Enhancing Global Network Monitoring with *Magnifier*

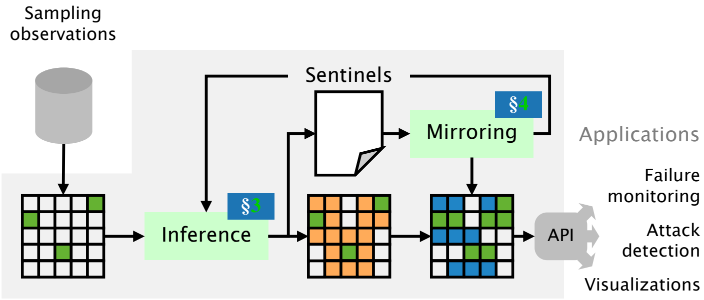

Tobias Bühler, Romain Jacob, Ingmar Poese (*), Laurent Vanbever

ETH Zürich and BENOCS (*)

NSDI 2023

April 19 2023

Where does my traffic enter and leave the network?

Where does my traffic enter and leave the network?

Easy?
Using control-plane data

Where does my traffic enter and leave the network?

Hard!
Requires data-plane data

Today, operators infer ingress and egress points using sampling or mirroring

Sampled flow statistics

For example with sFlow or NetFlow

Selected packet mirroring

Packet clones matching defined rules

Today, operators infer ingress and egress points using sampling or mirroring

	Sampled flow statistics	Selected packet mirroring
	For example with sFlow or NetFlow	Packet clones matching defined rules
Traffic coverage		
Inference accuracy		
Reporting speed		

Today, operators infer ingress and egress points using sampling or mirroring

	Sampled flow statistics	Selected packet mirroring
	For example with sFlow or NetFlow	Packet clones matching defined rules
Traffic coverage	*Any* flow can be sampled, but skewed towards long flows	Only packets matching defined rules are mirrored
Inference accuracy		
Reporting speed		

Today, operators infer ingress and egress points using sampling or mirroring

	Sampled flow statistics	Selected packet mirroring
	For example with sFlow or NetFlow	Packet clones matching defined rules
Traffic coverage	*Any* flow can be sampled, but skewed towards long flows	Only packets matching defined rules are mirrored
Inference accuracy	No guarantee that same flow is sampled over time or devices	*All* matching packets are mirrored, at the cost of twice the traffic amount
Reporting speed		

Today, operators infer ingress and egress points using sampling or mirroring

	Sampled flow statistics	Selected packet mirroring
	For example with sFlow or NetFlow	Packet clones matching defined rules
Traffic coverage	*Any* flow can be sampled, but skewed towards long flows	Only packets matching defined rules are mirrored
Inference accuracy	No guarantee that same flow is sampled over time or devices	*All* matching packets are mirrored, at the cost of twice the traffic amount
Reporting speed	Often delayed due to caches or aggregation steps	Nearly instantaneous feedback once a mirrored packet is generated

Today, operators infer ingress and egress points using sampling or mirroring, both of which are problematic

	Sampled flow statistics	Selected packet mirroring
	For example with sFlow or NetFlow	Packet clones matching defined rules
Traffic coverage	*Any* flow can be sampled, but skewed towards long flows	Only packets matching defined rules are mirrored
Inference accuracy	No guarantee that same flow is sampled over time or devices	*All* matching packets are mirrored, at the cost of twice the traffic amount
Reporting speed	Often delayed due to caches or aggregation steps	Nearly instantaneous feedback once a mirrored packet is generated

Magnifier combines the benefits of sampling and mirroring

without their drawbacks

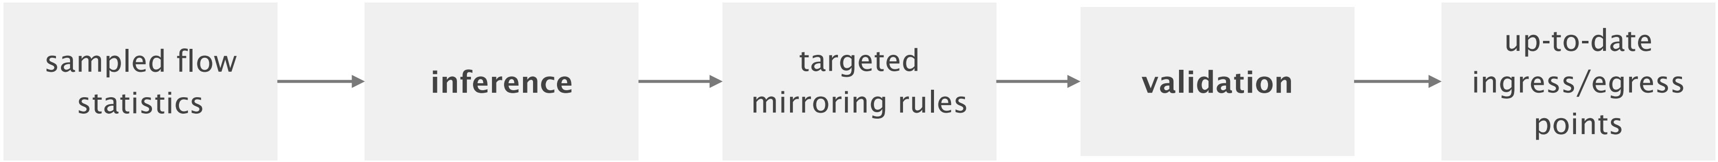
sampled flow statistics
inference
targeted mirroring rules
validation
up-to-date ingress/egress points

