

SHIP: A Scalable High-performance IPv6 Lookup Algorithm that Exploits Prefix Characteristics

Thibaut Stimpfling, Normand Belanger, J.M. Pierre Langlois *Member IEEE* and Yvon Savaria *Fellow IEEE*

Abstract—Due to the emergence of new network applications, current IP lookup engines must support high-bandwidth, low lookup latency and the ongoing growth of IPv6 networks. However, existing solutions are not designed to address jointly those three requirements. This paper introduces SHIP, an IPv6 lookup algorithm that exploits prefix characteristics to build a two-level data structure designed to meet future application requirements. Using both prefix length distribution and prefix density, SHIP first clusters prefixes into groups sharing similar characteristics, then it builds a hybrid trie-tree for each prefix group. The compact and scalable data structure built can be stored in on-chip low-latency memories, and allows the traversal process to be parallelized and pipelined at each level in order to support high packet bandwidth.

Evaluated on real and synthetic prefix tables holding up to 580 k IPv6 prefixes, SHIP has a logarithmic scaling factor in terms of the number of memory accesses, and a linear memory consumption scaling. Using the largest synthetic prefix table, simulations show that compared to other well-known approaches, SHIP uses at least 44% less memory per prefix, while reducing the memory latency by 61%.

Index Terms—Algorithm, Routing, IPv6 Lookup, Networking.

I. INTRODUCTION

GLOBAL IP traffic carried by networks is continuously growing, as around a zettabyte total traffic is expected for the whole of 2016, and it is envisioned to increase threefold between 2015 and 2019 [1]. To handle this increasing Internet traffic, network link working groups have ratified the 100-gigabit Ethernet standard (IEEE P802.3ba), and are studying the 400-gigabit Ethernet standard (IEEE P802.3bs). As a result, network nodes have to process packets at those line rates which puts pressure on IP address lookup engines used in the routing process. Indeed, less than 6 ns is available to determine the IP address lookup result for an IPv6 packet [2].

The IP lookup task consists of identifying the next hop information (NHI) to which a packet should be forwarded. The lookup process starts by extracting the destination IP field from the packet header, and then matching it against a list of entries stored in a lookup table, called the forwarding information base (FIB). Each entry in the lookup table represents a network defined by its prefix address. While a lookup key may match multiple entries in the FIB, only the longest prefix and its NHI are returned for result as IP lookup is based on the Longest Prefix Match (LPM) [3].

IP lookup algorithms and architectures that have been tailored for IPv4 technology are not performing well with

IPv6 [2], [4], due to the fourfold increase in the number of bits in IPv6 addresses over IPv4. Thus, dedicated IPv6 lookup methods are needed to support upcoming IPv6 traffic.

IP lookup engines must be optimized for high bandwidth, low latency, and scalability for two reasons. First, due to the convergence of wired and mobile networks, many future applications that require a high bandwidth and a low latency, such as virtual reality, remote object manipulation, eHealth, autonomous driving, and the Internet of Things, will be carried on both wired and mobile networks [5]. Second, the number of IPv6 networks is expected to grow, and so is the size of the IPv6 routing tables, as IPv6 technology is still being deployed in production networks [6], [7]. However, current solutions presented in the literature are not jointly addressing these three performance requirements.

In this paper, we introduce SHIP: a Scalable and High Performance IPv6 lookup algorithm designed to meet current and future performance requirements. SHIP is built around the analysis of prefix characteristics. Two main contributions are presented: 1) two-level prefix grouping, that clusters prefixes in groups sharing common properties, based on the prefix length distribution and the prefix density, 2) a hybrid trie-tree tailored to handle prefix distribution variations. SHIP builds a compact and scalable data structure that is suitable for on-chip low-latency memories, and allows the traversal process to be parallelized and pipelined at each level in order to support high packet bandwidth. SHIP stores 580 k prefixes and the associated NHI using less than 5.9 MB of memory, with a linear memory consumption scaling. SHIP achieves logarithmic latency scaling and requires in the worst case 10 memory accesses per lookup. For both metrics, SHIP outperforms known methods by over 44% for the memory footprint, and by over 61% for the memory latency.

The remainder of this paper is organized as follows. Section II introduces common approaches used for IP lookup and Section III gives an overview of SHIP. Then, two-level prefix grouping is presented in Section IV, while the proposed hybrid trie-tree is covered in Section V. Section VI introduces the method and metrics used for performance evaluation and Section VII presents the simulation results. Section VIII shows that SHIP fulfills the properties for hardware implementability, and compares SHIP performance with other methods. Lastly, we conclude the work by summarizing our main results in Section IX.

II. RELATED WORK

Many data structures have been proposed for the LPM operation applied to IP addresses. We can classify them in

four main types: hash tables, Bloom filters, tries and trees. Those data structures are encoding prefixes that are loosely structured. First, not only prefix length distribution is highly nonuniform, but it also varies with the prefix table used. Second, for any given prefix length, prefix density ranges from sparse to very dense. Thus, each of the four main data structures type comes with a different tradeoff between time and storage complexity.

Interest for IP lookup with hash table is twofold. First, a hash function aims at distributing uniformly a large number of entries over a number of bins, independently of the structure of the data stored. Second, a hash table provides $O(1)$ lookup time and $O(N)$ space complexity. However, a pure hash based LPM solution can require up to one hash table per IP prefix length. An alternative to reduce the number of hash tables is to use prefix expansion [8], but it increases memory consumption. Two main types of hash functions can be selected to build a hash table: perfect or non-perfect hash functions. A Hash table built with a perfect hash functions offers a fixed time complexity that is independent from the prefixes used as no collision is generated. Nevertheless, a perfect hash function cannot handle dynamic prefix tables, making it unattractive for a pure hash based LPM solution. On the other hand, a non-perfect hash function leads to collisions and cannot provide a fixed time complexity. Extra-matching sequences are required with collisions that drastically decrease performance [9], [10]. In addition, not only the number of collisions is determined after the creation of the hash table but it also depends on the prefix distribution characteristics. In order to reduce the number of collisions independently of the characteristics of the prefix table used, a method has been proposed that exploits multiple hash tables [8], [9]. This method divides the prefix table into groups of prefixes, and selects a hash function such that it minimizes the number of collisions within each prefix group [8], [9]. Still, the hash function selection for each prefix group requires to probe all the hash functions, making it unattractive for dynamic prefix tables. Finally, no scaling evaluation has been completed in recent publications [8], [11] making it unclear whether the proposed hash-based data structures can address forthcoming challenges.

Low-memory footprint hashing schemes known as Bloom filters have also been covered in the literature [10], [12]. Bloom filters are used to select a subgroup of prefixes that may match the input IP address. However, Bloom filters suffer from two drawbacks. First, by design, this data structure generates false positives independent of the configuration parameters used. Thus, a Bloom filter can improve the average lookup time, but it can also lead to poor performance in the worst case, as many sub-groups need to be matched. Second, the selection of a hash function that minimizes the number of false positives is highly dependent of the prefix distribution characteristics used. Hence, its complexity is similar to that of a hash function that minimizes the number of collisions in a regular hash table.

Tree solutions based on binary search trees (BST) or generalized B-trees have also been explored in [2], [4]. Such data structures are tailored to store loosely structured data such as prefixes, as their time complexity is independent from the

prefix distribution characteristics. Indeed, BST and 2-3 Trees have a time complexity of respectively $\log_2(N)$ and $\log_3(N)$, with N being the number of entries [2]. Nevertheless, such data structure provides a solution at the cost of a large memory consumption. Indeed, each node stores a full-size prefix, leading to memory waste. Hence, their memory footprint makes them unsuitable for the very large prefix tables that are anticipated in future networks.

At the other end of the tree spectrum, decision-trees (D-Trees) have been proposed in [13], [14] for the field of packet classification. D-Trees were found to offer a good tradeoff between memory footprint and the number of memory accesses. However, no work has been conducted yet on using this data structure for IPv6 lookup.

