

# Binary Search Schemes for Fast IP Lookups

Pronita Mehrotra, Paul D. Franzon

**Abstract**—IP route look up is the most time consuming operation of a router. Route lookup is becoming a very challenging problem due to the increasing size of routing tables. To reduce the number of entries, routing tables store variable length address prefixes. To determine the outgoing port for a given address, the longest matching prefix amongst all the prefixes, needs to be determined. This makes the task of searching in a large database even more difficult. Our paper describes binary search schemes that allow fast address lookups. Binary search can be performed on the number of entries or on the number of mutually disjoint prefixes, which for most practical routing tables is a little less than the number of entries. Lookups can therefore, be performed in  $O(N)$  time, where  $N$  is number of entries and the amount of memory required to store the binary database is also  $O(N)$ . These schemes scale very well with both large databases and for longer addresses (as in IPv6).

**Keywords**—Address Lookup, Binary Search, Longest Prefix Match

## I. INTRODUCTION

Traffic on the internet doubles every few months [1]. Additionally, the number of hosts on the internet have been increasing steadily, forcing a transition from the 32 bit address of IPv4 to the 128 bit addressing scheme of IPv6[2]. With bandwidth hungry applications like video conferencing becoming increasingly popular, the demand on the performance of routers has become very high. The most time consuming part in the design of a router is the route lookup. This paper deals with route lookup schemes that allow faster lookups when compared to other comparable schemes.

When an IP router receives a packet on one of its input port, it has to decide the outgoing port packet depending on the destination address of the packet. To make this decision, it has to look into a large database of destination networks and hosts. However, the problem with a lookup database is that if it were to store information for all possible destination addresses, then the database would become huge. For instance for a 32-bit IPv4 addresses, 4GB of memory would be required (assuming, output ports are stored in 8 bit fields). For 128-bit addressing of IPv6, an

The authors are with the Department of Electrical and Computer Engineering, North Carolina State University, Raleigh, NC 27695 email: {pmehrot,paulf}@eos.ncsu.edu, Ph:(919) 515 7351. This work was supported by ARDA under contract MDA904-00-C-2133 and NSF under contract EIA-9703090.

Prefix	Next Hop
10*	3
1011*	9
011*	8
010110*	5
001*	4
101101*	2
011010*	6
011100*	1
10111*	8
00101*	7

TABLE I

A SAMPLE ROUTING TABLE WITH PREFIXES AND NEXT HOPS

impossibly large amount of memory would be required to store all the information. To circumvent this problem, the classless interdomain routing (CIDR) scheme was introduced which allowed for the aggregation of networks to reduce the number of routing table entries. Current routing tables only store an address prefix which represents a group of addresses that can be reached from the output port. For instance, a routing table can have an entry of 128.\* with an associated output port of 2, which would mean that any destination address that begins with 128 has be directed to output port 2. There could also be another entry of 128.14.\* with an output port of 5. The rules now have to be interpreted as follows: if there is a destination address that begins with 128 it should be sent to output port 2, unless it begins with 128.14, in which case it should go to output port 5. The problem then becomes of finding the *longest matching prefix* to determine the output port or the next hop address. As an example, consider a sample routing table with prefixes as shown in Table I. For an incoming packet with a destination address beginning with 1011010\*, entries 10\*, 1011\*, 101101\* all give a match. The longest of these is 101101\* and so the correct next hop address in this case should be 2.

The rest of the paper is organized as follows. Section II discusses some of the previous work done in this field. Section III and Section IV then describe the details of our algorithms. We tested the performance of our schemes on

practical routing tables and the results of those are presented in Section V. Section VI discusses some of the key features and issues associated with the schemes and we finally conclude with Section VII.