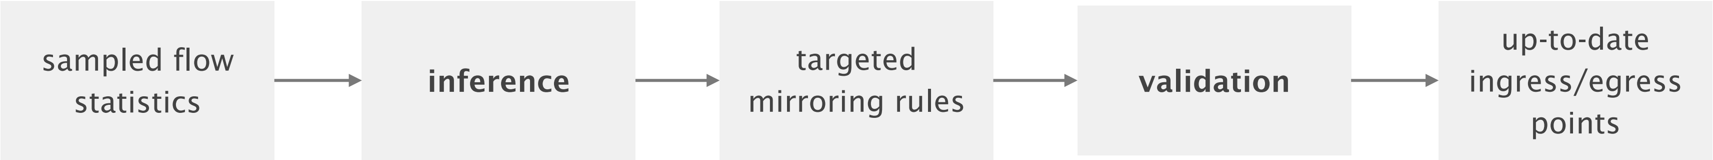

Infer largest subnets matching sampled IPs to ingress or egress

Mirror where we do *not* expect traffic to enter or leave

Magnifier uses sampled data to infer ingress and egress points

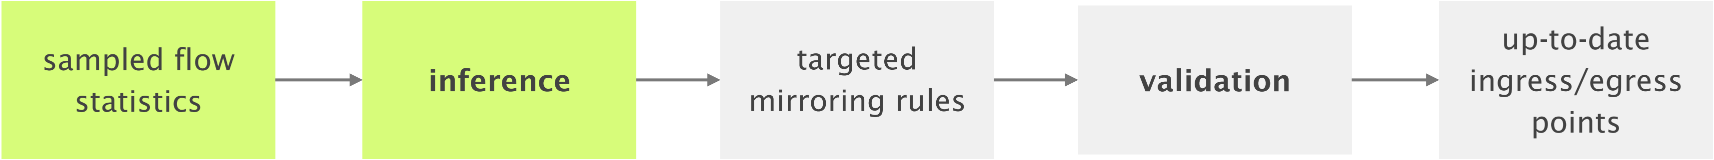

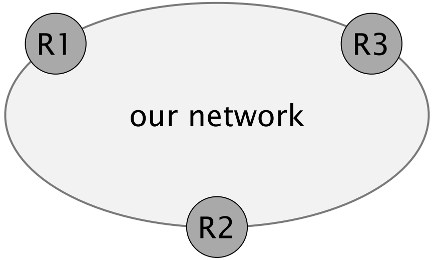

Magnifier uses sampled data to infer ingress and egress points

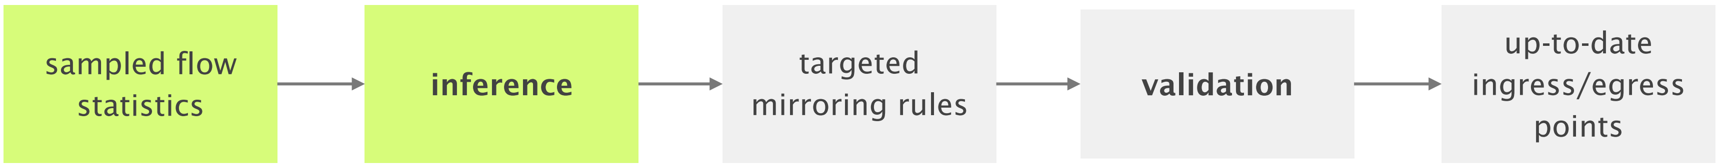

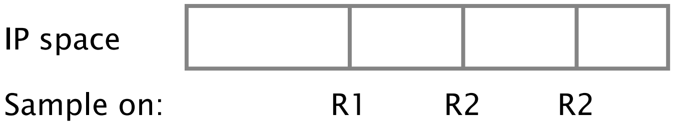

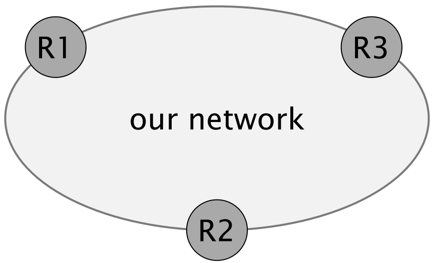

Magnifier uses sampled data to infer ingress and egress points

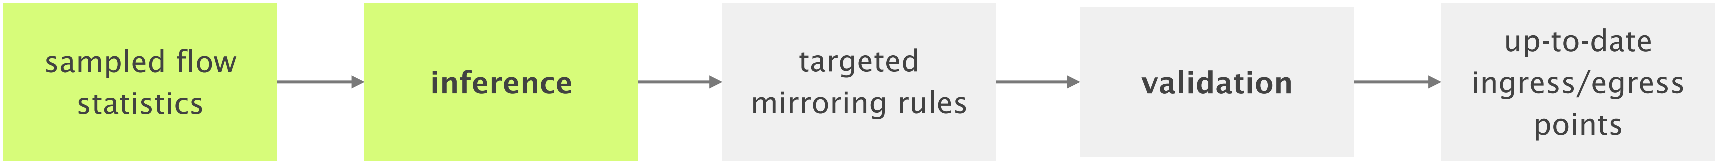

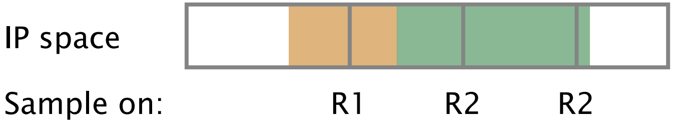

Goal: maximize the covered IP space

Subnet (orange) enters (leaves) on R1

Subnet (green) enters (leaves) on R2

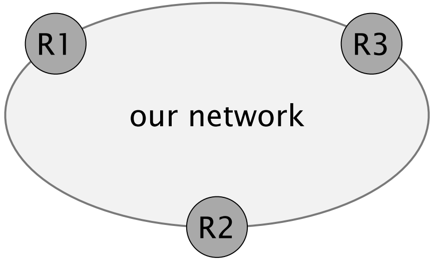

A *lack of* mirrored traffic validates Magnifier's inferences

sampled flow statistics → **inference** → targeted mirroring rules → **validation** → up-to-date ingress/egress points

Subnet ● on R1

Subnet ● on R2

Goal: high coverage

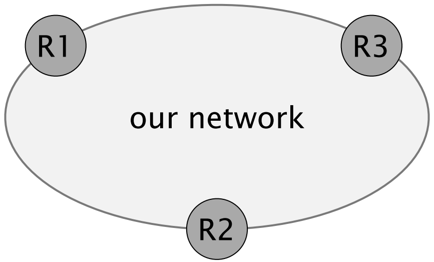

A *lack of* mirrored traffic validates Magnifier’s inferences

sampled flow statistics → **inference** → targeted mirroring rules → **validation** → up-to-date ingress/egress points

Subnet ● on R1

Subnet ● on R2

Goal: high coverage

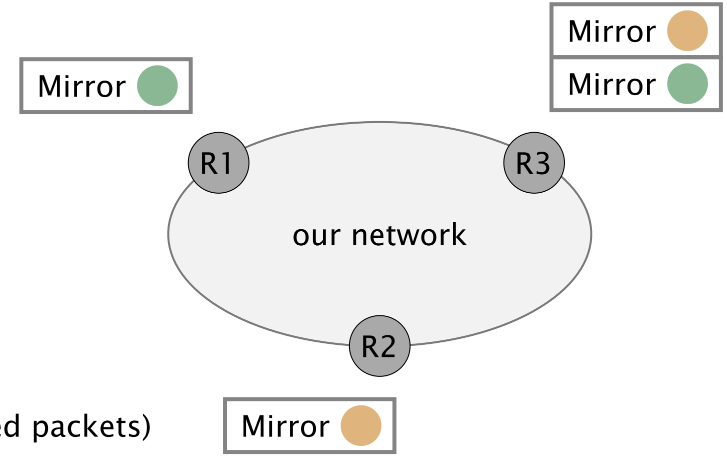

Goals: high accuracy

low overhead (few mirrored packets)