The trie data structure, also known as radix tree, has regained interest with tree bitmap [15]. Indeed, a k -bit trie requires k/W memory accesses, but has very poor memory efficiency when built with unevenly distributed prefixes. A tree bitmap improves the memory efficiency over a multi-bit trie, independently of the prefix distribution characteristics, by using a bitmap to encode each level of a multi-bit trie. However, tree bitmaps cannot be used with large strides, as the node size grows exponentially with the stride size, leading to multiple wide memory accesses to read a single node. An improved tree bitmap, the PC-trie, is proposed for the FlashTrie architecture [11]. A PC-Trie reduces the size of bitmap nodes using a multi-level leaf pushing method. This data structure is used jointly with a pre-processing hashing stage to reduce the total number of memory accesses. Nevertheless, the main shortcoming of the Flashtrie architecture lies in its pre-processing hashing module. First, similar to other hashing solutions, its performance highly depends on the distribution characteristics of the prefixes used. Second, the hashing module does not scale well with the number of prefixes used.

At the other end of the spectrum of algorithmic solutions, TCAMs have been proposed as a pure hardware solution, achieving $O(1)$ lookup time by matching the input key simultaneously against all prefixes, independently of their distribution characteristics. However, these solutions use a very large amount of hardware resources, leading to large power consumption and high cost, and making them unattractive for routers holding a large number of prefixes [11], [16].

Recently, information-theoretic and compressed data structures have been applied to IP lookup, yielding very compact data structures, that can handle a very large number of prefixes [17]. Even though this work is limited to IPv4 addresses, it is an important shift in terms of concepts. However, the hardware implementation of the architecture achieves 7 million lookups per second. In order to support a 100-Gbps bandwidth, this would require many lookup engines, leading to a memory consumption that is similar or higher than previous trie or tree algorithms [2], [8], [9], [11].

In summary, the existing data structures are not exploiting the full potential of the prefix distribution characteristics. In addition, none of the existing data structures were shown to optimize jointly the time complexity, the storage complexity, and the scalability.

III. SHIP OVERVIEW

SHIP consists of two procedures: the first one is used to build a two-level data structure, and the second one is used to traverse the two-level data structure. The procedure to build the data structure is called two-level prefix grouping, while the traversal procedure is called the lookup algorithm.

Two-level prefix grouping clusters prefixes upon their characteristics to build an efficient two-level data structure, presented in Fig 1. At the first level, SHIP leverages the low density of the IPv6 prefix MSBs to divide prefixes into M address block bins (ABBs). A pointer to each ABB is stored in an N -entry hash table. At the second level, SHIP uses the uneven prefix length distribution to sort prefixes held in each ABB into K Prefix Length Sorted (PLS) groups. For each non-empty $K \cdot M$ PLS groups, SHIP further exploits the prefix length distribution and the prefix density variation to encode prefixes into a hybrid trie-tree (HTT).

The lookup algorithm, which identifies the NHI associated to the longest prefix matched, is presented in Fig 1. First, the MSBs of the destination IP address are hashed to select an ABB pointer stored in an N entry hash table. The selected ABB pointer in this figure is held in the n -th entry of the hash table, represented with a dashed rectangle. This pointer identifies bin m , represented with a dashed rectangle. Second, the HTTs associated to each PLS group of the m -th bin, are traversed in parallel, using portions of the destination IP address. Each HTT can output a NHI, if a match occurs with its associated portion of the destination IP address. Thus, up to K HTT results can occur, and a priority resolution module is used to select the NHI associated to the longest prefix.

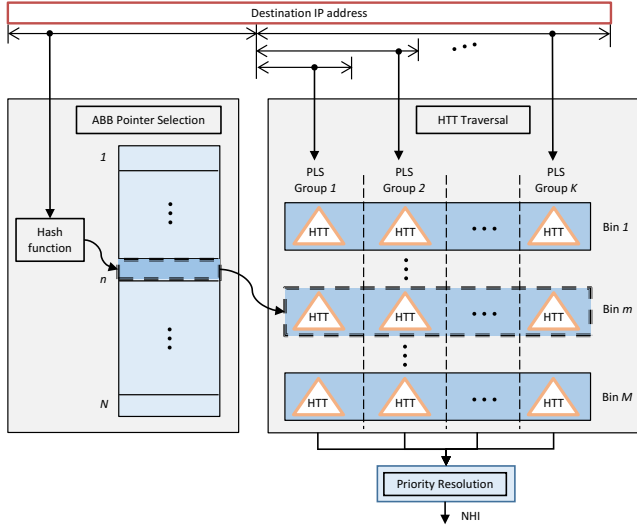


Fig. 1. SHIP two-level data structure organization and its lookup process with M address block bins and K prefix length sorting groups.

In the following section, we present the two-level prefix grouping procedure.

IV. TWO-LEVEL PREFIX GROUPING

This section introduces two-level prefix grouping, that clusters and sort prefixes into groups, and then builds the two-level

data structure. First, prefixes are binned and the first level of the two-level data structure is built with the address block binning method. Second, inside each bin, prefixes are sorted into groups, and then the HTTs are built with the prefix length sorting method.

A. Address block binning

The proposed address block binning method exploits both the structure of IP addresses and the low density of the IPv6 prefix MSBs to cluster the prefixes into bins, and then build the hash table used at the first level of SHIP data structure.

IPv6 addresses are structured into IP address blocks, managed by the Internet Assigned Numbers Authority (IANA) that assigns blocks of IPv6 addresses ranging in size from $/16$ to $/23$, that are then further divided into smaller address blocks. However, the prefix density on the first 23 bits is low, and the prefix distribution is sparse [18]. Therefore, the address block binning method bins prefixes based on their first 23 bits. Before prefixes are clustered, all prefixes with a prefix length that is less than $/23$ are converted into $/23$. The pseudo-code used for this binning method is presented in Algorithm 1. For each prefix held in the prefix table, this method checks whether a bin already exists for the first 23 bits. If none exists, a new bin is created, and the prefix is added to the new bin. Otherwise, the prefix is simply added to the existing bin. This method only keeps track of the M created bins. The prefixes held in each bin are further grouped by the prefix length sorting method that is presented next.

Address block binning utilizes a perfect hash function [19] to store a pointer to each valid bin. Let N be the size of the hash table used. The first 23 MSBs of the IP address represents the key of the perfect hash table. A perfect hashing function is chosen for three reasons. First, the number of valid keys on the first 23 bits is relatively small compared to the number of bits used, making hashing attractive. Second, a perfect hashing function is favoured when the data hashed is static, which is the case here, for the first 23 bits only, because it represents blocks of addresses allocated to regional internet registries that are unlikely to be updated on a short time scale. Finally, no resolution module is required because no collisions are generated.

The lookup procedure is presented in Algorithm 2. It uses the 23 MSBs of the destination IP address as a key for the hash table and returns a pointer to an address block bin. If no valid pointer exists in the hash table for the key, then a null pointer is returned.

The perfect hash table created with the address block binning method is an efficient data structure to perform a lookup on the first 23 MSBs of the IP address. However, within the ABBs the prefix length distribution can be highly uneven, which degrades the performance of the hybrid trie-trees at the second level. Therefore, the prefix length sorting method, described next, is proposed to address that problem.

B. Prefix length sorting

Prefix length sorting (PLS) aims at reducing the impact of the uneven prefix length distribution on the number of overlaps

Algorithm 1 Building the Address Block binning data structure

Input: Prefix table

Output: Address Block binning data structure

```

1: for each prefix held in the prefix table do
2:   Extract its 23 MSBs
3:   if no bin already exists for the extracted bits then
4:     Create a bin that is associated to the value of the
       extracted bits
5:     Add the current prefix to the bin
6:   else
7:     Select the bin that is associated to the value of the
       extracted bits
8:     Add the current prefix to the selected bin
9:   end if
10: end for
11: Build a hash table that stores a pointer to each address
    block bin, with the 23 MSBs of each created bin as a key.
    Empty entries are holding invalid pointers.
12: return the hash table
  
```

Algorithm 2 Lookup in the Address Block binning data structure

Input: Address Block binning hash table, destination IP Address

Output: Pointer to an address block bin

```

1: Extract the 23 MSBs of the destination IP address
2: Hash the extracted 23 MSBs
3: if the hash table entry pointed by the hashed key holds a
   valid pointer then
4:   return pointer to the address block bin
5: else
6:   return null pointer
7: end if
  
```

between prefixes held in each address block bin. By reducing the number of prefix overlaps, the performance of the HTTs is improved, as it will be shown later. The PLS method sorts the prefixes held in each address block bin by their length, into K PLS groups that cover disjoint prefix length ranges. Each range consists of contiguous prefix lengths that are associated to a large number of prefixes with respect to the prefix table size. For each PLS group, a hybrid trie-tree is built.

