Scalable Packet Classification using Distributed Crossproducting of Field Labels

David E. Taylor Jonathan S. Turner
{det3,jst}@arl.wustl.edu

Washington University in Saint Louis

Packet classification is an enabling function for a variety of applications including Quality of Service, security, monitoring, and multimedia communications. Such applications typically operate on packet flows; therefore, network nodes must classify individual packets traversing the node in order to assign a flow identifier, *FlowID*. Packet classification entails searching a set of filters for the highest priority filter or set of filters which match the packet. At minimum, filters contain multiple field values that specify an exact packet header or set of headers and the associated *FlowID* for packets matching all the field values. The type of field values are typically prefixes for IP address fields, an exact value or wildcard for the transport protocol number and flags, and ranges for port numbers.

Due to the complexity of the search, packet classification is often a performance bottleneck in network infrastructure; therefore, it has received much attention in the research community.The existing solutions explore various design tradeoffs to provide high search rates, power and space efficiency, fast incremental updates, and the ability to scale to large numbers of filters. There remains a need for techniques that achieve a favorable balance among these tradeoffs and scale to support classification on additional fields beyond the standard 5-tuple. We introduce *Distributed Crossproducting of Field Labels* (*DCFL*), a novel combination of new and existing packet classification techniques that can yield lookup performance comparable to Ternary Content Addressable Memory (TCAM) while scaling to additional filter fields and remaining memory, power, and update efficient. Like several existing approaches, we leverage observations of the structure of real filter sets in order to improve lookup performance and storage efficiency. Two key observations motivate our approach: the number of unique field values for a given field in the filter set is small relative to the number of filters in the filter set, and the number of unique field values matched by any packet is very small relative to the number of filters in the filter set. *DCFL* is also designed to take advantage of the millions of logic gates and hundreds of multi-port embedded memory blocks in the current generation of ASICs and FPGAs.

Using a high degree of parallelism, *DCFL* decomposes the multi-field packet classification problem and employs parallel search engines optimized for each filter field. Given that search techniques for single packet fields (longest prefix matching and range matching) are well-studied, the primary focus of this work is the development and analysis of a scalable technique for aggregating the results from each field search. Our approach is to perform a distributed set membership query using a network of aggregation nodes. Each query performs an intersection on the set of possible field combinations matched by the packet and the set of field combinations specified by filters in the filter set. We introduce several new concepts including field labeling, *Meta-Labeling* unique field combinations, *Field Splitting*, and optimized data structures such as *Bloom Filter Arrays* that minimize the number of memory accesses to perform set membership queries. By performing this aggregation in a distributed fashion, we avoid the exponential increase in the time or space incurred when performing this operation in a single step as in the seminal *Crossproducting* technique [2]. We also develop efficient encoding techniques for intermediate results allowing us to avoid the memory inefficiency suffered by similar techniques such as *Recursive Flow Classification* (RFC) [1].

Using a collection of 12 real filter sets and synthetic filter sets generated with the *ClassBench* tools [3], we provide an evaluation of *DCFL* performance and resource requirements for filter sets of various sizes and compositions. We show that an optimized implementation of *DCFL* can provide over 100 million searches per second and storage for over 200 thousand filters in a current generation FPGA or ASIC without the need for external memory devices.Furthermore, we show that *DCFL* retains its lookup performance and memory efficiency when the number of filters and number of fields in the filters increases. Scalability to classify on additional fields is a distinct advantage *DCFL* exhibits over existing decision tree algorithms and TCAM-based solutions. We continue to explore optimizations to improve the search rate and memory efficiency of *DCFL*. We also believe that *DCFL* has potential value for other searching tasks beyond traditional packet classification. Full technical report available at:
`http://www.arl.wustl.edu/~det3/`

SIGCOMM'04 Poster Session

1. REFERENCES

[1] P. Gupta and N. McKeown. Packet Classification on Multiple Fields. In *ACM Sigcomm*, August 1999.

[2] V. Srinivasan, S. Suri, G. Varghese, and M. Waldvogel. Fast and Scalable Layer Four Switching. In *ACM Sigcomm*, June 1998.

[3] D. E. Taylor and J. S. Turner. ClassBench: A Packet Classification Benchmark. Technical Report WUCSE-2004-28, Department of Computer Science & Engineering, Washington University in Saint Louis, May 2004.