

# Poster: Scaling Up Anomaly Detection Using In-DRAM Working Set of Active Flows Table

Rhongho Jang<sup>†‡</sup>, Seongkwang Moon<sup>†</sup>, Youngtae Noh<sup>†</sup>, Aziz Mohaisen<sup>‡</sup>, and Daehun Nyang<sup>†</sup>  
<sup>†</sup>INHA University      <sup>‡</sup> University of Central Florida

**Abstract**—In the zettabyte era, per-flow measurement becomes more challenging owing to the growth of both traffic volumes and the number of flows. Also, swiftness of detection of anomalies (e.g., DDoS attack, congestion, link failure, and so on) becomes paramount. For fast and accurate anomaly detection, managing an accurate working set of active flows (WSAF) from massive volumes of packet influxes at line rates is a key challenge. WSAF is usually located in a very fast but expensive memory, such as TCAM or SRAM, and thus the number of entries to be stored is quite limited. To cope with the scalability issue of WSAF, we propose to use In-DRAM WSAF with scales, and put a compact data structure called FlowRegulator in front of WSAF to compensate for DRAM’s slow access time by substantially reducing massive influxes to WSAF without compromising measurement accuracy. We prototype and evaluated our system in a large scale real-world experiment (connected to monitoring port of our campus main gateway router for 113 hours, and capturing 122.3 million flows). As one key application, FlowRegulator detected heavy hitters with 99.8% accuracy.

## I. INTRODUCTION

Due to the high speed network, traffic measurement now have to cope with enormous incoming data rates (*i.e.*, larger number of flows) with tight deadlines (*i.e.*, real-time). We stress that large-scale instant measurement is highly necessary for network anomaly detection. For example, if denial of service (DoS) attacks cause an influx of packets at 100 Gbps, detection delay of 100 ms will cause 1.2GB data to hit a server or a network. Therefore, to avoid large bandwidth penalties, instant anomaly detection is essential.

A working set of active flows (WSAF) is a type of cache of a full flow table, which can be found usually in TCAM (Ternary Content Addressable Memory), CAM, or sometimes SRAM for traffic monitoring. For instance, NetFlow uses TCAM for storing WSAF in which an entry consists of a flow ID and the counting value. However, the number of entries of the table cannot be large because those types of memories are quite expensive. For scalability, we can put WSAF in DRAM instead of the expensive memory (*i.e.*, incentive to cost-effectiveness). However, there is a speed issue for In-DRAM WSAF: a packet arrival rate is too fast to handle by In-DRAM WSAF, owing to the DRAM’s speed and WSAF table’s hash collision.

To cope with these issues, sketch-based techniques have been greatly enhanced over several decades because sketch-based counting algorithms only require a small amount of memory to *encode* a large volume of traffic in real-time. Due to their design, however, most of *decode* algorithms involve hundreds of hash calculations and memory accesses from statistically mixed random blocks to obtain meaningful

statistics [1]. For this reason, offline decoding in a high-performance server is commonly accepted in practice but inherently incurs huge network delay. Thus, online decoding is highly necessary for instant measurement and further timely detection.

Unfortunately, most sketch-based algorithms lack scalability and online decoding capabilities. Our approach to solve these two problems is 1) to use a counting algorithm that can perform online decoding and 2) to put a flow regulator before WSAF to slow down the incoming packet rate to WSAF. To realize both ideas, we designed a highly scalable counting and flow regulating algorithm called FlowRegulator. By design, instead of directly inserting or updating every packet of a flow into WSAF table, FlowRegulator (*i.e.*, a small cache buffer) retains a fraction of flow counts. By doing so, we can suppress frequent WSAF updates in DRAM; thereby FlowRegulator can support large-scale influx of flows with the use of cost-effective large DRAM. Consequently, FlowRegulator relaxes the necessity of expensive memories (TCAM or SRAM) for maintaining large WSAF, and further enables us to build a highly scalable and fast measurement system. We conducted a real-world campus network experiment for 113 hours by connecting our prototype of FlowRegulator to a mirroring port of a main gateway router, capturing 9.11 billion packets, 122.3 million flows, and 8.5TB bytes. FlowRegulator successfully measured the whole L4 flows with a standard error (0.65%) and detected heavy hitters with 99.8% accuracy.

