

# Small Forwarding Tables for Fast Routing Lookups

Paper review by Ehsun Darody  
810181050

This paper presents a forwarding data structure designed for quick routing lookups. Forwarding tables are small enough to fit in the cache of a conventional general-purpose processor. This approach allows fast IP routing lookups in software. Traditional implementation of routing tables use a version of Patricia trees, a data structure invented almost thirty years ago, with modifications for longest prefix match.

The forwarding engines or processing elements in a router use a local version of routing table, a forwarding table, downloaded from the network processor to make their routing decisions. The primary goal of designing a data structure used for these forwarding engines is to minimize lookup time. To reach that goal we simultaneously minimize two parameters: the number of memory accesses required during lookup and size of the data structure.

The forwarding table is essentially a binary tree with three levels. Searching one level requires one to four memory accesses in the worst case. The most likely number of accesses is eight or less. The forwarding table is a representation of the binary tree spanned by all routing entries. This is called prefix tree, which is a complete tree. Three levels of data structure, covers prefix tree down to depth 16, 17 to 24 and 25 to 32, respectively. A complete binary tree of height  $h$  will be represented by using only one bit per possible leaf at depth  $h$  (bit vector), plus one base index per 64 possible leafs, plus information stored in the leaves. The bit vector is divided into bit-masks of lengths 16.

For the first level, the data structure consists of an array of *code words* and an array of *base indices*. The first 12 bits of the IP address are an index to the proper code word, and the first 10 bits are an index to the array of base indices. Other levels use a similar data structure.

Forwarding tables constructed from this scheme for multiple real routing tables in the Internet varies in size from 28 KB to 160 KB. Lookup performances for two general purpose processors are 2.0 to 2.2 million routing lookups per second that show software routing is still possible for high speed internetworking.