The number of PLS groups, K , is chosen to maximize the HTT's performance. As will be shown experimentally in section VII, beyond a threshold value, increasing the value of K does not further improve performance. The prefix length range selection is based on the prefix length distribution and it is guided by two principles. First, to minimize prefix overlap, when a prefix length covers a large percentage of the total number of prefixes, this prefix length must be used as an upper bound of the considered group. Second, prefix lengths included in a group are selected such that group sizes are as balanced as possible.

To illustrate those two principles, an analysis of prefix length distribution using a real prefix table is presented in Fig. 2. The prefix table extracted from [7] holds approximately

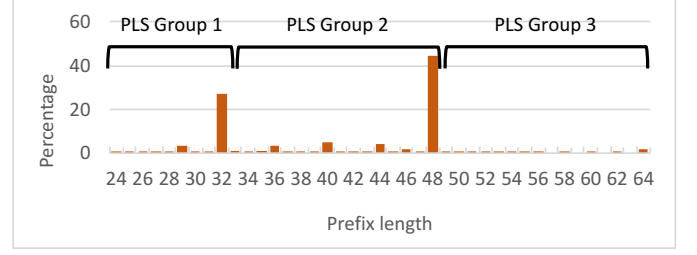


Fig. 2. The uneven prefix length distribution of real a prefix table used by the PLS method to create 3 PLS groups.

25 k prefixes. The first 23 prefix lengths are omitted in Fig. 2, as the address block binning method already bins prefixes on their 23 MSBs. It can be observed in Fig. 2 that the prefix lengths with the largest cardinality are /32 and /48 for this example. Applying the two principles of prefix length sorting to this example, the first group covers prefix lengths from /24 to /32, and the second group covers the second peak, from /33 to /48. Finally, all remaining prefix lengths, from /49 to /64 are left in the third prefix length sorting group.

For each of the K PLS group created, an HTT is built. Thus, the lookup step associated to the prefix length sorting method consists of traversing the K HTTs held in the selected address block bin.

To summarize, the created PLS groups cover disjoint prefix length ranges by construction. Therefore, the PLS method directly reduces prefix overlaps in each address block bin that increases the performance of HTT. However, within each PLS group, the prefix density variation remains uneven. Hence, a hybrid-trie tree is proposed that exploits the local prefix characteristics to build an efficient data structure.

V. HYBRID TRIE-TREE DATA STRUCTURE

The hybrid trie-tree proposed in this work is designed to leverage the prefix density variation. This hybrid data structure uses a density-adaptive trie, and a reduced D-Tree leaf when the number of prefixes covered by a density-adaptive trie node is below a fixed threshold value. A description of the two data structures is first presented, then the procedure to build the hybrid trie-tree is formulated, and finally the lookup procedure is introduced.

A. Density-Adaptive Trie

The proposed density-adaptive trie is a data structure that is built upon the prefix density distribution. A density-adaptive trie combines a trie data structure with the Selective Node Merge (SNM) method.

While a trie or multi-bit trie creates equi-sized regions whose size is independent of the prefix density distribution, the proposed SNM method adapts the size of the equi-sized regions to the prefix density. Low-density equi-sized regions created with a trie data structure are merged into variable region sizes by the SNM method. Two equi-sized regions are merged if the total number of prefixes after merging is equal to the largest number of prefixes held by the two equi-sized regions, or if it is less than a fixed threshold value. The

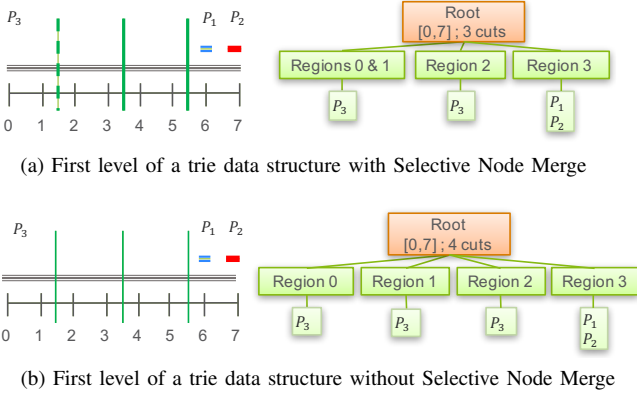


Fig. 3. Impact of Selective Node Merge on the replication factor for the first level of a trie data structure.

SNM method merges equi-sized regions both from the low indices to the highest ones and from the high indices to the lowest ones. For both directions, the SNM method first selects the lowest and highest index equi-sized regions, respectively. Second, it evaluates if each selected region can be merged with its next contiguous equi-sized region. The two steps are repeated until the selected region can no longer be merged with its next contiguous equi-sized region. Else, the two previous steps are repeated from the last equi-sized region that was left un-merged. The SNM method has two constraints with respect to the number of merged regions. First, each merged region covers a number of equi-sized regions that is restricted to powers of two, as the space covered by the merged region is described with the prefix notation. Second, the total number of merged regions is bounded by the size of a node header field of the adaptive trie. By merging equi-sized regions together, the SNM method reduces the number of regions that are stored in the data structure. As a result, the SNM method improves the memory efficiency of the data structure.

The benefit of the SNM method on the memory efficiency is presented in Fig. 3a for the first level of a multi-bit trie. As a reference, the first level of a multi-bit trie without the SNM method is also presented in Fig. 3b. In both figures, IP addresses are defined on 3 bits for the following prefix set $P_1 = 110/3$, $P_2 = 111/3$ and $P_3 = 0/0$. In both figures, the region is initially partitioned into four equi-sized regions, each corresponding to a different bit combination, called 0 to 3. In Fig. 3a, the SNM method merges the two leftmost equi-sized regions 0 and 1, separated by a dashed line, as they fulfill the constraints of SNM. In Fig. 3a, not only the SNM method reduces the number of nodes held in memory by 25% compared to the multi-bit trie presented in Fig. 3b but also prefix P_3 is replicated twice, that is a 33% reduction of the prefix replication factor. As a result, the SNM method increases the memory efficiency of the multi-bit trie data structures.

The regions that are traversed by the SNM method (merged or not) are stored in a SNM field of the adaptive-trie node header. The SNM field is divided into a *LtoH* and *HtoL* array. The *LtoH* and *HtoL* arrays hold the indices of the regions traversed respectively from low to high index values, and high to low index values. For each region traversed by the SNM

method, merged or equi-sized, one index is stored either in the *LtoH* or the *HtoL* array. Indeed, as a merged region holds two or more multiple contiguous equi-sized regions, a merged region can be described with the indices of the first and the last equi-sized region it holds. In addition, the SNM method traverses the equi-sized regions contiguously. Therefore, the index of the last equi-sized region held in a merged region can be determined implicitly using the index of the next region traversed by the SNM method. The index value of a non-merged region is sufficient to fully describe it.

B. Reduced D-Tree leaf

A reduced D-Tree leaf is built when the number of prefixes held in a region of the density-adaptive trie is below a fixed threshold value b . The proposed leaf is based on a D-tree leaf [13], [14] that is extended with the Leaf Size Reduction technique (LSR).

