

Detecting adversarial attacks on Bloom filters in P4 data plane systems

Dhiraj Saharia
IIT Hyderabad
dhiraj.saharia@cse.iith.ac.in

K.Shiv Kumar
IIT Hyderabad
cs21resch11003@iith.ac.in

ABSTRACT

The programmable data plane has made the network programmer define the intended packet processing behaviour inside the switch. P4, a high-level domain-specific language, let the programmers define the tables that a packet should hit, the actions applied to it and the conditionals to check different properties of the network traffic. P4 exposes a stateful construct, i.e. register, which can be used to implement memory-efficient hash-based probabilistic data structures like Bloom filter. The bloom filter is used for various network applications for efficiently tracking the network traffic for analysis, traffic engineering, security, load balancing etc. However, since the bloom filter is a probabilistic data structure, it can be polluted by inserting adversarial elements. Despite being memory efficient, bloom filter can become a target for the pollution and saturation attacks and result in corruption of network statistics, increased latency by making a path for DDoS attacks etc., by increasing false positive probabilities. An adversary can forge carefully chosen items that pollute/saturate the bloom filter, especially in such applications that cannot tolerate false positives. The adversary can also introduce delay by querying carefully chosen elements not present in the bloom filter making the query expensive. It is crucial to study the impact of such attacks on various data plane systems relying on the bloom filter and develop a system to detect such attacks.

1 INTRODUCTION

Bloom filter is a hash-based probabilistic data structure primarily used to check an element's membership in the set efficiently. It is implemented in the form of a bit array. It is one of the most popular data structures used in online security applications because it reduces the memory consumption of the applications. Recent advancements in programmable white box switches have made it possible to realize such data structures for various networking applications. However, despite being space-time efficient, the bloom filters are prone to False Positives, which can cause havoc in various network applications. The increase in the False Positive in the bloom filter will result in erroneous results, i.e., the bloom filter will show that the item is present in the bloom filter even though that item has never been inserted.

Different network applications can tolerate different values of FPR. For example, the FPR needs to be very small for security applications so that any malicious traffic cannot bypass the security policies. The optimal parameters have to be selected very carefully before implementation to use the bloom filter efficiently. However, the bloom filters can still be exploited by an adversary and can be polluted, resulting in the bloom filter giving erroneous results [1]. An adversary can send carefully crafted packets with an objective to increase the number of set bits in the bloom filter. Such attacks can make the bloom filter behave abnormally, for example, giving wrong per-flow loss in the case of network analysis applications like FlowRadar [3].

2 FUTURE STEPS

The main idea of this proposal is to explore the possibilities of detecting attacks on the bloom filter by utilizing the per-packet statistics collected by DBVal [2]. Using DBVal we can extract the path distribution of the traffic and differentiate the normal and adversarial traffic using statistical techniques such as Chi-squared test [4]. To realize this idea we have identified following high-level objectives for the future:

- To study and demonstrate the impact of possible adversarial attacks on bloom-filter based systems in P4 data-plane.
- To study the characteristics of normal and attack traffic and develop a framework for different P4 targets to detect adversarial attacks.

REFERENCES

- [1] Thomas Gerbet, Amrit Kumar, and Cédric Lauradoux. 2014. The Power of Evil Choices in Bloom Filters. In *IEEE IFIP*.
- [2] K Shiv Kumar, PS Prashanth, Mina Tahmasbi Arashloo, Venkanna U, and Praveen Tammana. 2021. DBVal: Validating P4 Data Plane Runtime Behavior. In *Proceedings of the ACM SIGCOMM Symposium on SDN Research (SOSR)*. 122–134.
- [3] Yuliang Li, Rui Miao, Changhoon Kim, and Minlan Yu. 2016. {FlowRadar}: A Better {NetFlow} for Data Centers. In *13th USENIX symposium on networked systems design and implementation (NSDI 16)*. 311–324.
- [4] Archit Sanghi, Krishna P Kadiyala, Praveen Tammana, and Saurabh Joshi. 2021. Anomaly Detection in Data Plane Systems using Packet Execution Paths. In *Proceedings of the ACM SIGCOMM 2021 Workshop on Secure Programmable network Infrastructure*. 9–15.