How much larger is SAV table compared to FIB? A study with real BGP data

Presenter: Lancheng Qin

July 27, 2023

Background

- SAV table
  - The table or data structure that implements SAV rules and is used for source address validation in the data plane
  - Three validation modes of SAV table are defined in [draft-huang-savnet-sav-table]

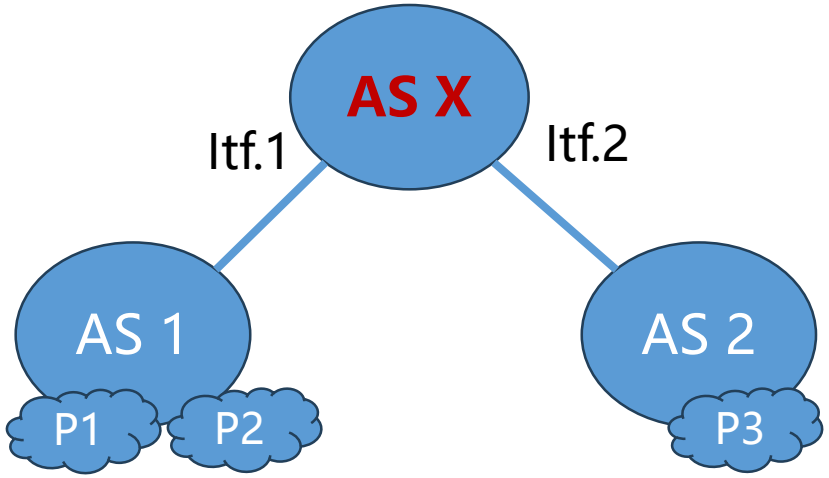

Mode 1

Interface-based prefix allowlist at Itf.1 of AS X
P1, P2

Interface-based prefix allowlist at Itf.2 of AS X
P3

Mode 2

Interface-based prefix blocklist at Itf.1 of AS X
P3

Interface-based prefix blocklist at Itf.2 of AS X
P1,P2

Mode 3

Prefix-based interface allowlist of AS X	
Src Prefix	Itf
P1	Itf.1
P2	Itf.1
P3	Itf.2

Background

- Many people are concerned about the scale of SAV tables with different validation modes
  - ◆We conduct simulations to compare the scale of FIB and SAV tables by using real RIB data provided by RouteViews and RIPE RIS

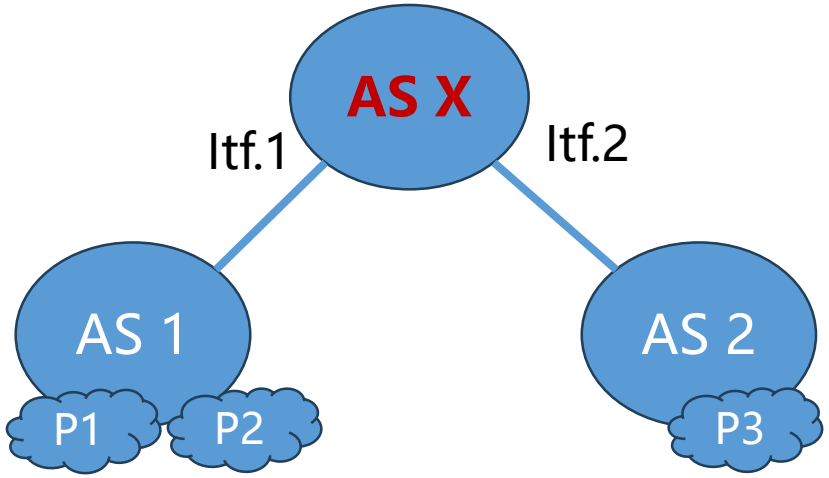

Mode 1

Interface-based prefix allowlist at Itf.1 of AS X
P1, P2

Interface-based prefix allowlist at Itf.2 of AS X
P3

Mode 2

Interface-based prefix blocklist at Itf.1 of AS X
P3

Interface-based prefix blocklist at Itf.2 of AS X
P1,P2

Mode 3

Prefix-based interface allowlist of AS X	
Src Prefix	Itf
P1	Itf.1
P2	Itf.1
P3	Itf.2