A D-Tree leaf is a bucket that stores the prefixes and their associated NHI held in a given region. A D-Tree leaf has a memory complexity and time complexity of $O(n)$ for n prefixes stored. A D-Tree leaf is used for the regions at the bottom of the density-adaptive trie because of its higher memory efficiency in those regions. Indeed, we observed that most of the bottom level regions of an density-adaptive trie hold highly unevenly distributed prefixes. Moreover, a D-Tree leaf has better memory efficiency with highly unevenly distributed prefixes over a density-adaptive trie. Whereas a density-adaptive trie can create prefix replication, which reduces the memory efficiency, no prefix replication is created with a D-Tree leaf. However, the D-Tree leaf comes at the cost of higher time complexity compared to a density-adaptive trie.

As a consequence, the LSR technique is introduced to reduce the time complexity of a D-Tree leaf by reducing the amount of information stored in a D-Tree leaf. In fact, a D-Tree leaf stores entirely each prefix even the bits that have already been matched by the density-adaptive trie before reaching the leaf. On the other hand, the LSR technique stores in the reduced D-Tree leaf only the prefix bits that are left unmatched. To specify the number of bits that are left unmatched, a new LSR leaf header field is added, coded on 6 bits. The LSR technique reduces the amount of information that is stored in each reduced D-Tree leaf. As a result, not only does the reduced D-Tree leaf requires fewer memory accesses but it also has a better memory efficiency over a D-Tree leaf.

C. HTT build procedure

The hybrid trie-tree build procedure is presented in Algorithm 3, starting with the root region that holds all the prefixes (line 1). If the number of prefixes stored in the root region is below a fixed threshold b , then a reduced D-Tree leaf is built (line 2 – 3). Else, the algorithm iteratively partitions this region into equi-sized regions (lines 4 – 10). The SNM method is then applied on the equi-sized regions (line 11). Next, for each region, if the number of prefixes is below the threshold value (line 12), a reduced D-Tree leaf is built (line 13), else a density-adaptive trie node is built (line 14 – 16) and the region is again partitioned.

Algorithm 3 Hybrid Trie-Tree build procedure**Input:** Prefix Table, stack Q**Output:** Hybrid Trie-Tree

```

1: Create a root region covering all prefixes
2: if the number of prefixes held in that region is below the
   threshold value  $b$  then
3:   Create a reduced D-Tree leaf for those prefixes
4: else
5:   Push the root region onto Q
6: end if
7: while Q is not empty do
8:   Remove the top node in Q and use it as the reference
   region
9:   Compute the number of partitions in the reference
   region
10:  Partition the reference region according to the previous
   step
11:  Apply the SNM method on the partitioned reference
   regions
12:  for each partitioned reference region do
13:    if it holds a number of prefixes that is below the
    threshold value then
14:      Create a reduced D-Tree leaf for those prefixes
15:    else
16:      Build an adaptive-density trie node for those
      prefixes
17:      Push this region onto Q
18:    end if
19:  end for
20: end while
21: return the Hybrid Trie-Tree

```

The number of partitions in a region (line 9 of Algorithm 3) is computed by a greedy heuristic proposed in [13]. The heuristic uses the prefix distribution to adapt the number of partitions, as expressed in Algorithm 4. An objective function, the *Space Measurement* (Sm) is evaluated at each iteration (lines 4 and 5) and compared to a threshold value, the *Space Measurement Factor* ($Smpf$) evaluated in the first step (line 1). The number of partitions increases by a factor of two at each iteration (line 3), until the value of the objective function Sm (line 4) becomes greater than the threshold value (line 5). The objective function estimates the memory usage efficiency with the prefix replication factor by summing the number of prefixes held in each j equi-sized region created $\sum_{j=0}^{N_p} Num_{Prefixes}(equi-sizedregion_j)$ (line 4). The prefix replication factor is impacted by the prefix distribution. If prefixes are evenly distributed, the replication factor remains very low until the equi-sized regions become smaller than the average prefix size. Then, the prefix replication factor increases exponentially. Thus, to avoid over-partitioning a region if the replication factor remains low for many iterations, the number of partitions N_p and the result of the previous iterations $Sm(N_{p-1})$ are used as a penalty term that is added to the objective function (line 4). On the other hand, if prefixes are unevenly distributed, the prefix replication factor increases linearly until the largest prefixes in the region partitioned

become slightly smaller compared to an equi-sized region. Passed this point, an exponential growth of the replication factor is observed. The heuristic creates fine-grained partition size in a dense region, and coarse-grained partition size in a sparse region.

Algorithm 4 Heuristic used to compute the number of partitions in a region**Input:** Region to be cut**Output:** Number of partitions (N_p)

```

1:  $N_p = 1; Smpf = Num_{Prefixes} \cdot 8; Sm(N_p) = 0;$ 
2: do
3:    $N_p = N_p \cdot 2;$ 
4:    $Sm(N_p) = \sum_{j=0}^{N_p} Num_{Prefixes}(equi-sizedregion_j) + N_p + Sm(N_{p-1});$ 
5: while  $Sm(N_p) \leq Smpf$ 
6: return  $N_p$ 

```

The number of partitions in a region is a power of two. Thus, the base-2 logarithm of the number of partitions represents the number of bits from the IP address used to select the equi-sized region covering this IP address.

D. HTT lookup procedure

The hybrid trie-tree lookup algorithm starts with a traversal of the density-adaptive trie until a reduced D-Tree leaf is reached. Next, the reduced D-Tree leaf is traversed to identify the matching prefix and its NHI.

The traversal of the density-adaptive trie consists in computing the memory address of the child node that matches the destination IP address, calculated with Algorithm 5. This algorithm uses as input parameters the memory base address, the destination-IP bit-sequence, the *LtoH* and the *HtoL* arrays that are extracted from the node header. The SNM method can merge multiple equi-sized nodes into a single node in memory, and thus the destination-IP bit-sequences cannot be used directly as an index to the child node. Therefore, Algorithm 5 computes for each destination-IP bit-sequence the number of equi-sized nodes that are skipped in memory based on the characteristics of the merged regions described in the *LtoH* and the *HtoL* arrays. The value of the destination-IP bit-sequence can point to a region that is either included 1) in a merged region described in the *LtoH* array (line 1), or 2) in a merged region described in the *HtoL* array (line 4), or 3) in a equi-sized region that has not been traversed by the SNM method (line 7).

The following notation is introduced: L represents the size of the *HtoL* and *LtoH* arrays, $LtoH[i]$ and $HtoL[i]$ are respectively the i -th entry of the *LtoH* and the *HtoL* arrays. In the first case, each entry of the *LtoH* array is traversed to find the closest $LtoH[i]$ that is less than or equal to the destination-IP bit-sequence (line 1). The index of the matched child node is equal to $index_{LtoH}$ (line 2), where $index_{LtoH}$ is the index of the *LtoH* array that fulfills this condition. In the second case, each entry of the *HtoL* array is similarly traversed to find the closest $HtoL[i]$ that is greater than or equal to the destination-IP bit-sequence

(line 4). The $index_{HtoL}$ in the $LtoH$ array that fulfills this condition is combined with the characteristics of the $LtoH$ and $HtoL$ arrays to compute the index of the selected child node (line 5). In the third case, the algorithm evaluates only the number of equi-sized nodes that are skipped in memory based on the characteristics of the $LtoH$ array and the destination IP address bit sequence of the matched child node (line 7). Finally, the index that is added to the base address points in memory to the matched child node.