Magnifier continuously updates validated ingress and egress points

sampled flow statistics → **inference** → targeted mirroring rules → **validation** → up-to-date ingress/egress points

Subnet (orange) on R1

Subnet (green) on R2

Goal: high coverage

Mirror (orange) on R2 & R3

Mirror (green) on R1 & R3

Goals: high accuracy

low overhead

Magnifier continuously updates validated ingress and egress points

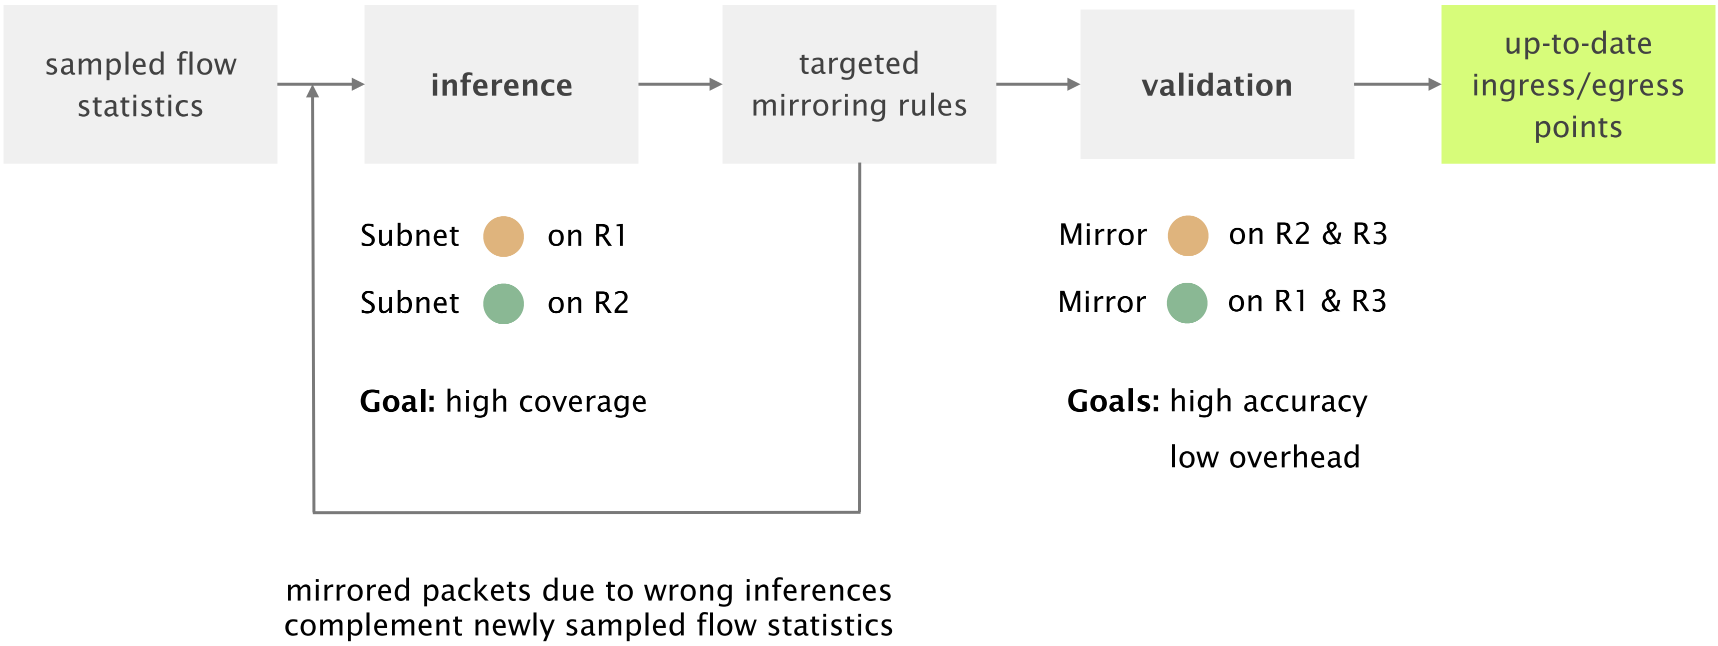

Sounds great, what is the catch?
Today's router resources are a limiting factor

Memory limitations	Router memory is limited Mirroring rules share resources with other router features
Deployment time	Adding/removing large number of rules takes time Especially if they are actively used
Controller placement	Rule deactivation can be slow in large networks Magnifier could generate a lot of mirrored traffic

Magnifier mitigates the mirroring overhead by...

Memory limitations	Prioritization of rules according to custom metrics
Deployment time	Activation of pre-deployed rules in batches
Controller placement	Deployment of sub-controllers close to the border

We performed various simulations and lab experiments

Simulations	Using simulated sampling and mirroring operations We assume unlimited resources
Lab experiments	Using Cisco switches in our lab at ETH We allow at most 500 mirroring rules on one device
CAIDA traces	Using CAIDA packet traces as input We get full insight (ground-truth information)

We need to assign CAIDA packets (IPs) to ingress points

We need to assign CAIDA packets (IPs) to ingress points

the evaluation focuses
on ingress observations

We need to assign CAIDA packets (IPs) to ingress points

Random	Ingress of one IP can change randomly over time *No* continuity over time *or* IP space
Static	Every IP is statically assigned to one ingress point *No* continuity over the IP space
Permuted	IP subnets are persistently permuted to different ingresses Continuity over time *and* IP space

We need to assign CAIDA packets (IPs) to ingress points

Random — Ingress of one IP can change randomly over time
No continuity over time *or* IP space

Static — Every IP is statically assigned to one ingress point
No continuity over the IP space