Assumptions

- Every AS is treated as a node to generate AS-level FIB, Interface-level SAV table with Mode 1 and Mode 2, and AS-level SAV table with Mode 3
  - FIB generation
    - Shortest path policy is used to select the best forwarding path among multiple routes
  - SAV table generation
    - SAV rules are generated based on the best forwarding paths. For example, after knowing the forwarding paths from AS Y to AS X, AS X identifies legitimate incoming directions for source prefixes of AS Y
- AS relationships are not considered
  - For each AS, we generate a SAV table with Mode 1 and Mode 2 at each AS-level interface, regardless of whether the connected AS is a customer, provider, or peer.
- To measure the scale, we calculate the number of prefixes for each table
- Only IPv4 is considered

Simulation Experiment

- Step 1:
  - We extract the full RIB of 319 ASes by using all public data provided by RouteViews[1] and RIPE RIS[2], and generate the FIB of each AS using shortest path policy
- Step 2:
  - For each of the 319 ASes, we generate its SAV rules for prefixes of the other 318 ASes, and organize SAV tables with different validation modes
- Step 3:
  - For each of the 319 ASes, we calculate the scale of FIB containing only prefixes of the other 318 ASes and the scale of SAV tables with different validation modes

[1] RouteViews. https://www.routeviews.org/routeviews/index.php/archive/
[2] RIPE RIS. https://www.ripe.net/analyse/internet-measurements/routing-information-service-ris/archive/ris-raw-data

Result (1)

The number of prefixes of FIB and SAV tables

	Mean	Max	Min
AS-level FIB	20603	26860	8598
Interface-level SAV table (Mode 1)	1073	32744	1
Interface-level SAV table (Mode 2)	31501	32743	0
AS-level SAV table (Mode 3)	32613 (1.6X of FIB)	32744	29172

Result (2)

