Optimization of packet classification for FPGA and off-chip memory

Orsák, Michal

Packet classification is a crucial operation which has an influence on the final throughput of a network device. For networks of speed $10Gb/s$ and faster it is required to use a hardware acceleration. Many hardware accelerators for this task were developed but utilization of off-chip memories for $Tbps$ networks is still a challenging task due update time limitations, memory access and size limitation. Packet classification corresponds to a N-dimensional point location problem. Decision-tree and decomposition are the most well-known algorithmic approaches. Decision-tree based algorithms suffers from long update time and rule replication problem which is causing memory insufficiency. Decomposition methods requires cross production tables to merge results from sub-classifiers. Cross production tables tends to be too large with increasing number of dimension, thus causing memory insufficiency.

In this paper we focus on classifiers for OpenFlow using FPGA with off-chip or HBM memory. OpenFlow uses more than 40 dimensions and requires fast update. Current OpenFlow switches usually using decomposition based algorithms like Tuple Splace Search (TSS)[2]. There are decision-tree based algorithms with better memory utilization, some of them are using heterogeneous trees which can be improved further to satisfy fast average update time.

The EffiCuts algorithm partially solves rule replication problem by separation of overlapping rules into separate decision trees. This approach greatly improves memory efficiency but in the cost of additional memory access[3]. The SmartSplit algorithm improves memory efficiency and classification time by utilizing multiple type of operations in the node of a decision-tree[1]. The PartitionSort algorithm introduces the concept of ruleset sortability which improves update time by separation of the sub-tree which requires update when there is a change in ruleset[4].

In this paper we present a method for design of hybrid hardware accelerators for packet classification which combines multiple existing algorithms into a single accelerator and an algorithm for ruleset mapping. Goal is use of off-chip memories, high utilization of on-chip memories and low update time.

Core of our accelerator is based on decision-trees and hashing with bloom filters. The ruleset mapping algorithm is based on PartitionSort algorithm. The modified algorithm stores parts of tree into hash table, thus reducing size of tree and update complexity. To ensure small update time there is an Tuple Space Search based classifier which is used as a buffer to amortize large updates of main classification tree. Architecture is optimized for specific memory resources and statistical features of ruleset. Entropy based heuristic approach is used for detection of decision-tree sub graphs suitable for specific component type.

References

[1] P. He, G. Xie, K. Salamatian, and L. Mathy. Meta-algorithms for software-based packet classification. In *2014 IEEE 22nd International Conference on Network Protocols*, pages 308–319, Oct 2014.

[2] Ben Pfaff, Justin Pettit, Teemu Koponen, Ethan J. Jackson, Andy Zhou, Jarno Rajahalme, Jesse Gross, Alex Wang, Jonathan Stringer, Pravin Shelar, Keith Amidon, and Martín Casado. The design and implementation of open vswitch. In *Proceedings of the 12th USENIX Conference on Networked Systems Design and Implementation*, NSDI'15, pages 117–130, Berkeley, CA, USA, 2015. USENIX Association.

[3] Balajee Vamanan, Gwendolyn Voskuilen, and T. N. Vijaykumar. Efficuts: optimizing packet classification for memory and throughput. *SIGCOMM Comput. Commun. Rev.*, 41(4):–, August 2010.

[4] S. Yingchareonthawornchai, J. Daly, A. X. Liu, and E. Torng. A sorted-partitioning approach to fast and scalable dynamic packet classification. *IEEE/ACM Transactions on Networking*, 26(4):1907–1920, Aug 2018.