Permuted — IP subnets are persistently permuted to different ingresses
Continuity over time *and* IP space

Easier for Magnifier

A *realistic* assignment is a combination of these extreme cases

Magnifier validates the ingress of a large amount of packets, while generating few mirrored packets

simulation results: 32 ingresses; sampling rate 1/1024; CAIDA trace replayed at full speed

Magnifier validates the ingress of a large amount of packets, while generating few mirrored packets

simulation results: 32 ingresses; sampling rate 1/1024; CAIDA trace replayed at full speed

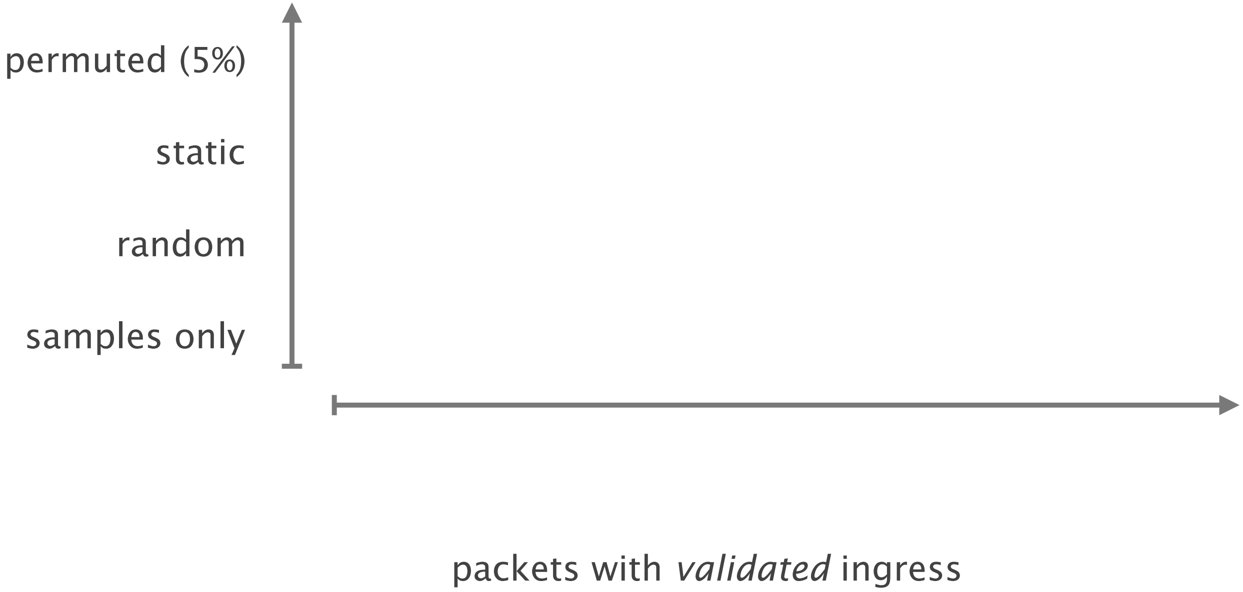

Magnifier validates the ingress of a large amount of packets, while generating few mirrored packets

simulation results: 32 ingresses; sampling rate 1/1024; CAIDA trace replayed at full speed

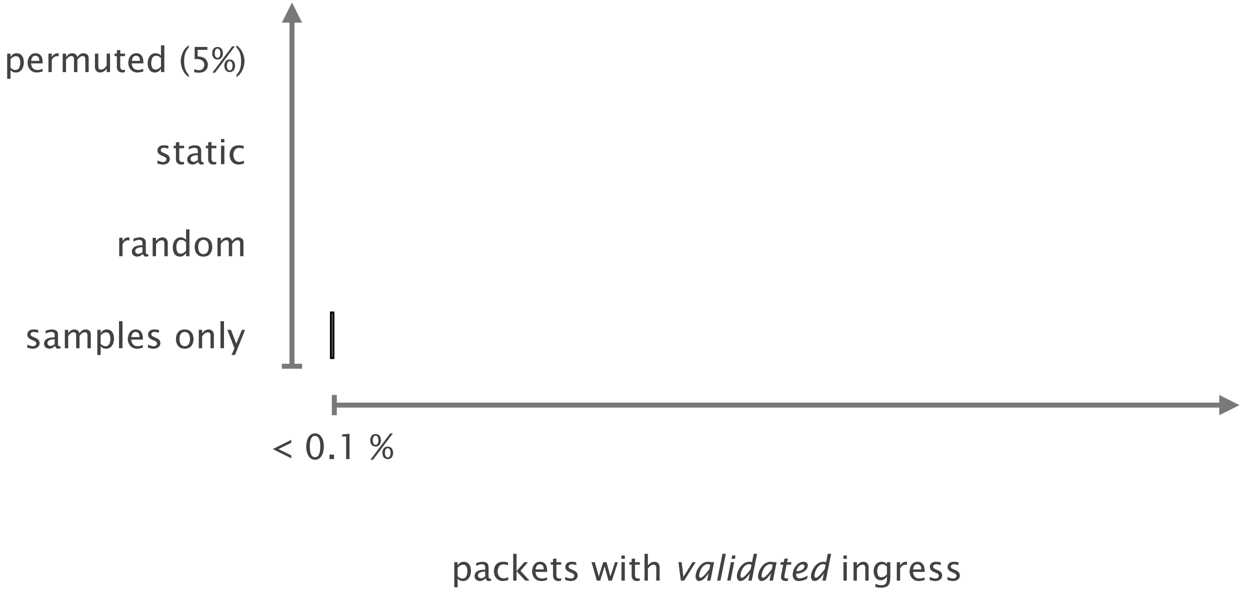

Magnifier validates the ingress of a large amount of packets, while generating few mirrored packets

simulation results: 32 ingresses; sampling rate 1/1024; CAIDA trace replayed at full speed