## II. PREVIOUS WORK

Most of the approaches used to solve the longest matching prefix problem, fall under one of the two categories [3]. The first is the "thumb indexing" approach where each bit of the address is checked to see which part of the data space the search has to be directed to. The performance of these schemes depends on the size of the address space. These schemes perform reasonably well for IPv4 where the address size is 32 bits. However, with IPv6, these schemes are expected to perform poorly. The other methods use binary searches on either the number of routing entries or the number of possible prefix lengths. These schemes would scale better with the higher address size of IPv6. Our paper deals with binary search methods where the search is performed on the routing table entries.

Gupta et al [4] presented a hardware implementation based on an indirect lookup scheme where the number of memory lookups required to determine the output port is quite small (1-2). However, their scheme exploited a feature of the IPv4 routing tables that most of the prefixes are smaller than 24 bits. This may not necessarily be true for IPv6 tables. Also, the amount of memory required to store the forwarding tables was around 33MB. A much larger memory would be required for IPv6.

Binary tries use lesser memory to store the forwarding data bases, however the number of memory accesses required to evaluate the next hop is much higher. A number of optimizations have been suggested to improve the performance of binary tries. The NetBSD implementation uses a Patricia Trie [5], where a count of bits skipped is maintained for one-way branches. This reduces the average depth of the trie to some extent. The average length of the search in the Patricia implementation is  $1.44 \log(N)$  where  $N$  is the number of entries in the routing table. For large databases ( $>30,000$ ), 22-23 memory accesses are still required. Level Compression (LC) can be used to further reduce the average depth of the trie [6].

Some approaches have used compression techniques to make the forwarding databases small enough to fit on an on-chip SRAM or a cache. This avoids some expensive DRAM accesses. The scheme by Degermark et al [7] makes a forwarding table which uses only 5-6 bytes per routing table entry. The hardware implementation by Huang et al [8] compacts a large forwarding table of 40,000 entries into 450-470KB. Still other schemes [9], [10] have used caching to improve the performance of route lookup.

Caching relies on the temporal locality of data and this may not be very useful for core routers which exhibit very little temporal locality.

Other approaches are variations of the binary search. Binary search by itself can work only with numbers. Prefixes in the routing tables represent range of numbers and not discrete numbers and a naive implementation of the binary search on prefixes would fail. A few modifications to the binary search have been proposed. The first of these is the scheme by Lampson et al [11], where each prefix is expanded into two entries. The set of these entries is then processed to compute pointers to help in the search process. Yazdani et al [12] defined a sorting scheme to sort prefixes of different lengths and then applied the binary search to the sorted list. However, their scheme leads to an unbalanced tree and the number of memory accesses for a search becomes variable. The above schemes performed a binary search on the number of entries in the routing table. Waldvogel et al [13] suggested a hash based scheme where a binary search is performed on the number of possible prefix lengths. Their scheme scales well with the size of the routing table and at most 5 hash lookups (for IPv4) are required to determine the next hop address. However, as pointed out in [11], this scheme would not scale well with longer addresses as in IPv6.

Our work focuses on binary searches performed on the number of routing table entries. In that sense, it is closest to the work in [11] and their scheme has been used to compare our results with in the rest of the paper.

## III. DESCRIPTION OF THE ALGORITHM

To understand the algorithm, we first look at a variable-degree tree constructed from the prefixes in Table I as shown in Figure 1. To place prefixes in their relative positions in the tree, two conditions were used. For two prefixes,  $A = a_1a_2\dots a_m$  and  $B = b_1b_2\dots b_n$ ,

1. If  $A \subset B$ , then  $A$  is a parent of  $B$  (where the parent could be any node along the path from the node to the root of the tree)
2. If  $A \not\subset B$  and  $A < B$ , then  $A$  lies on the left of  $B$ . To compare  $A$  and  $B$ , if the prefix lengths of  $A$  and  $B$  are equal, i.e.  $n = m$ , then the prefixes can be compared by taking their numerical values. However, if  $n \neq m$ , then the longer prefix is chopped to the length of the shorter prefix and the numerical values compared.