## II. FLOWREGULATOR DESIGN

Our large WSAF in DRAM is in contrast to the small WSAF in TCAM (*i.e.*, industry practice). In DRAM, we can store much more flows, thereby, we do not need a remote collector for decoding. However, the downside is that we cannot evade the “sluggishness” of DRAM.

**FlowRegulator to relax the {ips = pps} constraint:** Instead of directly inserting or updating every flow packet into the table, we put a small buffer called FlowRegulator to retain a fraction of flow counts before WSAF. FlowRegulator has a memory block (or a virtual vector initialized to all 0’s) for every single flow, and whenever a packet comes in, the corresponding block is updated by setting a random bit of the block. When the block saturates (or a portion of block has set to 1’s), the resulting counting fraction (we note that this is not the total size of a flow) is added up to the WSAF (*i.e.*, a hash table in DRAM). Because FlowRegulator retains mice flows whose sizes are lower than the saturation condition, not all the packets are fed into WSAF, but only the packets that trigger

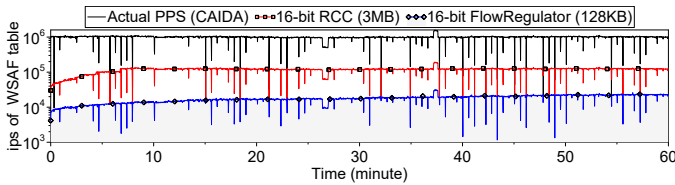


Fig. 1. WASF relaxation: FlowRegulator (FR) and RCC ips of CAIDA dataset

the saturation condition are given to WSAF. This design will greatly reduce insert per second (ips) even under a high packet per second (pps) condition.

To develop FlowRegulator, we utilize the recyclable counter with confinement (RCC) [4] that already has online decoding capability, and proven to be useful for measurement in the wireless SDN environment [2], [3]. To investigate its feasibility, we have tested RCC for its rate regulation (defined as Output ips/Input pps). Given that access time of SRAM is 10-20 times faster than DRAM's (and even faster with TCAM), RCC's rate regulation should be less than 5%. However, its regulation and retention capacity (the maximum number of packets in a virtual vector) are not operationally sufficient. Thus, it is impossible to work with RCC for building FlowRegulator. One way to increase the rate regulation is to give RCC a larger virtual vector, but that does not expand the retention capacity.

**Two-layer design for higher rate regulation:** Here, our observation is that enlarging the virtual vector size increases the retention capacity just in an additive manner, and thus, this is not a viable (*i.e.*, scalable) option. Instead, we designed a new counting algorithm for FlowRegulator, which has two layers of probabilistic counters to achieve the higher rate regulation. Our FlowRegulator plays a key role of retaining flows (from feeding into WSAF) for a while as well as counting flows. In the two-layer design, the second (higher) layer's one bit encodes multiple packets of a flow from a saturated sketch of the first (lower) layer. This design has substantially improved the rate regulation in a multiplicative manner. It enables higher rate regulation while not being detrimental to the accuracy and speed, while being scalable.

### III. EVALUATION

**WSAF ips relaxation.** In Fig. 1, the x-axis represents the time line of our CAIDA dataset, and the black solid line on the top represents the actual pps of the trace. Below the pps line, RCC's and FlowRegulator's regulation rates are shown in red squares and blue diamonds, respectively. The figure shows that RCC relaxes ips to feed packets to WSAF table at the speed of 112 kips (thousand ips), which corresponds to 12% regulation rate. FlowRegulator effectively regulated flows to pass only 1.02% with 128KB DRAM memory. As results, FlowRegulator has sufficient margin, while RCC does not have as can be seen in Fig. 1.

