

Vinay Bharadwaj

The lecture covered getaround solutions to IP lookup. Since IP lookup is expensive, programmers usually try to find workarounds to avoid performing IP lookup as much as possible.

Some of the workarounds proposed to avoid/reduce IP lookup costs are as follows:

1) VPI/VCI

One of the workarounds is using a Virtual Path Identifier/Virtual Channel Identifier as in ATM networks. VPI/VCI are usually 24 bit ID fields. In this method, IP packets are fragmented into cells and are viewed as IP over connection oriented ATM network. Using this method, one can avoid IP lookup.

2) Threaded Index

In this method, each router's routing table has a threaded index to the next hop. Each record of IP address has two parts - the Address of the next router and the Index in the next router's table. The Address of the next router determines the output port and the index in the next router's table indicates exactly where the next route is stored. Hence, it is not required to search the entire routing table to find the next hop's address. One can directly jump to the specified index. This reduces the lookup complexity from $O(n)$ to a constant $O(1)$.

This method is similar to MPLS (Multi protocol Label switching). Each entry in the routing table is basically a label consisting of the next hop address and index.

3) Flow switching

In flow switching, packet routing is viewed as an aggregated flow. Instead viewing the routing process packet-wise, we classify multiple packets into flows depending on some factors and route them. For example, packets belonging to the same destination may constitute a single flow. This method does not completely avoid IP lookup but reduces the cost of IP lookup significantly. As an example, instead of performing IP lookup for each packet, we look at the flow the packet belongs to and route them. With flow switching, the router performs the following functions:

- Detect flows
- Fast path switching

4) Caching

Caching is a non-algorithmic method to reduce IP lookup cost. Already performed IP lookup is cached and the future packets are just matched with the cache entry. If there is a cache hit, the packet can be directly routed to the destination. If there is a cache miss, a new IP lookup should be performed and cached. Therefore, this method is a tradeoff between cache space and IP lookup complexity. Some factors to be considered are the cache size and cache replacement policies. If there are a lot of IP addresses, caching is expensive.

5) TCAM

TCAM consists of 1 rule for each entry and requires a lot of space. It also consumes a lot of power, gets hot and requires significant cooling solutions. Basically, the TCAM consists of a table with two columns - Content and Block. Content consists of the IP prefix - 0, 1 or *.

6) Packet Classification

Packet classification is a method where packets are classified according to some criteria. They are majorly used in the following:

- Firewalls: To decide which packets to allow and which ones to block for example.
- Policing: Policing here refers to the policing done for access control. For example, it is used in BGP to determine whether the packet conforms to some export/import policy. It is more of a business rule/policy.
- VPN
- Priority of rules for classification. No longest prefix matching.