By applying these conditions to all the prefixes, the tree in Figure 1 can be constructed.

The problem with performing a binary search on variable length prefixes can now be seen. By simply sorting the prefixes in some fashion and performing a binary search, it would not be possible to determine the longest




Fig. 1. Binary Tree Constructed from Table I

Prefix	Next Hop
00101*	7
001*	4
010110*	5
011010*	6
011100*	1
011*	8
101101*	2
10111*	8
1011*	9
10*	3

TABLE II

PREFIXES FROM TABLE I AFTER SORTING

matching prefixes. For instance consider a sorting defined as follows:

For the two prefixes  $A = a_1a_2\dots a_m$  and  $B = b_1b_2\dots b_n$ ,

1. If the prefixes are of equal lengths i.e. if  $n = m$ , then the numerical value of the prefixes determines which is the larger of the two.
2. If the prefixes are of unequal lengths i.e. if  $n \neq m$ , then the longer prefix is chopped to the length of the shorter prefix and the numerical values are compared. If after chopping the prefixes are equal, then the shorter prefix is considered to be the larger of the two.

Using the above definition, the prefixes in Table I can be sorted in ascending order as shown in Table II. This is equivalent to having performed a post-order depth-first-search on the binary tree in Figure 1.

Performing a simple binary search for a given address on the sorted prefixes, would not necessarily lead to the longest matching prefix. For instance, prefixes 10110\* and 100\* both lie between the entries 011\* and 101101\* but both of them have different longest matching prefixes. In general, the prefix could be any of the parent nodes or the root (default) node. Therefore, to determine the correct next hop, additional (and a variable number of) steps

would need to be performed. This problem arose because the node 101101\* did not carry any additional information about its parent nodes. We can avoid this problem by storing an additional field at all nodes that gives information about all the parent nodes. This additional field, which we call the Path Information Field, is a 32 bit entry (for IPv4) where a 1 in any bit position in the field means that there is a parent node with a prefix till that bit position. For example, for the leaf node of 101101\* the Path Information field would look like 0...101010, i.e. the second bit (corresponding to 10\*), the fourth bit (corresponding to 1011\*) and the sixth bit (corresponding to the leaf node itself) are set to 1. In addition, the node would also contain a pointer to a list of next hop addresses for the corresponding bits in the path information field. In this case, the list of next hop addresses would be 2,9,3. Now by looking at the path information field the longest matching prefix can be determined and the correct next hop address obtained from the list. The data structure used at the nodes is shown in Figure 2. In our implementation, we store the following information at each of the leaf nodes: the prefix, prefix mask, the next hop address corresponding to the leaf node, the number of internal nodes in the path and a pointer to the list of next hop addresses corresponding to the internal nodes. Table III shows the relevant information stored with each of the entries in Table II. Using the same example, if the address 10110\* were to be searched for in the list, the search would point to between the upper and lower entries of 011\* and 101101\*, respectively. A match between the address 10110\* and the lower entry 101101\* results in a match of upto 5 bits. By looking at the path information field, bit 5 is not set to 1. The next lower bit that is set, is bit 4. Therefore, the longest matching prefix for the given address is 1011\* and the corresponding next hop address in the next hop list is 9. In this case, the address needs to be compared with only the lower entry resulting from the search failure. To see why this is true, we look at different cases a search can end up in:

1. Between two leaf nodes with a common parent node. In this case either of the nodes can be used to determine the next hop address.
2. Between a node and its parent node. In this case, the parent node is the longest matching prefix and the parent node is the lower (larger) entry.
3. Between a node connected to the root node and a leaf node of the next branch (as in 011\* and 101101\*). In this case, nothing is gained by comparing with the smaller entry, since it can only lead to the root node (default next hop).

In all possible cases, it suffices to compare the address against only the lower entry to look for partial matches.




