

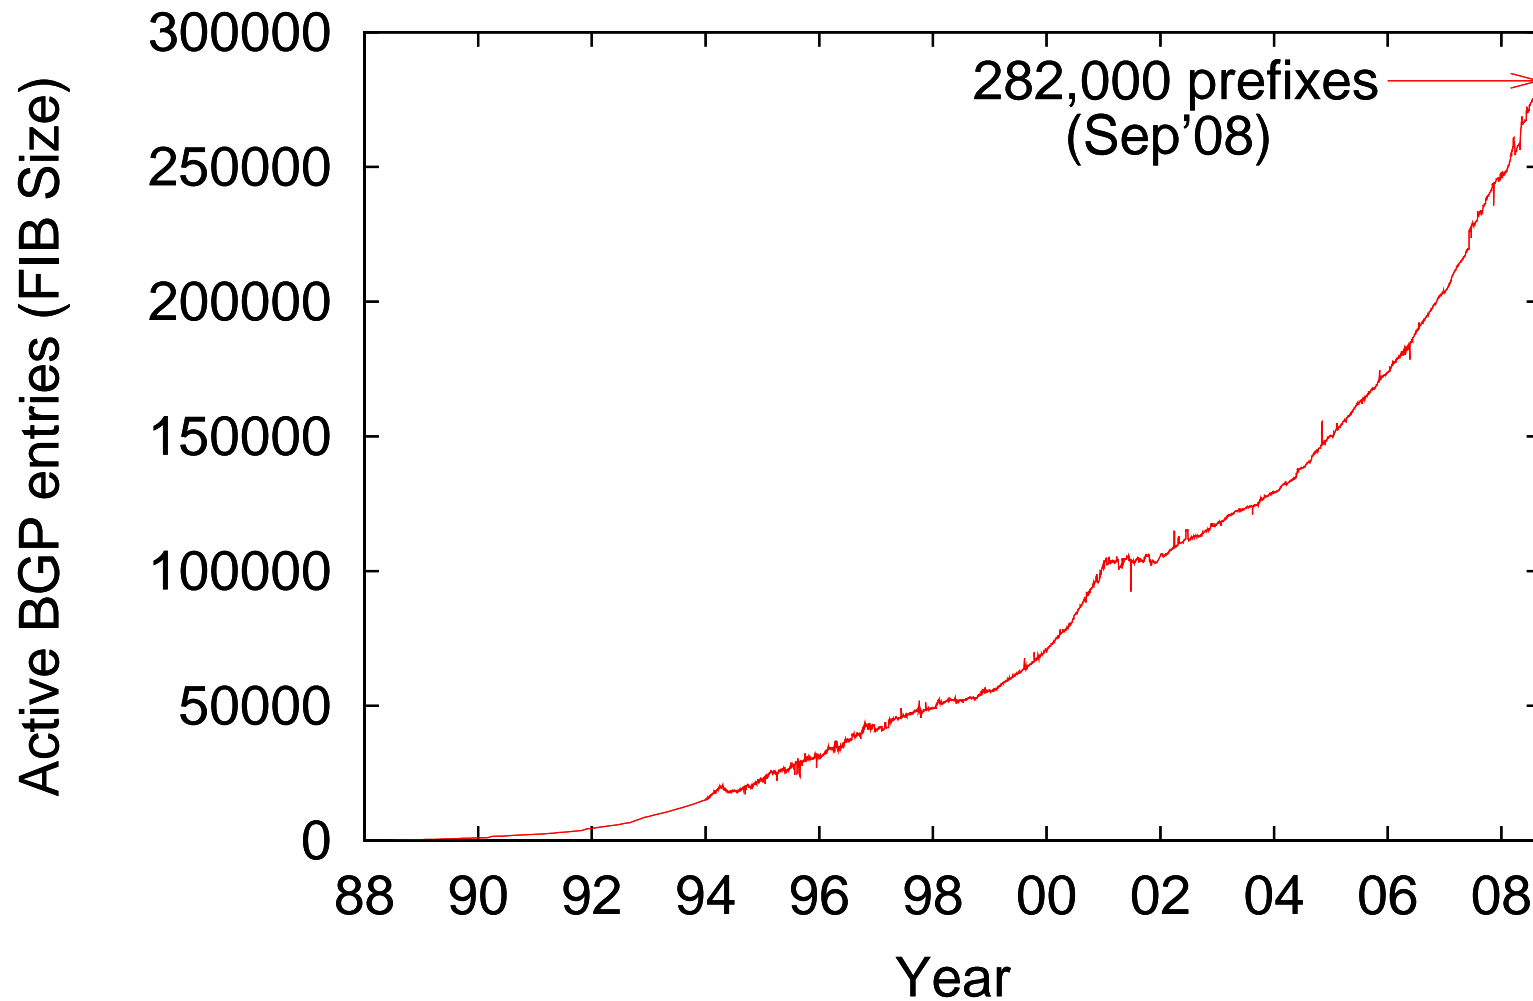
ViAggre: Making Routers Last Longer!

