Fast Incremental FIB Aggregation (FIFA)

Yaoqing Liu
yliu6@memphis.edu
The University of Memphis
Dr. Lan Wang
lanwang@memphis.edu
The University of Memphis
Dr. Beichuan Zhang
bzhang@cs.arizona.edu
The University of Arizona

Outline

- **Motivation**
- FIB aggregation overview
- Update handling
- Evaluation
- Operational impact
- Conclusion

Motivation (1)

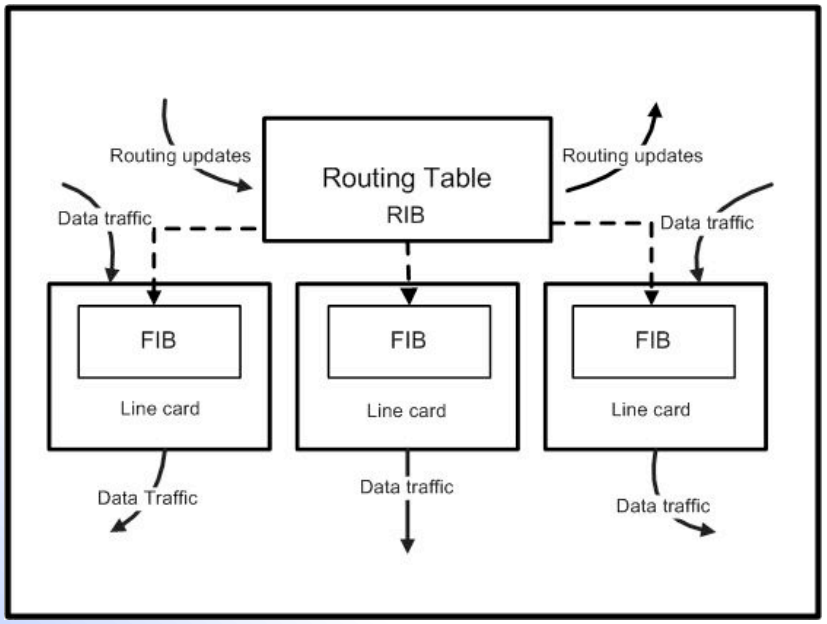

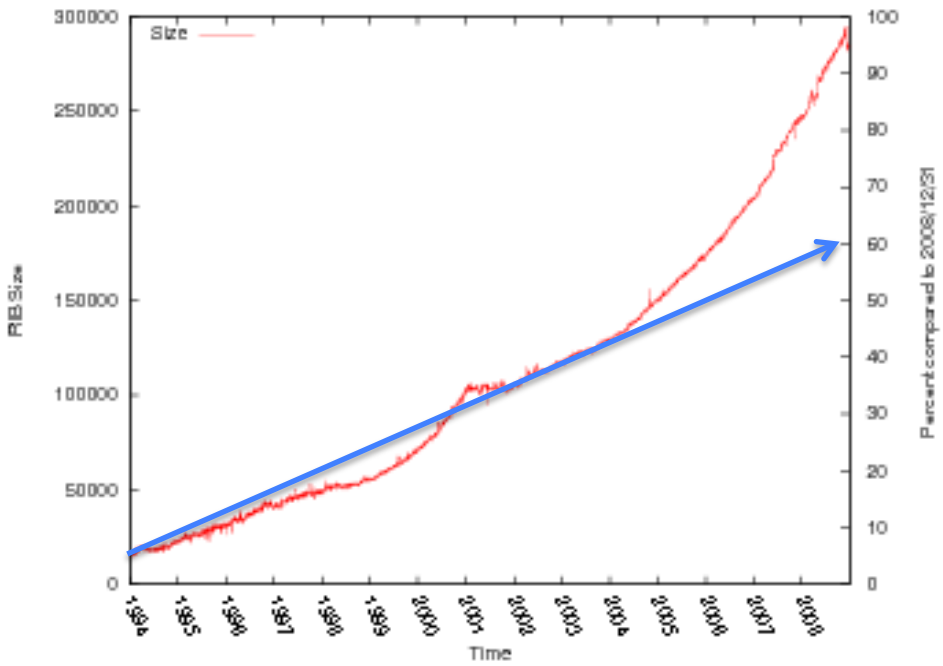

- Fast growth of RIB=>Large FIB (Forwarding Information Base)
- Large FIB may not fit line-card memory
  - Expensive FIB memory
  - High cost for operators and customers
- FIB Aggregation is a local, low cost, and software solution to reduce FIB size [2]

Motivation (2)

- Existing works focus on reducing table size (as much as 70%) [1, 3, 4, 5]
- We also focus on reducing the aggregation overhead
- Most challenging problem: update handling
  - Quickly apply updates to aggregated table
  - Still maintain good compression ratio

- **Objectives:**
  - Extend line-card life-time by more than 5 years
  - Reduce the overall processing time
  - Reduce total number of FIB changes
  - Mitigate single burst of FIB changes

Outline

- Motivation
- **FIB aggregation overview**
- Update handling
- Evaluation
- Operational impact
- Conclusion

What is FIB Aggregation?

- Merge multiple FIB entries into one

(a) Original FIB Entries

Label	Prefix	Next Hop
A	141.225.0.0/16	1
B	141.225.64.0/18	1
C	141.225.32.0/19	1
D	141.225.96.0/19	2
E	141.225.48.0/20	2

(b) Aggregated FIB Entries

Label	Prefix	Next Hop
A	141.225.0.0/16	1
D	141.225.96.0/19	2
E	141.225.48.0/20	2

- 100% same forwarding behaviors before and after aggregation

FIFA Model

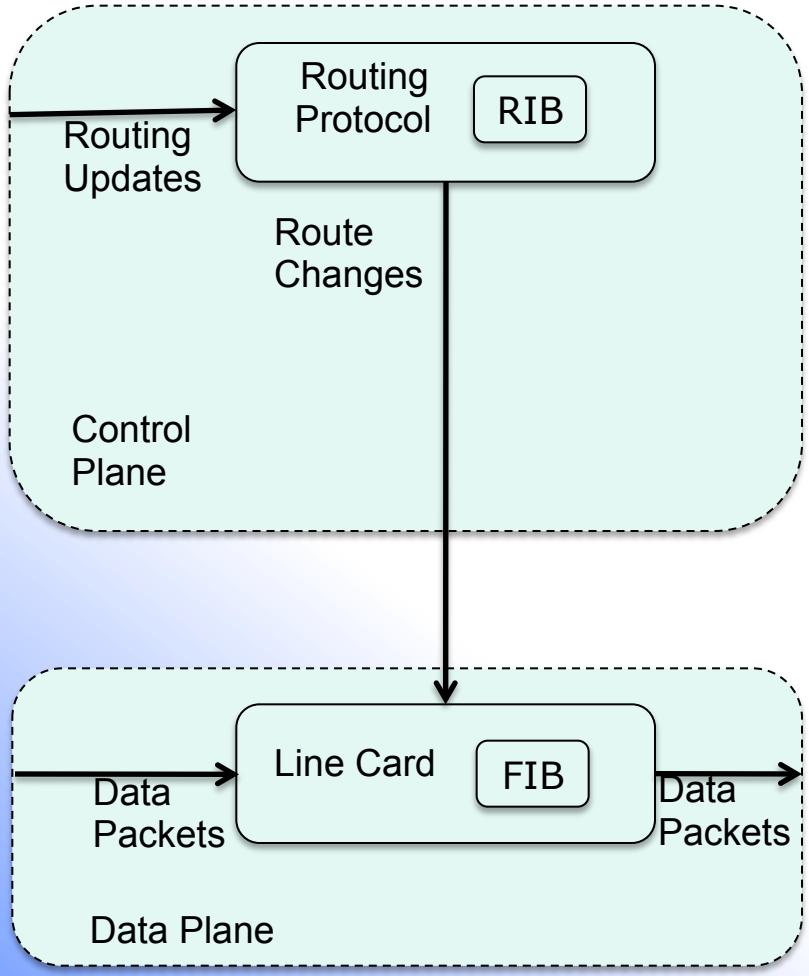

Normal Router Operations

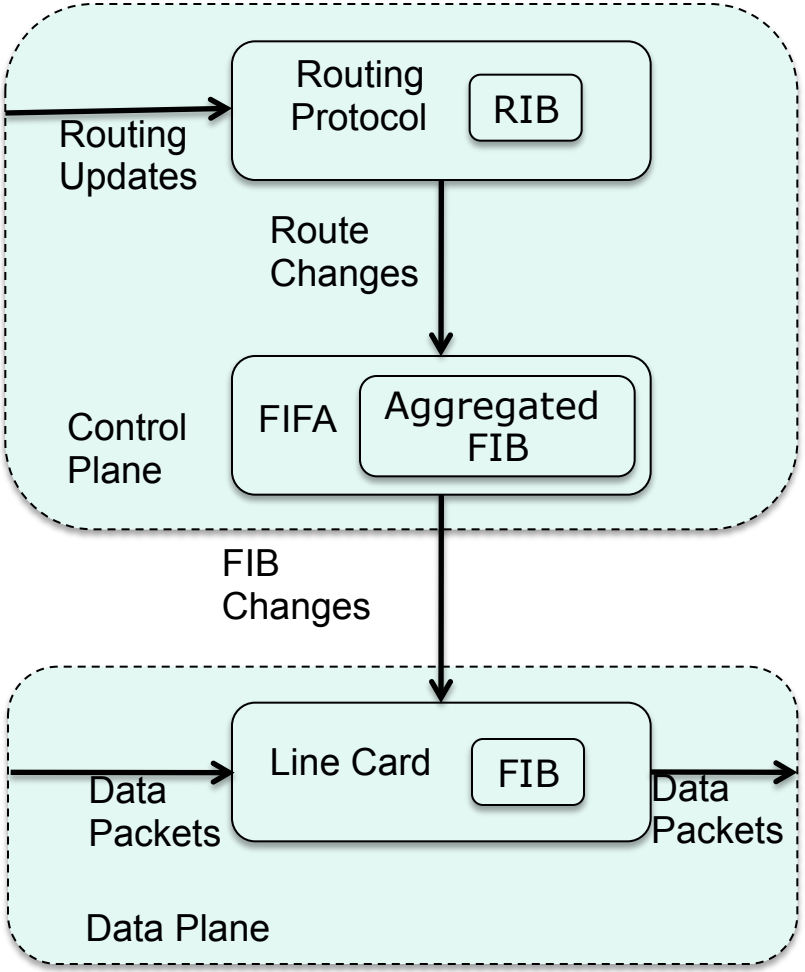

Router Operations with FIFA

Our Proposed Algorithms-FIFA

- FIFA: Three fast incremental FIB aggregation algorithms based on Improved ORTC [1]
  - [+] Patricia Trie (Two passes, 2.5 faster, 44% memory)
  - [+] Strong forwarding correctness
- FIFA-S: Optimal Size Update (Global Optimality)

  [+] Always smallest FIB size, no re-aggregation
- FIFA-T: Minimal Time Update (Local Optimality)

  [+] Fastest
- FIFA-H: Hybrid Update (Regional Optimality)

  [+] No re-aggregation, small FIB changes

ORTC-Binary Tree Implementation

(a) Original FIB Entries

Label	Prefix	Next Hop
A	141.225.0.0/16	1
B	141.225.64.0/18	1
C	141.225.32.0/19	1
D	141.225.96.0/19	2
E	141.225.48.0/20	2