**Monitoring in the wild.** We implemented our FlowRegulator in an off-the-shelf device and measured up-link traffics (1 Gbps bandwidth) at the backbone gateway (Juniper EX9208 switch) of our campus for 113 hours in total. During 113 hours, 9.1 billion packets of 122.3 billion L4 flows were measured simultaneously both in packets and in bytes.

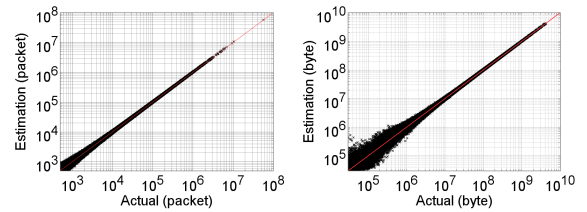


Fig. 2. Accuracy of packet counting (left) and byte counting (right).

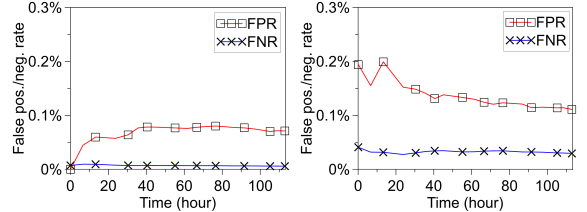


Fig. 3. False positive and false negative rates of packet heavy hitter detection (left) and byte volume heavy hitter detection (right).

**Accuracy.** FlowRegulator used 128KB for sketch, and 33MB for the WSAF table. Sketches and WSAF table are all in DRAM. Fig. 2 shows the estimation accuracy by standard error for the real-world experiment. For packet counting, we report 0.54% standard error over 350 flows of which size is 1000K+, 1.61% over 11,047 flows for 100K packets, 3.46% over 104292 flows for 10K+ packets. For byte counting, we report 0.63% over 414 flows of which byte size is 1G+, 1.74% over 12,125 flows of 100MB+, 3.65% over 107,726 flows of 10MB+. This accuracy matches the accuracy observed in the lab experiment with the CAIDA dataset.

**Heavy hitter detection.** Fig. 3 shows FlowRegulator's heavy hitter detection accuracy in terms of false positive/negative rate. Owing to FlowRegulator capability of counting both in packets and in bytes, it can detect both packet heavy hitters and byte heavy hitters. False negative rates in both cases are negligible, and the false positive rates of packet/byte heavy hitters are less than 0.1% and 0.2%, respectively.

### IV. CONCLUSION

In this work, we have developed FlowRegulator for instant flow monitoring. Our approach is different from conventional measurement frameworks by introducing a new notion of very large In-DRAM working set of active flows. In the future work, we plan to demonstrate FlowRegulator's performance and feasibility through an extensive analyses.

**Acknowledgement.** This work was supported by NRF grant number 2016K1A1A2912757 (Global Research Lab Initiative).

### REFERENCES

- [1] Q. Huang, X. Jin, P. P. C. Lee, R. Li, L. Tang, Y. Chen, and G. Zhang. Sketchvisor: Robust network measurement for software packet processing. In *Proc. of ACM SIGCOMM*, pages 113–126, 2017.
- [2] R. Jang, D. Cho, A. Mohaisen, Y. Noh, and D. Nyang. Two-level network monitoring and management in WLAN using software-defined networking: poster. In *In Proc. of ACM WiSec*, pages 279–280, 2017.
- [3] R. Jang, D. Cho, Y. Noh, and D. Nyang. Rflow+: An sdn-based wlan monitoring and management framework. in *Proc. of INFOCOM*, 2017.
- [4] D. Nyang and D. Shin. Recyclable counter with confinement for real-time per-flow measurement. *IEEE/ACM Trans. Netw.*, PP(99):1–1, 2016.