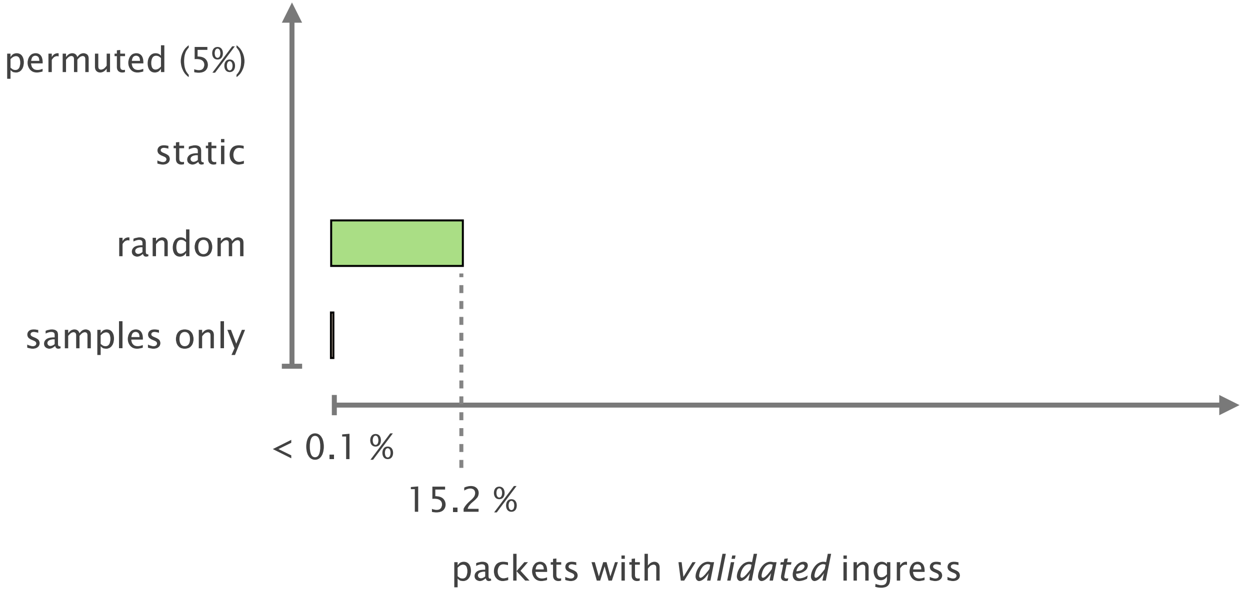

Magnifier validates the ingress of a large amount of packets, while generating few mirrored packets

simulation results: 32 ingresses; sampling rate 1/1024; CAIDA trace replayed at full speed

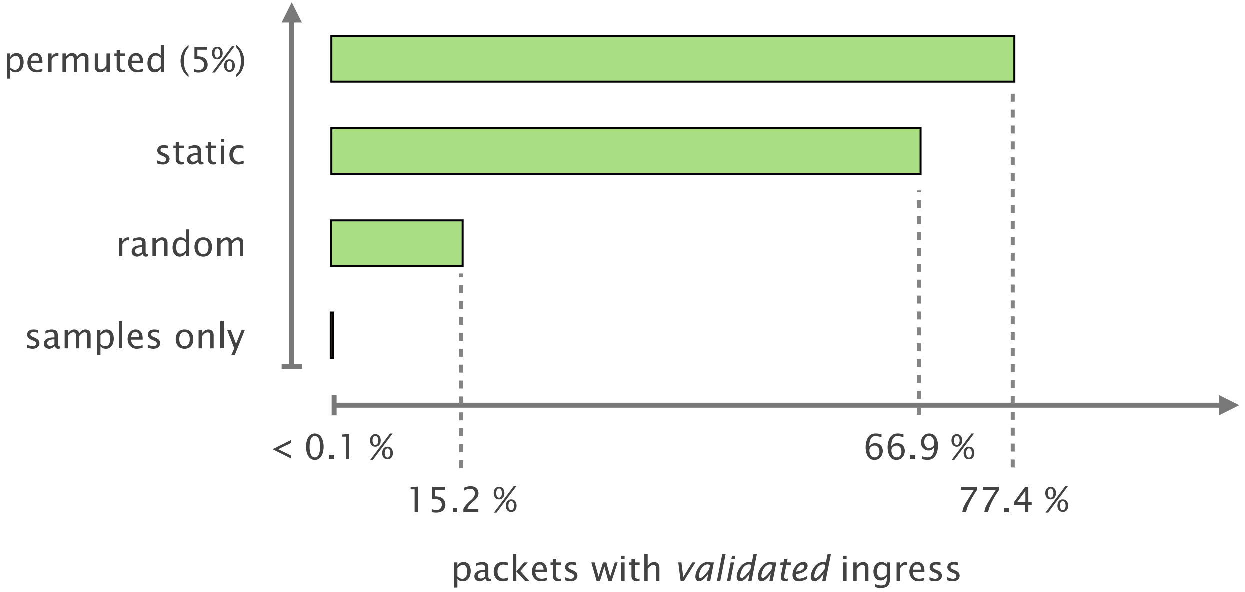

Magnifier validates the ingress of a large amount of packets, while generating few mirrored packets

simulation results: 32 ingresses; sampling rate 1/1024; CAIDA trace replayed at full speed