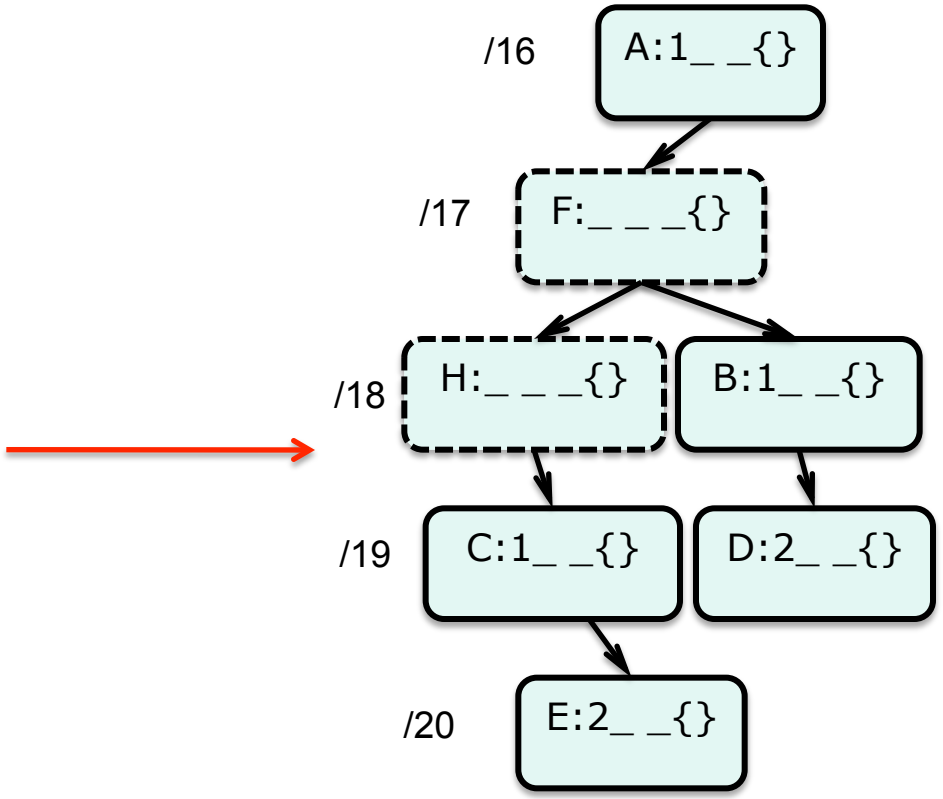

Original Tree

ORTC-Pass 1: Normalize Binary Tree

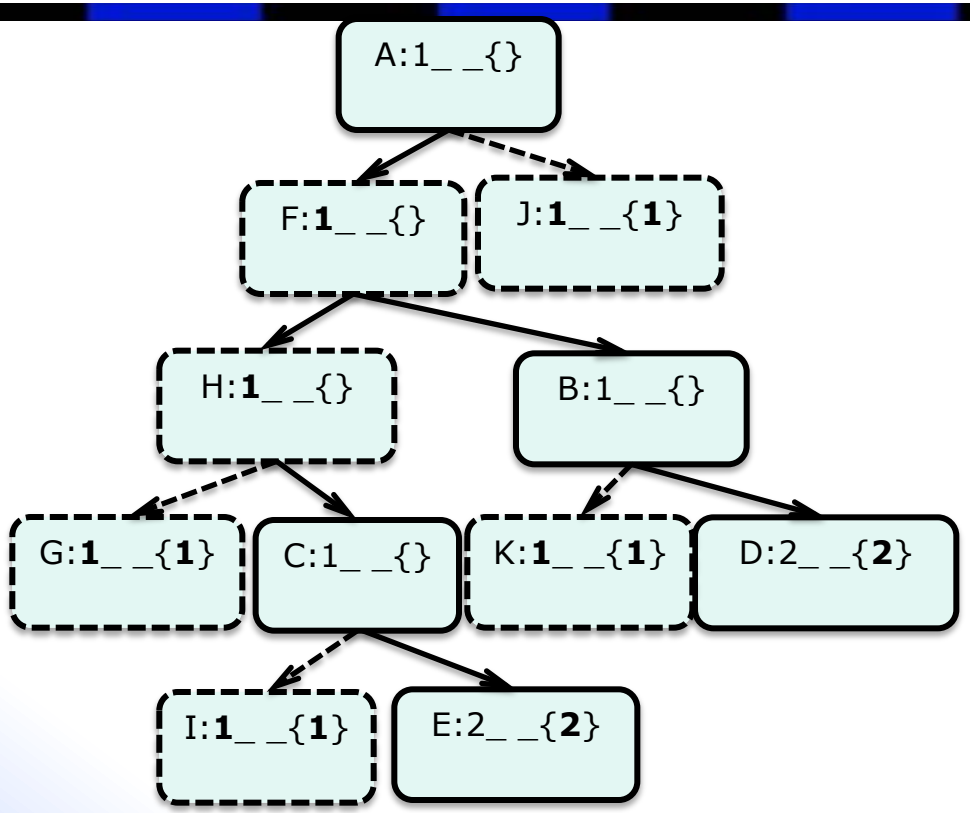

- Expand the binary tree to a complete binary tree in which a node has zero or two children
- Each expanded node has the same next hop as its nearest ancestor's

ORTC-Pass 2: Merge Next Hops

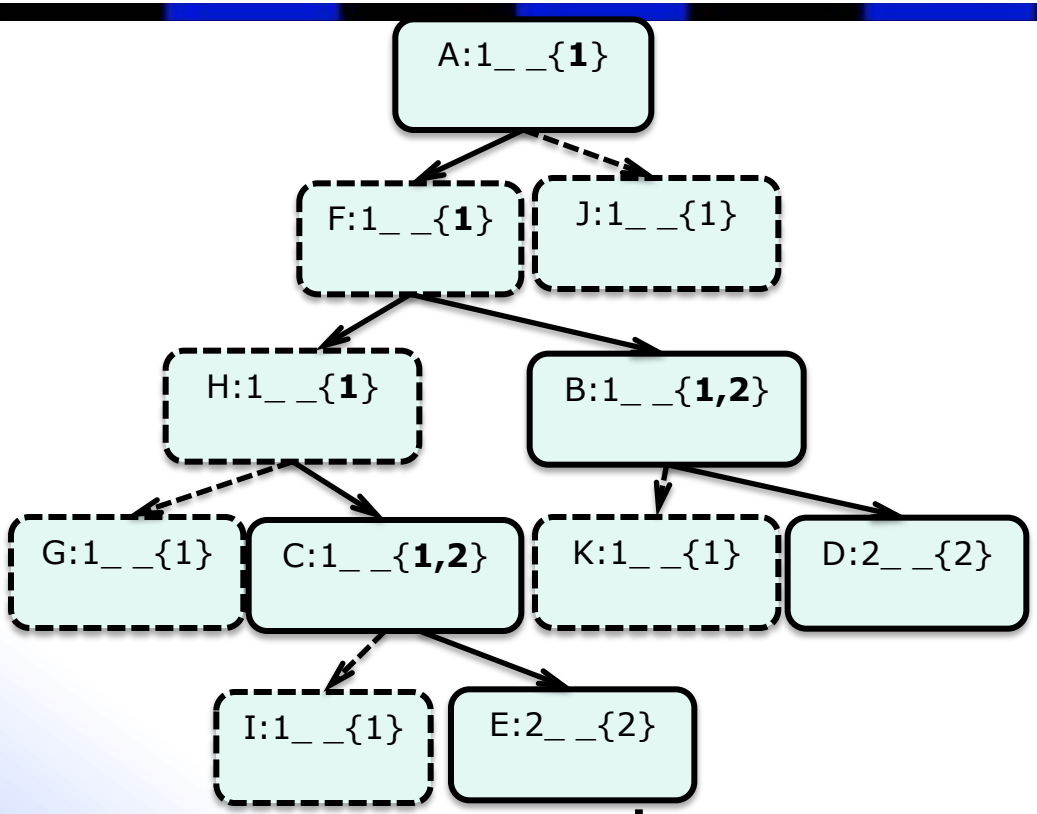

Post order, merge children's nexthops to their parent node

- If two children have overlapped nexthops,
  do intersection operation
- If two children have no overlapped nexthops,
  do union operation

ORTC - Pass 3: Select Next Hop

From root of tree, select nexthop from the merged nexthops till reaching all leaves

- If descendant selected nexthop is the same as its ancestor, then don't put it in FIB
- Otherwise, keep it in FIB

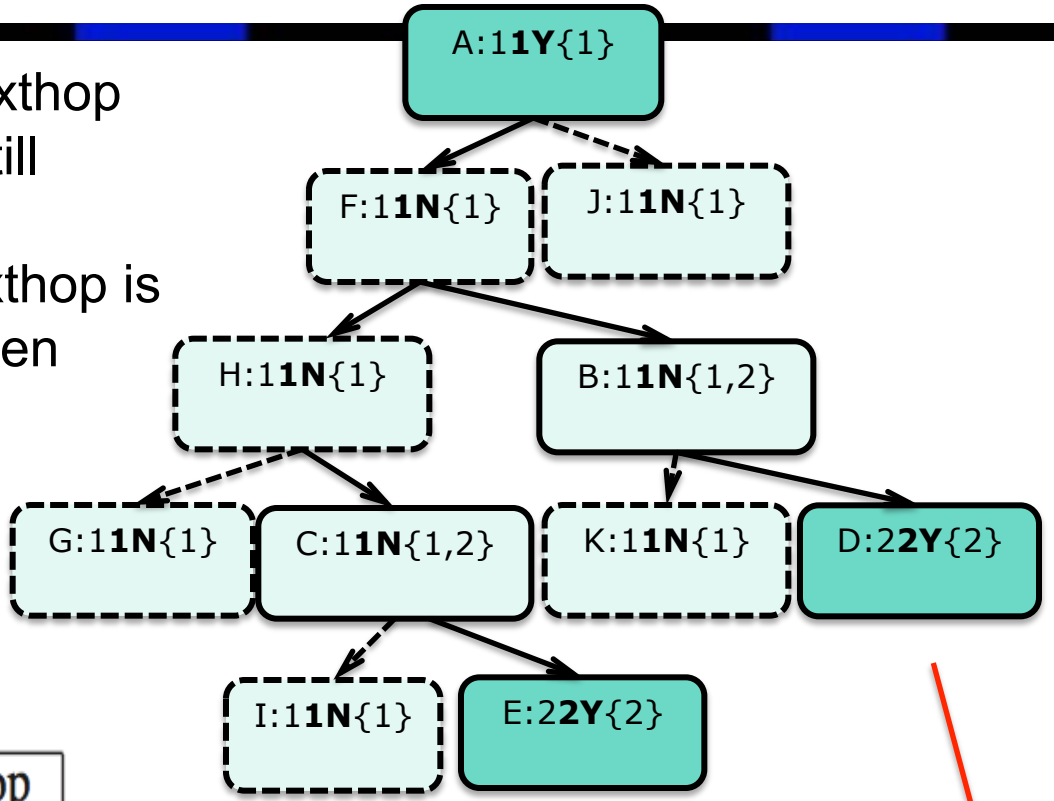

(a) Original FIB Entries

Label	Prefix	Next Hop
A	141.225.0.0/16	1
B	141.225.64.0/18	1
C	141.225.32.0/19	1
D	141.225.96.0/19	2
E	141.225.48.0/20	2