Hitesh Ballani


Paul Francis, Tuan Cao and Jia Wang
Cornell University and AT&T Labs – Research

HotNets 2008

Motivation: Rapid Routing Table Growth



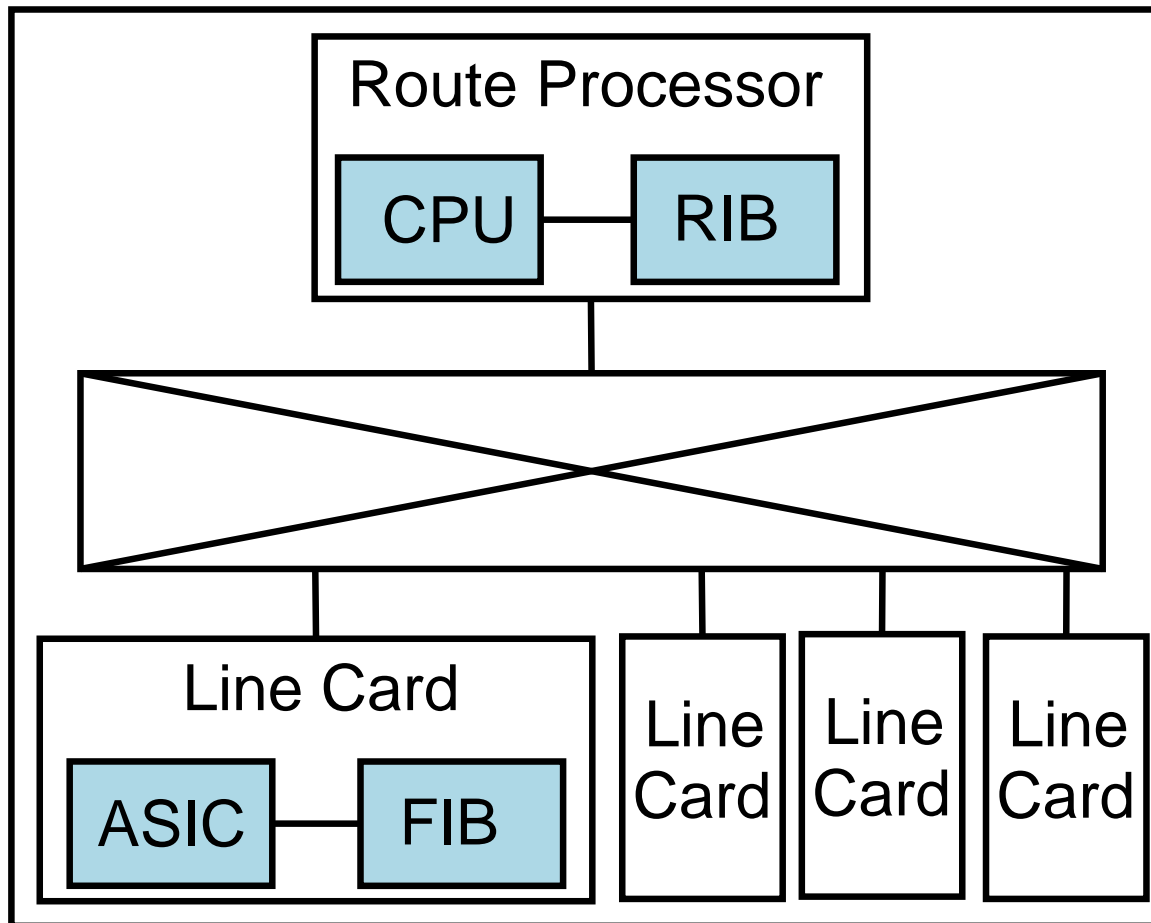
Motivation: Rapid Routing Table Growth



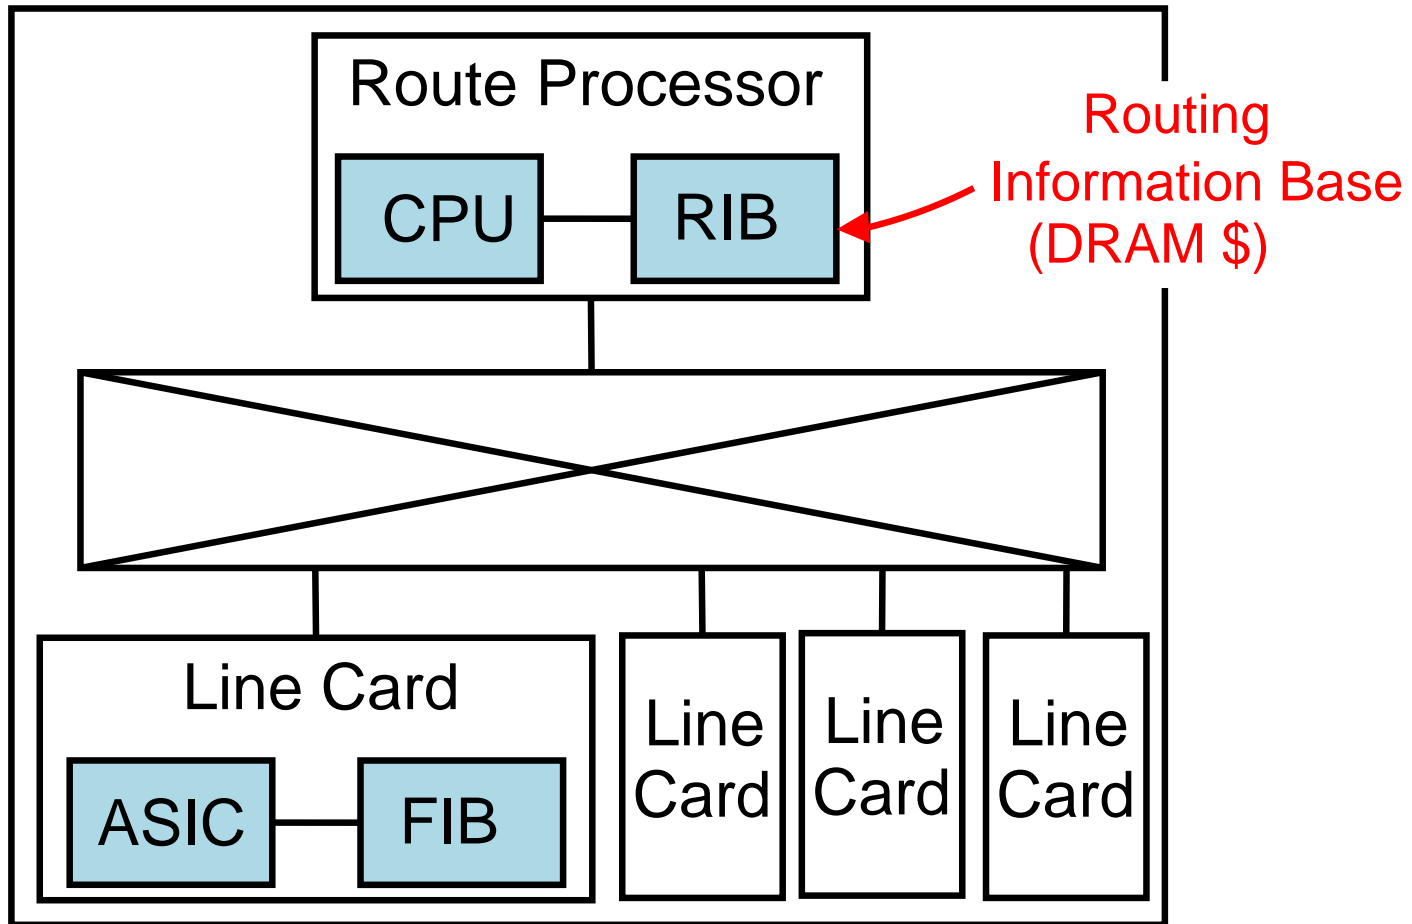
Rapid future growth

- ▶ IPv4 exhaustion
- ▶ IPv6 deployment

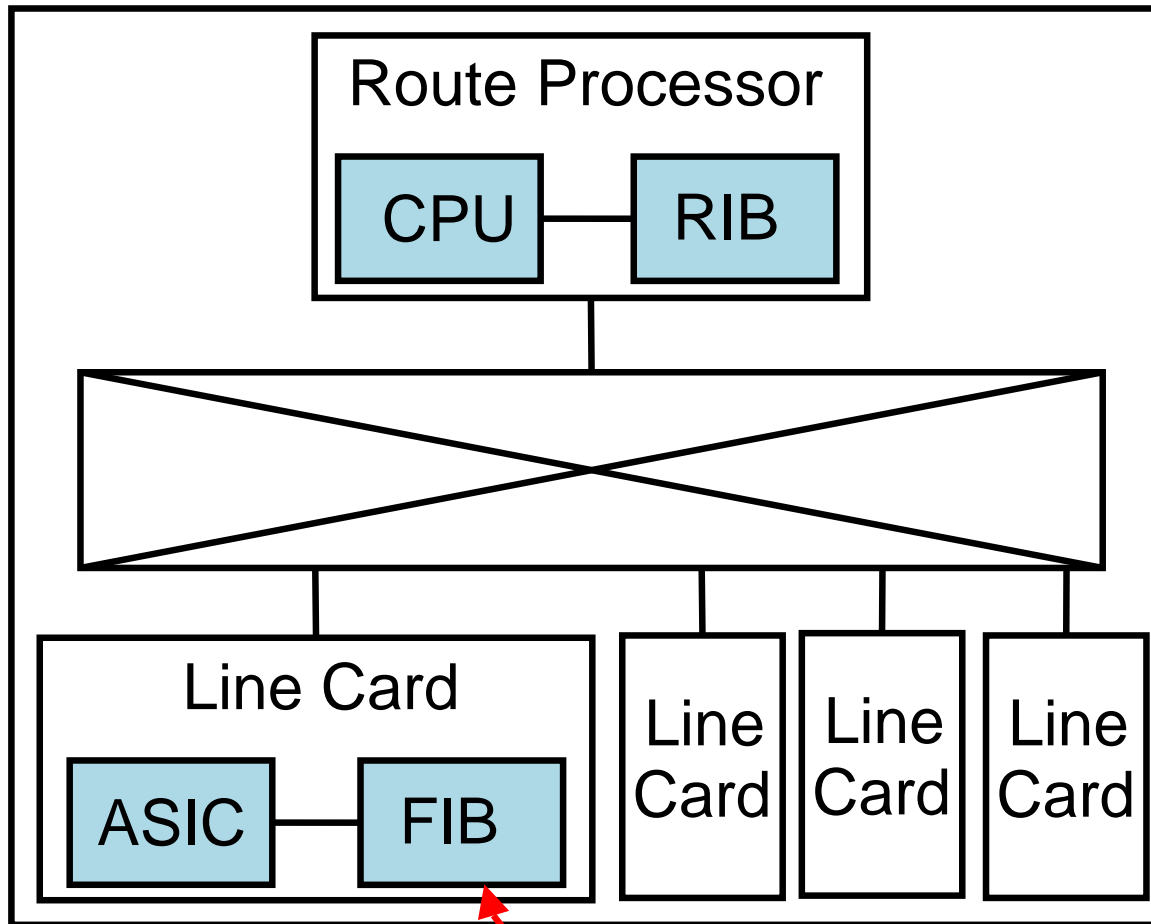
Typical Router Innards



Typical Router Innards



Typical Router Innards



Forwarding Information Base
(SRAM \$\$\$)

Does (FIB) Size Matter?