Algorithm 5 Memory address of the matched child node using SNM method

Input: Children base address, destination IP address bit sequence, $LtoH$ and $HtoL$ arrays

Output: Child node address

```

    ▷ Index included in a region using SNM method
1: if destination IP address bit sequence  $\leq LtoH[L-1]$  then
    ▷ In  $LtoH$  array
2:   Index =  $Index_{LtoH}$ 
3: else
4:   if destination IP address bit sequence  $\geq HtoL[0]$  then
    ▷ In  $HtoL$  array
5:     Index =  $index_{HtoL} + HtoL[0] - LtoH[L-1] + L - 1$ 
6:   else ▷ destination IP address bit sequence included
    in an equi-sized region that has not been traversed by the
    SNM method
7:     Index = destination IP address bit sequence
        -  $LtoH[L-1] + L - 1$ 
8:   end if
9: end if
10: return Child node address = base address + Index

```

Algorithm 5 is illustrated with Figures 4 in which $L = 3$ and the destination IP address bit sequence is arbitrarily set to 10. Based on Fig. 4, the destination IP address bit sequence 10 matches the equi-sized region with the index 10 before the SNM method is applied. However, after the SNM method is applied, the destination IP address bit sequence matches a merged node with the index 9. Based on the SNM header, the destination IP address bit sequence 10 is greater than both $LtoH[L-1] = 3$ and $HtoL[0] = 9$. Thus, we must identify the number of equi-sized nodes that are skipped in memory with the $LtoH$ and $HtoL$ arrays. Because $LtoH[L-1] = 3$, two equi-sized nodes have been merged. As one node is skipped in memory, any child index greater than 3 is stored at offset $index - 1$ in memory. Moreover, the destination IP address bit sequence is greater than $HtoL[0] = 9$. However, $HtoL[1] = 10$, meaning that indices 9 and 10 are not merged, and no entry is skipped in memory for the first two regions held in the $HtoL$ array. As a consequence, only one node is skipped in memory, and thus the child node index is $10 - 1 = 9$.

The density-adaptive trie is traversed until a reduced D-Tree leaf is reached. The lookup procedure of a reduced D-Tree leaf is presented in Algorithm 6. The leaf header is first parsed, and then prefixes are read (lines 1 to 2). Next, all prefixes are matched against the destination IP address, and their prefix length is recorded if matches are positive (lines 3 to 6). When

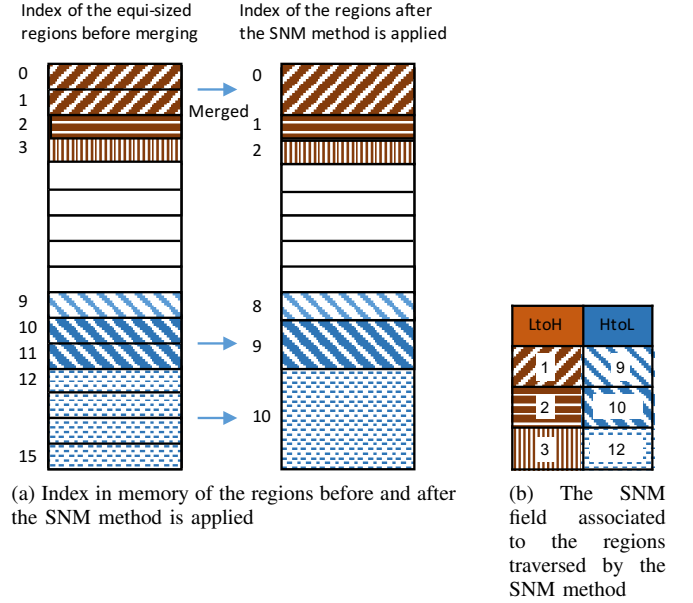


Fig. 4. SNM method applied to a region that holds 11 nodes after merging, and its associated SNM field

Algorithm 6 Lookup in the Reduced D-tree leaf

Input: Reduced D-Tree leaf, destination IP Address

Output: LPM and its NHI

```

1: Parse the leaf header
2: Read the prefixes held in the leaf
3: for each prefix held in the leaf do
4:   Match the destination IP address against the selected
     prefix
5:   if Positive Match then
6:     Record the prefix length of the matched prefix
7:   end if
8: end for
9: Identify the longest prefix match amongst all positive
   matches
10: return the longest prefix match and its NHI

```

all the prefixes are matched, only the longest prefix match is returned with its NHI (lines 7 – 8).

VI. PERFORMANCE MEASUREMENT METHODOLOGY

This section describes the methodology used to evaluate SHIP performance using both real and synthetic prefix tables. Eleven real prefix tables were extracted using the RIS remote route collectors [7], and each one holds approximately 25 k prefixes. Each scenario, noted *rrc* followed by a two-digit number, characterizes the location in the network of the remote route collector used. For prefix tables holding up to 580 k entries, synthetic prefixes were generated with a method that uses IPv4 prefixes to generate IPv6 prefixes, in a one-to-one-mapping [20]. The IPv4 prefixes used were also extracted from [7]. Using the IPv6 prefix table holding 580 k prefixes, four smaller prefix tables were created, with a similar prefix length distribution, holding respectively 290 k, 116 k, 58 k and 29 k prefixes.

The performance of SHIP was evaluated using two metrics: the number of memory accesses to traverse its data structure and its memory consumption. For the two metrics, the performance is reported separately for the hash table used by the address block binning method, and the HTTs built by the prefix length sorting method. SHIP performance is characterized using 1 to 6 groups for two-level prefix grouping, and as a reference the performance of a single HTT without grouping is also presented. The number of groups is limited to six, as we have observed with simulations that increasing it further does not improve the performance.

For the evaluation of the number of memory accesses, it is assumed that the selected hybrid trie-trees within an address block bin are traversed in parallel, using dedicated traversal engines. Therefore, the reported number of memory accesses is the largest number of memory accesses of all the hybrid trie-trees amongst all address block bins. It is also assumed that the memory bus width is equal to a node size, in order to transfer one node per memory clock cycle.

The memory consumption is evaluated as the sum of all nodes held in the hybrid trie-tree for the prefix length sorting method, and of the size of the perfect hash table used for the address block binning method. In order to evaluate the data structure overhead, this metric is given in bytes per byte of prefix. This metric is evaluated as the size of the data structure divided by the size of the prefixes held in the prefix table. The format and size of a non-terminal node and a leaf header used in a hybrid trie-tree are detailed respectively in Table I and in Table II. The node type field, coded with 1 bit, specifies whether the node is a leaf or a non-terminal node. The following fields are used only for non-terminal nodes. Up to 10 bits can be matched at each node, corresponding to a node header field coded with 4 bits. The fourth field is used for SNM, to store the index value of the traversed regions. Each index is restricted to 10 bits, while the *HtoL* and *LtoH* arrays each store up to 5 indices. The third field, coded in 16 bits, stores the base address of the first child node associated with its parent's node.

TABLE I
NON-TERMINAL NODE HEADER FIELD SIZES IN BITS

Header Field	Size
Node type	1
Number of cuts	4
Pointer to child node	16
Size of selective node merge array	$5 \cdot 10 + 5 \cdot 10$

The leaf node format is presented in Table II. A leaf can be split over multiple nodes to store all its prefixes. Therefore, two bits are used in the leaf header to specify whether the current leaf node is a terminal leaf or not. The next field gives the number of prefixes stored in the leaf. It is coded with 4 bits because in this work, the largest number of prefixes held in a leaf is set to 12 for each hybrid trie-tree. The LSR field stores the number of bits that need to be matched, using 6 bits. If a leaf is split over multiple nodes, a pointer coded with 16 bits points at the remaining nodes that are part of the leaf. Inside a leaf, prefixes are stored alongside their prefix length and with

their NHI. The prefix length is coded with the number of bits specified by the LSR field while the NHI is coded with 8 bits.

TABLE II
LEAF HEADER FIELD SIZES IN BITS

Header Field	Size
Node type	2
Number of prefixes stored	4
LSR field	6
Pointer to remaining leaf entries	16
Prefix and NHI	Value specified in the LSR field + 8

VII. RESULTS

SHIP performance is first evaluated using real prefixes, and then with synthetic prefixes, for both the number of memory accesses and the memory consumption.

A. Real Prefixes

The performance analysis is first made for the perfect hash table used by the address block binning method. In Table III, the memory consumption and the number of memory accesses for the hash table are shown. The ABB method uses between 19 kB and 24 kB, that is between 0.7 and 0.9 bytes per prefix byte for the real prefix tables evaluated. The memory consumption is similar across all the scenarios tested as prefixes share most of the 23 MSBs. On the other hand, the number of memory accesses is by construction independent of the number of prefixes used, and constant to 2.