Fig. 2. Data Structure used at a Node

Prefix	Next Hop	Path Information	Next Hop List
00101*	7	0...010100	4
001*	4	0...000100	-
010110*	5	0...100000	-
011010*	6	0...100100	8
011100*	1	0...100100	8
011*	8	0...000100	-
101101*	2	0...101010	9,3
10111*	8	0...011010	9,3
1011*	9	0...001010	3
10*	3	0...000010	-

TABLE III

SEARCH SPACE OF THE PREFIXES IN TABLE II

### A. Building the Data Structure

Building the data structure is fairly simple and the steps involved are listed below:

Step 1 Each entry is read from the routing table and stored in an array.

Step 2 The entries are then sorted from the smallest to the highest. To compare prefixes, the rules listed previously are used. If prefixes are of equal lengths then the numerical values of the prefixes are used to compare them. If prefixes are of unequal lengths, then the longer prefix is chopped to the length of the smaller prefix and then compared. If after chopping, the prefixes are equal then the shorter prefix is considered to be the larger of the two. Doing this ensures that in the following step, the child nodes get processed before the parent nodes.

Step 3 Entries from the sorted list are then processed and added in an array, one at a time. Each entry is also tested with the last array entry to see if it is a subset of the array entry.

Step 3a If it is a subset, then the next hop information corresponding to the array entry is added to the next hop list of the array entry. The path information field of the array entry is updated and so is the field containing the number

of nodes. This step is then repeated for previous entries till the test for subset fails. To see why this is necessary, consider the last prefix in Table II. When prefix 10\* is compared with the last array entry (1011\*), it is added in the next hop list of 1011\*. However, 10\* also lies in the path of 10111\* and 101101\* and the corresponding entries need to be updated as well.

### B. Searching the Data Structure

To search for the longest matching prefix, a binary search is performed on the entries. The search algorithm is summarized below:

Step 1 Binary search of the address is performed on the array entries which leads to a value between two consecutive array entries.

Step 2 The address is then matched with the lower entry, and checked against the path information field. The longest matching prefix and the corresponding next hop address from the next hop list is picked. If no match is found, the default next hop address is returned.

### C. Updating the Data Structure

Inserting or Deleting entries from the data space is equivalent to adding or deleting an entry from an array. To add an entry, a binary search is performed as outlined in the previous section, to find the location of the insertion. The entry is then added into the array, which is an  $O(N)$  process. The entry is also checked against entries above it to see if it is a subset or not, and the corresponding path information field and next hop list is updated. Deleting an entry follows a similar procedure. Updating the data structure, therefore, does not require the entire data structure to be built from scratch.

## IV. USING DISJOINT PREFIXES FOR BINARY SEARCH

The search space used in the previous scheme can be reduced further by using only mutually disjoint prefixes. Two prefixes are considered disjoint, if none of them is a prefix of the other. It is easy to see that these correspond to the leaf nodes of the tree shown in Figure 1. All internal nodes can be removed from the search space since the information corresponding to them is already contained in the path information field and the next hop list of the leaf nodes. The search space can then be shortened as shown in Table IV

From Table III and Table IV it might appear that a considerable amount of memory might be wasted in storing internal nodes a multiple number of times. For instance, next hop addresses corresponding to nodes 1011\* and 10\* are stored in the list of next hops for both 101101\* and 10111\*. This is not necessarily true. An examination of

Prefix	Next Hop	Path Information	Next Hop List
00101*	7	0...010100	4
010110*	5	0...100000	-
011010*	6	0...100100	8
011100*	1	0...100100	8
101101*	2	0...101010	9,3
10111*	8	0...011010	9,3

TABLE IV  
SEARCH SPACE OF THE PREFIXES IN TABLE II

practical routing tables from [14] shows that most of the nodes in fact do not have any internal nodes and next hop lists to store. Figure 3 shows the number of internal nodes for all the leaf nodes for various routing tables. From the figure, it can be seen that more than 93% of the leaf nodes do not have any internal nodes in the path to the root node. Therefore, the overhead in memory to store internal nodes multiple number of times is actually quite small.

