## **ETH** zürich

#### Network Security Group Department of Computer Science

# ALBUS: a Probabilistic Monitoring Algorithm to Counter Burst-Flood Attacks

Simon Scherrer, Adrian Perrig ETH Zürich

Jo Vliegen, Arish Sateesan, Nele Mentens\* KU Leuven, \*Leiden University

Hsu-Chun Hsiao National Taiwan University

SRDS 2023, Marrakech





Goal: Guarantee availability of network resources under malicious traffic from many sources

PoseidonRippleJaqenCOLIBRIACC-Turbo(Zhang et al. 2020)(Xing et al. 2021)(Liu et al. 2021)(Giuliari et al. 2021)(Alcoz et al. 2022)



Goal: Guarantee availability of network resources under malicious traffic from many sources

PoseidonRippleJaqenCOLIBRIACC-Turbo(Zhang et al. 2020)(Xing et al. 2021)(Liu et al. 2021)(Giuliari et al. 2021)(Alcoz et al. 2022)

**DDoS Defense System** 



























DDoS defense systems mostly use the algorithms CountMin-Sketch and Count-Sketch for detection



DDoS defense systems mostly use the algorithms CountMin-Sketch and Count-Sketch for detection

+ Limited memory + Efficient processing

- Limited accuracy



DDoS defense systems mostly use the algorithms CountMin-Sketch and Count-Sketch for detection





DDoS defense systems mostly use the algorithms CountMin-Sketch and Count-Sketch for detection





DDoS defense systems mostly use the algorithms CountMin-Sketch and Count-Sketch for detection



DDoS defense systems mostly use the algorithms CountMin-Sketch and Count-Sketch for detection



DDoS defense systems mostly use the algorithms CountMin-Sketch and Count-Sketch for detection



DDoS defense systems mostly use the algorithms CountMin-Sketch and Count-Sketch for detection



DDoS defense systems mostly use the algorithms CountMin-Sketch and Count-Sketch for detection



DDoS defense systems mostly use the algorithms CountMin-Sketch and Count-Sketch for detection + Limited memory + Efficient processing - Limited accuracy Link capacity Flow sending rate Time Reset Rese For each flow: Reset period Estimated Expected flow volume flow volume >in period so far over period Suspicious flow!

**ETH** zürich

DDoS defense systems mostly use the algorithms CountMin-Sketch and Count-Sketch for detection



DDoS defense systems mostly use the algorithms CountMin-Sketch and Count-Sketch for detection



**ETH** zürich

DDoS defense systems mostly use the algorithms CountMin-Sketch and Count-Sketch for detection



Suspicious flow!

DDoS defense systems mostly use the algorithms CountMin-Sketch and Count-Sketch for detection



DDoS defense systems mostly use the algorithms CountMin-Sketch and Count-Sketch for detection

+ Limited memory + Efficient processing - Limited accuracy



**ETH** zürich



Scenario: Malicious bursts last for 200 milliseconds and are 20% larger than allowed

 $\text{Recall} = \frac{\text{Reported malicious bursts}}{\text{Malicious bursts}}$ 

 $\label{eq:Precision} \mbox{Precision} = \frac{\mbox{Reported malicious bursts}}{\mbox{Reported bursts}}$ 



$\text{Recall} = \frac{\text{Reported malicious bursts}}{\text{Malicious bursts}}$	$Precision = \frac{Precision}{Precision} + Pr$
Recall 🛧	Precision 🔥
100%	100%
*	+



Scenario: Malicious bursts last for 200 milliseconds and are 20% larger than allowed



3/19

Scenario: Malicious bursts last for 200 milliseconds and are 20% larger than allowed



3/19









Scenario: Malicious bursts last for 200 milliseconds and are 20% larger than allowed



3/19

#### How Can We Better Detect Malicious Bursts?

To withstand burst-flood attacks, a monitoring algorithm must satisfy the following requirements:




























ALBUS: Adaptive Leaky-Bucket Undulation Sensor



#### ALBUS: Adaptive Leaky-Bucket Undulation Sensor

Flows:  $f_1$   $f_2$   $f_3$   $f_4$   $f_5$   $f_6$   $f_7$   $f_8$   $\cdots$   $f_{1475286}$ 



ALBUS: Adaptive Leaky-Bucket Undulation Sensor

Flows:  $f_1$   $f_2$   $f_3$   $f_4$   $f_5$   $f_6$   $f_7$   $f_8$   $\cdots$   $f_{1475286}$ 





ALBUS: Adaptive Leaky-Bucket Undulation Sensor

Flows:  $f_1$   $f_2$   $f_3$   $f_4$   $f_5$   $f_6$   $f_7$   $f_8$   $\cdots$   $f_{1475286}$ 





