(b) Aggregated FIB Entries

Label	Prefix	Next Hop
A	141.225.0.0/16	1
D	141.225.96.0/19	2
E	141.225.48.0/20	2

Outline

- Motivation
- FIB aggregation overview
- **Update handling**
- Evaluation
- Operational impact
- Conclusion

Normal Update Without Aggregation

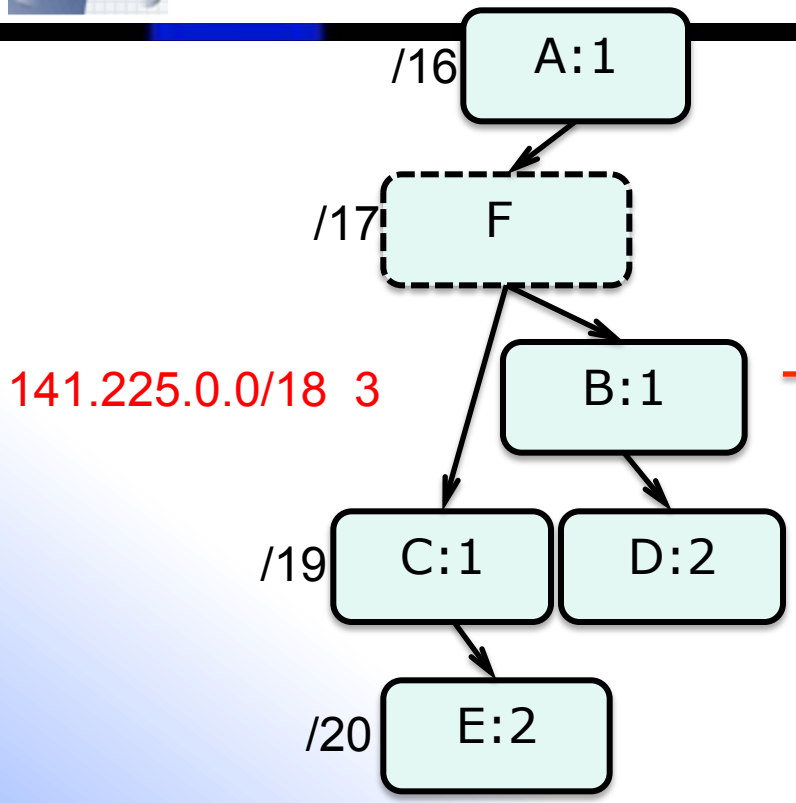

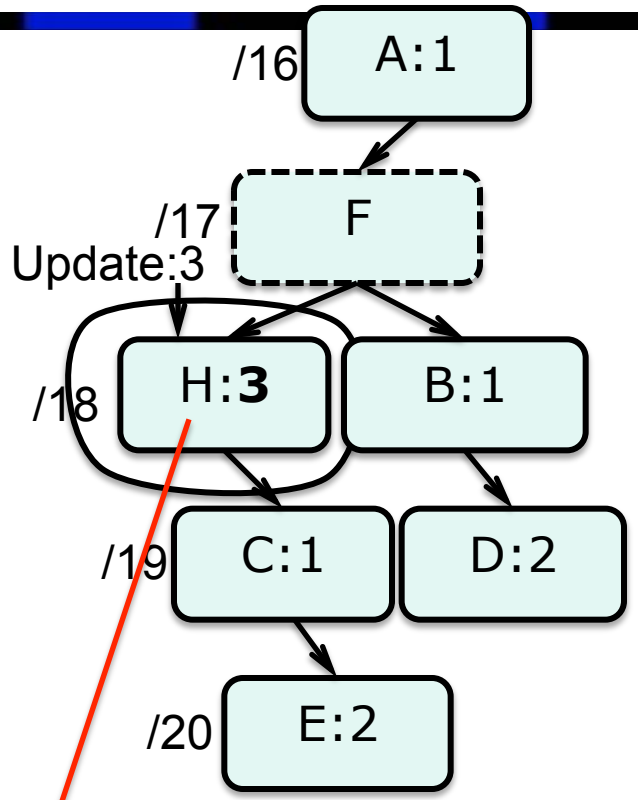

(a) Original FIB Entries

Label	Prefix	Next Hop
A	141.225.0.0/16	1
B	141.225.64.0/18	1
C	141.225.32.0/19	1
D	141.225.96.0/19	2
E	141.225.48.0/20	2

(a) Original FIB Entries

Label	Prefix	Next Hop
A	141.225.0.0/16	1
B	141.225.64.0/18	1
C	141.225.32.0/19	1
D	141.225.96.0/19	2
E	141.225.48.0/20	2
H	141.225.0.0/18	3

FIFA-S Update Handling

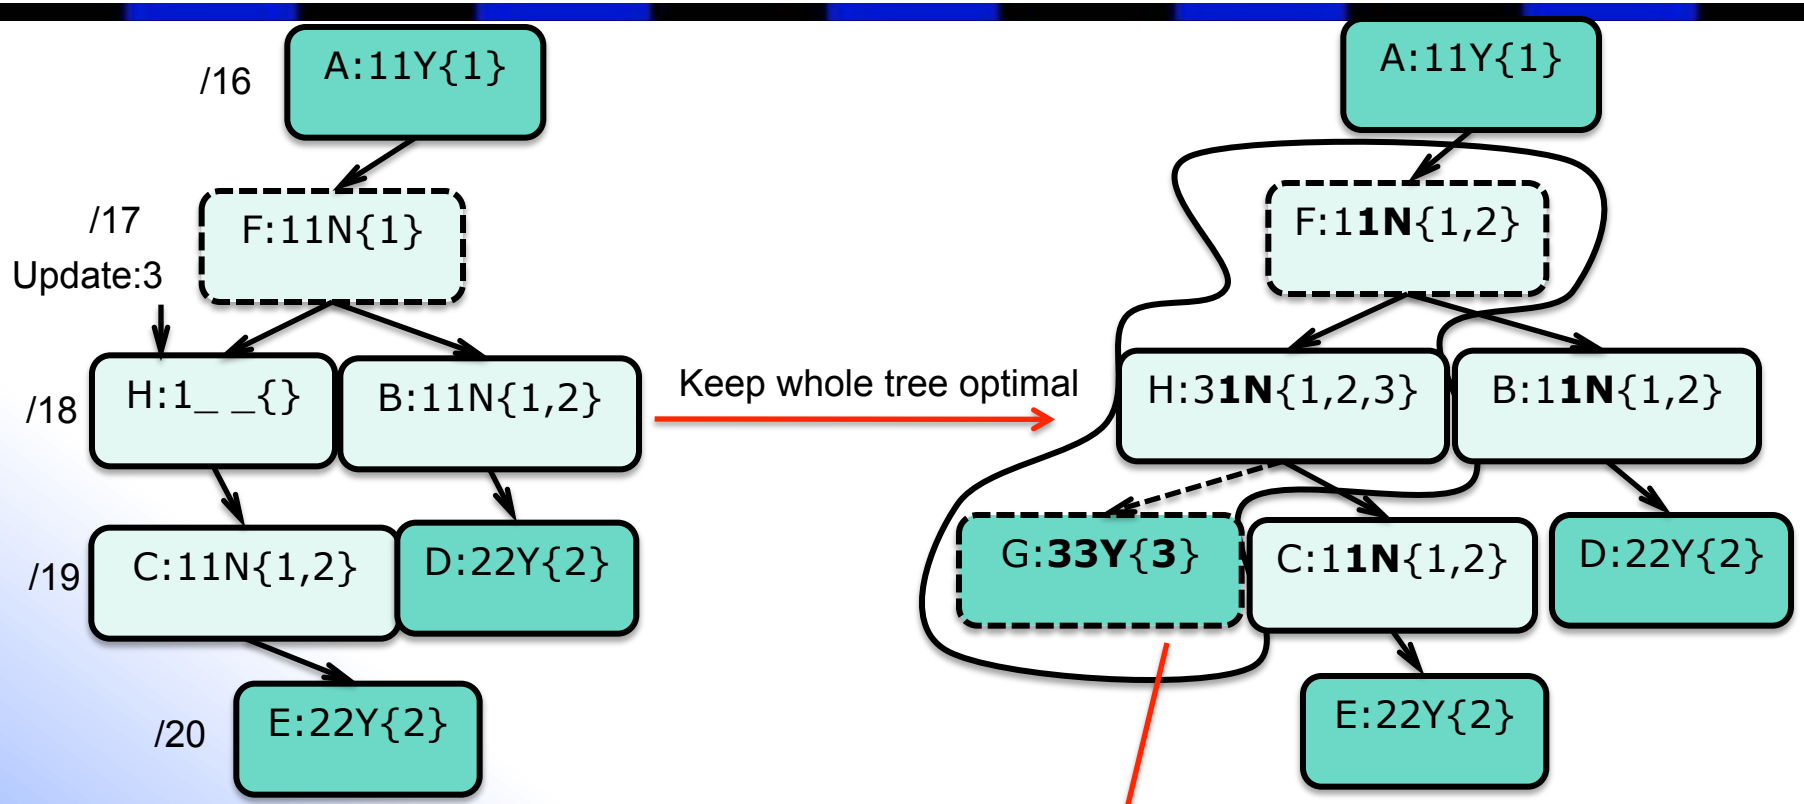

(a) Original FIB Entries

Label	Prefix	Next Hop
A	141.225.0.0/16	1
B	141.225.64.0/18	1
C	141.225.32.0/19	1
D	141.225.96.0/19	2
E	141.225.48.0/20	2
H	141.225.0.0/18	3

(b) Aggregated FIB Entries

Label	Prefix	Next Hop
A	141.225.0.0/16	1
D	141.225.96.0/19	2
E	141.225.48.0/20	2
G	141.225.0.0/19	3