TABLE III
MEMORY CONSUMPTION OF THE ADDRESS BLOCK BINNING METHOD FOR REAL PREFIX TABLES

Scenario	Hashing Table size (kB)	Memory Accesses
<i>rrc00</i>	20	2
<i>rrc01</i>	19	2
<i>rrc04</i>	24	2
<i>rrc05</i>	19	2
<i>rrc06</i>	19	2
<i>rrc07</i>	20	2
<i>rrc10</i>	21	2
<i>rrc11</i>	20	2
<i>rrc12</i>	20	2
<i>rrc13</i>	22	2
<i>rrc14</i>	21	2

In Figures 5a and 5b, the performance of the HTTs is evaluated respectively on the memory consumption and the number of memory accesses. In both figures, 1 to 6 groups are used for two-level prefix grouping. As a reference, the performance of the HTT without grouping is also presented.

In Fig. 5a, the memory consumption of the HTTs ranges from 1.36 to 1.60 bytes per prefix byte for all scenarios, while it ranges between 1.22 up to 3.15 bytes per byte of prefix for a single HTT. Thus, using two-level prefix grouping, the overhead of the HTTs ranges from 0.36 to 0.6 byte per byte

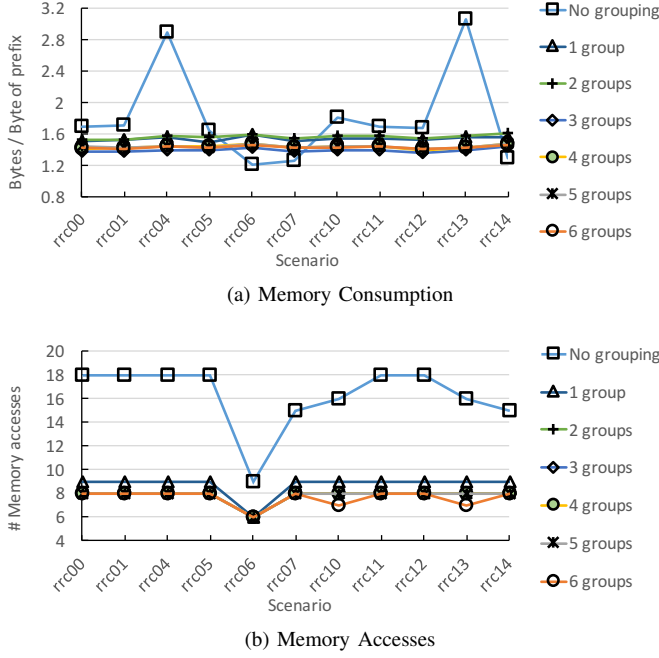


Fig. 5. Real prefix tables: impact of the number of groups on the memory consumption (a) and the number of memory accesses (b) of the HTTs.

of prefix. However, a single HTT leads to an overhead of 0.85 on average, and up to 3.15 bytes per byte of prefix for scenario *rrc13*. Thus, two-level grouping reduces the memory consumption and smooths its variability, but it also reduces the hybrid trie-tree overhead.

Fig. 5a shows that increasing the number K of groups up to three reduces the memory consumption. However, using more groups does not improve the memory consumption, and even worsens it. Indeed, it was observed experimentally that when increasing the value of K most groups hold very few prefixes, leading to a hybrid trie-tree holding a single leaf with part of the allocated node memory left unused. Thus, using too many groups increases memory consumption.

It can be observed in Fig. 5b that the number of memory accesses to traverse the HTTs ranges from 6 to 9 with two-level prefix grouping, whereas it varies between 9 and 18 with a single HTT. So, two-level prefix grouping smooths the number of memory accesses variability, but it also reduces on average the number of memory accesses approximatively by a factor 2.

However, increasing the number K of groups used by two-level prefix grouping, from 1 to 6, yields little gain on the number of memory accesses, as seen in Fig. 5b. Indeed, for most scenarios, one memory access is saved, and up to two memory accesses are saved in two scenarios, by increasing the number K of groups from 1 to 6. Indeed, for each scenario, the performance is limited by a prefix length that cannot be divided in smaller sets by increasing the number of groups. Still, using two or more groups, in the worst case, 8 memory accesses are required for all scenarios. The performance is similar across all scenarios evaluated, as few variations exist between the prefix groups created using two-level grouping for those scenarios.

B. Synthetic Prefixes

The complexity of the perfect hash table used for the address block binning method is presented in Table IV with synthetic prefix tables. It requires on average 2.7 bytes per prefix byte for the 5 scenarios tested, holding from 29 k up to 580 k prefixes. The perfect hash table used shows linear memory consumption scaling. For the number of memory accesses, its value is independent of the prefix table, and is equal to 2.

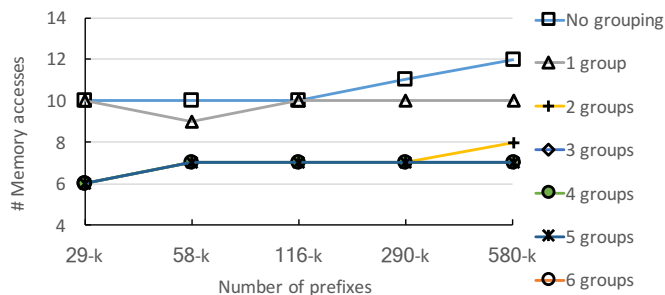
TABLE IV
COST OF BINNING ON THE FIRST 23 BITS FOR SYNTHETIC PREFIX TABLES

Prefix Table Size	Hashing Table size (kB)	Memory Accesses
580 k	1282	2
290 k	642	2
110 k	322	2
50 k	162	2
29 k	82	2

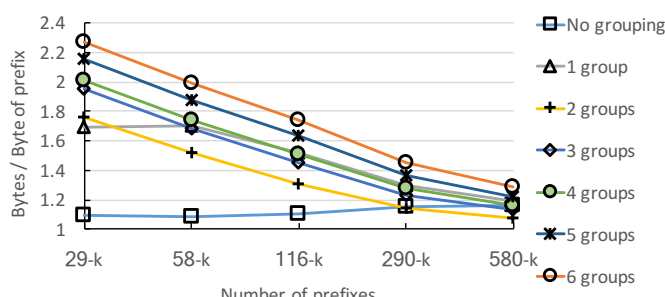
The performance of the HTTs with synthetic prefixes is evaluated for the number of memory accesses, the memory consumption, and the memory consumption scaling, respectively in Fig. 6a, 6b, and 6c. For each of the three figures, 1 to 6 groups are used for two-level prefix grouping. The performance of the HTT without grouping is also presented in the three figures, and is used as a reference.

Two behaviors can be observed for the memory consumption in Fig. 6b. First, for prefix tables with 290 k prefixes and more, it can be seen that two-level prefix grouping used with 2 groups slightly decreases the memory consumption over a single HTT. Using this method with two groups, the HTTs consumes between 1.18 and 1.09 byte per byte of prefix, whereas the memory consumption for a single HTT lies between 1.18 and 1.20 byte per byte of prefix. However, increasing the number of groups to more than two does not improve memory efficiency, as it was observed that most prefix length sorting groups hold very few prefixes, leading to hybrid trie-tree holding a single leaf, with part of the allocated node memory that is left unused. Even though the memory consumption reduction brought by two-level prefix grouping over a single HTT is small for large synthetic prefix tables, it will be shown in this paper that the memory consumption remains lower when compared to other solutions. Moreover, it will be demonstrated that two-level prefix grouping reduces the number of memory accesses to traverse the HTT with the worst case performance over a single HTT, for all synthetic prefix table sizes. Second, for smaller prefix tables with up to 116 k prefixes, a lower memory consumption is achieved using only a single HTT for two reasons. First, the synthetic prefixes used have fewer overlaps and are more distributed than real prefixes for small to medium size prefix tables, making two-level prefix grouping less advantageous in terms of memory consumption. Indeed, a larger number M of address block bins has been observed compared to real prefix tables with respect to the number of prefixes held the prefix tables, for small and medium prefix tables. Thus, on average, each bin holds fewer prefixes compared to real prefix tables. As a consequence, we observe that the average and maximum number of prefixes

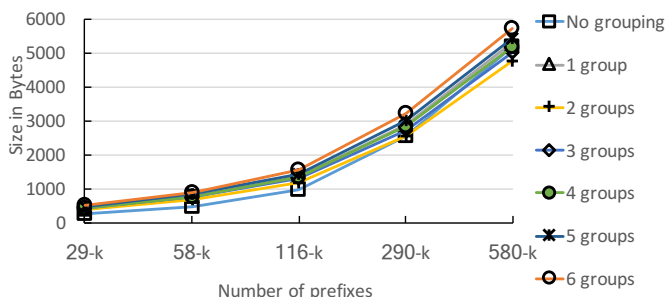
held in each PLS group is smaller for prefix tables holding up to 116 k prefixes. It then leads to hybrid trie-trees where the allocated leaf memory is less utilized, achieving lower memory efficiency and lower memory consumption.