**ETH** zürich



A leaky bucket reports a flow that sends more than  $\gamma w + \beta$  during a time window with arbitrary width w



A leaky bucket reports a flow that sends more than  $\gamma w + \beta$ during a time window with arbitrary width w

 $\gamma$  = Leakage rate = Flow base rate



A leaky bucket reports a flow that sends more than  $\gamma w + \beta$ 



- during a time window with arbitrary width w
- $\gamma$  = Leakage rate = Flow base rate
- $\beta$  = Bucket volume = Burstiness allowance



A leaky bucket reports a flow that sends more than  $\gamma w + \beta$ 

- $\gamma$  = Leakage rate = Flow base rate
- $\beta$  = Bucket volume = Burstiness allowance



A leaky bucket reports a flow that sends more than  $\gamma w + \beta$ 

- $\gamma$  = Leakage rate = Flow base rate
- $\beta$  = Bucket volume = Burstiness allowance



A leaky bucket reports a flow that sends more than  $\gamma w + \beta$ 

- $\gamma$  = Leakage rate = Flow base rate
- $\beta$  = Bucket volume = Burstiness allowance



A leaky bucket reports a flow that sends more than  $\gamma w + \beta$ 

- $\gamma$  = Leakage rate = Flow base rate
- $\beta$  = Bucket volume = Burstiness allowance







But: A leaky bucket can only monitor a single flow at a time!





nent of Computer Science







Flows: 
$$f_1$$
  $f_2$   $f_3$   $f_4$   $f_5$   $f_6$   $f_7$   $f_8$   $\cdots$   $f_{1475286}$ 

	LB 1		LB 2		LB 3		LB 4		LB 256	
ucket ers:										

Leaky-bucket counters:


















How Can We Apply the LB Algorithm in a Memory-Efficient Fashion?





How Can We Apply the LB Algorithm in a Memory-Efficient Fashion?



Eviction criterion:

If LB net inflow turns negative



How Can We Apply the LB Algorithm in a Memory-Efficient Fashion?



ALBUS: Adaptive Leaky-Bucket Undulation Sensor

Flows:  $f_1$   $f_2$   $f_3$   $f_4$   $f_5$   $f_6$   $f_7$   $f_8$   $\cdots$   $f_{1475286}$ 







ALBUS: Adaptive Leaky-Bucket Undulation Sensor



ALBUS: Adaptive Leaky-Bucket Undulation Sensor



































#### Question:

How to find the most bursty flow among the background flows of a leaky-bucket counter?



#### Question: How to find the most bursty flow among the background flows of a leaky-bucket counter?





#### Question: How to find the most bursty flow among the background flows of a leaky-bucket counter?





#### Question: How to find the most bursty flow among the background flows of a leaky-bucket counter?

Probabilistic decay (Yang et al. 2019)

















Department of Computer Science

10/19







ment of Computer Science



ment of Computer Science

















**ETH** zürich


















































## How Does ALBUS Work?





**ETH** zürich

## How Does ALBUS Work?





**ETH** zürich

## How Does ALBUS Perform Under Burst-Flood Attacks?



#### How Does AI BUS Perform Under Burst-Flood Attacks?

Scenario: Malicious bursts last for 200 milliseconds and are 20% larger than allowed



#### How Does AI BUS Perform Under Burst-Flood Attacks?

Scenario: Malicious bursts last for 200 milliseconds and are 20% larger than allowed





#### How Does AI BUS Perform Under Burst-Flood Attacks?

Scenario: Malicious bursts last for 200 milliseconds and are 20% larger than allowed





#### How Does ALBUS Perform Under Burst-Flood Attacks?

Scenario: Malicious bursts last for 200 milliseconds and are 20% larger than allowed



Department of Computer Science

To withstand burst-flood attacks, a monitoring algorithm must satisfy the following requirements:









To withstand burst-flood attacks, a monitoring algorithm must satisfy the following requirements:







ALBUS



Network Security Group Department of Computer Science

To withstand burst-flood attacks, a monitoring algorithm must satisfy the following requirements:





To withstand burst-flood attacks, a monitoring algorithm must satisfy the following requirements:







Yes. ALBUS has low computational complexity:



Yes. ALBUS has low computational complexity:

Single hash computation



Yes. ALBUS has low computational complexity:

Single hash computation

No counter-array iterations



Yes. ALBUS has low computational complexity:

Single hash computation

No counter-array iterations

No associative arrays



Yes. ALBUS has low computational complexity:

Single hash computation

No counter-array iterations

No associative arrays

⇒ ALBUS is hardware-friendly



Yes. ALBUS has low computational complexity:

Single hash computation

No counter-array iterations

No associative arrays

 $\implies$  ALBUS is hardware-friendly

FPGA implementation for Xilinx Virtex UltraScale+ FPGA:



Hardware design of a LB-BC combination



Yes. ALBUS has low computational complexity:

Single hash computation

No counter-array iterations

No associative arrays

⇒ ALBUS is hardware-friendly

FPGA implementation for Xilinx Virtex UltraScale+ FPGA:



 $\begin{array}{l} \textit{Throughput:}\\ \texttt{200 million packets}\\ \textit{per second}\\ \sim \texttt{560 Gbps!} \end{array}$ 

Hardware design of a LB-BC combination





ALBUS is considerably more effective under burst-flood attacks than previous monitoring algorithms.