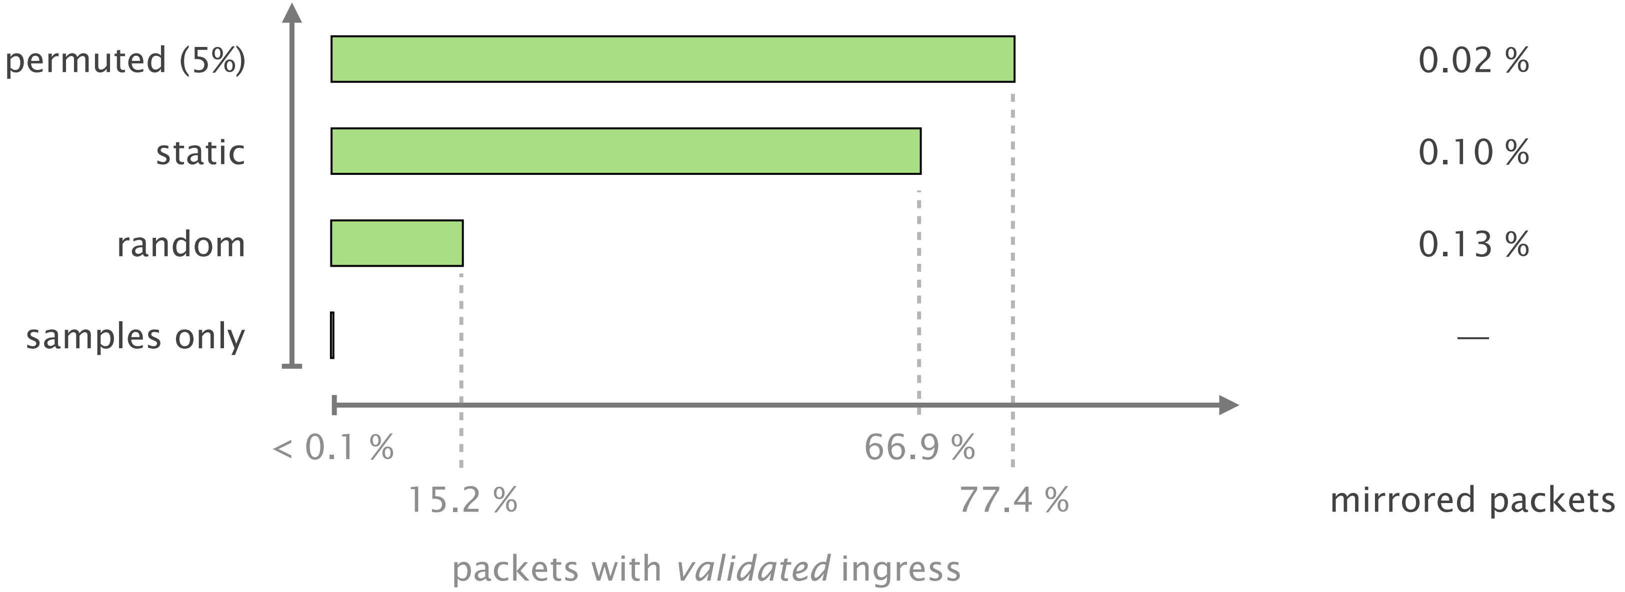

Lab experiments with 500 mirroring rules *only*, confirm simulation observations (*random* IP to ingress assignment)

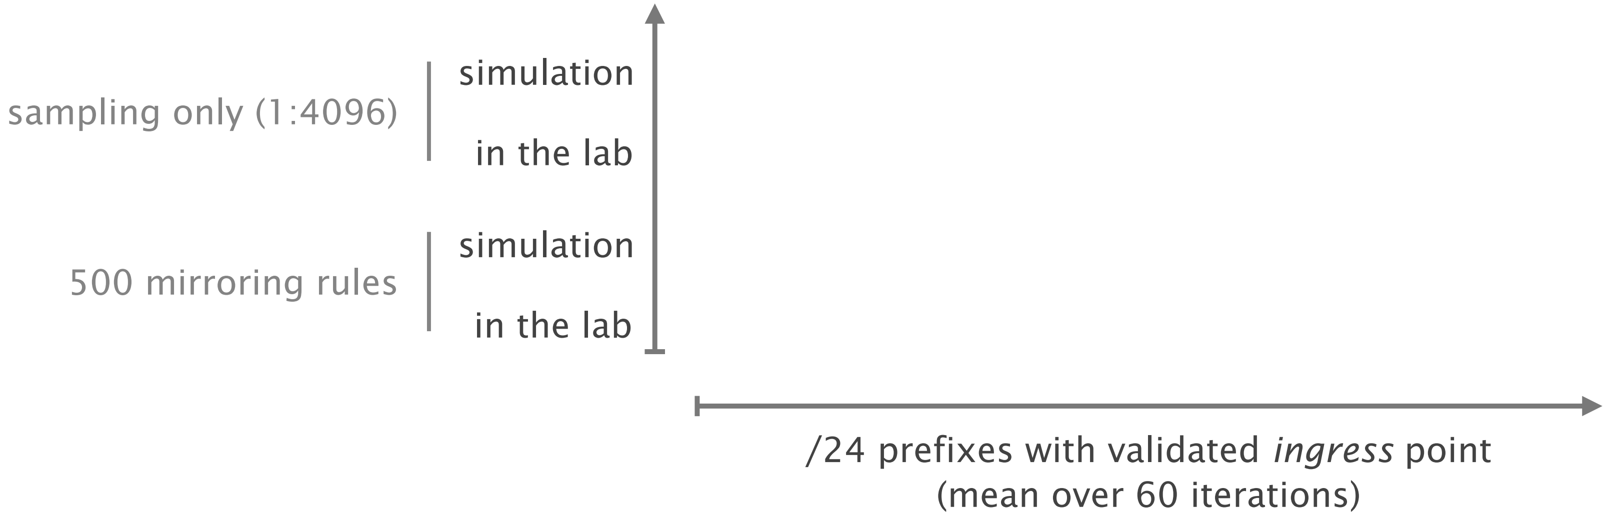

Lab experiments with 500 mirroring rules *only*, confirm simulation observations (*random* IP to ingress assignment)