Distribution of the number of prefixes of FIB and SAV tables

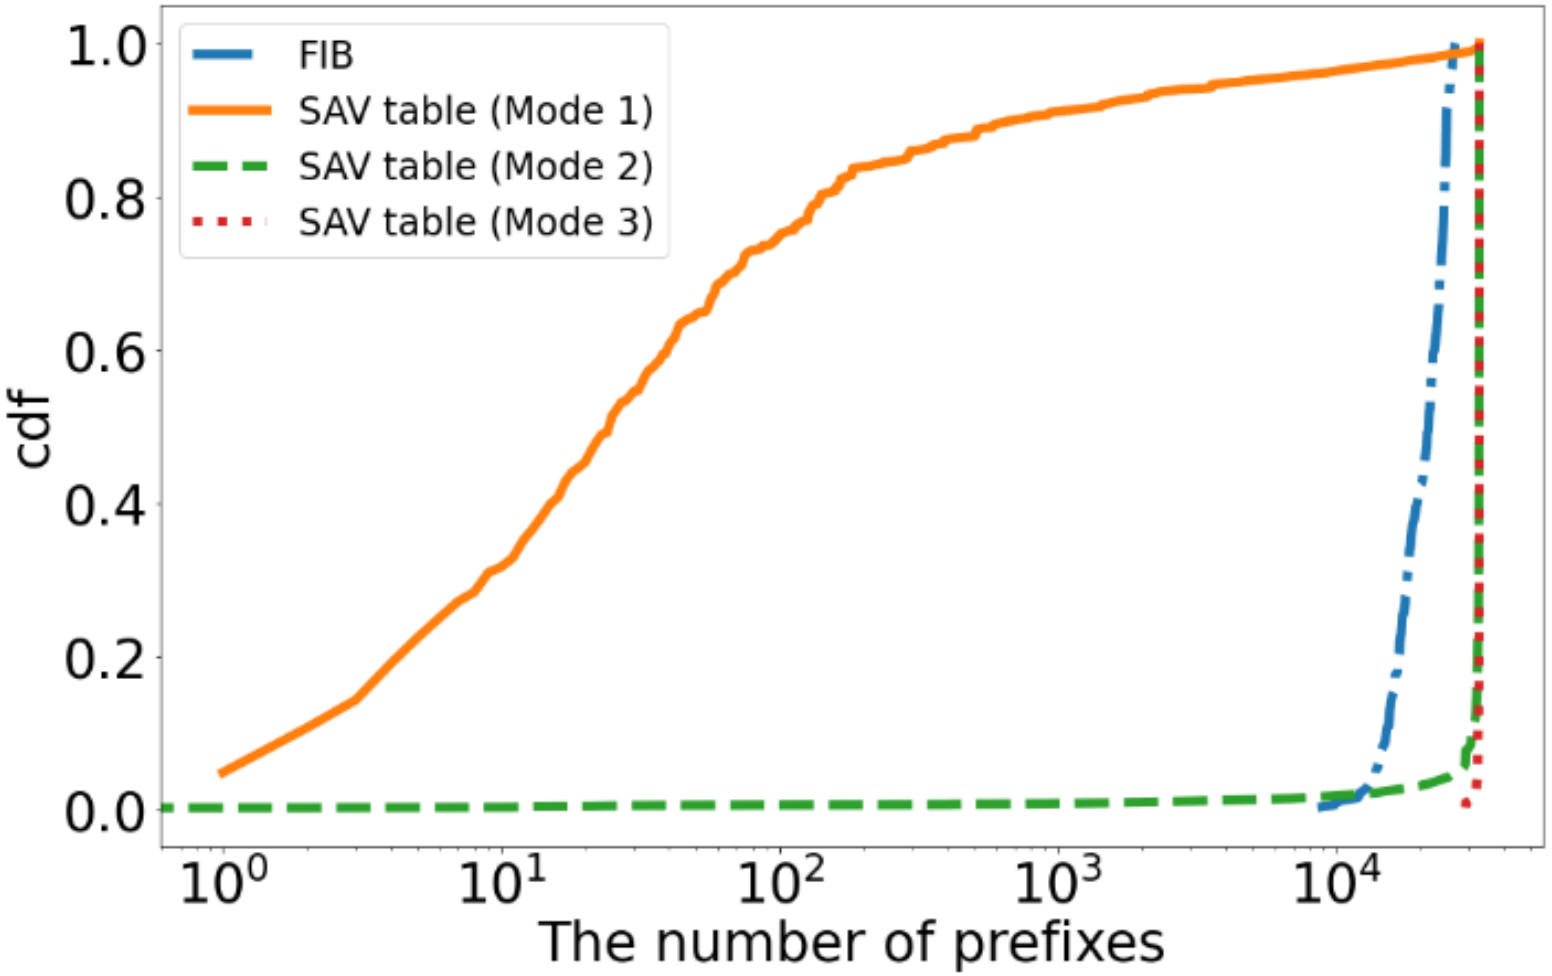

Result (3)

- Distribution of the ratio of the scale of SAV table (Mode 3) versus that of FIB
  - For 100% of ASes, the scale of SAV table (Mode 3) is larger than FIB
  - For 17.6% of ASes, the scale of SAV table (Mode 3) is 2X larger than FIB
  - For 0.9% of ASes, the scale of SAV table (Mode 3) is 3X larger than FIB