Technical concerns

- ▶ More Memory
- ▶ More Processing
- ▶ Power and Heat dissipation problems

Does (FIB) Size Matter?

Technical concerns

- ▶ More Memory
- ▶ More Processing
- ▶ Power and Heat dissipation problems

Business concerns

- ▶ Less cost-effective networks
 - ▶ Price per byte forwarded increases
- ▶ Router memory upgrades

Does (FIB) Size Matter?

Technical concerns


- ▶ More Memory
- ▶ More Processing
- ▶ Power and Heat dissipation problems

Business concerns

- ▶ Less cost-effective networks
 - ▶ Price per byte forwarded increases
- ▶ Router memory upgrades

ISPs are willing to undergo some pain to reduce FIB size

Routing Scalability Problem Space



[Deering, '96]

[O'Dell, ID'97]

[Zhang et. al., ICNP'06]

[Farinacci, ID'07]

[Massey et. al., ID'07]

[Jen et. al., HotNets'08]


[Francis, CNIS'94]

[Deering et. al., ID'00]

[Hain, ID'02]

[Krioukov et. al., Arxiv'05]

Routing Scalability Problem Space



[Deering, '96]

[O'Dell, ID'97]

[Zhang et. al., ICNP'06]

[Farinacci, ID'07]

[Massey et. al., ID'07]

[Jen et. al., HotNets'08]

[Francis, CNIS'94]


[Deering et. al., ID'00]

[Hain, ID'02]

[Krioukov et. al., Arxiv'05]

Separate edge from the core

Routing Scalability Problem Space



[Deering, '96]

[O'Dell, ID'97]

[Zhang et. al., ICNP'06]

[Farinacci, ID'07]

[Massey et. al., ID'07]

[Jen et. al., HotNets'08]

[Francis, CNIS'94]


[Deering et. al., ID'00]

[Hain, ID'02]

[Krioukov et. al., Arxiv'05]

Geographical routing

Routing Scalability Problem Space



[Deering, '96]

[O'Dell, ID'97]

[Zhang et. al., ICNP'06]

[Farinacci, ID'07]

[Massey et. al., ID'07]

[Jen et. al., HotNets'08]

[Francis, CNIS'94]


[Deering et. al., ID'00]

[Hain, ID'02]

[Krioukov et. al., Arxiv'05]

Compact routing

Routing Scalability Problem Space



[Deering, '96]

[O'Dell, ID'97]

[Zhang et. al., ICNP'06]

[Farinacci, ID'07]

[Massey et. al., ID'07]

[Jen et. al., HotNets'08]

[Francis, CNIS'94]


[Deering et. al., ID'00]

[Hain, ID'02]

[Krioukov et. al., Arxiv'05]

All require architectural change
So many ideas, so little impact!

Routing Scalability Problem Space



[Deering, '96]

[O'Dell, ID'97]

[Zhang et. al., ICNP'06]

[Farinacci, ID'07]

[Massey et. al., ID'07]

[Jen et. al., HotNets'08]

[Francis, CNIS'94]

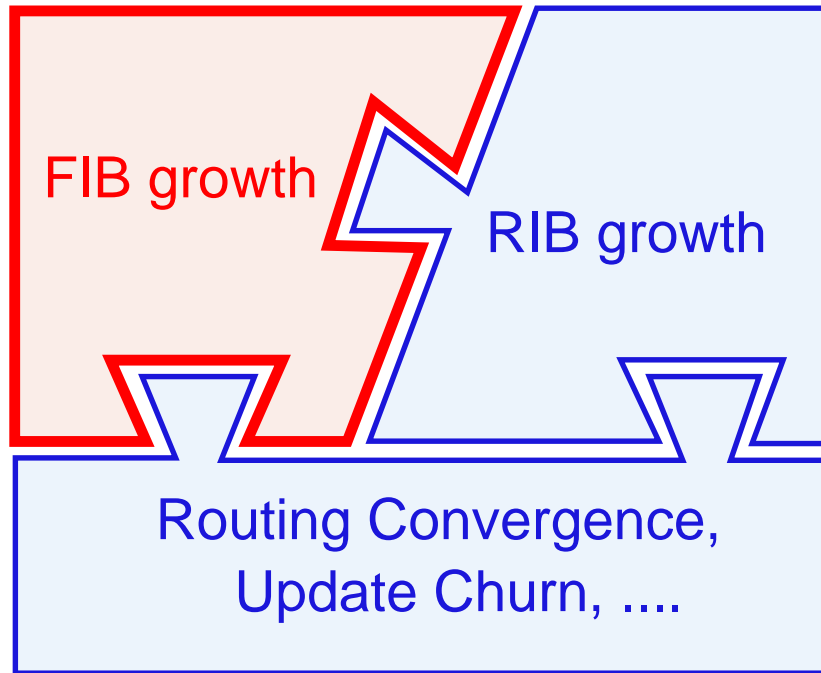
[Deering et. al., ID'00]

[Hain, ID'02]

[Krioukov et. al., Arxiv'05]

Tackle routing scalability through a series of incremental, individually cost-effective upgrades

Routing Scalability Problem Space



[Deering, '96]

[O'Dell, ID'97]

[Zhang et. al., ICNP'06]

[Farinacci, ID'07]

[Massey et. al., ID'07]

[Jen et. al., HotNets'08]

[Francis, CNIS'94]

[Deering et. al., ID'00]

[Hain, ID'02]

[Krioukov et. al., Arxiv'05]

