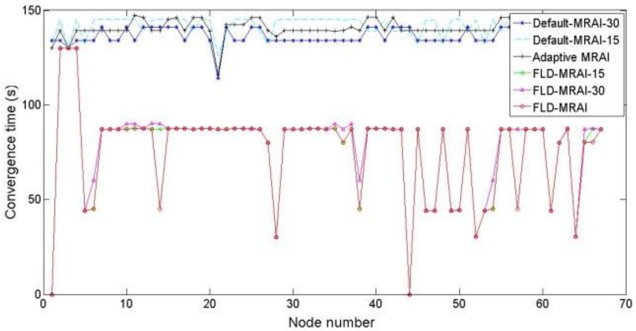

Tlong event

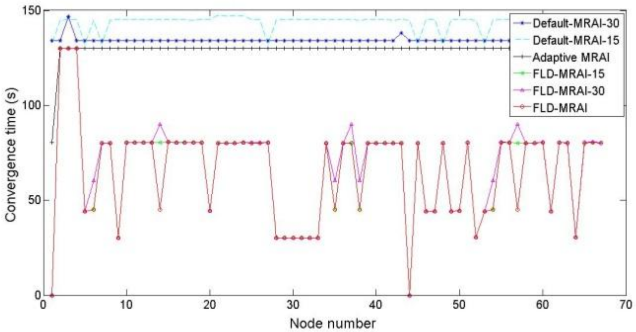

Tdown event

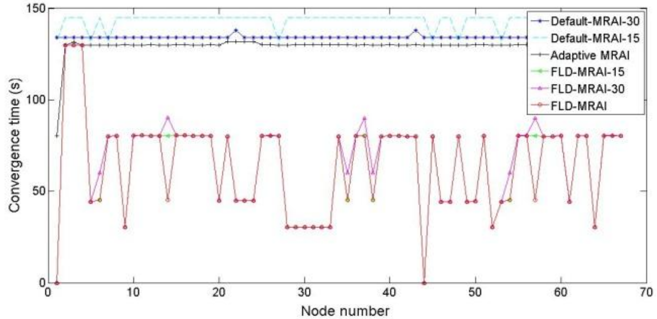

Tup event

Performance evaluation: network Topology 1

- FLD-MRAI algorithm: overall number of update messages for network Topology 1 for various BGP options

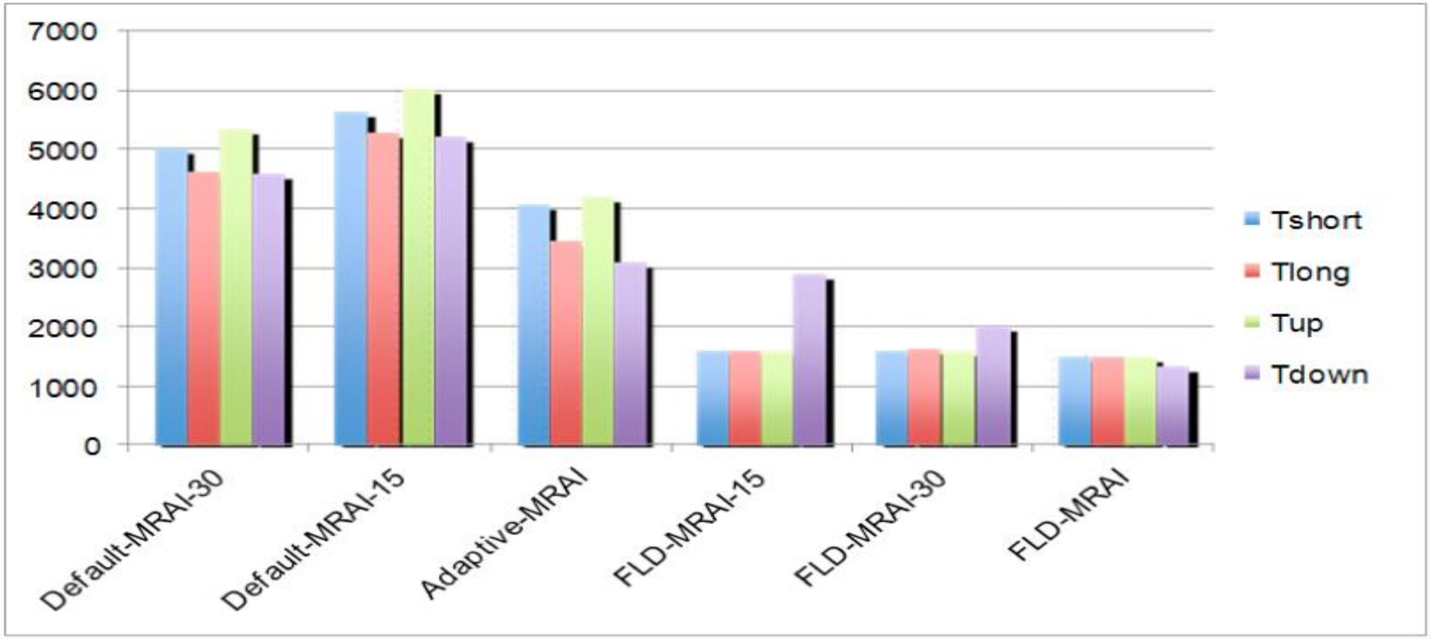

Performance evaluation: network Topology 1

- FLD-MRAI algorithm: convergence time (s) of the high load scenario for network Topology 1 for various BGP options