The build and search algorithms need to be modified slightly to accommodate the changes and the modified algorithms are described in the following subsections.

#### A. Building the Data Structure

The main difference in building the data structure is that not all entries get added to the search space. If an entry is found to be a subset of the previous entry, then it is not added to the search space. The modified algorithm is described below:

Step 1 Each entry is read from the routing table and stored in an array.

Step 2 The entries are then sorted in an ascending order.

Step 3 Entries from the sorted list are then processed one at a time. Each entry is tested with the last of the leaf node entry added to see if it is a subset of the leaf node.

Step 3a If it is a subset, then the corresponding next hop information is added to the next hop list of the leaf node. The path information field of the leaf node is updated and so is the field containing the number of nodes. This step is then repeated for previous leaf node entries till the test for subset fails.

Step 3b If it is not a subset, then a new leaf node is added.

#### B. Searching the Data Structure

A binary search is performed on the leaf nodes, but the address now needs to be checked against both the upper and lower entries returned by the search. The steps involved in the search are listed below:

Step 1 Binary search of the address is performed on the leaf nodes which leads to a value between two consecutive leaf node entries.

Step 2 The address is then matched with the two leaf nodes to see which of the leaf nodes is a better match.

Step 3 The better match from Step 2 is picked and checked against the path information field of the corresponding leaf node. The longest matching prefix and the corresponding next hop address from the next hop list is picked. If no match is found, the default next hop address is returned.

#### C. Updating the Data Structure

Inserting or Deleting entries from the data structure does not require the search space to be built from scratch, as before. To add an entry, a binary search is performed to find the location of the insertion. If the entry turns out to be an internal node, only the next hop lists and the parent information fields of the corresponding leaf nodes get updated. If the entry to be added turns out to be a leaf node, then the leaf node array is updated which is an  $O(N)$  process. Deleting an entry follows a similar procedure.

## V. RESULTS AND COMPARISONS

We ran our algorithms on a Sun Ultra 5 with a 333 MHz processor and 512MB of RAM. The programs were written in C and compiled with gcc with the compiler optimization level 3. We also compared our results against the binary search implementation of Lampson et al [11]. The binary search part of all algorithms were identical. We used the search time, build time and memory consumption to evaluate the performance of the schemes. Practical routing tables from [14] were used in the experiments.

*Average Search Time* Random IP addresses were generated and a lookup was performed using both schemes. Table V lists the average lookup times for different routing tables. From Table V it can be seen that our scheme which uses all nodes in the search space results in over 10% improvement in lookup speeds as opposed to the scheme from [11]. With internal nodes eliminated from the search space, an improvement of 15-20% can be obtained. The difference starts getting larger as the size of the routing table increases as seen from Figure 4.

*Build Time of Data Structure* The time required to build the searchable structure from the routing tables is shown in Table VI. The time shown does not include the initial time taken to read entries from a file and store them in an array. This means that for all algorithms time starts from sorting the entries and ends when the searchable structure is built. As seen from the table the time taken to build the searchable space using Lampson et al's scheme [11] is more than two times the time taken to build ours. A pro-




Fig. 3. Profile of routing tables from

Routing Table	No of Entries	Binary Search (All nodes)	Binary Search (Only Leaves)	Binary Scheme from [11]
MaeEast	23113	662ns	610ns	761ns
MaeWast	35752	742ns	652ns	845ns
PacBell	27491	703ns	656ns	761ns
Paix	17641	640ns	634ns	739ns
AADS	31958	700ns	640ns	777ns

TABLE V

AVERAGE SEARCH TIMES FOR DIFFERENT ROUTING TABLES




Fig. 4. Average Lookup Time for Different Schemes

file of the times taken, using gprof, shows that most of the difference can be accounted for by the initial sorting of entries. Since the number of entries in their scheme is more than twice ours, sorting becomes a fairly expensive operation. Build time for the scheme that uses only leaves in its

searchable space is not significantly different from the one that uses all the nodes. This is due to the fact that most of the entries in these routing tables, end up as leaves and the overhead in adding internal nodes is very small.