This Paper: Focuses on reducing FIB size

Virtual Aggregation, aka ViAggre

A “configuration-only” approach to shrinking router FIBs

- ▶ Applies to legacy routers
- ▶ Can be adopted independently by any ISP

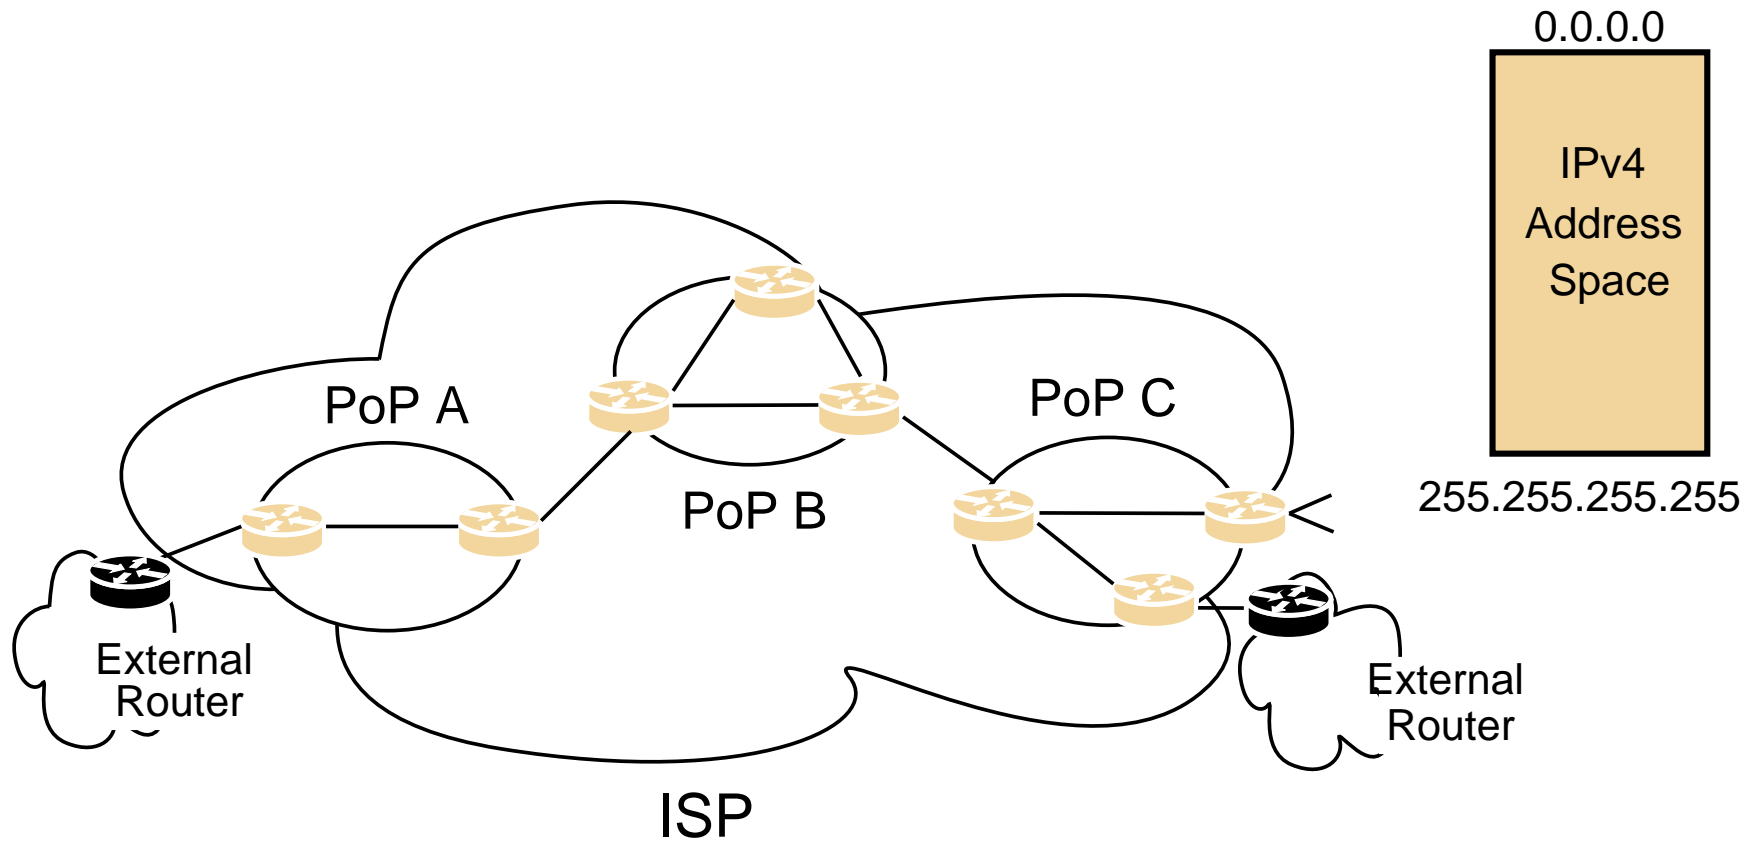
Key Insight: Divide the routing burden

A router only needs to keep routes for a fraction of the address space

Talk Outline


- ▶ Introduction
- ▶ **ViAggre: Basic Idea**
- ▶ ViAggre Design
- ▶ Evaluation
- ▶ Deployment
- ▶ Conclusions

ViAggre: Basic Idea




Today: All routers have routes to all destinations

ViAggre: Basic Idea



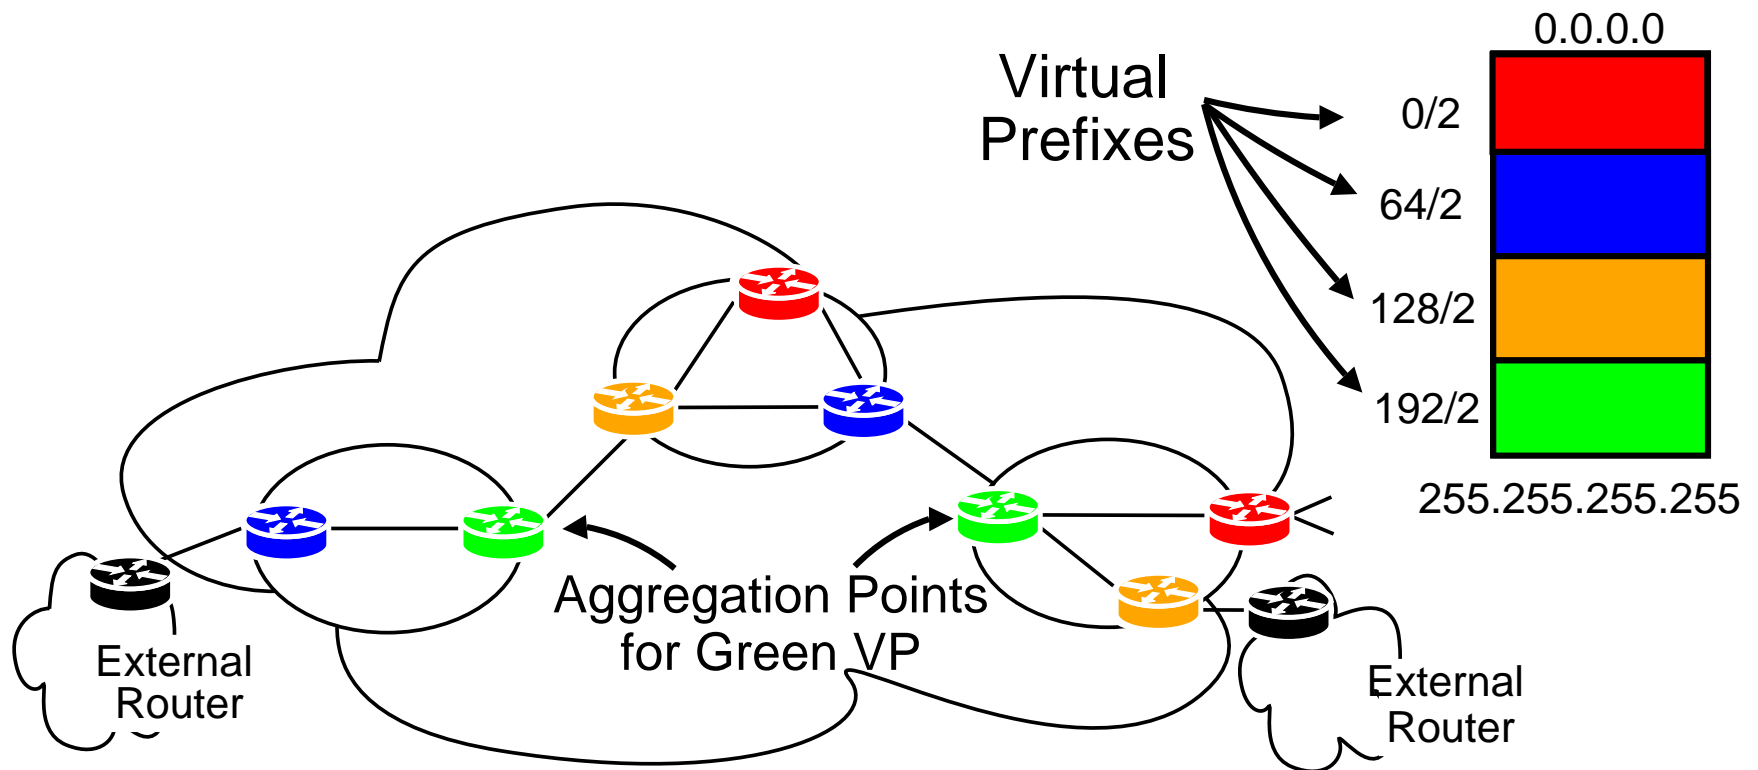
Divide address space into Virtual Prefixes (VPs)

ViAggre: Basic Idea



Assign Virtual Prefixes to the routers
Routers only have routes to a fraction of the address space

ViAggre: Basic Idea



How to achieve such division of the routing table?