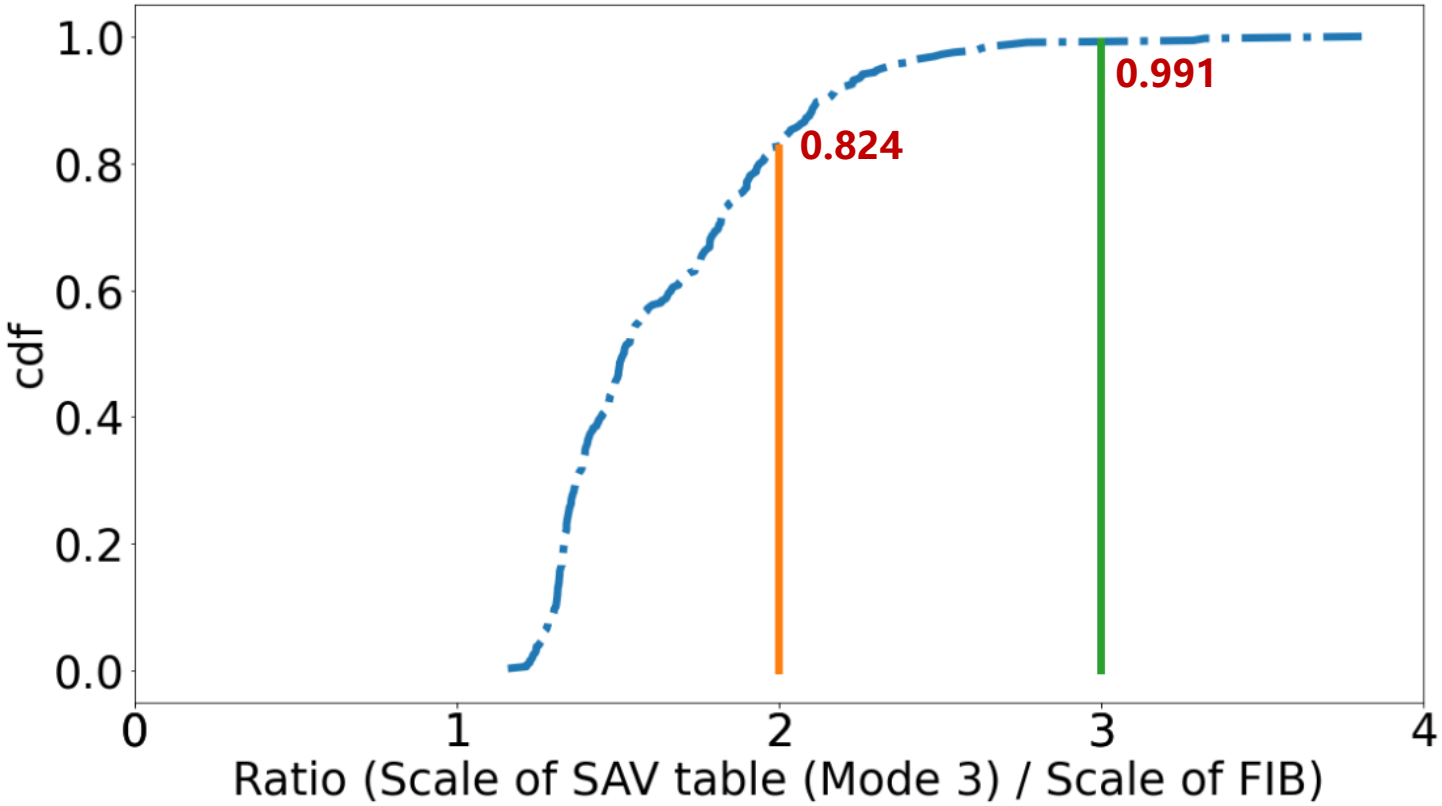

Result (3)

- Distribution of the ratio of the scale of SAV table (Mode 3) versus that of FIB
  - For 100% of ASes, the scale of SAV table (Mode 3) is larger than FIB
  - For 17.6% of ASes, the scale of SAV table (Mode 3) is 2X larger than FIB
  - For 0.9% of ASes, the scale of SAV table (Mode 3) is 3X larger than FIB