# + '37.&LFM#"5)%FJ6FN)' / "-8FO.7.#75')FJ35)%F0)DO9NKF P '&\$5)%FM.7F'EFN#754.F:-';3FC",.-.

%J\*" . J\*(R3" . 90\*" . , )3" . '\*\*\*" S\*G" .130 N\*J &T-T(' \* J3-/0" \$30UG"(N?3" .  
+JW3" . X , "-. J1/VG42V0<G / , B\*\* X /0463>V-" J3J34V , + ?1"\* JX -" J3J34V , + B\* J3-/0" XG42V0<G " ?3" . X -" J3J34V , +

## !NO% I \$PCO! I N

!"#\$%&'()\*+,-./01(\*2(341-50(26\*)/(78&9:;  
!"#\$%&'()\*+ / B( B? / ?0? )' / C (3 . / B@ (4A, 2-40/7)K-2/  
'.)B-.B &/-)+4/&/ .B ') ?0?2% . / K( ) -4% 3(4 B? / . / B@ (4A  
- . ( & -2% C/B/KB\* ( . P  
+&' , - / 9(4 & ( . \*B(4\* . 0 B4-33\*K) , ! / B92( @ + / ) > 1GQ 3(4)B(4\* . 0  
R#G9\* . @?K? - . / . B4% K( . )B (3 - 32( @ || " - . C B? / K( + . B\* . 0  
N-2+ / P S( @ / N/4, B? / . + & : / 4 (3 / . B4\* / ) (3 B? / B- : 2/ K- . . (B : /  
2-40/ : / K- + ) / B? ( ) / B% ' / ) (3 & / & (4\* / ) -4/ T+\*B/ / U' / . )N/P  
0\* . " 9(4)K-2- : \*2B% , @ / K- . ' + B R#G9\* . " DGQ\* . )B/ - C (3 B? /  
/ U' / . )N/ & / & (4% I\*P/P , \* . K / . B\*N/ B( K( ) B7/33/KB\*N/ . / ) )OP  
12"- . )% . 3 S( @ / N/4, B? / 4\* ) - ' / / C \* ) + / 3(4 || . 7" DGQ R#G9V  
- ' - KA/B - 44\*N-2 4-B/ \* ) B( ( 3- ) B B( ? - . C2/ : % || . 7" DGQ R#G9,  
( @ \* . 0 B( B? / " DGQ || ) ' / / C - . C R#G9 B- : 2/ || ) ? - ) K( 22\* ) ( . P  
+&. 45' (3 "66&" "#2.3 !A/BK?7: - ) / C B/K? . \*T+ / ) ? - N / : / . + ) / C  
( . 2% 4/T+\*4/ - ) & - 22 - & ( + . B (3 & / & (4% B( / . K(C/ - 2-40/  
N(2+& / (3 B4-33\*K\* . 4/ - 27B\* & / P X . 3(4B+ . - B/2% , & ( ) B ) A/BK?7: - ) / C  
- 20(4\*B? & ) 2- KA 3#"- , 5-578 - . C \* "6-" 0 < 04\* < - . K- ' - : \*2B\* / ) P