(a) Memory accesses



(b) Memory consumption



(c) Memory consumption scaling

Fig. 6. Synthetic prefix tables: impact of the number of groups on the number of memory accesses (a), the memory consumption (b) and scaling (c) of the HTTs.

In order to observe the memory consumption scaling of the HTTs, Fig. 6c shows the total size of the HTTs using synthetic prefix tables, two-level prefix grouping, and a number K of groups that ranges from 1 to 6. The memory consumption of the HTTs with and without two-level prefix grouping grows exponentially for prefix tables larger than 116 k. However, because the abscissa uses a logarithmic scale, the memory consumption scaling of the proposed HTT is linear with and without two-level prefix grouping. In addition, the memory consumption of the HTTs is 4,753 kB for the largest scenario with 580 k prefixes, two-level prefix grouping, and $K = 2$.

Next, we analyze in Fig. 6a the number of memory accesses required to traverse the HTT leading to the worst-case performance, for synthetic prefix tables, with two-level prefix grouping using 1 to 6 groups. It can be observed that two-level prefix grouping reduces the number of memory accesses

over a single HTT, for all the number of groups and all prefix table sizes. The impact of two-level prefix grouping is more pronounced when using two groups or more, as the number of memory accesses is reduced by 40% over a single HTT. Using more than 3 groups does not further reduce the number of memory accesses as the group leading to the worst-case scenario cannot be reduced in size by increasing the number of groups. Finally, it can be observed in Fig. 6a that the increase in the number of memory accesses for a search is at most logarithmic with the number of prefixes, since each curve is approximately linear and the x-axis is logarithmic.

The performance analysis presented for synthetic prefixes has shown that two-level prefix grouping improves the performance over a single HTT for the two metrics evaluated. Although the performance improvement of the memory consumption is limited to large prefix tables, using few groups, the number of memory accesses is reduced for all prefix table sizes and for all numbers of groups. In addition, it has been observed experimentally that the HTTs used with two-level prefix grouping have a linear memory consumption scaling, and a logarithmic scaling for the number of memory accesses. The hash table used in the address block binning method has shown to offer a linear memory consumption scaling and a fixed number of memory accesses. Thus, SHIP has a linear memory consumption scaling, and a logarithm scaling for the number of memory accesses.

VIII. DISCUSSION

This section first demonstrates that SHIP is optimized for a fast hardware implementation. Then, the performance of SHIP is compared with previously reported results.

A. SHIP hardware implementability

We demonstrate that SHIP is optimized for a fast hardware implementation, as it complies with the following two properties; 1) pipeline-able and parallelizable processing to maximize the number of packets forwarded per second, 2) use of a data structure that can fit within on-chip memory to minimize the total memory latency.

A data structure traversal can be pipelined if it can be decomposed into a fixed number of stages, and for each stage both the information read from memory and the processing are fixed. First, the HTT traversal can be decomposed into a pipeline, where each pipeline stage is associated to a HTT level. Indeed, the next node to be traversed in the HTT depends only on the current node selected and the value of the packet header. Second, for each pipeline stage of the HTT both the information read from memory and the processing are fixed. Indeed, the information is stored in memory using a fixed node size for both the adaptive-density trie and the reduced D-Tree leaf. In addition, the processing of a node is constant for each data structure and depends only on its type, as presented in Section V-D. As a result, the HTT traversal is pipeline-able. Moreover, the HTTs within the K PLS groups are independent, thus their traversal is by nature parallelizable. As a consequence, by combining a parallel traversal of the HTTs with a pipelined traversal of each HTT, property 1 is

fulfilled. The hash table data structure used for the address block binning technique has been implemented in hardware in previous work [4], and thus it already complies with property 1.

For the second property, SHIP uses 5.9 MB of memory for 580 k prefixes, with two-level prefix grouping, and $K = 2$. Therefore, SHIP data structure can fit within on-chip memory of the current generation of FPGAs and ASICs [21]–[23]. Hence, SHIP fulfills property 2. As both the hash table used by two-level prefix grouping, and the hybrid trie-tree comply with properties 1 and 2 required for a fast hardware implementation, SHIP is optimized for a fast hardware implementation.

B. Comparison with previously reported results

Table V compares the performance of SHIP and previous work in terms of memory consumption and worst case memory latency. If available, the time and space complexity are also shown. In order to use a common metric between all reported results, the memory consumption is expressed in bytes per prefix, obtained by dividing the size of the data structure by the number of prefixes used. The memory latency is based on the worst-case number of memory accesses to traverse a data structure. For the following comparison, it is assumed that on-chip SRAM memory running at 322 MHz [2] is used, and off-chip DDR3-1600 memory running at 200 MHz is used.

Using both synthetic and real benchmarks, SHIP requires in the worst case 10 memory accesses, and consumes 5.9 MB of memory for the largest prefix table, with 2 groups for two-level prefix grouping. Hence, the memory latency to complete a lookup with on-chip memory is equal to $10 \cdot 3.1 = 31$ ns.

FlashTrie has a high memory consumption, as reported in Table V. The results presented were reevaluated using the node size equation presented in [9] due to incoherence with equations shown in [11]. This algorithm leads to a memory consumption per prefix that is around $11\times$ higher than the SHIP method, as multiple copies of the data structure have to be spread over DDR3 DRAM banks. In terms of latency, in the worst case, two on-chip memory accesses are required, followed by three DDR3 memory bursts. However, DRAM memory access comes at a high cost in terms of latency for the FlashTrie method. First, independently of the algorithm, a delay is incurred to send the address off-chip to be read by the DDR3 memory controller. Second, the latency to complete a burst access for a given bank, added to the maximum number of bank-activate commands that can be issued in a given period of time, limits the memory latency to 80 ns and reduces the maximum lookup frequency to 84 MHz. Thus, FlashTrie memory latency is $2.5\times$ higher than SHIP.

The FlashLook [8] architecture uses multiple copies of data structures in order to sustain a bandwidth of 100 Gbps, leading to a very large memory consumption compared to SHIP. Moreover, the memory consumption of this architecture is highly sensitive to the prefix distribution used. For the memory latency, in the worst case, when a collision is detected, two on-chip memory accesses are required, followed by three memory bursts pipelined in a single off-chip DRAM, leading to a total latency of 80 ns. The observed latency of SHIP is

61% smaller. Finally, no scaling study is presented, making it difficult to appreciate the performance of FlashLook for future applications.

The method proposed in [2] uses a tree-based solution that requires 19 bytes per prefix, which is 78% larger than the proposed SHIP algorithm. Regarding the memory accesses, in the worst case, using a prefix table holding 580 k prefixes, 22 memory accesses are required, which is more than twice the number of memory accesses required by SHIP. In terms of latency, their implementation leads to a total latency of 90 ns for a prefix table holding 580 k prefixes, that is $2.9\times$ higher than the proposed SHIP solution. Nevertheless, similar to SHIP, this solution has a logarithmic scaling factor in terms of memory accesses, and scales linearly in terms of memory consumption.

Finally, Tong et al. [4] present the CLIPS architecture [24] extended to IPv6. Their method uses 27.6 bytes per prefix, which is about $2.5\times$ larger than SHIP. The data structure is stored in both on-chip and off-chip memory, but the number of memory accesses per module is not presented by the authors, making it impossible to give an estimate of the memory latency. Finally, the scalability of this architecture has not been discussed by the authors.