*Memory Consumption* Table VII shows the memory required in storing the searchable structure for all schemes. Memory required for both binary schemes is close to half of that required in [11]. This is because each prefix in their algorithm gives rise to two entries in the binary search table. Memory requirement for the scheme using only leaves in its searchable space is not significantly different from the one using all the nodes, for the reason pointed before.

## VI. DISCUSSION

We have compared the performance of our scheme with the scheme by Lampson et al [11] because that is the only other binary search scheme that we are aware of, which searches on the number of entries in the routing table. In the search scheme by Lampson et al [11], each entry in the routing table is expanded to two entries in the binary

Routing Table	No of Entries	Binary Search (All Nodes)	Binary Search (Only Leaves)	Binary Scheme from [11]
MaeEast	23113	80ms	80ms	210ms
MaeWast	35752	130ms	120ms	330ms
PacBell	27491	90ms	90ms	260ms
Paix	17641	60ms	50ms	150ms
AADS	31958	120ms	100ms	300ms

TABLE VI  
TIME TAKEN TO BUILD SEARCHABLE STRUCTURE

Routing Table	No of Entries	Binary Search (All Nodes)	Binary Search (Only Leaves)	Binary Scheme from [11]
MaeEast	23113	0.62MB	0.58MB	1.06MB
MaeWast	35752	0.96 MB	0.9MB	1.64MB
PacBell	27491	0.74MB	0.7MB	1.26MB
Paix	17641	0.48MB	0.45MB	0.81MB
AADS	31958	0.86MB	0.81MB	1.46MB

TABLE VII  
MEMORY REQUIREMENT FOR DIFFERENT ROUTING TABLES

search space. The time taken to search for an entry is of the order of  $\log(2N)$  where  $N$  is the number of entries in the routing table. Once the binary search is performed, an additional memory lookup is required (approximately half the time) to obtain the next hop address. In comparison, the search space in our schemes is  $\leq N$ . Our schemes therefore, result in 1-2 fewer memory accesses in just the binary search. After the search is narrowed down, our schemes, for most of the time does not require any additional memory lookup to determine the next hop address. For these reasons, the average lookup time in our scheme is lower than the lookup time in Lampson et al's scheme.

The other main point of difference is the build time of the searchable structure. As pointed before, the build time of our schemes is also much smaller than the build time of Lampson et al's scheme (see Table VI). The main reason for this difference is the sorting step which is required by all schemes. Sorting becomes more and more expensive as the number of entries to be sorted increases. This accounts for majority of the time difference. The build time will be particularly crucial for larger routing tables, because updating entries requires the entire searchable space to be built from scratch in Lampson et al's scheme. For large routing tables ( $>100,000$  entries), this could be serious problem in their scheme.

One potential problem with our schemes is that the path

information field has to be equal to the size of the address. For IPv4 this means that the path information field has to be 32 bits. For IPv6, this would mean storing a 128 bit field, leading to higher memory consumption and longer times to process instructions using this 128 bit field. In comparison, the additional information that Lampson et al's scheme uses are pointers, the size of which would depend on the number of entries. For a routing table of 100,000 entries, the pointers need to be only 18 bits wide. Storing the high and low pointers in their scheme, would make the extra memory consumption to 36 bits as compared to 128 bits in our case. However, this could be improved somewhat as we show next. If very long or very short prefixes do not exist in the routing tables, then the number of bits used in storing the path information can be reduced. Figure 5 shows the various possible prefix lengths and the corresponding number of entries for the MaeEast routing table from [11]. This may not be clear from the figure, but, no prefixes exist for prefix lengths smaller than 8 bits and larger than 30 bits. The path information field, therefore, only needs to have 23 bits to store the relevant information. We do not have any data for how core routing tables would look like for IPv6, but we expect to be able to considerably reduce the number of bits used in storing the path information.