## I G+(3AA+\*34J

@. 8F5\* . " ` +B\$ - \$32 ( @ \$ / O + 2 - B ( 4\$ \* . \$34 ( . B\$ ( 3\$ R # G9\$B ( \$K ( & ' / . ) - B / \$3 ( 4\$  
" DGQ || ) \$ 2 ( @ \$ - KK / ) ) \$B\* & / \$ : % \$ ) + : ) B - . B\* - 22% \$ 4 / C + K\* . O\$ & - ) \* N / \$\* . 32 + U / ) \$  
B ( \$ R # G9\$  
%063@ (1J0D-A/ E(AA/F(4\* " /1+3-" 1K  
! R'B? ( +B\$92 ( @ D / O + 2 - B ( 4, \$ \* - KA / B\$ / 4\$ ) / K ( . C\$ I ' ' ) O\$ ( 3\$B4 - 33\*K ) \$ / T + - 2\$  
B ( \$\* . ) / 4B\$ ' / 4\$ ) / K ( . C\$ I ' ' ) O\$ ( 3\$ R G # 9\$ I "#\$#&#\* ' ) a\$ ' ' ) OP  
! R'B? \$92 ( @ D / O + 2 - B ( 4, \$ @ / \$ 4 / B\* - . \$ - \$34 - KB\* ( . \$ ( 3\$32 ( @ \$K ( + . B ) \$ : / 3 ( 4 / \$  
R # G9\$ R ? / . \$B? / \$32 ( @ \$K ( + . B ) \$ ) - B + 4 - B / ) , \$B? / \$K ( + . B\* . O\$N - 2 + / \$\* ) \$  
- CC / C\$ + ' \$B ( \$B? / \$ R # G9\$ I "#\$#&#\* ' ) b\$ ' ' ) OP  
: 031G+0/ (\*2( : 6\* ) %0 . G631\* +K  
! # & - 22\$ & / & ( 4% \$ + ) - O / \$ I "#\$#&#\* ? + . C4 / C ) \$ ( 3\$F \ O\$B ( \$3\$B\$B? / \$L - \$K - K? / P  
! #K - 2 - : 2 / \$K ( + . B\* . O\$K - ' - K'B\$ I "#\$#&#\* 4 / B\* - . \* . O\$ ' ( 4B\* ( . \$ ( 3\$2 - 40 / \$32 ( @ ) ) OP  
! # & - 22\$32 ( @ \$ ) - & ' 2\* . O\$ I "#\$#&#\* ' 4 ( : - : \*2\* ) B\* K - 22% OP\$  
! c . 2\* . / \$C / K ( C\* . O\$K - ' - K'B\$ I "#\$#&#\* / UB4 - KB\$K ( + . B\* . O\$N - 2 + / ) P\$  
!BA60B0" 131- \* "  
! > @ ( 72 - % / 4\$C ) \* O . \$ ( 3\$D / K%K2 - : 2 / \$ 1 ( + . B / 4\$ @ B? \$ 1 ( . 3\* . / & / . B\$ I D 1 1 , \$  
' + : 2\* ) ? / C\$\* . \$ > ( ! O\$3 ( 4\$ ) K - 2 - : 2 / \$ \* ) 4 / C + KB\* ( . \$ I "#\$#&#\* ' ) bb\$ ' ' ) OP\$

## Q536G31-\* "

=3>\*+31\*+?(0@A0+-B0"1(7C&!\$&(<313/01;  
8&9: DAA/ E-A/F+063@31-\* "

! 911 < &' ' ) Y < < 5 A' ' ) I < 5 Z @ B? [ Q \ & / & ( 4% OP  
! : - ' ; 9 . % ( - ' 7 ' < &' ' ) Y < ; A' ' ) I < 2 ; 5 Z @ B? < 5 ] F \ & / & ( 4% OP  
&44G+34? 7/13 "<3+<0++\*";

! + "#\$.7#'( )7.& <=>? @ ! / . / AB ; P5 ] I < ; ; F^ & ( @ ) 0 , ; P == Z I < ; ; F^ & ( @ ) 0 ,  
! 187.#'( )7.& <=>? @ ! / . / AB ; P58 Z I < J \ ^ & ( @ ) 0 , <? ; ; Z I < ; ; Q \ \ ^ & ( @ ) 0  
! C'60@5\*. )755#75' ) D / K - 22 \_ = HZ 18( ' < &' 22' ( . & ( @ ) OP

## H+\*1\*1?A0

Hardware configuration Scenario (campus gateway) Multi-core design (DPDK)

%036#) \*+6<(0@A0+-B0"1(7LLM#J\*G+(43BAG/(.310)3?;  
&44G+34?