Without changes to routers and routing protocols

Without cooperation from external networks

Talk Outline

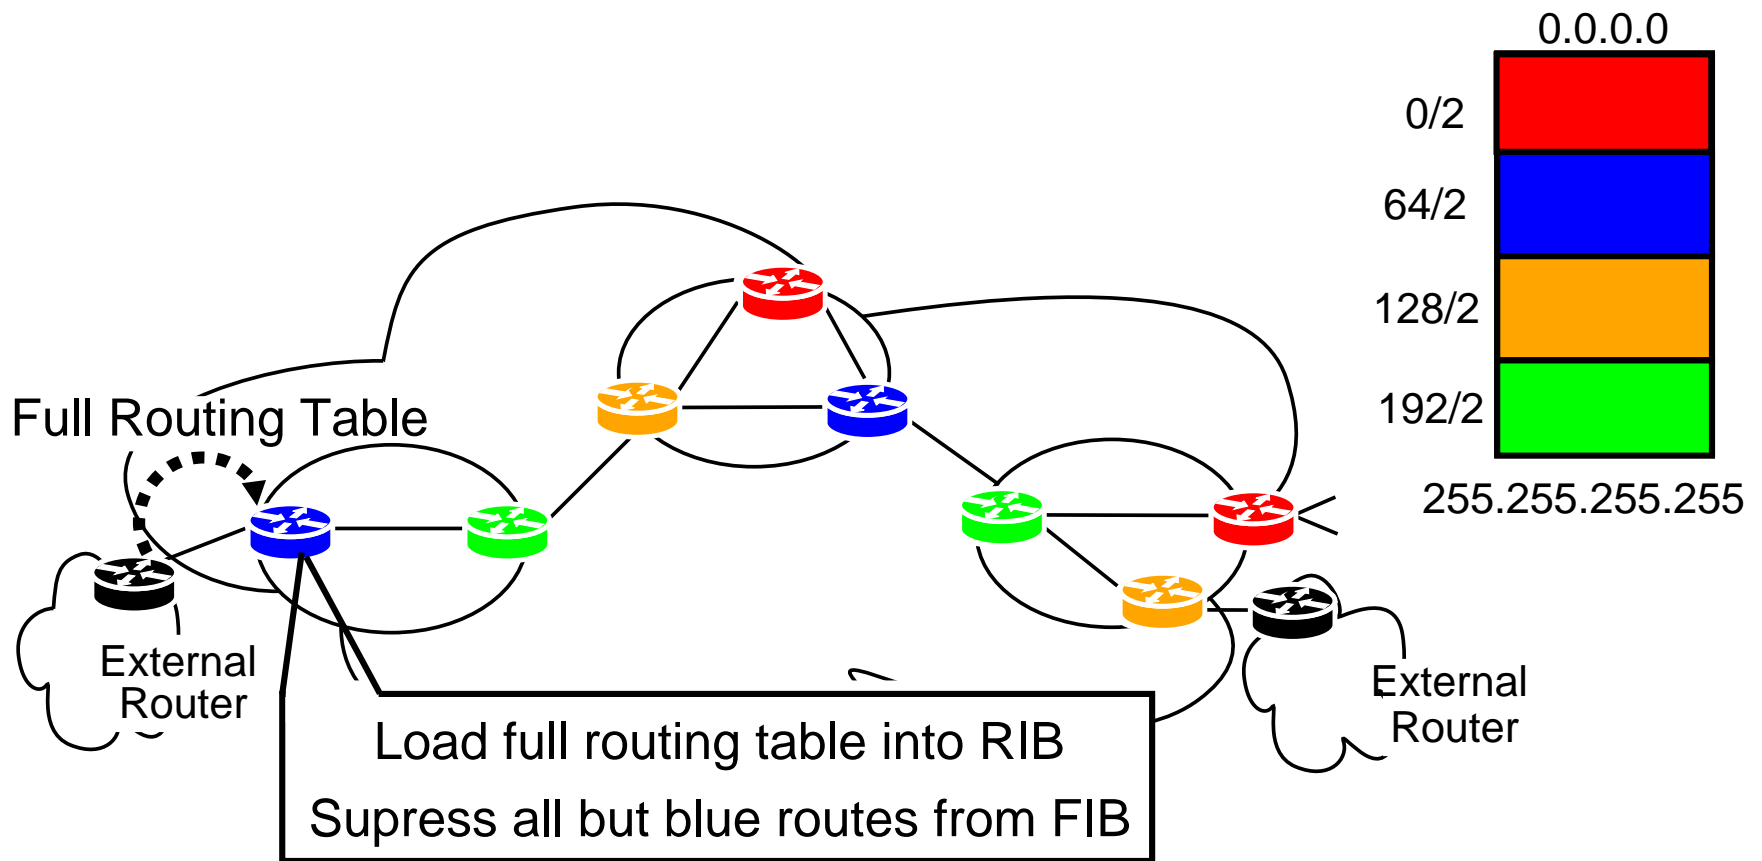
- ▶ Introduction
- ▶ ViAggre: Basic Idea
- ▶ **ViAggre Design**
- ▶ Evaluation
- ▶ Deployment
- ▶ Conclusions

ViAggre Control-Plane



eBGP Peers may advertise full routing table


ViAggre Control-Plane



FIB Suppression


Blue routers only load blue routes into their FIB

Data-Plane paths



Consider packets destined to a prefix in the red VP


Data-Plane paths



ViAggre path

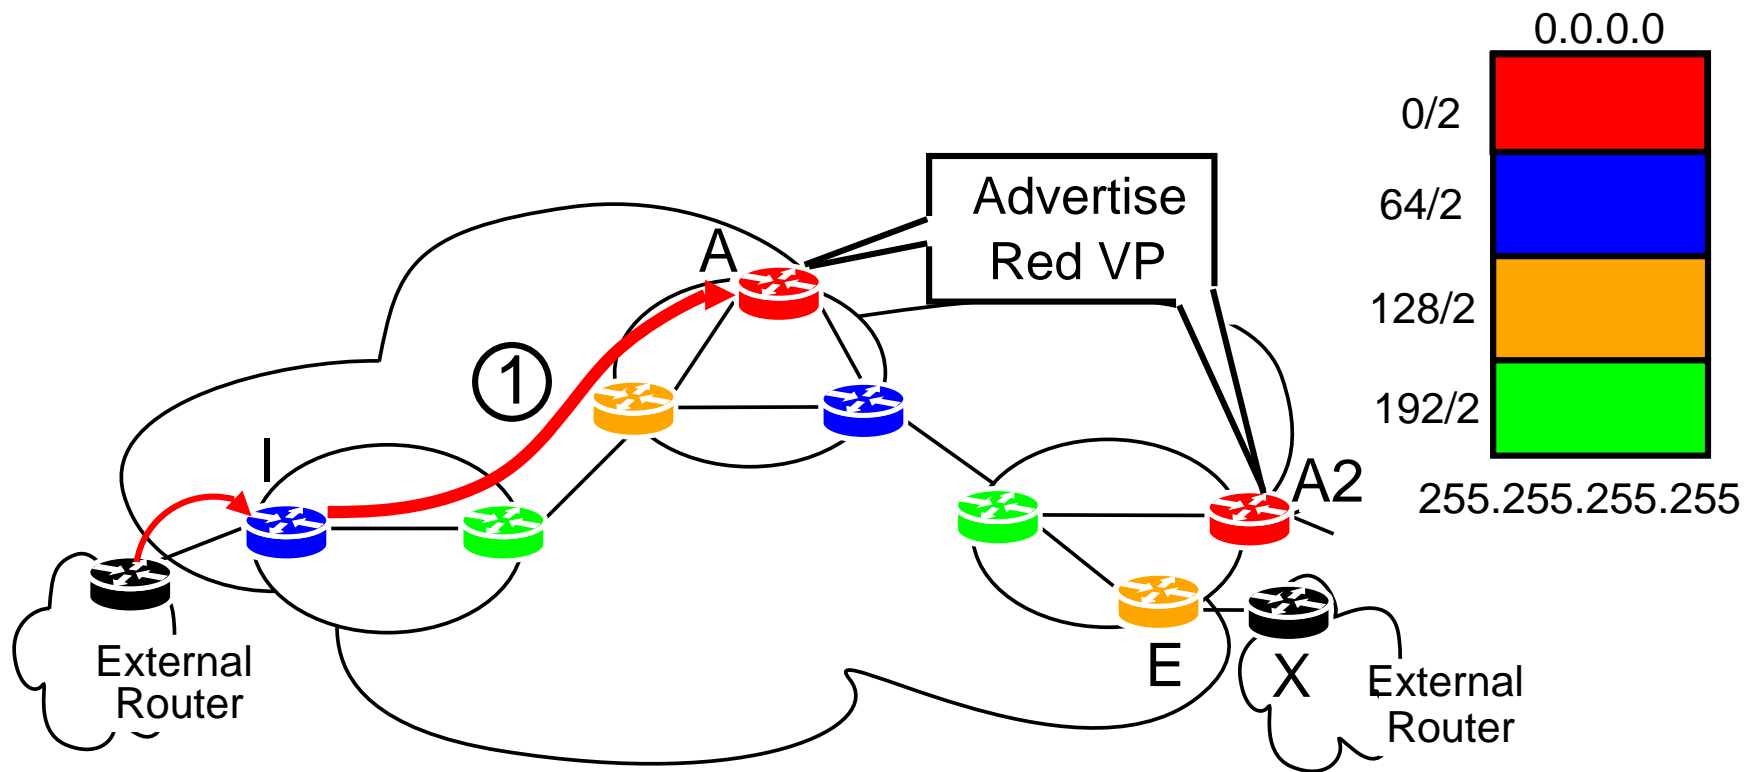
Ingress (I) → Aggregation Pt (A) → Egress (E)