FIB Size under FIFA

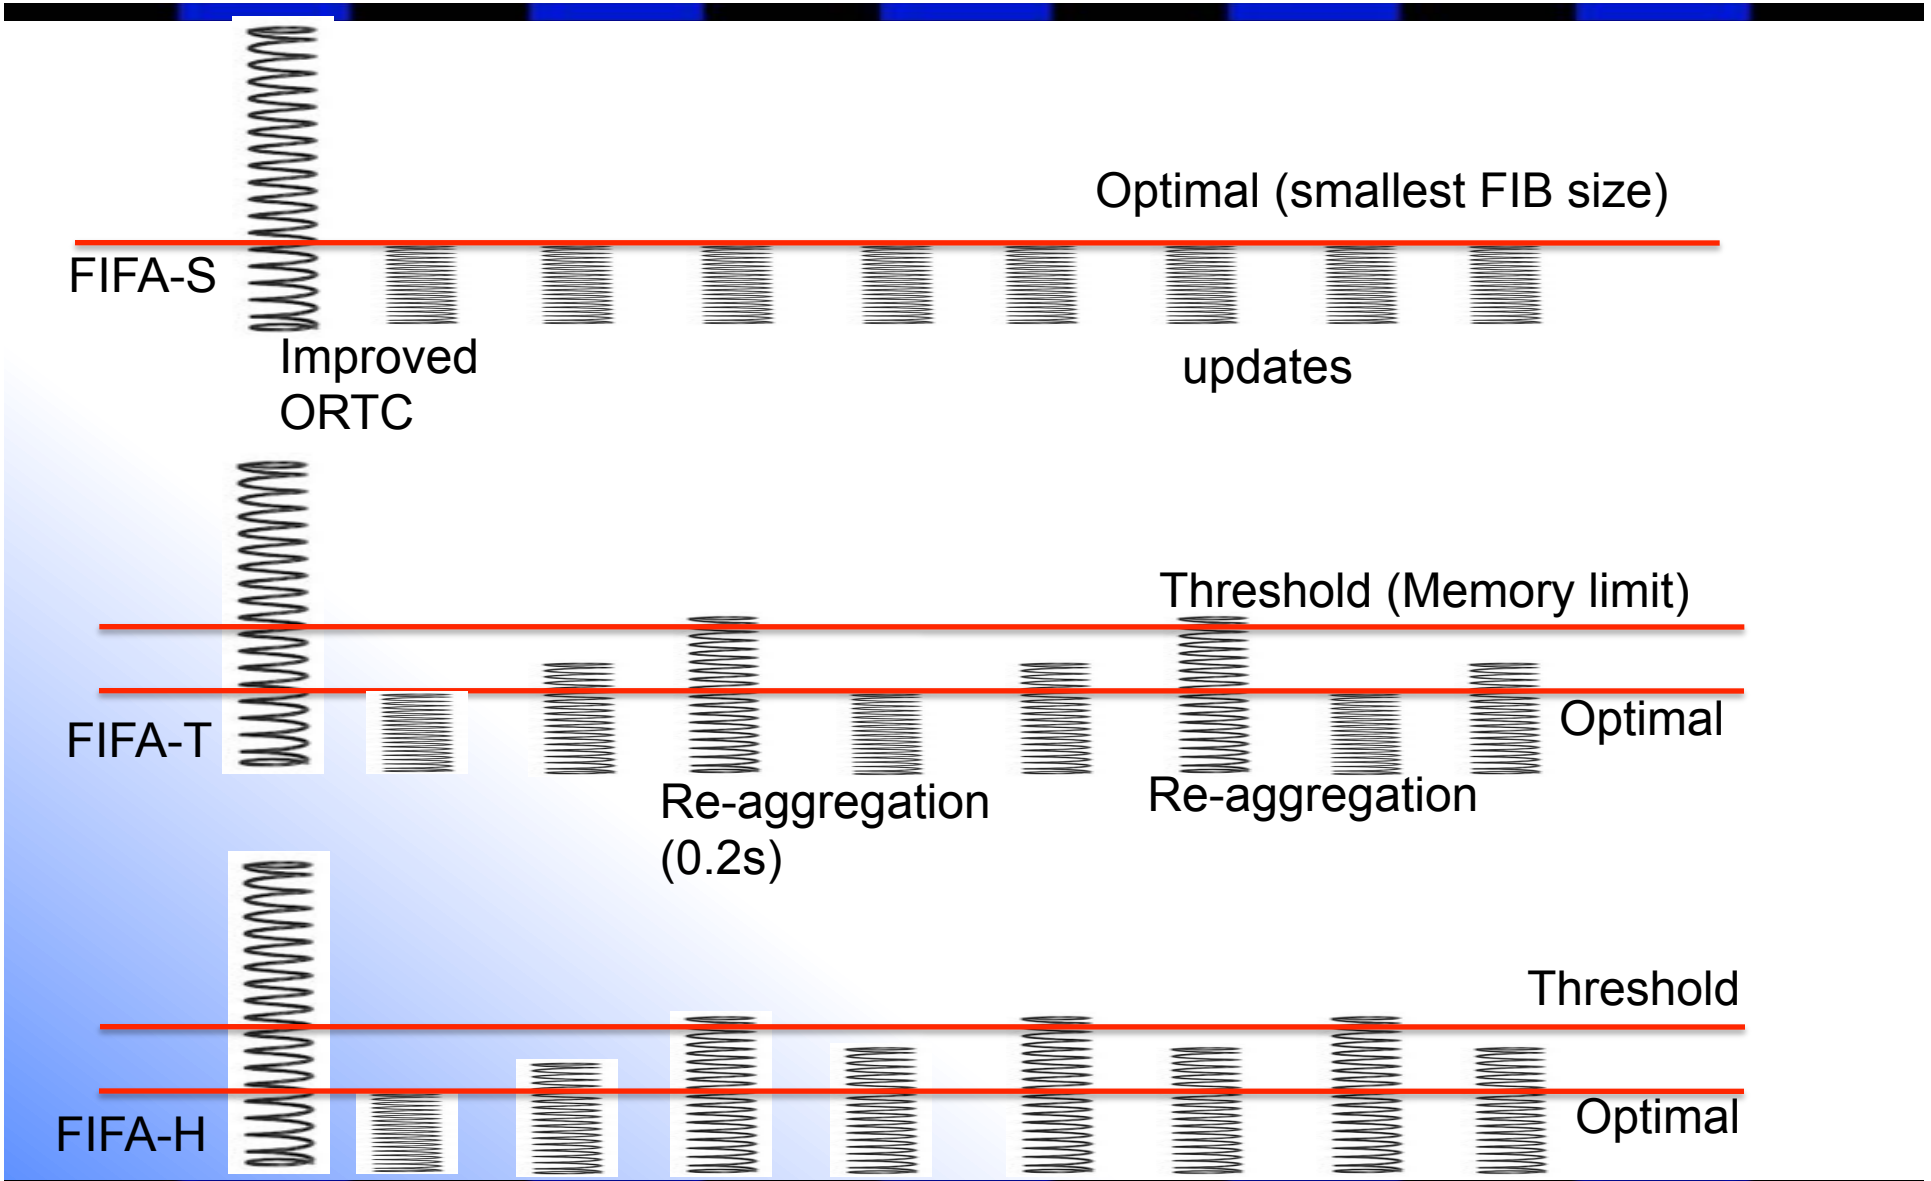

Outline

- Motivation
- FIB aggregation overview
- Update handling
- **Evaluation**
- Operational impact
- Conclusion

Evaluation Methodology (1)

- Data source: route-views ribs and updates
  - Data period: 01/01/2011-12/31/2011
  - Peer: **4.69.184.193 has most next AS hops among 36 peers: ~3000(estimate worst case)**
- Data pre-processing: filter out all duplicate RIB updates
  - Total # updates after filtering: 54,095,965
- Use Next AS Hop as nexthop
  - Prefixes sharing the same next-AS-hop are likely to share the same iBGP neighbor and thus the same next- hop router [3]

Evaluation Methodology (2)

- Experiment on a normal desktop of Intel Core 2 Quad 2.83GHz CPU
- Metrics
  - **FIB size**: total number of entries in aggregated FIB
  - **Time cost**: total update handling time
  - **FIB changes**: total number of FIB updates caused by all RIB updates
  - **FIB burst**: number of FIB updates caused by one individual RIB update, this can be 0