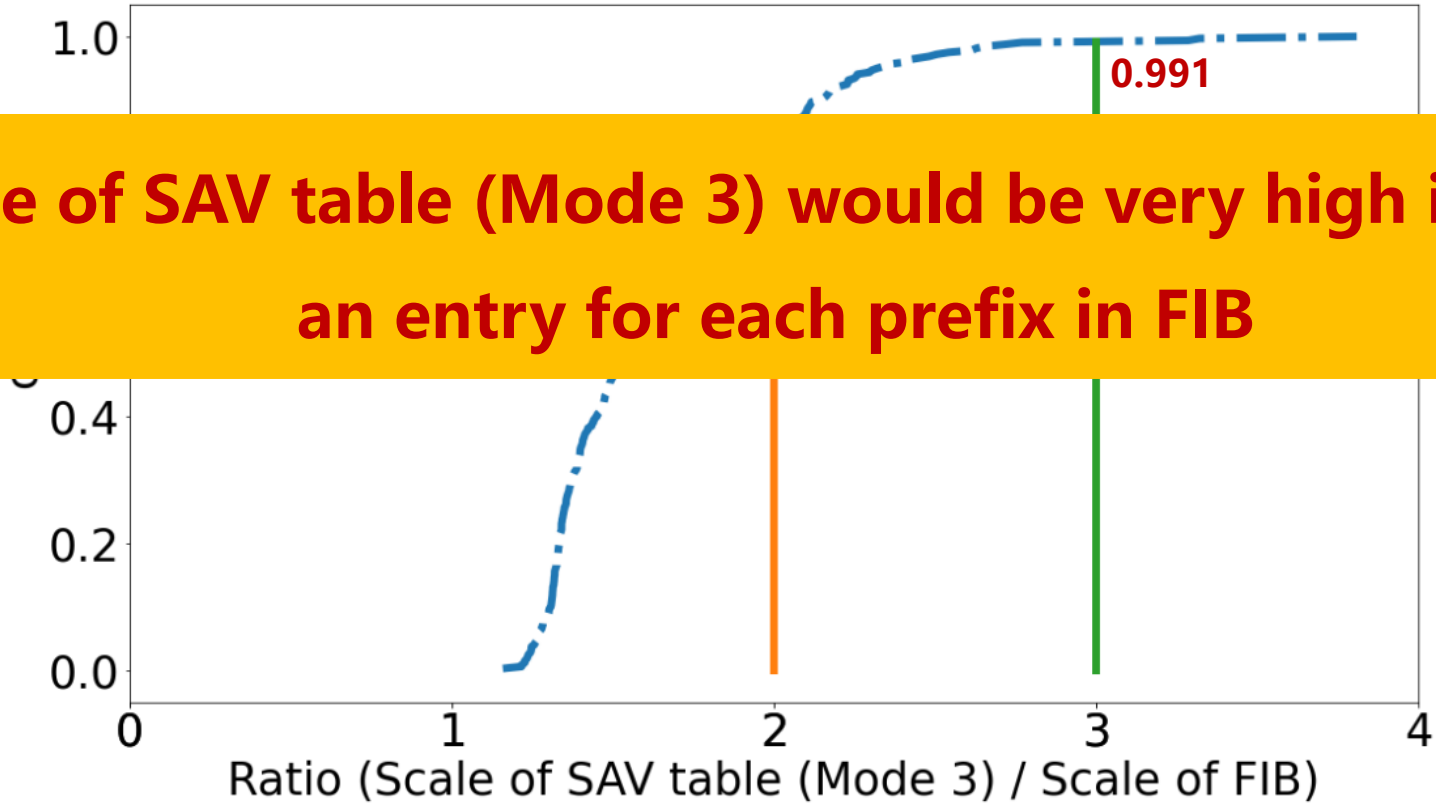

The scale of SAV table (Mode 3) would be very high if we record an entry for each prefix in FIB

How to Reduce the Scale of SAV Tables?

- Source prefix aggregation

Assume Prefix 2 is the sub prefix of Prefix 1

Interface-level SAV table (Mode 1)
Prefix 1, Prefix 2

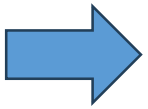

Interface-level SAV table (Mode 1)
Prefix 1

Interface-level SAV table (Mode 2)
Prefix 1, Prefix 2

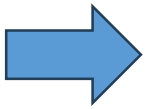

Interface-level SAV table (Mode 2)
Prefix 1

AS-level SAV table (Mode 3)	
Src Prefix	Itf
Prefix 1	Interface 1
Prefix 2	Interface 1

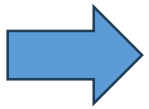

AS-level SAV table (Mode 3)	
Src Prefix	Itf
Prefix 1	Interface 1

Source Prefix Aggregation: Result (1)

Before source prefix aggregation

	Mean	Max	Min
AS-level FIB	20603	26860	8598
Interface-level SAV table (Mode 1)	1073	32744	1
Interface-level SAV table (Mode 2)	31501	32743	0
AS-level SAV table (Mode 3)	32613 (1.6X of FIB)	32744	29172

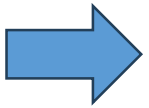

After source prefix aggregation

	Mean	Max	Min
AS-level FIB	20603	26860	8598
Interface-level SAV table (Mode 1)	520	16103	1
Interface-level SAV table (Mode 2)	15464	16102	0
AS-level SAV table (Mode 3)	15945 (0.77X of FIB)	16103	14437