Ingress → Aggregation Point



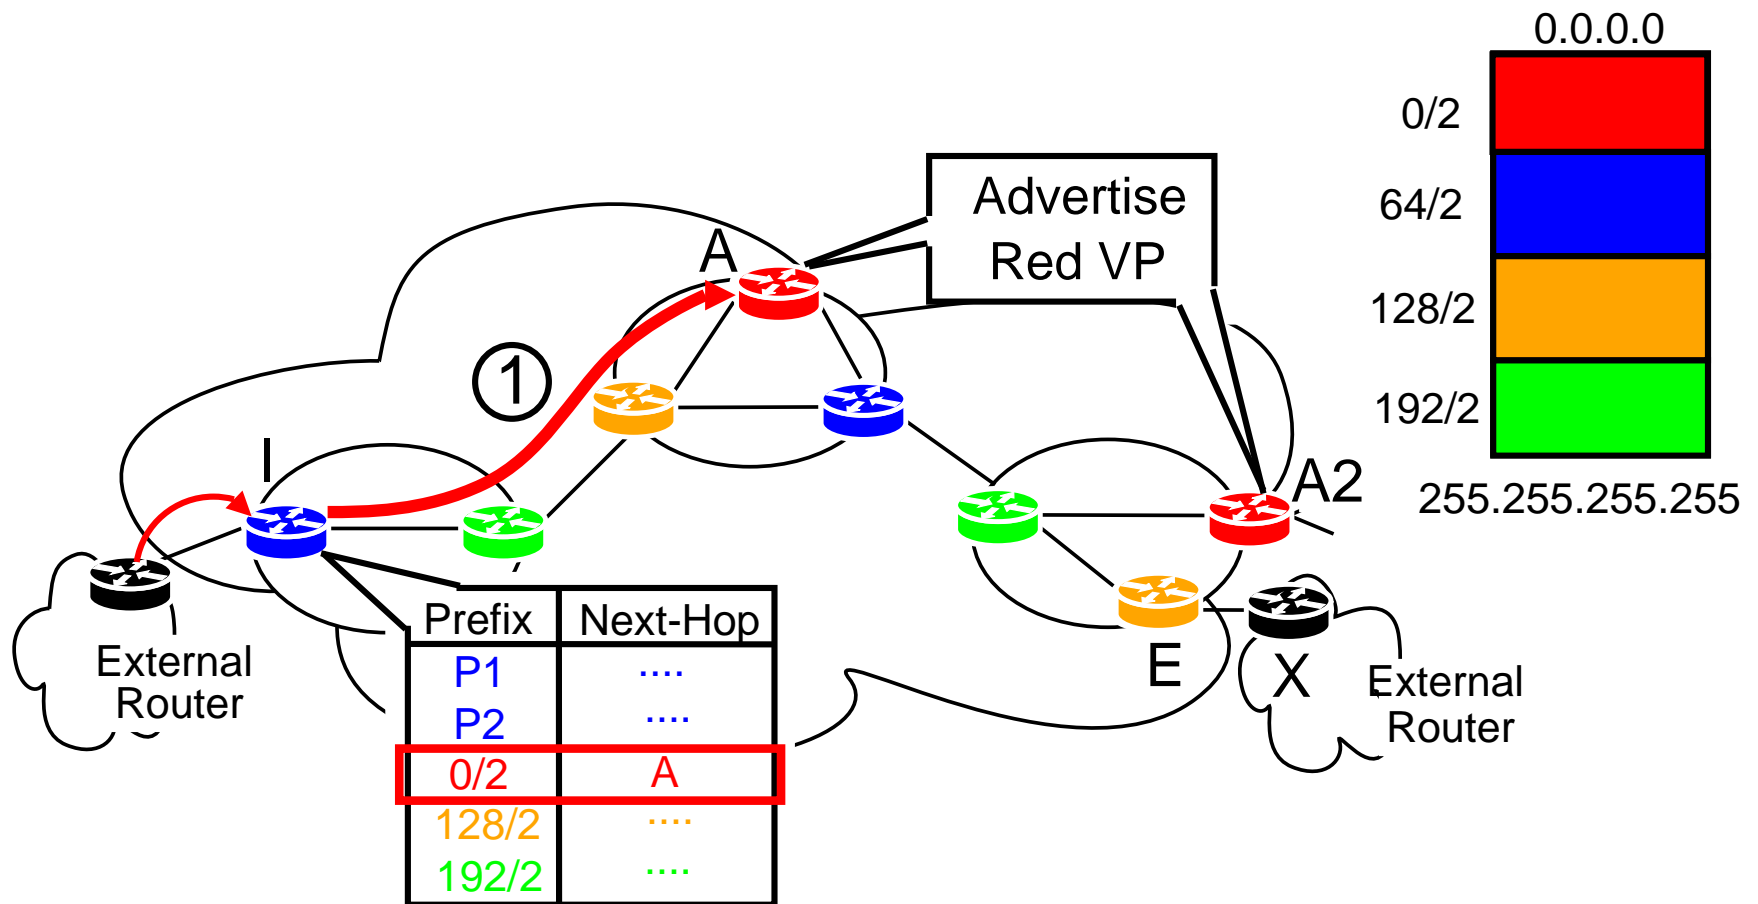
Router I doesn't have a route for destination prefix