These results show that SHIP reduces the memory consumption over other solutions and decreases the total memory latency to perform a lookup. It also offers a logarithmic scaling factor for the number of memory accesses, and it has a linear memory consumption scaling.

IX. CONCLUSION

In this paper, SHIP, a scalable and high performance IPv6 lookup algorithm, has been proposed to address current and future application performance requirements. SHIP exploits prefix characteristics to create a shallow and compact data structure. First, two-level prefix grouping leverages the prefix length distribution and prefix density to cluster prefixes into groups that share common characteristics. Then, for each prefix group, a hybrid trie-tree is built. The proposed hybrid trie-tree is tailored to handle local prefix density variations using a density-adaptive trie and a reduced D-Tree leaf structure.

Evaluated with real and synthetic prefix tables holding up to 580 k IPv6 prefixes, SHIP builds a compact data structure that can fit within current on-chip memory, with very low memory lookup latency. Even for the largest prefix table, the memory consumption per prefix is 10.64 bytes, with a maximum number of 10 on-chip memory accesses. Moreover, SHIP provides a logarithmic scaling factor in terms of the number of memory accesses and a linear memory consumption scaling. Compared to other approaches, SHIP uses at least 44% less memory per prefix, while reducing the memory latency by 61%.

ACKNOWLEDGMENTS

The authors would like to thank the Natural Sciences and Engineering Research Council of Canada (NSERC), Prompt, and Ericsson Canada for financial support to this research.

TABLE V
COMPARISON RESULTS

Method	Memory Consumption (in bytes per prefix)	Latency (ns)	Complexity	
			Memory Consumption	Memory latency
Tree-based [2]	19.0	90	$O(N)$	$O(\log_2(N)) \leq \text{Latency} \leq 2 \cdot O(\log_3(N))$
CLIPS [4]	27.6	N/A	N/A	N/A
FlashTrie [11]	124.2	80	N/A	N/A
FlashLook [8]	1010.0	90	N/A	N/A
SHIP	10.64	31	$O(N)$	$O(\log(N))$

REFERENCES

- [1] Cisco. (2016, Feb.) Cisco visual networking index: Global mobile data traffic forecast update, 2015–2020 white paper. [Online]. Available: <http://www.cisco.com/c/en/us/solutions/collateral/service-provider/visual-networking-index-vni/mobile-white-paper-c11-520862.html>
- [2] H. Le and V. K. Prasanna, “Scalable tree-based architectures for IPv4/v6 lookup using prefix partitioning,” *IEEE Transactions on Computers*, vol. 61, no. 7, pp. 1026–1039, July 2012.
- [3] H. J. Chao and B. Liu, *High Performance Switches and Routers*. Wiley, 2007.
- [4] D. Tong, Y. H. E. Yang, and V. K. Prasanna, “A memory efficient IPv6 lookup engine on FPGA,” in *2012 International Conference on Reconfigurable Computing and FPGAs*, Dec 2012, pp. 1–6.
- [5] N. Alliance, “5G white paper,” Tech. Rep., 2015.
- [6] (2016, April) IPv6 BGP Table Data. [Online]. Available: <http://bgp.potaroo.net/v6/as2.0/index.html>
- [7] (2015, October) RIS Raw Data Set. [Online]. Available: <https://www.ripe.net/analyse/internet-measurements/routing-information-service-ris/ris-raw-data>
- [8] M. Bando, N. S. Artan, and H. J. Chao, “FlashLook: 100-Gbps hash-tuned route lookup architecture,” in *International Conference on High Performance Switching and Routing, 2009.*, June 2009, pp. 1–8.
- [9] M. Bando and H. J. Chao, “Flashtrie: Hash-based prefix-compressed trie for IP route lookup beyond 100Gbps,” in *2010 Proceedings IEEE INFOCOM*, 2010, pp. 1–9.
- [10] H. Song, F. Hao, M. Kodialam, and T. Lakshman, “IPv6 lookups using distributed and load balanced bloom filters for 100gbps core router line cards,” in *INFOCOM 2009, IEEE*. IEEE, 2009, pp. 2518–2526.
- [11] M. Bando, Y.-L. Lin, and H. J. Chao, “Flashtrie: beyond 100-Gb/s IP route lookup using hash-based prefix-compressed trie,” *IEEE/ACM Transactions on Networking*, vol. 20, no. 4, pp. 1262–1275, 2012.
- [12] H. Yu and R. Mahapatra, “A Power and Throughput-Efficient Packet Classifier with n Bloom Filters,” *IEEE Transactions on Computers*, vol. 60, no. 8, pp. 1182–1193, 2011.
- [13] P. Gupta and N. McKeown, “Classifying packets with hierarchical intelligent cuttings,” *IEEE Micro*, vol. 20, no. 1, pp. 34–41, Jan 2000.
- [14] D. E. Taylor and J. S. Turner, “Scalable packet classification using distributed crossproducing of field labels,” in *INFOCOM 2005. 24th Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings IEEE*, vol. 1, March 2005, pp. 269–280 vol. 1.
- [15] W. Eatherton, G. Varghese, and Z. Dittia, “Tree bitmap: hardware/software IP lookups with incremental updates,” *SIGCOMM Comput. Commun. Rev.*, vol. 34, no. 2, pp. 97–122, Apr. 2004. [Online]. Available: <http://doi.acm.org/10.1145/997150.997160>
- [16] K. Zheng, C. Hu, H. Lu, and B. Liu, “A TCAM-based distributed parallel IP lookup scheme and performance analysis,” *IEEE/ACM Transactions on Networking*, vol. 14, no. 4, pp. 863–875, Aug. 2006.
- [17] G. Rétvári, J. Tapolcai, A. Kőrösi, A. Majdán, and Z. Heszberger, “Compressing IP Forwarding Tables: Towards Entropy Bounds and Beyond,” *IEEE/ACM Transactions on Networking*, vol. 24, no. 1, pp. 149–162, Feb. 2016.
- [18] (2016, January) IPv6 global unicast address assignments. [Online]. Available: <http://www.iana.org/assignments/ipv6-unicast-address-assignments/ipv6-unicast-address-assignments.xhtml>
- [19] Z. J. Czech, G. Havas, and B. S. Majewski, “Perfect hashing,” *Theoretical Computer Science*, vol. 182, no. 1-2, pp. 1–143, Aug. 1997. [Online]. Available: [http://dx.doi.org/10.1016/S0304-3975\(96\)00146-6](http://dx.doi.org/10.1016/S0304-3975(96)00146-6)
- [20] M. Wang, S. Deering, T. Hain, and L. Dunn, “Non-random generator for IPv6 tables,” in *12th Annual IEEE Symposium Proceedings on High Performance Interconnects, 2004. Proceedings*. Washington, DC, USA: IEEE Computer Society, 2004, pp. 35–40. [Online]. Available: <http://dx.doi.org/10.1109/CONECT.2004.1375198>
- [21] Altera. Overview of the Startix-10 family. [Online]. Available: <https://www.altera.com/products/fpga/stratix-series/stratix-10/overview.html>
- [22] Xilinx. Ultrascale architecture and product overview. [Online]. Available: https://www.xilinx.com/support/documentation/data_sheets/ds890-ultrascale-overview.pdf
- [23] P. Bosshart, G. Gibb, H.-S. Kim, G. Varghese, N. McKeown, M. Izard, F. Mujica, and M. Horowitz, “Forwarding metamorphosis: Fast programmable match-action processing in hardware for SDN,” in *ACM SIGCOMM Computer Communication Review*, vol. 43, no. 4. ACM, 2013, pp. 99–110.
- [24] Y. H. E. Yang, O. Erdem, and V. K. Prasanna, “High Performance IP Lookup on FPGA with Combined Length-Infix Pipelined Search,” in *2011 IEEE 19th Annual International Symposium on Field-Programmable Custom Computing Machines*, May 2011, pp. 77–80.