On average, the scale of SAV table is reduced by more than 50%

Source Prefix Aggregation: Result (2)

Before source prefix aggregation

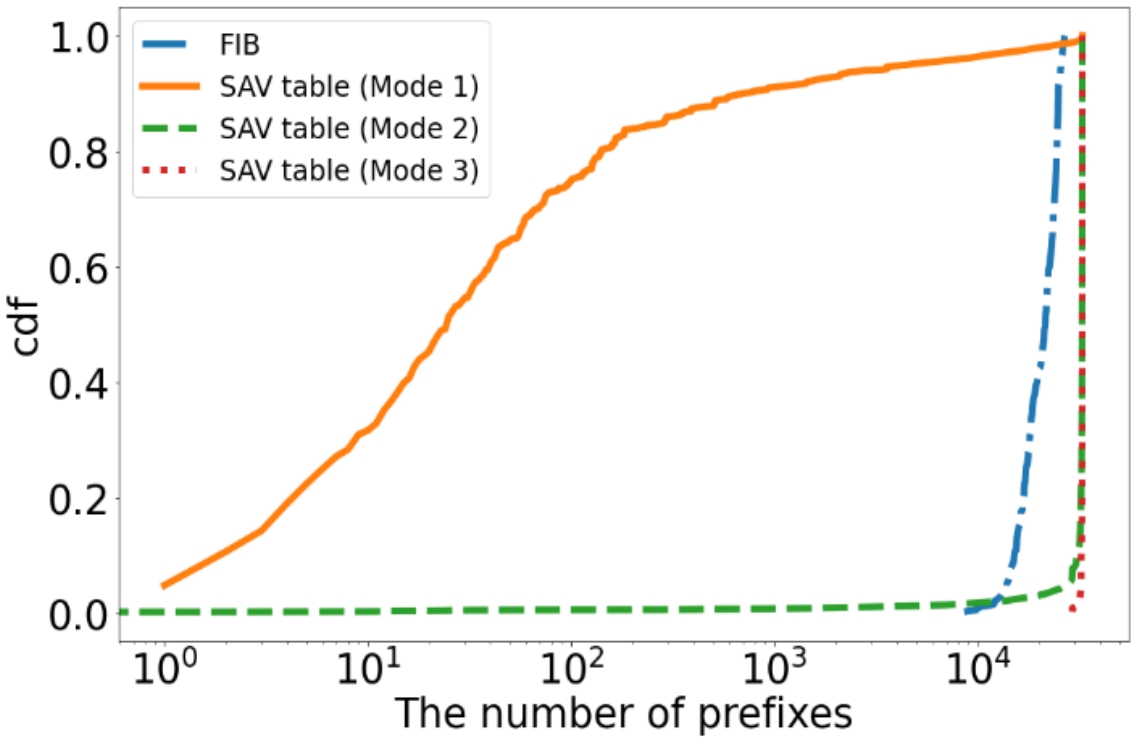

After source prefix aggregation

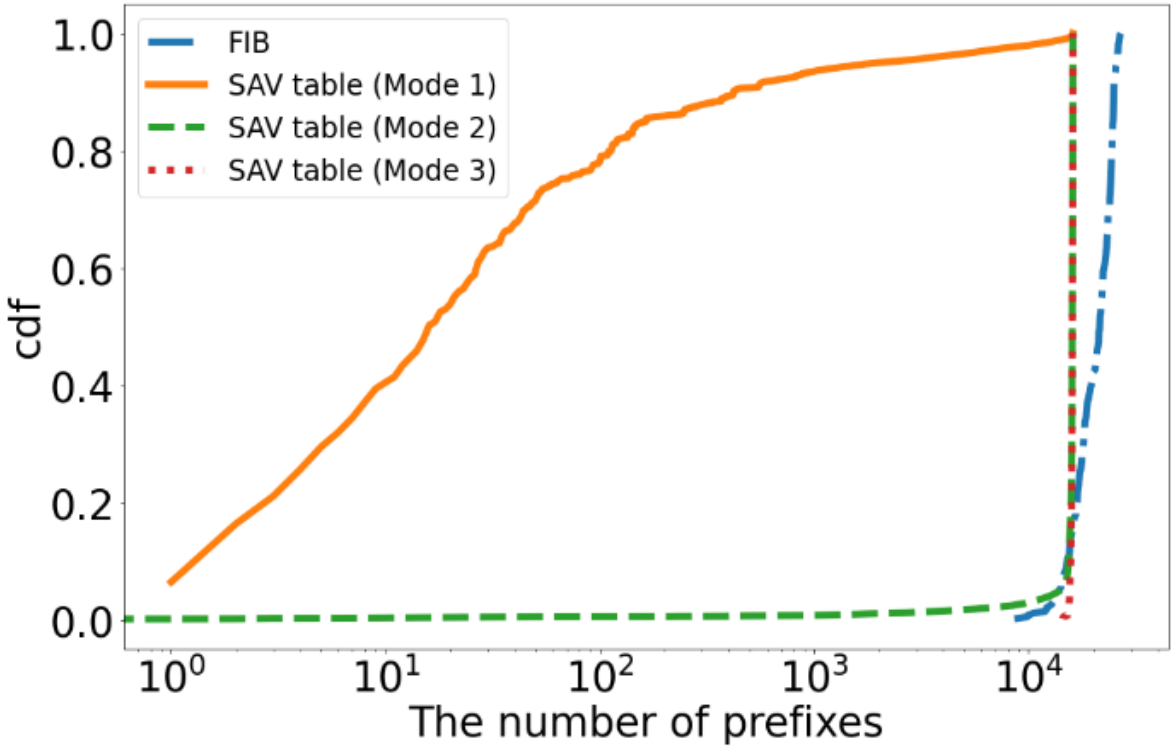

On average, the scale of SAV table is reduced by more than 50%

Source Prefix Aggregation: Result (3)