FIB Size

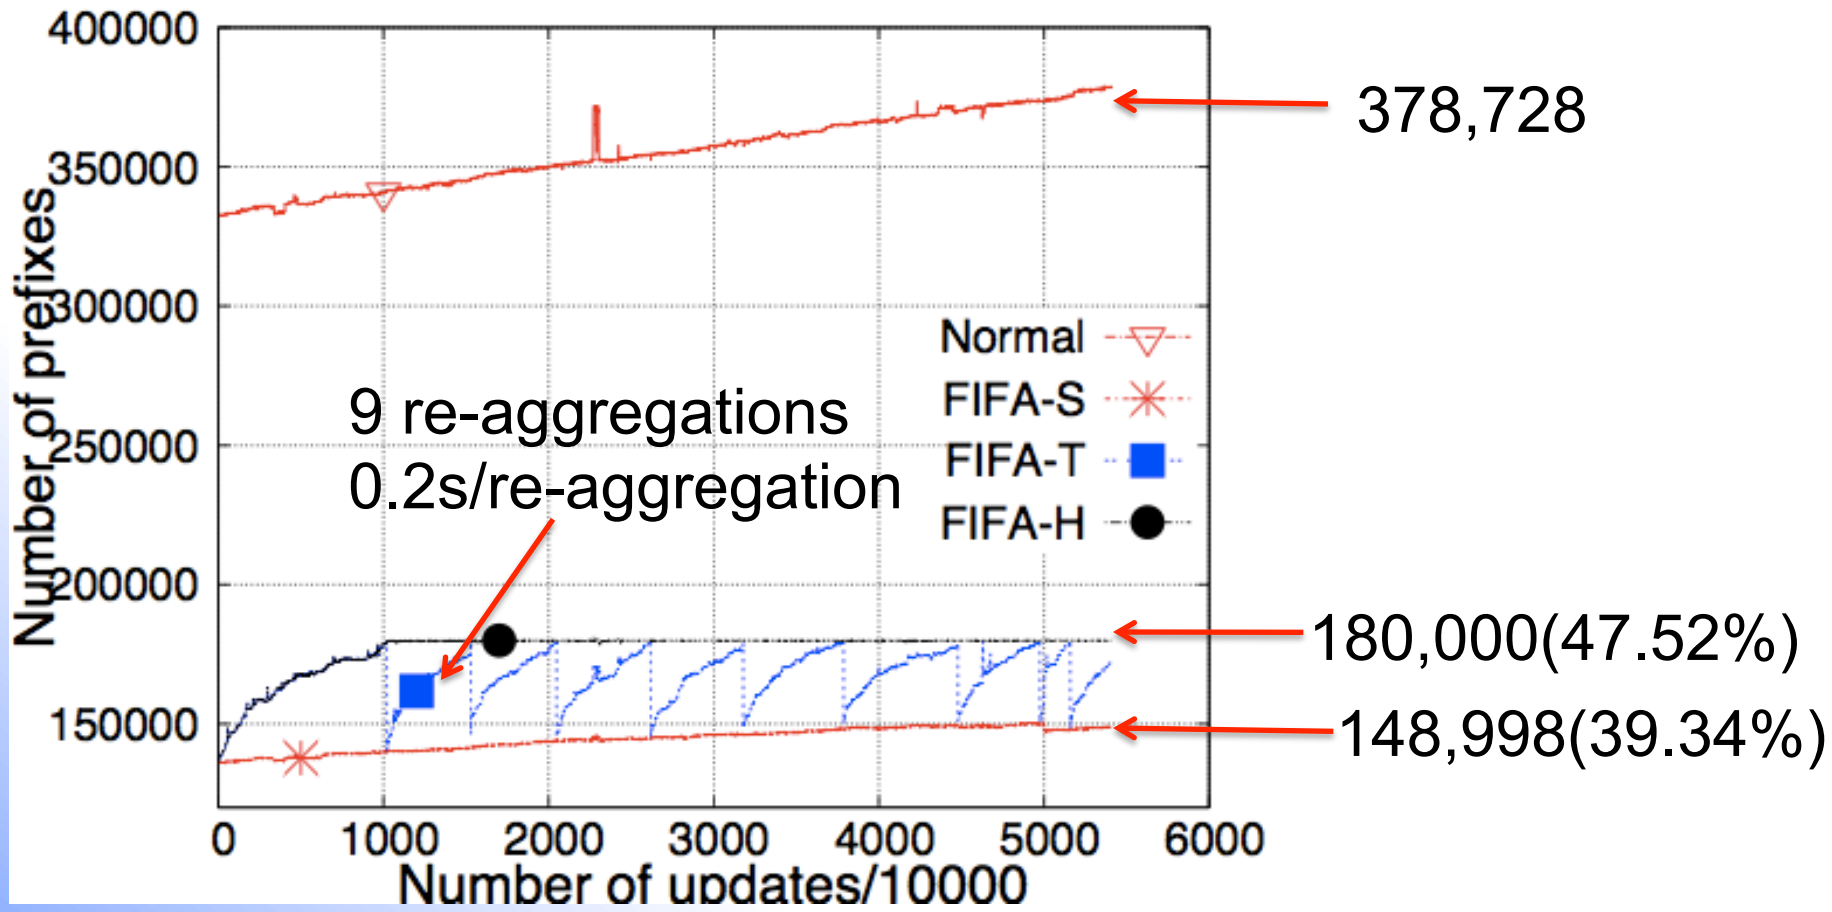

- Aggregate more than 50%

Time

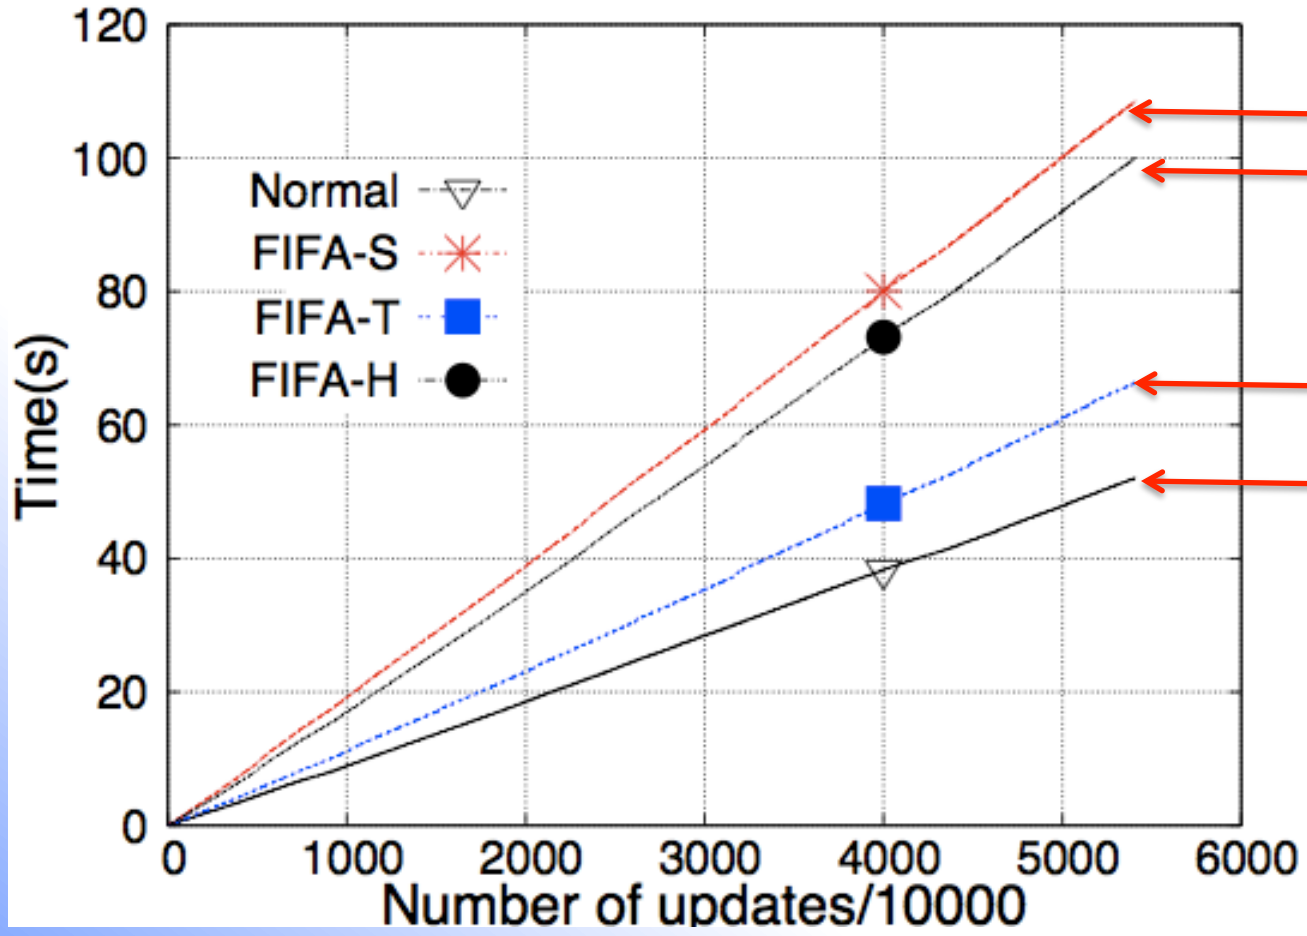

2us/update

1.9us/update

1.2us/update

1us/update

- All affordable

FIB Changes

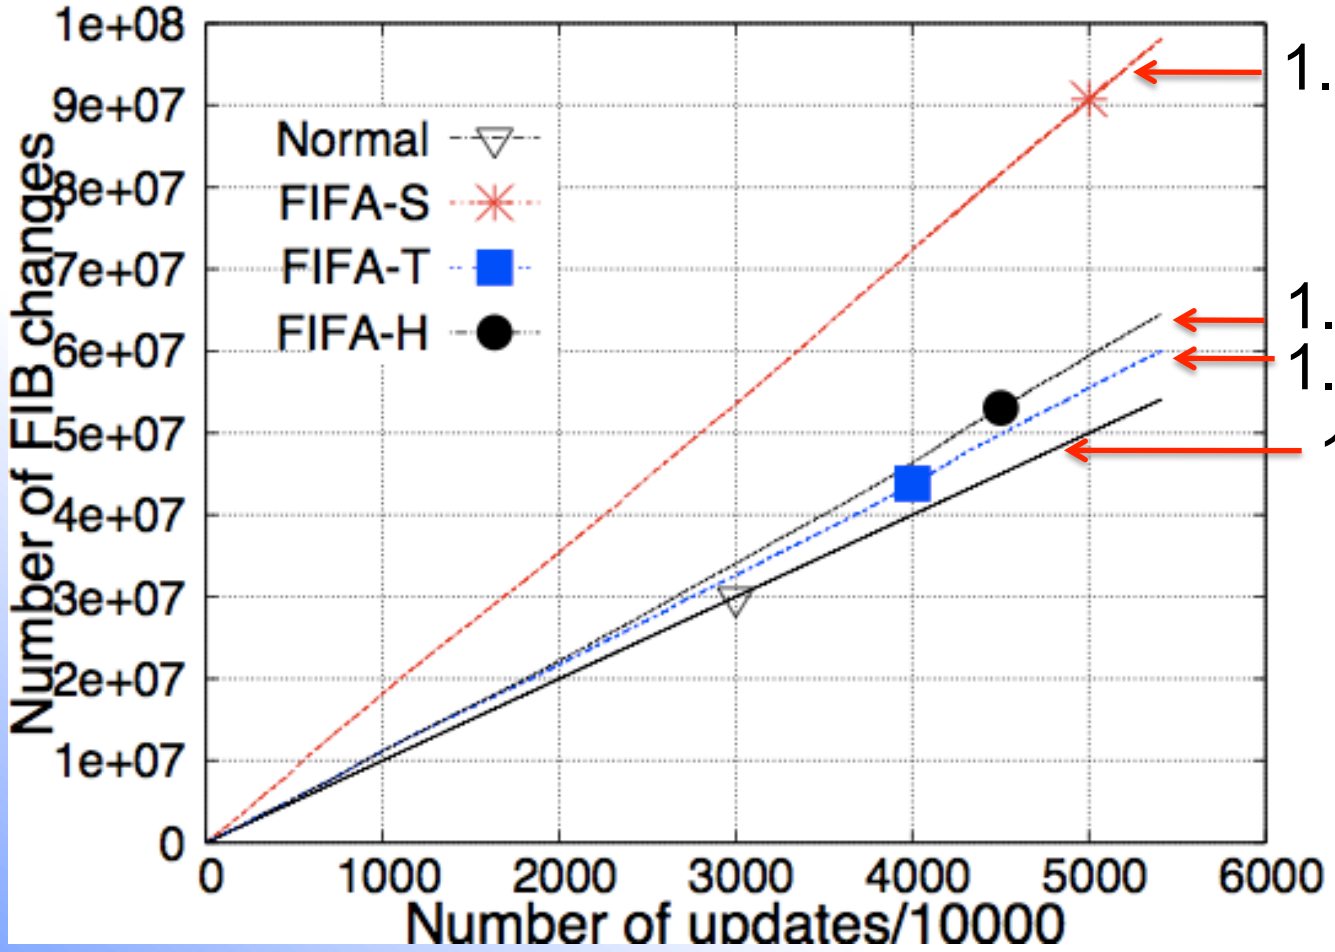

1.8 changes/update

1.2 changes/update

1.1 changes/update

1 change/update

- All acceptable

FIB Burst

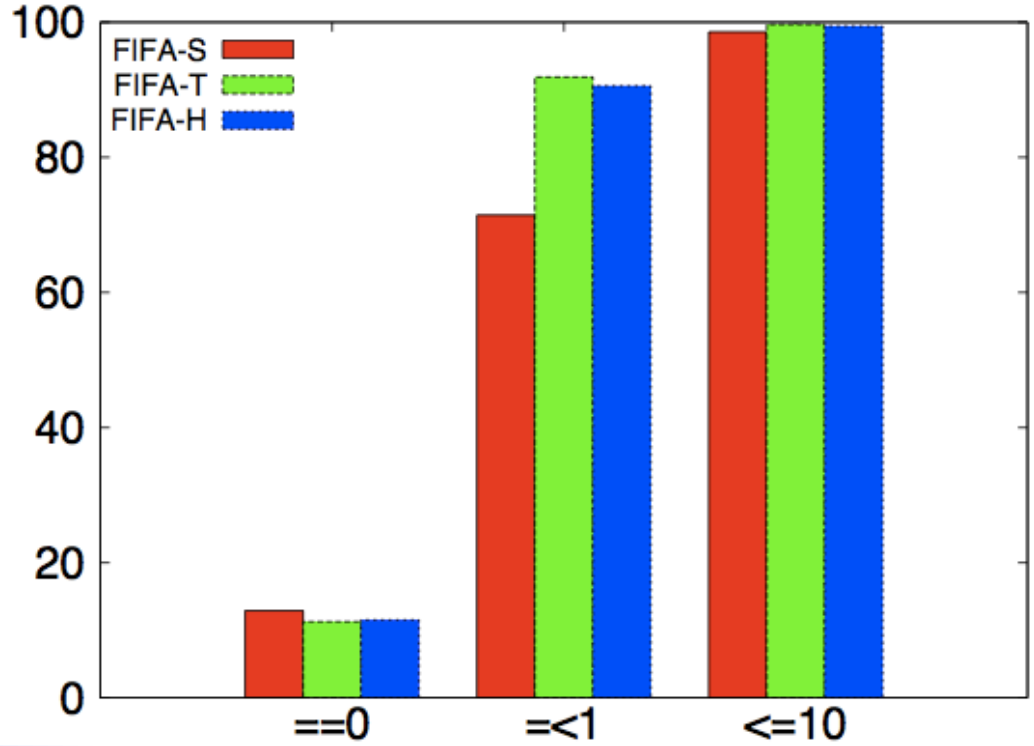

- Most RIB updates cause zero or one FIB change.
- About 99% FIB bursts are less than 10 FIB changes.
- With FIFA-S and FIFA-H, the heaviest FIB bursts have **568** and **1182** FIB changes, respectively
- FIFA-T usually has small FIB bursts, but the heaviest bursts can get very large (**69,526**) due to re-aggregation
- **Pick the one fit for you**