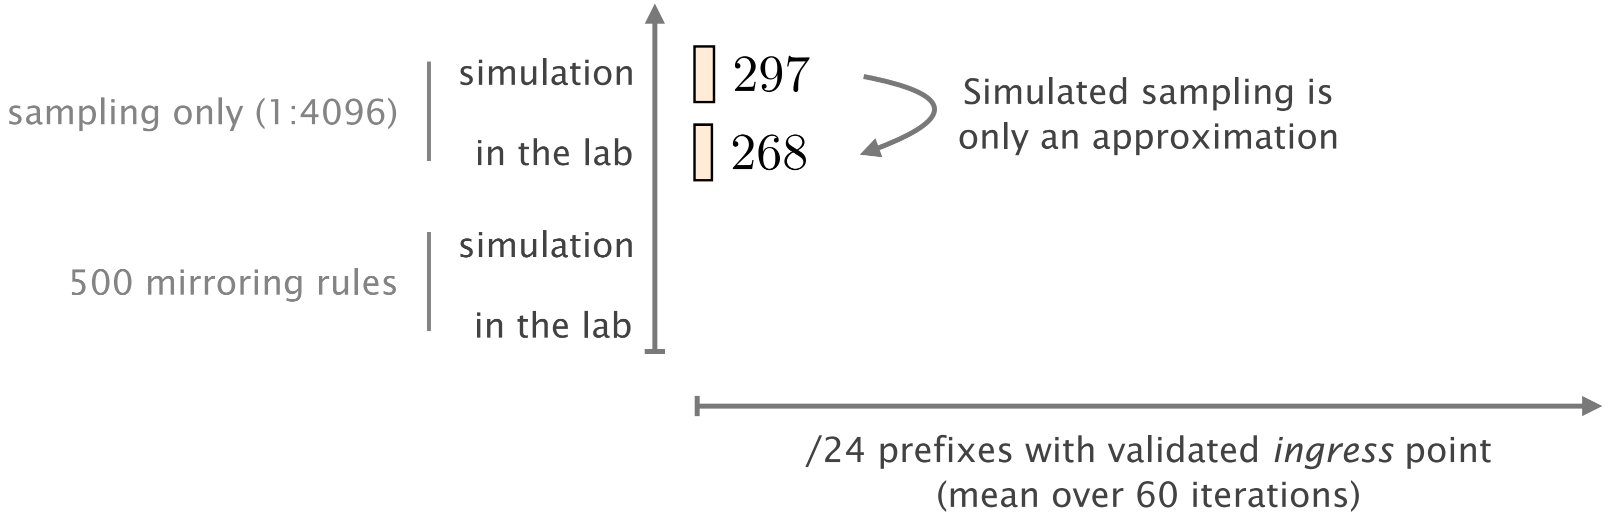

Lab experiments with 500 mirroring rules *only*,
confirm simulation observations (*random* IP to ingress assignment)

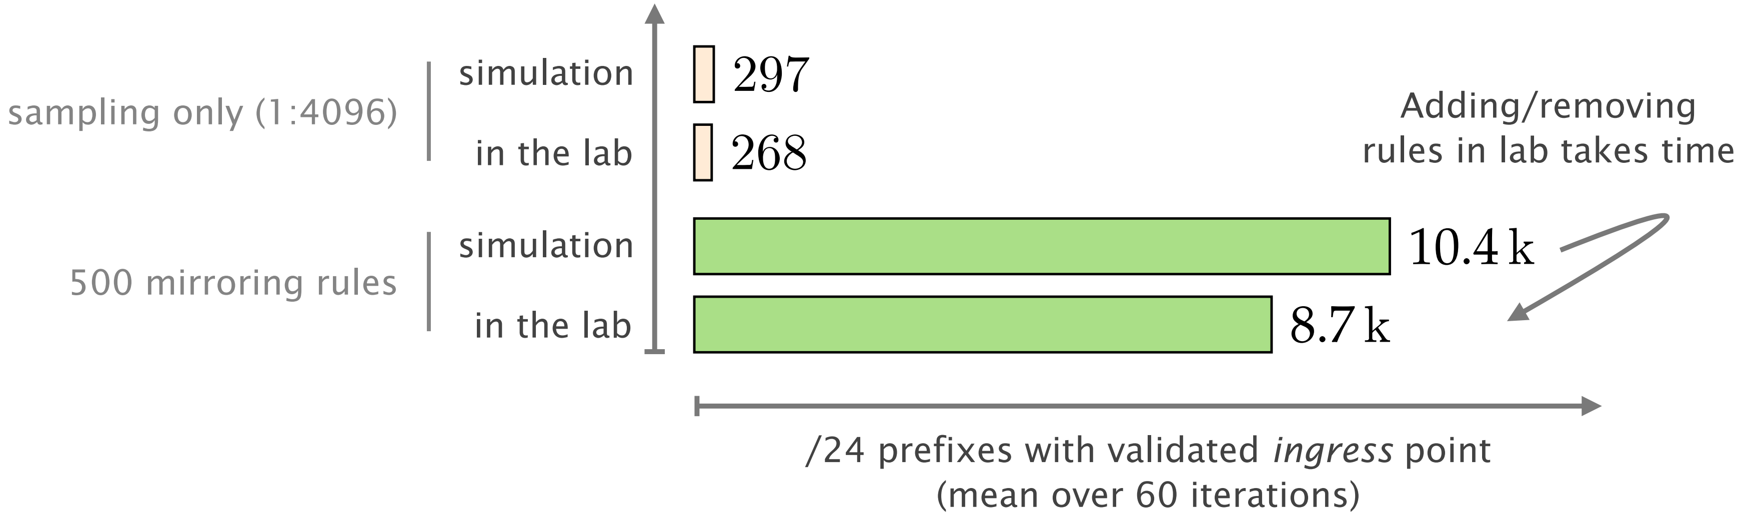

Lab experiments with 500 mirroring rules *only*, confirm simulation observations (*random* IP to ingress assignment)