! + "#\$.7#'( )7.&F=H@ 1F / . / AB R / \$4 / ' ( 4B\$ ; PH6 Z \$ ) B - . C - 4C\$ / 44 ( 4\$ N / 4\$ [ H ; \$32 ( @ ) \$ ( 3\$ @ ? \* K ? \$  
) d / \$ ) \$ ; ; F^ < P E < Z \$ ( N / 4\$ < ; ; 68\$32 ( @ ) \$3 ( 4\$ < ; ; F^ & [ P6E Z \$ ( N / 4\$ < ; ; 6,5 = 5\$32 ( @ ) \$3 ( 4\$ < ; ; F^ & \$  
! 187.##'( )7.&F=H@ 1F / . / AB R / \$4 / ' ( 4B\$ ; PE [ Z \$ ( N / 4\$ 6 < 6\$32 ( @ ) \$ ( 3\$ @ ? \* K ? \$ ) d / \$ ) \$ < J \ ^ , \$  
< P6E Z \$ ( N / 4\$ < ; 5H\$32 ( @ ) \$ ( 3\$ < ; ; Q \ \ ^ [ PEH Z \$ ( N / 4\$ < ; ; 8,85E\$32 ( @ ) \$ ( 3\$ < ; ; Q \ \ ^

H0+2\*B3"40(7U035?(J-110+(Y(CHP(Y' Q' ;

False positive and false negative rates of packet heavy hitter detection (left) and byte volume heavy hitter detection (right).

! I . "48f2577. & 9-2) / \$ . / O - B\* N / \$4 - B / ) \$\* . \$ : ( B? \$K - ) / ) \$ - 4 / \$ . / O? O' : 2 / , \$ - . C\$B? / \$3-2) / \$' ( ) B\* N / \$  
4 - B / ) \$ ( 3\$\* - KA / B\$ : \$B / \$ ? / - N% \$ ? \* B\$ / 4 ) \$ - 4 / \$ / ) ) \$B? - . \$ ; P < Z \$ - . C\$ ; P5 Z , \$4 / ) ' / KB\* N / 2% P\$  
! 1+Jf>? / \$K ( 4 / I ) \$ @ ( 4A ( - C\$ & - BK? / ) ) \$B? / \$B4 - 33\*K\$ - BB / 4 . \$ - . C\$B? / \$K ( 4 / \$ + ) - O / C\$ C\$ . ( B\$O ( \$ ( N / 4\$  
6 ; Z \$ - \$B - . % \$ ( ' . B\$P\$  
! K . / ' & 8F\$ ) \$3 ( 4\$B? / \$T + / + / \$1 : 2 - KA\$C\* - & ( . C ) \$\* . \$B? / \$3\* O + 4 / 0 , \$1\$B\$C\$ . ( B\$O4 ( @ \$ . ( B\* K / - : 2% P\$  
> ? / \$4 / ) + B\$ ) K ( . 3\*4 & / C\$B? - B\$FlowReguloar \* & ' 2 / & / . B / C\$ ( . \$G\$ ( & \$ : ( - 4C\$ @ ( 4A / C\$ @ / 22\$3 ( 4\$B? / \$ < \$  
J : ' ) \$ . / B @ ( 4A\$ & ( . B ( 4\* . O\$ . - . C\$3 ( 4\$ - \$T + \* B / \$ ( . O\$B\$ / P

!"#\$%&'()\*+&,\$  
#-. \$\*/O(,\$1-2\*3(4.- \$  
56758\$9/ : 4+ - 4%\$5 ; < =

>?\*)\$@ (4A\$@- )\$) + ' ' (4B / C\$ : %  
! D9\$04 - . B\$ . + & : / 4\$5 ; < EF < G < G5 = < 58H8\$  
! J2 ( : - 2\$D / ) / - 4K? \$L - : \$I . B\* - B\* N / 0