Summary

Attribute/ Option	FIFA-S	FIFA-T	FIFA-H
FIB Size	**Optimal**	**Threshold**	**Threshold**
Time Cost	**Medium**	**Very low**	**Low**
FIB Changes	**High**	**Low**	**Medium**
FIB Burst	**Light**	**Heavy**	**Medium**
Re-Aggregation	**No**	**Fast**	**No**

Benefit

- FIB aggregation: extend a router's life up to 7 years

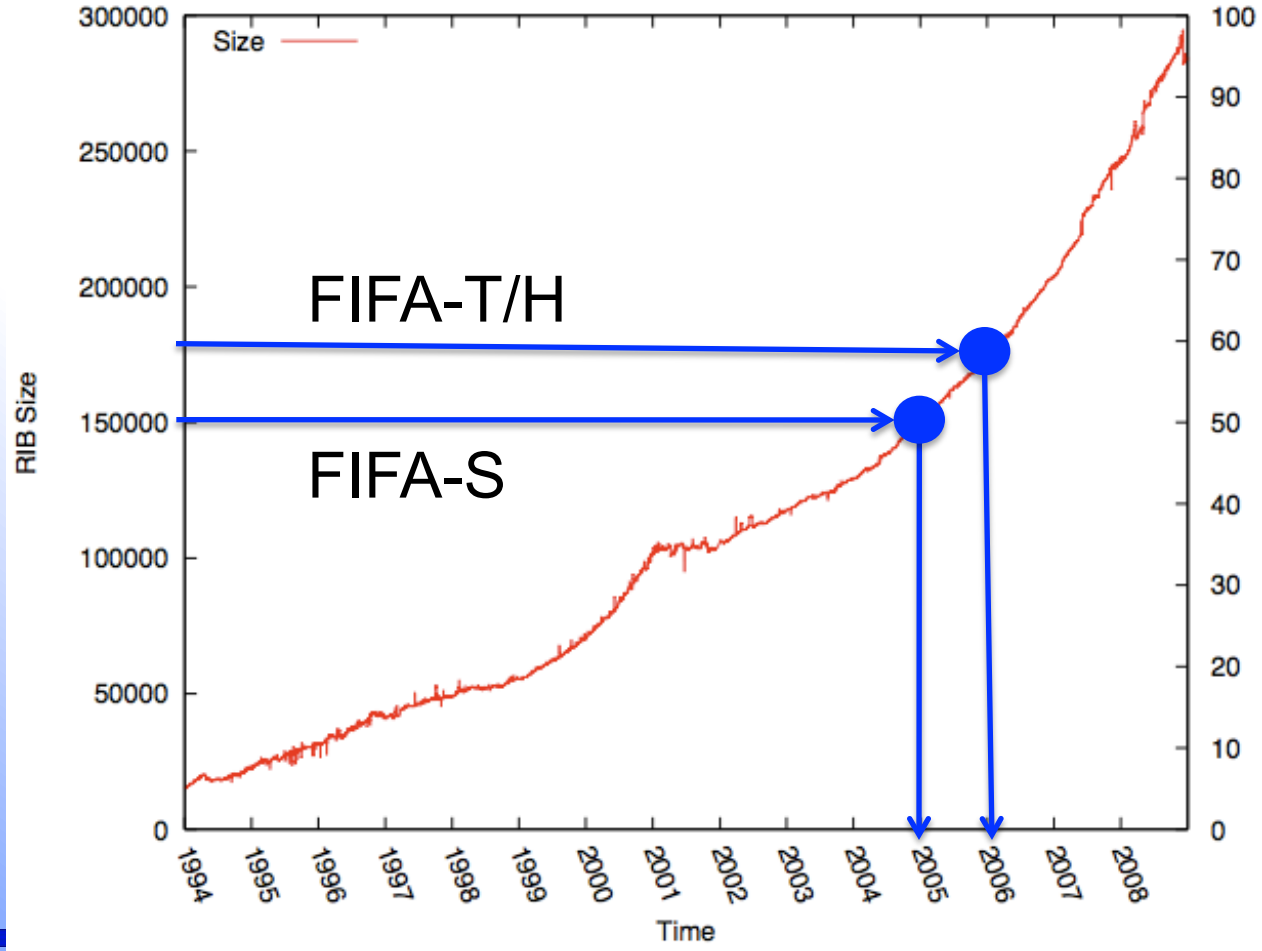

Operational impact

- Need to maintain aggregated FIB table in main memory
- FIFA-T
  - 1.1 times of RIB updates
  - 1.2 times longer processing time for each update to maintain forwarding correctness
  - periodical fast re-aggregations (0.2s)
- FIFA-S
  - 1.8 times of RIB updates
  - 2 times longer processing time to maintain **optimal aggregated FIB size**
  - No re-aggregation
- FIFA-H
  - 1.2 times of RIB updates
  - 1.9 times longer processing time to maintain forwarding correctness
  - No re-aggregation

Conclusion

- Benefits
  - Extend router lifetime for more than five years
  - Small overhead (time, FIB changes)
- Costs (for all FIB aggregations)
  - Additional memory in control plane
  - Handle FIB bursts

References

1. R. Draves, C. King, S. Venkatachary, and B. D. Zill, "**Constructing Optimal IP Routing Tables**," in *Proc. IEEE INFOCOM, 1999.*
2. X. Zhao, D. J. Pacella, and J. Schiller, "**Routing scalability: an operator's view**," *IEEE Journal on Selected Areas in Communications, vol. 28, no. 8, pp. 1262–1270, Oct. 2010.*
3. X. Zhao, Y. Liu, L. Wang, and B. Zhang, "**On the Aggregatability of Router Forwarding Tables**," in *Proc. IEEE INFOCOM, 2010.*
4. Y. Liu, X. Zhao, K. Nam, L. Wang, and B. Zhang, "**Incremental forwarding table aggregation**," in *Proc. IEEE Globecom, 2010.*
5. Z. A. Uzmi, M. Nebel, A. Tariq, S. Jawad, R. Chen, A. Shaikh, J. Wang, and P. Francis, "**SMALTA: practical and near-optimal FIB aggregation**," in *Proc. CoNEXT, 2011.*

Acknowledgement

- NSF grants support
  - Lan Wang: 0721645
  - Beichuan Zhang: 0721863
- Lixia Zhang-UCLA

Thanks!

yliu6@memphis.edu
Yaoqing Liu

Aggregated FIB

ORTC: three passes over the tree
Improved ORTC: two rounds

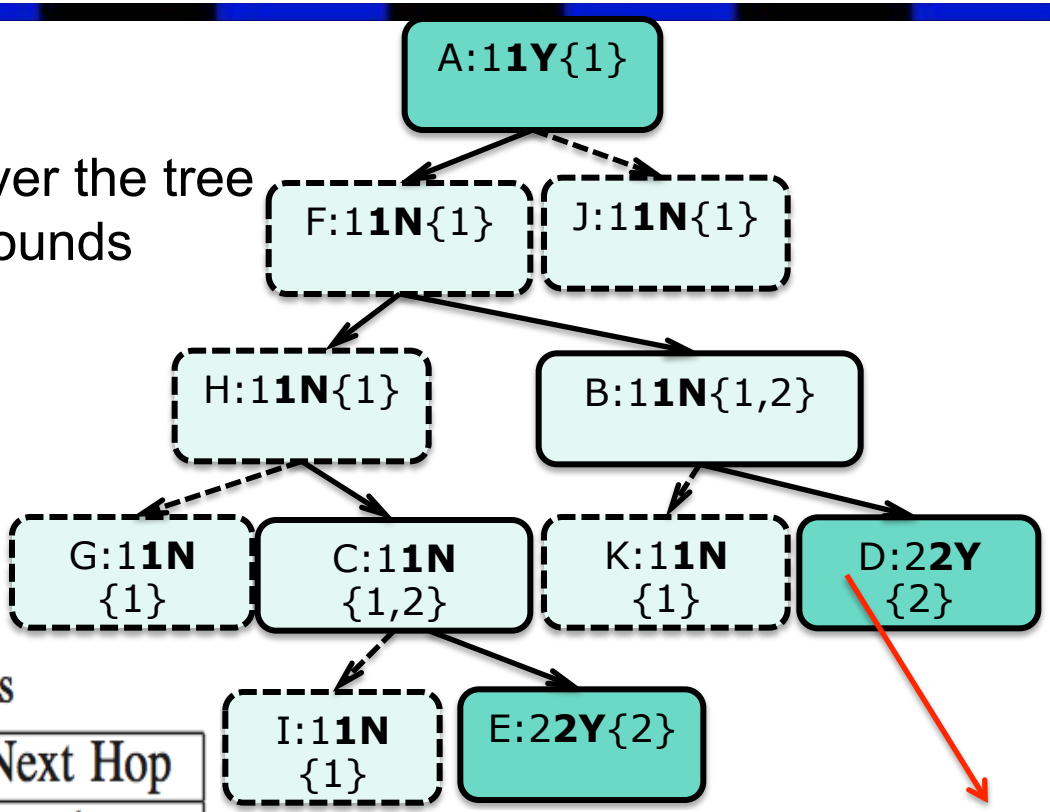

(a) Original FIB Entries

Label	Prefix	Next Hop
A	141.225.0.0/16	1
B	141.225.64.0/18	1
C	141.225.32.0/19	1
D	141.225.96.0/19	2
E	141.225.48.0/20	2

