

High-Throughput and Memory-Efficient Multimatch Packet Classification Based on Distributed and Pipelined Hash Tables

Abstract:

The emergence of new **network** applications, such as the **network** intrusion detection system and packet-level accounting, requires packet classification to report all matched rules instead of only the best matched rule. Although several schemes have been proposed recently to address the multimatch packet classification problem, most of them require either huge memory or expensive ternary content addressable memory (TCAM) to store the intermediate data structure, or they suffer from steep performance degradation under certain types of classifiers. In this paper, we decompose the operation of multimatch packet classification from the complicated multidimensional search to several single-dimensional searches, and present an asynchronous pipeline architecture based on a signature tree structure to combine the intermediate results returned from single-dimensional searches. By spreading edges of the signature tree across multiple hash tables at different stages, the pipeline can achieve a high throughput via the interstage parallel access to hash tables. To exploit further intrastage parallelism, two edge-grouping algorithms are designed to evenly divide the edges associated with each stage into multiple work-conserving hash tables. To avoid collisions involved in hash table lookup, a hybrid perfect hash table construction scheme is proposed. Extensive simulation using realistic classifiers and traffic traces shows that the proposed pipeline architecture outperforms HyperCuts and B2PC schemes in classification speed by at least one order of magnitude, while having a similar storage requirement. Particularly, with different types of classifiers of 4K rules, the proposed pipeline architecture is able to achieve a throughput using perfect hash tables.