

Hardware Implementation of Layer 3 and 4 Policy-based Packet Forwarding

Mark J. Boyd

Department of Computer Engineering
University of California, Santa Cruz
Santa Cruz, California 95064 USA

Abstract

The ability to rapidly differentiate service to users requires real-time operation of filters which match a large number of fields and wildcards. In the worst case, a lookup would need to continuously provide a result for gigabit speed forwarding and packets as small as a few bytes. This means the lookup must be complete within less than 50ns, using 20,000 filters, and matching fields of up to 128 bits.

We implement a combination of various schemes using FPGA hardware. We necessarily implement a scheme which provides slow insertions and deletions of filters in order to provide faster lookups. Although the implementation does not meet the full requirements, we implement a scalable design which can be implemented fully on a PCB with additional memory resources.

1 Introduction

We implement various filtering schemes to achieve a hardware implementation of layer-4 forwarding. Our approach is the design of a scalable implementation which sacrifices insertion and deletion speed for sustainable rapid lookups.

2 Background

Two techniques proposed for fast lookups are multi-level data structures and grid-of-tries. Multi-level data structures provide very fast lookups, moderate redundant storage, but fairly slow insertions and deletions.

A grid-of-tries with switch pointers reduces the amount of storage required while preserving a quick lookup time, while again having slow insertions and deletions.

3 Top-Level Design

The design necessarily must implement several approaches to accommodate lookups quickly. Multi-level data structures provide quick performance for the IP address portions, while other fields (which may have

ranges, for example) require much more creative approaches to provide quick lookup.

We deal with each field distinctly, as some lend themselves well to a high degree of parallelism to reduce complexity and improve speed (the range problem), and others are best handled by pointers (IP lookups and wildcards). For particularly difficult lookups, parallel memory accesses and redundant memory use is implemented as a last resort to improve speed.

4 Implementation

The design will be implemented using Xilinx Foundation Series 1.4 software, which allows full design and simulation of a downloadable FPGA design of several thousand gates. The implementation will not meet all the requirements of a full scale router, but will demonstrate a scalable design which would provide fast lookups when married with an appropriate PCB and memory.

5 Bibliography

Doeringer, W., Karjoth, G. and Nassehi, M. Routing on longest-matching prefixes, IEEE/ACM Transactions on Networking, vol.4, (no.1), IEEE; ACM, Feb. 1996. (Corrections available in IEEE/ACM Transactions on Networking, vol. 5, (no.4), IEEE; ACM, August 1997)

Waldvogel, M., Varghese, G., Turner, J. and Platner, B. Scalable high-speed IP routing lookups, Computer Communication Review, vol.27, (no.4), (ACM SIGCOMM 97 Conference. Applications, Technologies, Architectures, and Protocols for Computer Communication, Cannes, France, 14-18 Sept. 1997.)

Degermark, M., Brodnik, A., Carlsson, S. and Pink, S. Small forwarding tables for fast routing lookups, Computer Communication Review, vol.27, (no.4), (ACM SIGCOMM 97 Conference. Applications, Technologies, Architectures, and Protocols for Computer

Communication, Cannes, France, 14-18 Sept. 1997.)
ACM, Oct. 1997. p.3-14.

Srinivasan, V. and Varghese, G. Faster IP lookups using controlled prefix expansion, Performance Evaluation Review, vol.26, (no.1), (SIGMETRICS '98/PERFORMANCE'98. Joint International Conference on Measurement and Modeling of Computer Systems, Madison, WI, USA, 22-26 June 1998.) ACM, June 1998

Gupta, P., Lin, S. and McKeown, N. Routing lookups in hardware at memory access speeds, Proceedings of the IEEE INFOCOM '98.

Srinivasan, V., Varghese, G., Suri, S. and Waldvogel, M. Fast and scalable layer four switching, Computer Communication Review, vol.28, (no.4), (ACM SIGCOMM'98 Conference. Applications, Technologies, Architectures, and Protocols for Computer Communication, Vancouver, BC, Canada, 2-4 Sept. 1998.) ACM, Oct. 1998. p.191-202.

Lakshman, T.V. and Stiliadis, D. High-speed policy-based packet forwarding using efficient multi-dimensional range matching, Computer Communication Review, vol.28, (no.4), (ACM SIGCOMM'98 Conference. Applications, Technologies, Architectures, and Protocols for Computer Communication, Vancouver, BC, Canada, 2-4 Sept. 1998.) ACM, Oct. 1998. p.203-14.