(b) Aggregated FIB Entries

Label	Prefix	Next Hop
A	141.225.0.0/16	1
D	141.225.96.0/19	2
E	141.225.48.0/20	2

FIFA-T/FIFA-H

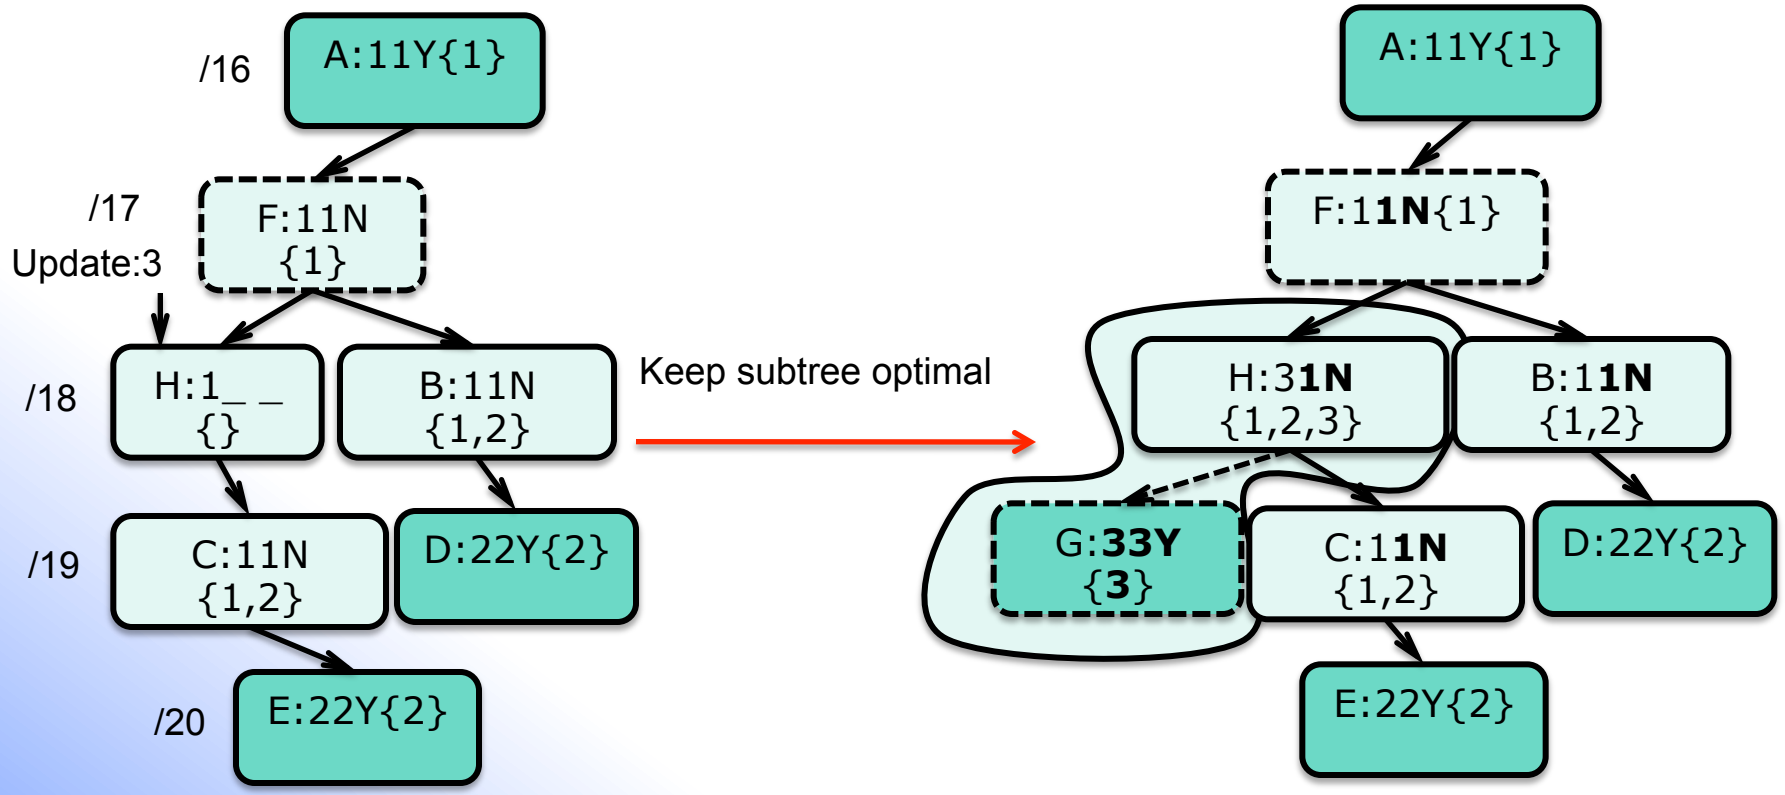

FIFA-T keeps the updated subtree optimal

FIFA-H keeps the updated subtree optimal before threshold

FIFA-H keeps a bigger subtree (under CAP) optimal after threshold

SMALTA

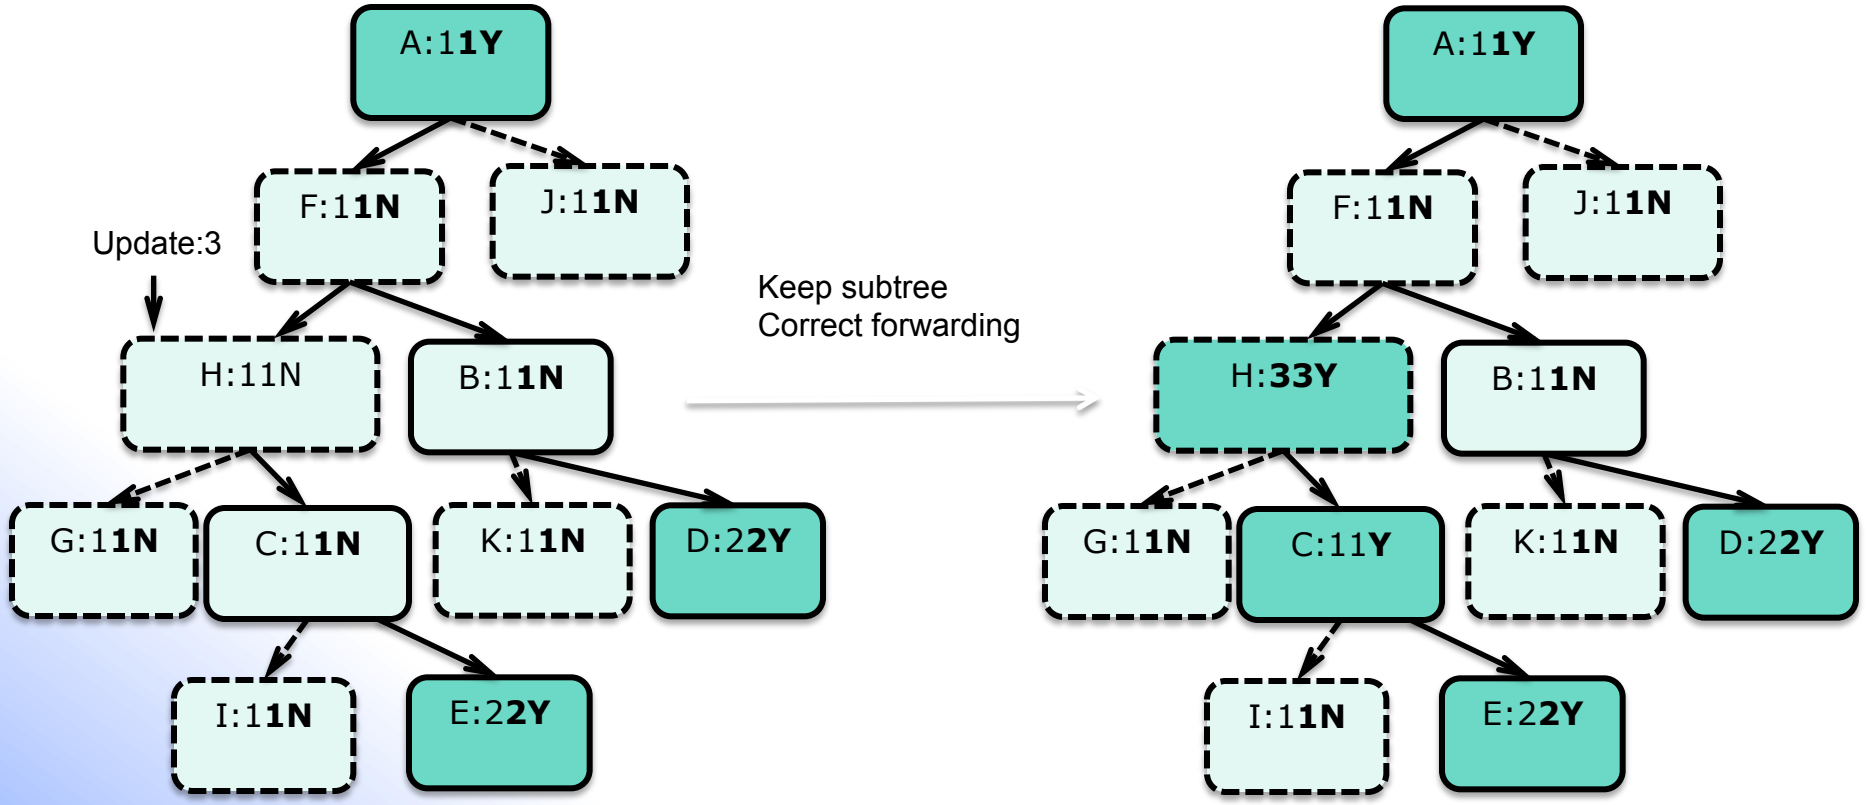

SMALTA keeps the updated subtree correct forwarding for each update

FIB Size

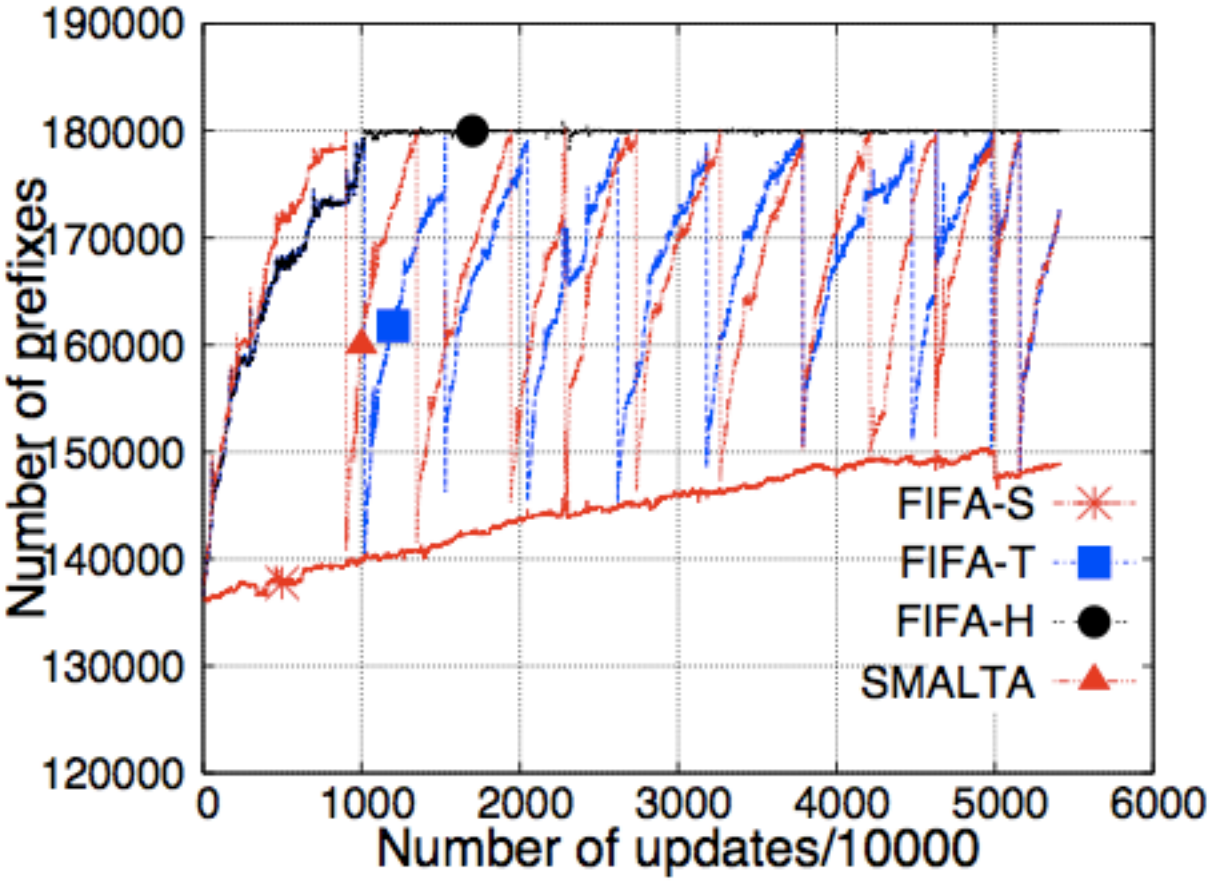

- FIFA-S always has optimal FIB size and never does FIB re-aggregation
- FIFA-T does fast re-aggregations when threshold is reached (180,000).
- FIFA-H keeps close to the threshold size when reaching it and never does FIB re-aggregation
- SMALTA has more re-aggregations