Ingress → Aggregation Point




Aggregation Points advertise corresponding Virtual Prefixes

Ingress → Aggregation Point




Blue router has a route for the red Virtual Prefix

Aggregation Point → Egress




Aggregation Pt. A tunnels packet to external router

Aggregation Point → Egress



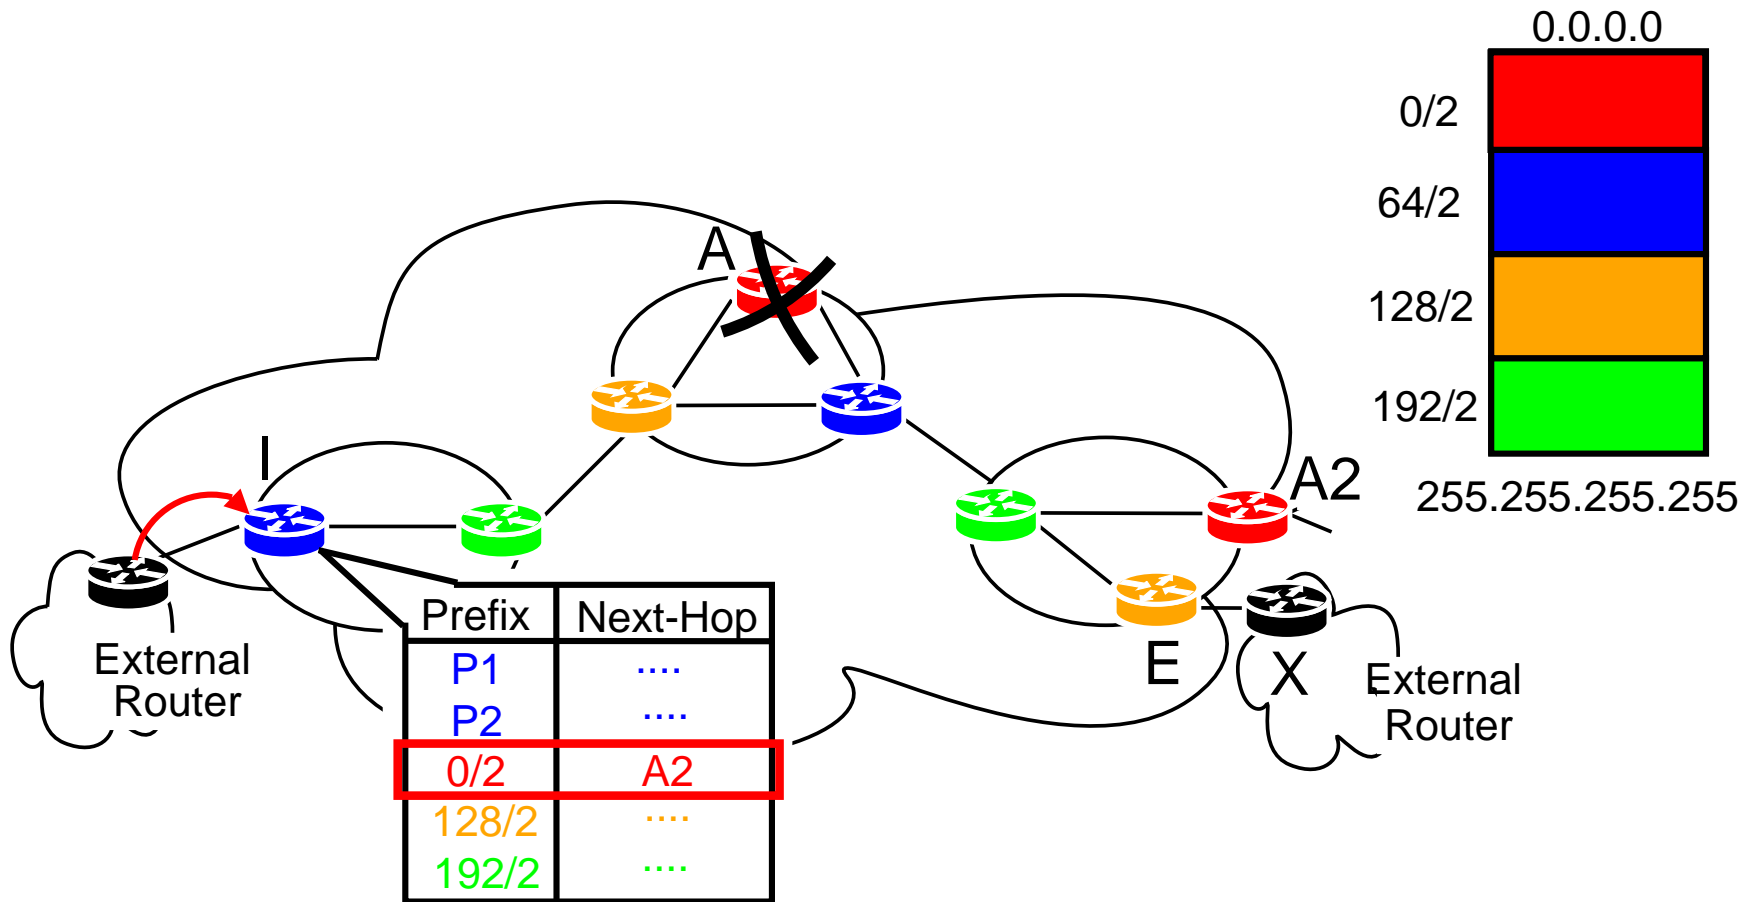
Egress Router strips the tunnel header off outgoing packets

Failure of Aggregation Point




What if Aggregation Pt. A fails?

Failure of Aggregation Point




Router I installs the route advertised by A2

Failure of Aggregation Point



Packets are re-routed appropriately

ViAggre's impact on ISP's traffic



ViAggre paths can be longer than native paths
Traffic stretch, increased router and link load, etc.

Popular Prefixes

Traffic volume follows power-law distribution

- ▶ 95% of the traffic goes to 5% of prefixes
- ▶ Has held up for years

Install “Popular Prefixes” in routers

- ▶ Stable over weeks
- ▶ Mitigates ViAggre’s impact on the ISP’s traffic

Talk Outline

- ▶ Introduction
- ▶ ViAggre: Basic Idea
- ▶ ViAggre Design
- ▶ **Evaluation**
- ▶ Deployment
- ▶ Conclusions

Stretch Vs FIB Size

Assigning more routers to aggregate a virtual prefix

- ▶ Reduces Stretch imposed on Traffic
- ▶ Increases FIB size

Aggregation Point Assignment Problem

- ▶ Minimize Worst FIB size, subject to constraint on Worst stretch
- ▶ NP-complete problem
- ▶ Implemented a greedy approximation