Before source prefix aggregation

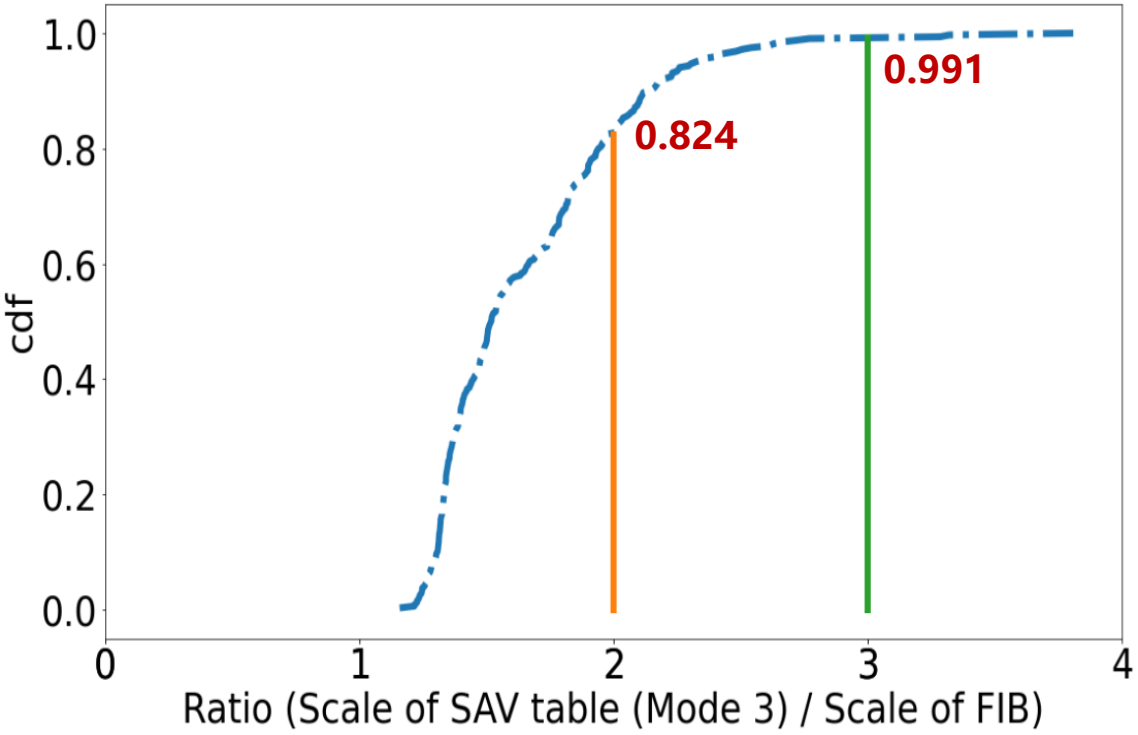

After source prefix aggregation

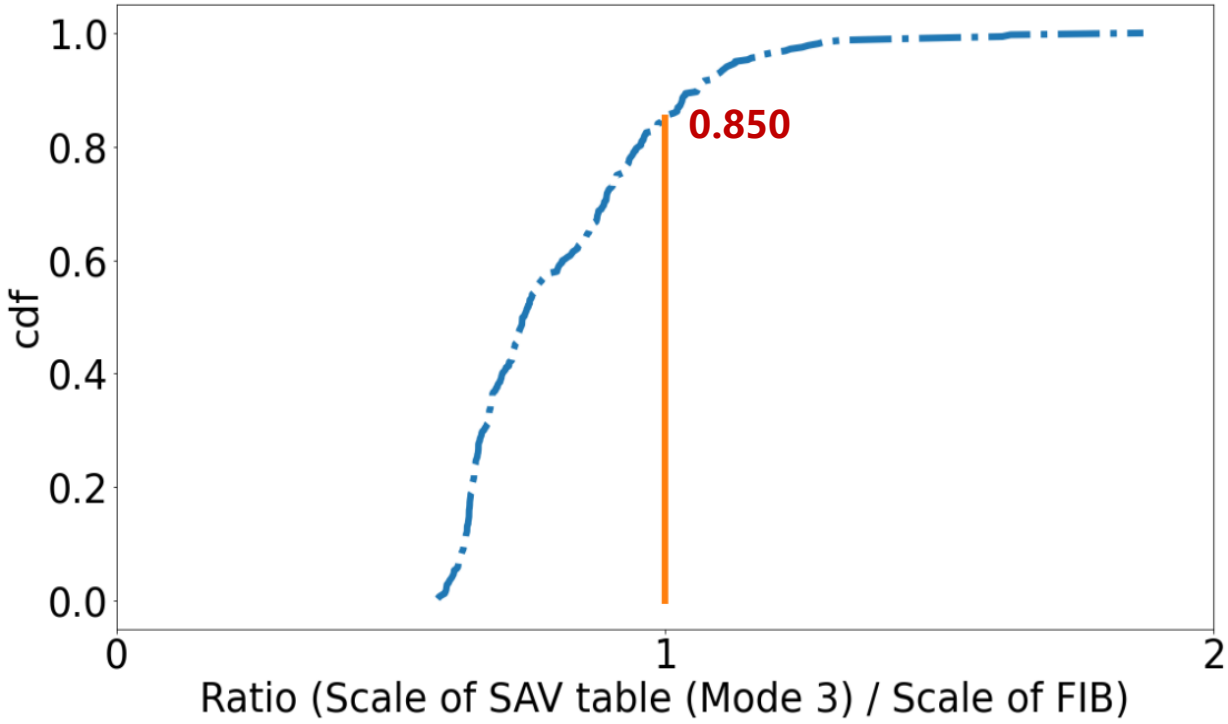

On average, the scale of SAV table is reduced by more than 50%

Summary

- The scale of Interface-level SAV table (Mode 1 or Mode 2) is **smaller** than FIB
  - For an AS, the whole scale of interface-level SAV tables at all AS-level interfaces may be larger than the AS-level FIB, but different interface-level SAV tables are usually stored in different border routers
- For **85%** of ASes, the scale of AS-level SAV table (Mode 3) is **smaller** than FIB
  - On average, the scale of AS-level SAV table (Mode 3) is **77% of that of FIB**
- The scale of SAV table can be further reduced by using advanced compression techniques
  - Such as CLP[3], Rétvári, et al.[4]

[3] Rodrigues, Kirk, Yu Luo, and Ding Yuan. CLP: Efficient and Scalable Search on Compressed Text Logs. OSDI 2021.
[4] Rétvári, Gábor, et al. Compressing IP forwarding tables: Towards entropy bounds and beyond. SIGCOMM 2013.

Thanks!