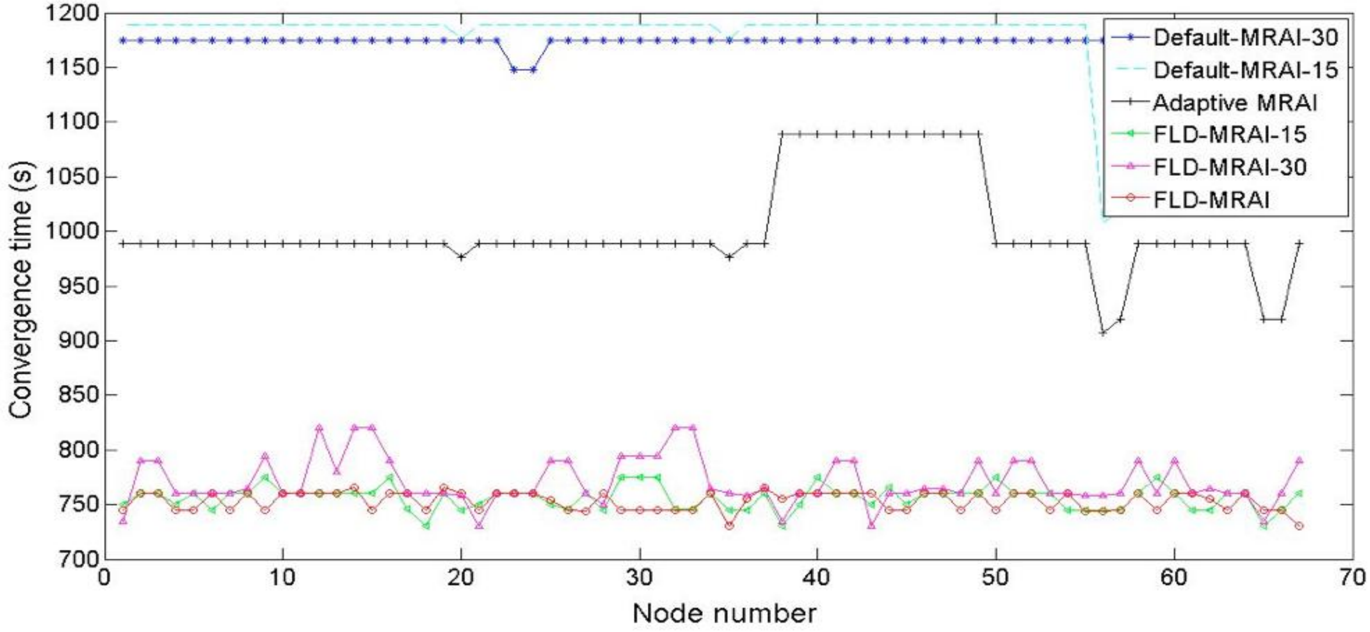

Performance evaluation: network Topology 1

- FLD-MRAI algorithm: overall number of update messages for the high load scenario