Time

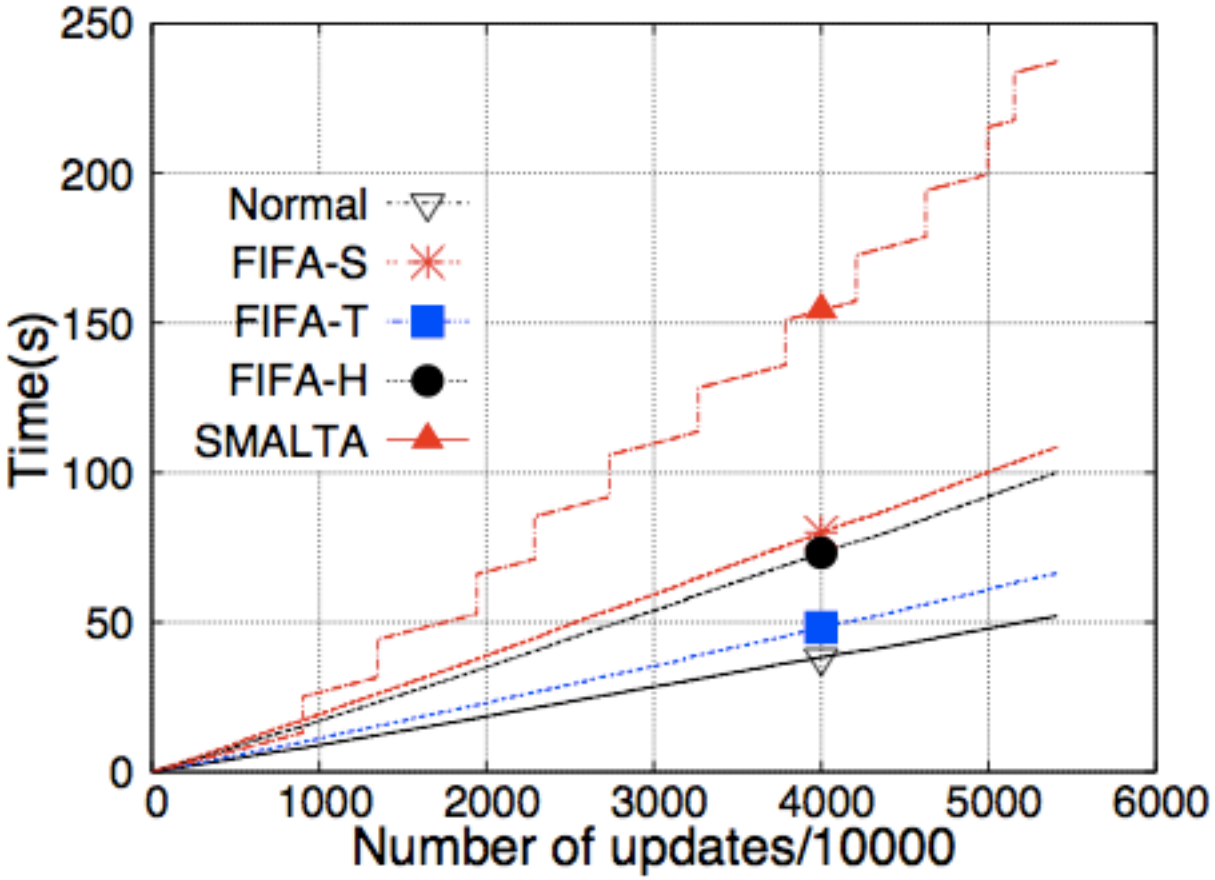

- FIFA-S takes 108s, about 2 times of normal update, 2us/update
- FIFA-T takes 66s, about 1.2 times of normal update,1.2us/update (0.2s/fast re-aggregations)
- FIFA-H takes 100s, about 1.9 times of normal update 1.9us/ update
- SMALTA takes 237s, about 4.5 times of normal update 4.5us/ update(15s/slow re-aggregations)

FIB Changes

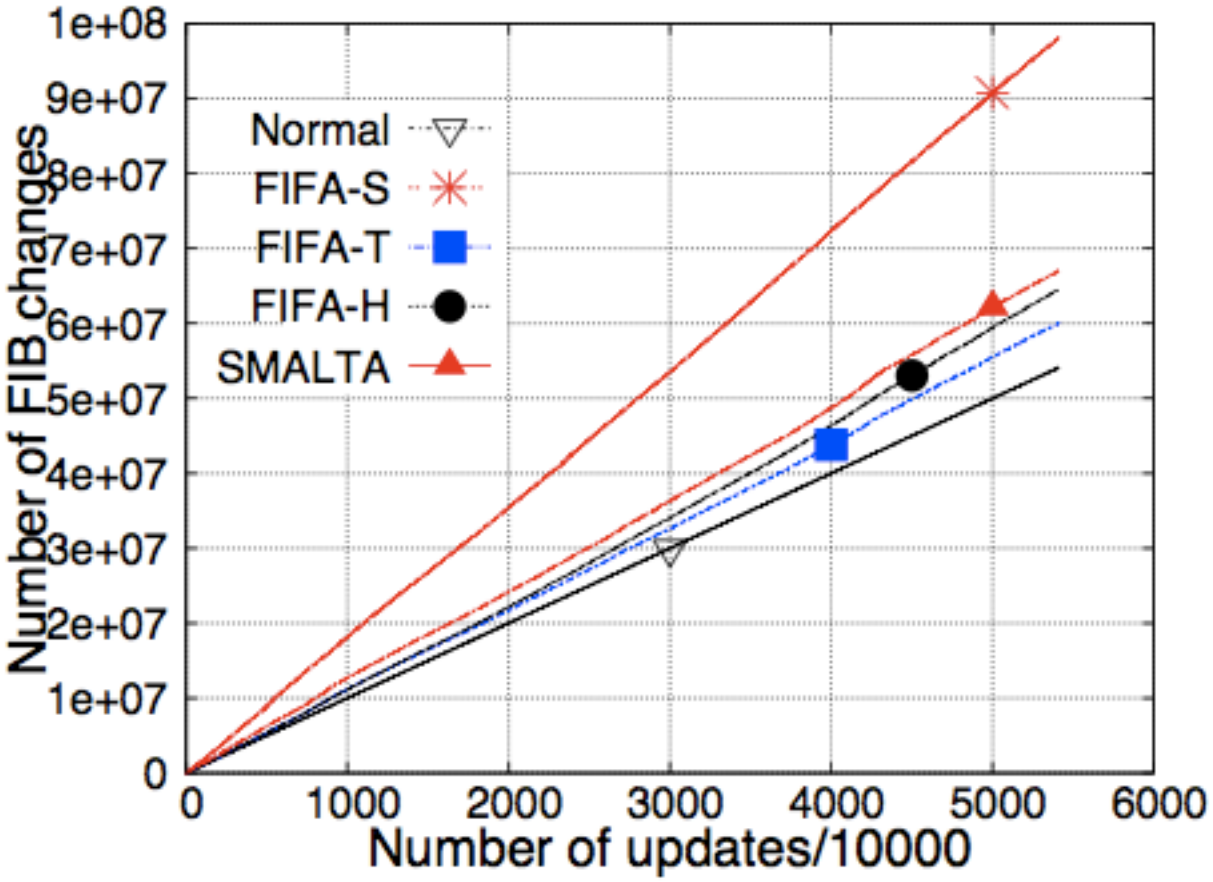

- FIFA-S has 1.8 times of total FIB changes compared to normal update to maintain optimality of the aggregated FIB size (1.8 changes to FIB/update)
- FIFA-T has 1.1 times of total FIB changes compared to normal update
- FIFA-H has 1.2 time of total FIB changes compared to normal update
- SMALTA has 1.2 time of total FIB changes compared to normal update

FIB Burst

FIB Changes	Min	Max	Med	==0	<=1	<=10	Total
FIFA-S	0	568	1	6,961,449	38,645,578	53,318,607	54,095,965
	-	-	-	12.87%	71.43%	98.56%	100%
FIFA-T	0	69,526	1	6,150,664	49,704,177	53,919,736	54,095,965
	-	-	-	11.36%	91.88%	99.67%	100%
FIFA-H	0	1,182	1	6,232,328	48,997,278	53,784,161	54,095,965
	-	-	-	11.52%	90.57%	99.42%	100%

- Most RIB updates cause zero or one FIB change.
- About 99% FIB bursts are less than 10 FIB changes.
- With FIFA-S and FIFA-H, the heaviest FIB bursts have 568 and 1182 FIB changes, respectively
- FIFA-T usually has small FIB bursts, but the heaviest bursts can get very large (re-aggregation)
- **Pick the one fit for you**

FIB Burst

	Min	Max	Med	==0	<=1	<=10	Total
FIFA-S	0	568	1	6,961,449	38,645,578	53,318,607	54,095,965
	-	-	-	12.87%	71.43%	98.56%	100%
FIFA-T	0	69,526	1	6,150,664	49,704,177	53,919,736	54,095,965
	-	-	-	11.36%	91.88%	99.67%	100%
FIFA-H	0	1,182	1	6,232,328	48,997,278	53,784,161	54,095,965
	-	-	-	11.52%	90.57%	99.42%	100%
SMALTA	0	72,856	1	4,456,410	48,297,973	53,873,603	**504,095,976**
	-	-	-	8.23%	89.28%	99.58%	100%

- FIFA-T usually has small FIB bursts, but the heaviest bursts can get very large
- SMALTA has the heaviest FIB burst.

Summary

Attribute/ Option	FIFA-S	FIFA-T	FIFA-H	SMALTA
FIB Size	**Optimal**	**Threshold**	**Threshold**	Threshold
Time Cost	**Medium**	**Very low**	**Low**	High
FIB Changes	**High**	**Low**	**Medium**	Medium
FIB Burst	**Light**	**Heavy**	**Medium**	Heavy
Re-Aggregation	**No**	**Fast**	**No**	Very Slow