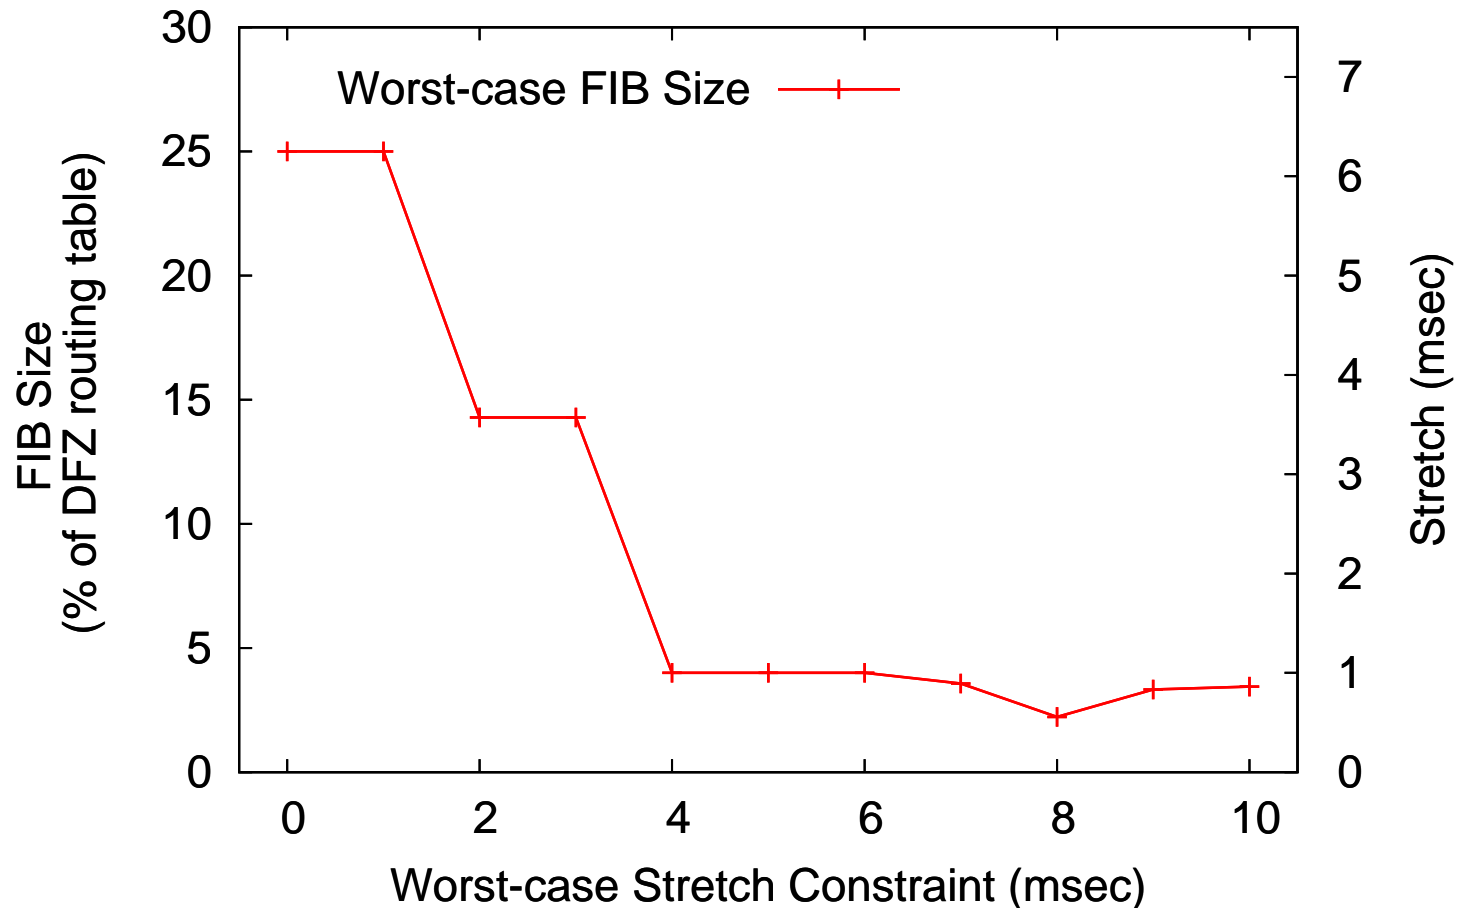
Performance Study

Data from tier-1 ISP

- ▶ Topology, Routing tables, Traffic matrix

Used out algorithm with varying stretch constraints

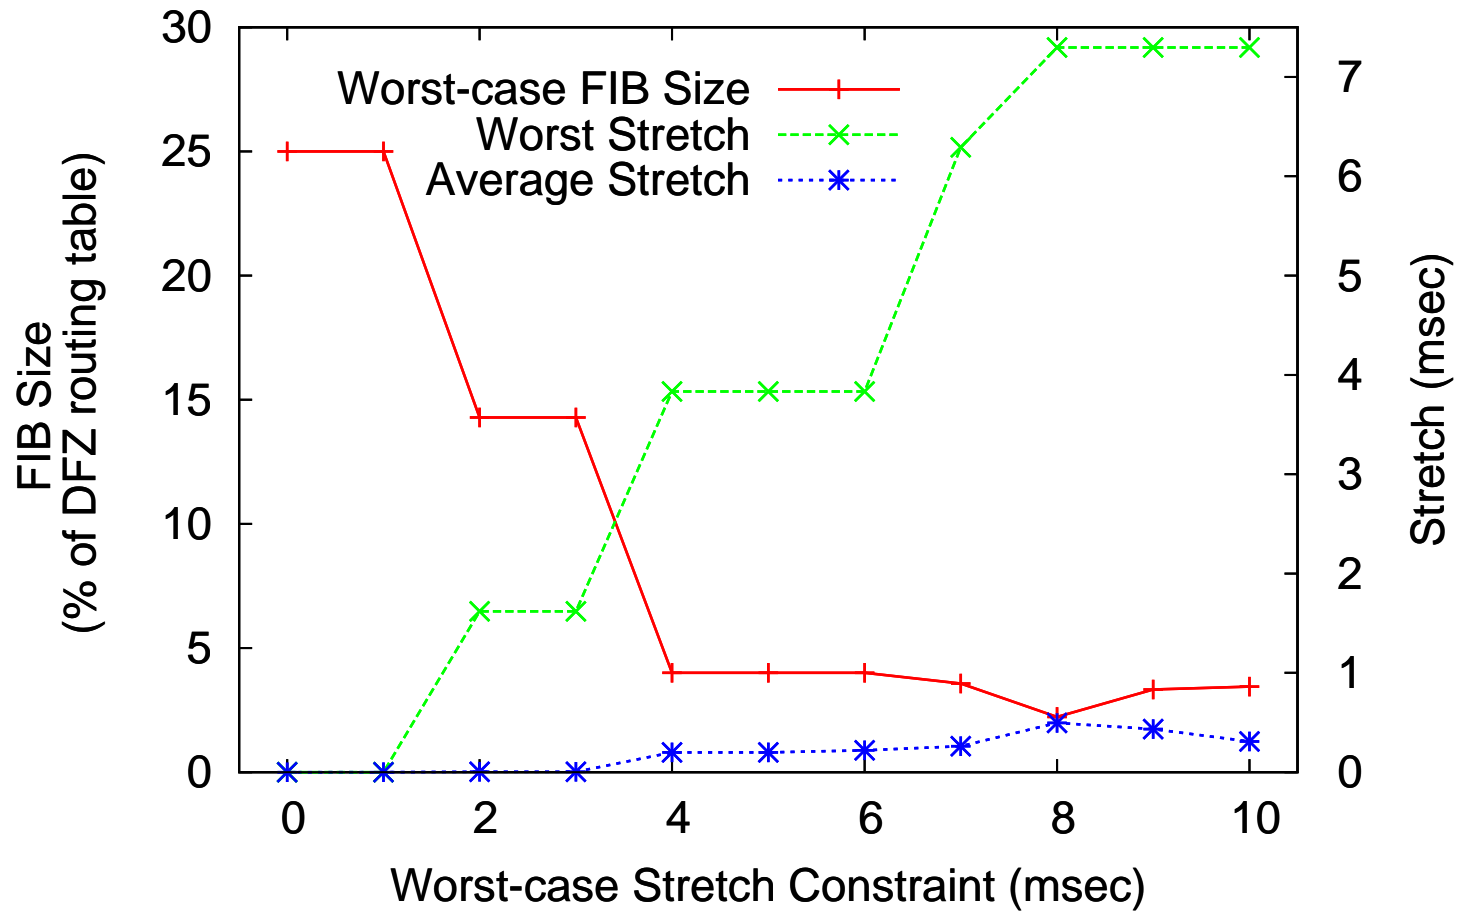
Constraining Worst Stretch



FIB Size reduces as Stretch constraint is relaxed

Worst-case Stretch $\leq 4\text{ms} \Rightarrow$ Worst FIB = 10,226 prefixes
(4% of global routing table)

Constraining Worst Stretch



Average Stretch is negligible

Worst-case Stretch $\leq 4\text{ms} \Rightarrow$ Average Stretch = 0.2msec

Router Load

Deployment with Worst-case Stretch $\leq 4\text{msec}$

- ▶ Shrinks FIB by more than 20x
- ▶ Median router load increases by 31.3%

Using popular prefixes

- ▶ 5% popular prefixes carry 96.7% of traffic
- ▶ Median and Worst-case router load increase $\approx 1\%$

Talk Outline

- ▶ Introduction
- ▶ ViAggre: Basic Idea
- ▶ ViAggre Design
- ▶ Evaluation
- ▶ **Deployment**
- ▶ **Conclusions**

ViAggre Pros

- ▶ Shrinks router FIB substantially
- ▶ Can be incrementally deployed
- ▶ Can be deployed on a limited-scale
- ▶ Incentive for deployment
- ▶ No change to ISP's routing setup
 - ▶ Does not affect convergence times
 - ▶ Does not affect routes advertised to neighbors
 - ▶ Does not restrict routing policies
 - ▶ ...

Can it be deployed?

Configuration overhead of a configuration-only solution

- ▶ Configuring FIB suppression on routers
- ▶ Configuring LSP advertisements on edge routers

Planning Overhead

- ▶ Choosing virtual prefixes
- ▶ Assigning aggregation points
- ▶ Assuring network robustness
- ▶ ...

ViAggre management overhead

Deployed ViAggre on WAIL

- ▶ Cisco 7300 routers
- ▶ Developed Configuration Tool
 - ▶ ~330 line python script
 - ▶ Extracts information from existing configuration files
 - ▶ Generates ViAggre configuration files
- ▶ Planning tool in the works

Working with a router vendor (Huawei)

- ▶ Implement ViAggre natively
- ▶ IETF Draft

Conclusion

ViAggre shrinks the FIB on routers

- ▶ Can extend the lifetime of installed routers

Is this a “complete” solution? **No**

- ▶ A simple and effective first step
- ▶ Next Step: Inter-domain ViAggre

Thank You!