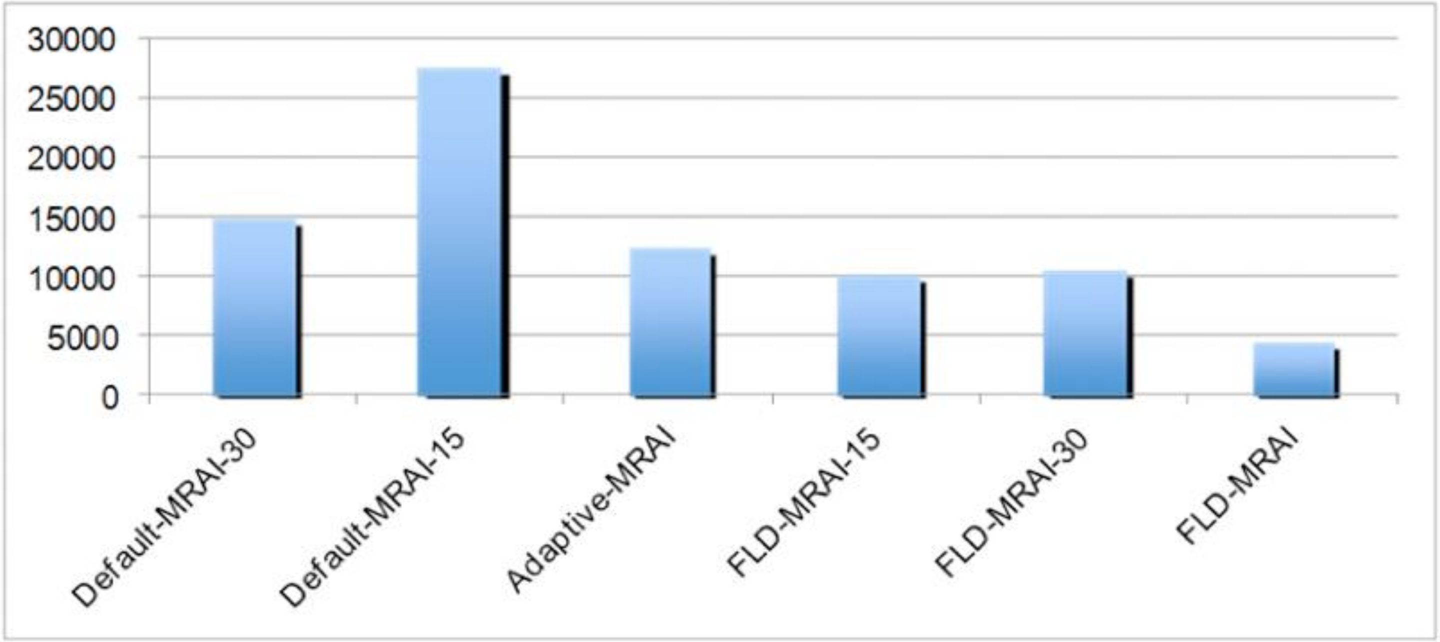

Performance evaluation: network Topology 5

- FLD-MRAI algorithm: average convergence time (s) for network Topology 5 for various BGP options

Scenarios	Default-MRAI-30	Default-MRAI-15	Adaptive MRAI	FLD-MRAI-15	FLD-MRAI-30	FLD-MRAI
Tshort	772.91	775.71	782.42	659.90	792.67	601.50
Tlong	795.61	778.60	783.26	794.67	660.66	608.66
Tup	773.03	775.65	782.34	659.66	793.05	602.51
Tdown	796.13	779.60	784.71	794.46	661.09	609.33
High load	918.02	909.48	906.42	930.95	951.70	530.39

Performance evaluation: network Topology 5

- FLD-MRAI algorithm: overall number of update messages for network Topology 5 for various BGP options

Scenarios	Default-MRAI-30	Default-MRAI-15	Adaptive MRAI	FLD-MRAI-15	FLD-MRAI-30	FLD-MRAI
Tshort	8,330	12,298	6,526	6,342	7,755	2,579
Tlong	8,349	12,286	6,514	6,315	7,721	2,564
Tup	8,323	12,292	6,520	6,331	7,734	2,523
Tdown	8,326	12,286	6,526	6,315	7,721	2,565
High load	13,353	13,422	10,466	6,141	6,256	2,672

Performance evaluation

- The average percentage of improvement of FLD-MRAI over default MRAI (30 s) based on different network topologies

Events	Convergence time (s)	Overall number of updates
Tshort	24%	47%
Tlong	26%	51%
Tup	27%	47%
Tdown	25%	46%

Roadmap

- Introduction
- BGP convergence
- Flexible Load Dispersing (FLD)-MRAI algorithm:
  - CPU utilization and modified DoP
  - modified reusable timers
  - duration of MRAI
  - space and time complexity
- Implementation of FLD-MRAI
- Performance evaluation
- Conclusions, future work, references

Conclusions

- Proposed FLD-MRAI algorithm employs:
  - BGP modifications to reduce the convergence time and number of update messages exchanged during normal and high traffic loads
  - modified DoP that depends on the calculation of available CPU
  - separate durations of MRAI for different events that occur during BGP advertisements
  - modified reusable MRAI timers

Conclusions

- Simulation results show that FLD-MRAI performs better than other BGP options at the cost of computing available CPU of neighboring routers
- The CPU processing capability and duration of MRAI timers greatly affect the BGP convergence time
- FLD-MRAI exhibit the best performance in networks with large diameter and may help improve performance of today's Internet

Future work

- Analyse the effect of iBGP on the convergence time and the number of update messages
- Explore the effect of routing policies on the BGP convergence time along with the MRAI
- Implement various simulation scenarios and settings
- Test the FLD-MRAI algorithm in a real test-bed

iBGP : Interior Border Gateway Protocol

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: MRAI

- B. Wang, “The research of BGP convergence time,” in *Proc. ITAIC*, Chongqing, China, Aug. 2011, vol. 2, pp. 354–357.
- A. Fabrikant, U. Syed, and J. Rexford, “There's something about MRAI: timing diversity can exponentially worsen BGP convergence,” in *Proc. INFOCOMM*, Shanghai, China, Apr. 2011, pp. 2975–2983.
- 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.
- S. Deshpande and B. Sikdar, “On the impact of route processing and MRAI timers on BGP convergence times,” in *Proc. GLOBECOM*, Dallas, TX, USA, Nov. 2004, vol. 2, pp. 1147–1151.
- A. Feldmann, H. Kong, O. Maennel, and A. Tudor, “Measuring BGP pass-through times,” in *Proc. PAM*, Antibes Juan-les-Pins, France, Apr. 2004, pp. 267–277.
- B. Premore, An analysis of convergence properties of the border gateway protocol using discrete event simulation, Ph. D. Thesis, Dartmouth College, 2003.
- 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.

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

Thank you!