ALBUS is considerably more effective under burst-flood attacks than previous monitoring algorithms.

#### CountMin-Sketch Count-Sketch

And the operation of th	Fixed tim	e-windows
Suspecieus ferul	Child advance spatial standing of the spatial standing	nn Countrille Statut au Count Statut for dimension or promoting - Longe Statut (Countril St



ALBUS is considerably more effective under burst-flood attacks than previous monitoring algorithms.

#### CountMin-Sketch Count-Sketch

	5-001100005
Etild allevar spanne nanty se tra slightlen a unad anny a filter for and the start of the star	Countral Status and Count Status for denoting processing - united Status of Burst Faced Attack Burst Status and Barst Rous Burst and process Burst and process Rout and process Rout and process Rout and process Rout and process





ALBUS is considerably more effective under burst-flood attacks than previous monitoring algorithms.

CountMin-Sketch Count-Sketch

Fixed time-	windows
The rank for a spectrum of the	American and Cauce Starts for detection and the starts of the starts of the starts the starts of the starts of the starts Bear View Free Area: Bear View Fre
FILADA DISTRICT.	**

Accuracy trade-off	$\vdash$
Additional program and an additional program and additional program	









ALBUS is considerably more effective under burst-flood attacks than previous monitoring algorithms.





**ETH** zürich

ALBUS is considerably more effective under burst-flood attacks than previous monitoring algorithms.



17/19

# Additional Material



#### How Do Background Counters Help Detection Accuracy?

Probabilistic decay identifies large background flows



#### How Do Background Counters Help Detection Accuracy?

Probabilistic decay identifies large background flows p =Count-modification probability



#### How Do Background Counters Help Detection Accuracy?

Probabilistic decay identifies large background flows

p = Count-modification probability

p = 1: Majority algorithm (Boyer et al. 1981)


Probabilistic decay identifies large background flows

p =Count-modification probability

p = 1: Majority algorithm (Boyer et al. 1981)

BC contains flow f with probability  $q_f(1) = \min\left(1, \frac{\text{Volume of flow } f}{\text{Volume of all flows } \neq f}\right)$ 

Probabilistic decay identifies large background flows

p =Count-modification probability

p = 1: Majority algorithm (Boyer et al. 1981)

BC contains flow f with probability  $q_f(1) = \min\left(1, \frac{\text{Volume of flow } f}{\text{Volume of all flows } \neq f}\right)$ 



Probabilistic decay identifies large background flows

p =Count-modification probability

p = 1: Majority algorithm (Boyer et al. 1981)

BC contains flow f with probability  $q_f(1) = \min\left(1, \frac{\text{Volume of flow } f}{\text{Volume of all flows } \neq f}\right)$ 



Probabilistic decay identifies large background flows

p =Count-modification probability

p = 1: Majority algorithm (Boyer et al. 1981)

BC contains flow f with probability  $q_f(1) = \min\left(1, \frac{\text{Volume of flow } f}{\text{Volume of all flows } \neq f}\right)$ 





Probabilistic decay identifies large background flows p =Count-modification probability

p = 1: Majority algorithm (Boyer et al. 1981)

BC contains flow f with probability  $q_f(1) = \min\left(1, \frac{\text{Volume of flow } f}{\text{Volume of all flows} \neq f}\right)$ 





Probabilistic decay identifies large background flows

p =Count-modification probability

p = 1: Majority algorithm (Boyer et al. 1981)

BC contains flow f with probability  $q_f(1) = \min\left(1, \frac{\text{Volume of flow } f}{\text{Volume of all flows} \neq f}\right)$ 





Probabilistic decay identifies large background flows p =Count-modification probability

p = 1: Majority algorithm (Boyer et al. 1981)

BC contains flow f with probability  $q_f(1) = \min\left(1, \frac{\text{Volume of flow } f}{\text{Volume of all flows} \neq f}\right)$ 





Probabilistic decay identifies large background flows p =Count-modification probability

p = 1: Majority algorithm (Boyer et al. 1981)

BC contains flow f with probability  $q_f(1) = \min\left(1, \frac{\text{Volume of flow } f}{\text{Volume of all flows} \neq f}\right)$ 