Since the main component of both schemes is the binary




Fig. 5. Prefix Length Distribution for the MaeEast Router

search on the entries in the searchable structure, the performance of both schemes can be enhanced by going from a binary search to a multiway search. This has already been shown in [11] and therefore is not the focus of this paper.

## VII. CONCLUSIONS

We have described methods to adapt binary search to work with variable length prefixes. By storing information corresponding to parent nodes, we can reduce the time taken to search for a next hop address. Further improvement in performance can be obtained by using only mutually disjoint prefixes in the searchable space. 10-20% improvement in average lookup time over other schemes can be obtained. In addition, reducing the search space to less than half, the amount of memory taken by the searchable space is also reduced to about half. The time required to build the searchable structure is similarly reduced considerably due to the fewer number of entries.

## REFERENCES

- [1] "Internet Host and Traffic Growth." (<http://www.cs.columbia.edu/hgs/internet/growth.html>).
- [2] "Internet Growth Summary." (<http://www.mit.edu/people/mkgray/net/internet-growth-summary.html>).
- [3] N. McKeown and B. Prabhakar, "High Performance Switches and Routers: Theory and Practice," in *Hot Interconnects Tutorial Slides* (<http://tiny-tera.stanford.edu/nickm/talks/index.html>), Aug. 1999.
- [4] P. Gupta, S. Lin, and N. McKeown, "Routing Lookups in Hardware at Memory Access Speeds," in *Proc. IEEE INFOCOM'98*, (San Francisco, CA), pp. 1382–1391, 1998.
- [5] K. Sklower, "A Tree-Based Routing Table for Berkeley Unix," in *Technical Report*, (University of California, Berkeley).
- [6] S. Nilsson and G. Karlsson, "IP-Address Lookup Using LC-Tries," *IEEE Journal on Selected Areas in Communications*, vol. 17, pp. 1083–1092, June 1999.
- [7] M. Degermark, A. Brodnik, S. Carlsson, and S. Pink, "Small Forwarding tables for fast routing lookups," in *Proc. ACM SIGCOMM*, vol. 27, pp. 3–14, Oct. 1997.
- [8] N.-F. Huang and S.-M. Zhao, "A Novel IP-Routing Lookup Scheme and Hardware Architecture for Multigigabit Switching Routers," *IEEE Journal on Selected Areas in Communications*, vol. 17, pp. 1093–1104, June 1999.
- [9] T. cker Chiueh and P. Pradhan, "High-Performance IP Routing Table Lookup Using CPU Caching," in *Proc. IEEE INFOCOM'99*, pp. 1421–1428, 1999.
- [10] T. cker Chiueh and P. Pradhan, "Cache Memory Design for Network Processors," in *Proceedings of Sxth International Symposium on High-Performance Computer Architecture, 2000*, vol. HPCA-6, pp. 409–418, 2000.
- [11] B. Lampson and G. V. V. Srinivasan, "IP Lookups using Multiway and Multicolumn Search," in *Proc. IEEE INFOCOM'98*, vol. 3, (San Francisco, CA), pp. 1248–1256, 1998.
- [12] N. Yazdani and P. S. Min, "Fast and Scalable schemes for the IP address Lookup Problem," in *Proc. IEEE Conference on High Performance Switching and Routing*, pp. 83–92, 2000.
- [13] M. Waldvogel, G. Varghese, J. Turner, and B. Plattner, "Scalable High Speed IP Routing Lookups," in *Proc. ACM SIGCOMM*, vol. 27, pp. 25–36, Oct. 1997.
- [14] "Michigan University and Merit Network. Internet Performance Management and Analysis (IPMA) Project." (<http://nic.merit.edu/ipma>).