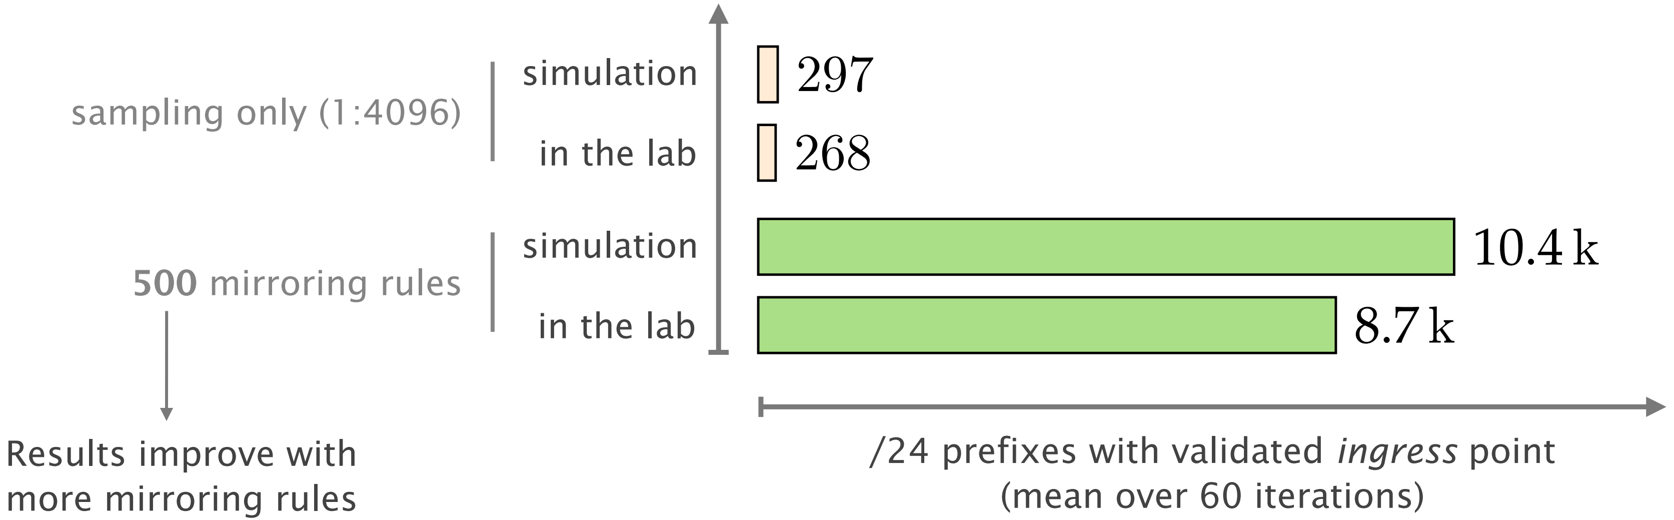

All of that while mirroring fewer than 1.4 % of all packets (in the lab)

Compared to related work, Magnifier generates few mirrored packets and does not require end-host support

[Everflow]	Mirrors packets of *every* flow (based on TCP flags) Up to **5 %** of all packets in our simulations
[Flowyager]	Uses Flowtrees to store flow information efficiently Limited by the available information in *sampled* data
[Pingmesh]	Performs active pings between data center end hosts Infeasible in an ISP setting

Infer largest subnets matching sampled IPs to ingress or egress

High coverage

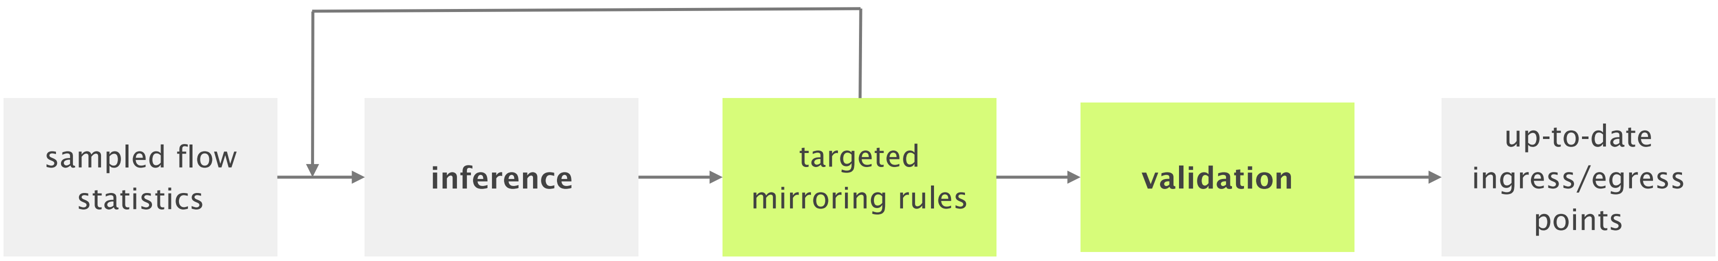

Mirror where we do *not*
expect traffic to enter or leave

High accuracy

Low overhead

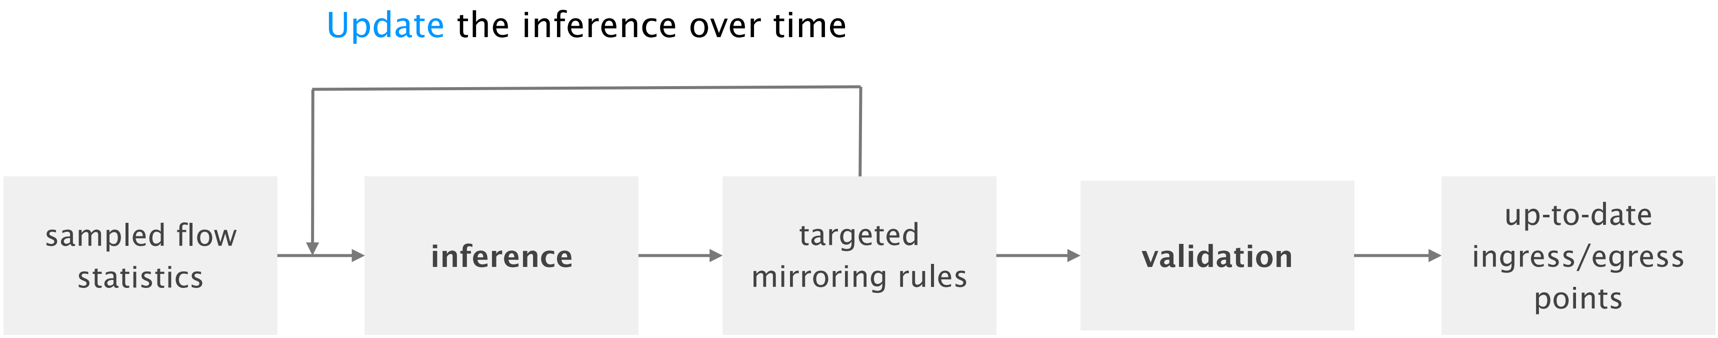

Infer largest subnets matching sampled IPs to ingress or egress

High coverage

Mirror where we do *not* expect traffic to enter or leave

High accuracy

Low overhead