
THE UBIQUITOUS SKIPLIST: A SURVEY OF WHAT CANNOT BE SKIPPED ABOUT THE SKIPLIST AND ITS APPLICATIONS IN BIG DATA SYSTEMS

A PREPRINT

Venkata Sai Pavan Kumar Vadrevu

Department of Computer Science
Purdue University
West Lafayette, IN 47906
vvadrevu@purdue.edu

Lu Xing*

Department of Computer Science
Purdue University
West Lafayette, IN 47906
xingl@purdue.edu

Walid G. Aref

Department of Computer Science
Purdue University
West Lafayette, IN 47906
aref@purdue.edu

ABSTRACT

Skiplists have become prevalent in systems. The main advantages of skiplists are their simplicity and ease of implementation, and the ability to support operations in the same asymptotic complexities as their tree-based counterparts. In this survey, we explore skiplists and their many variants. We highlight many scenarios of how skiplists are useful and fit well in these usage scenarios. We study several extensions to skiplists to make them fit for more applications, e.g., their use in the multi-dimensional space, network overlaying algorithms, as well as serving as indexes in database systems. Besides, we also discuss systems that adopt the idea of skiplists and apply the probabilistic skip pattern into their designs.

Keywords Skiplist, skip list, index

1 Introduction

Skiplists, or skip lists, have been first introduced by William Pugh in [76]. Pugh’s paper coins the term “skip list”, but with its popularity, it has become one word and has been termed the “skiplist” in later papers. Since then, the skiplist has become ubiquitous, not only as a probabilistic data structure that is of theoretical interest, but rather as a practical data structure that has been used widely in many systems. The main advantages of skiplists are their simplicity and ease of implementation, and the ability to support operations in the same asymptotic complexities as their tree-based counterparts.

In this paper, we survey the skiplist, its variants, and its applications in various systems including big data systems. In Section 2, we present basics of the skiplist, its structure, and its corresponding algorithms for insert, delete, and search. Section 3 compares and contrasts the skiplist with the B/B⁺-tree. Section 4 presents skiplist variants that support concurrent operations. Section 5 introduces a taxonomy for skiplists and the versatile uses of skiplists in various big data systems. Sections 6-8 overview deterministic skiplists, skiplist adaptivity to data access patterns, and several skiplist optimizations, respectively. Section 9 presents skiplists as customized for specific hardware platforms. Sections 10-12 survey skiplists for the multi-dimensional space, interval skiplists, and skiplist realizations of priority queues, respectively. Finally, Section 14 concludes the paper.

*The first two authors contributed equally to this research.

2 Skiplist basics

In this section, we present the skiplist basics. First, we review the structure of a skiplist. Then, we discuss how insert, delete, and search operations are performed in the skiplist and the cost analysis associated with each operation. Next, we present common extensions and optimizations associated with the basic skiplist. In Section 5, we explain these extensions in more detail when we discuss variants of the skiplist.

2.1 Basic Structure

Refer to Figure 1a for illustration. The skiplist is a probabilistic data structure that consists of multiple levels, where each level contains a linked list of nodes and pointers. Levels are numbered bottom up from one to i ($i > 1$). Level-1 is a linked list connecting all the data items in sorted order. In level- i ($i > 1$), a fraction p ($p \in (0, 1)$) of the nodes in level- i also appear in level- $(i + 1)$ pointers. In [76], Pugh uses $p = 1/2$. Each node has at most one forward pointer per level pointing to its successor (right sibling) in the same level. The header of the skiplist has forward pointers; one per level. When the skiplist is empty, all forward pointers of the header point to NIL. A node with k forward pointers is referred to as a *Level k node*. Given a node, its level is randomly chosen when it is inserted into the skiplist. The number of levels is capped by a constant *MaxLevel* that Pugh [76] suggests to use $\log_{1/p} N$ (N is an upper bound of the number of elements in the skiplist). Figure 1a gives an example of a skiplist containing seven elements with $p = 1/2$. We show how operations are conducted over skiplists in the next section.

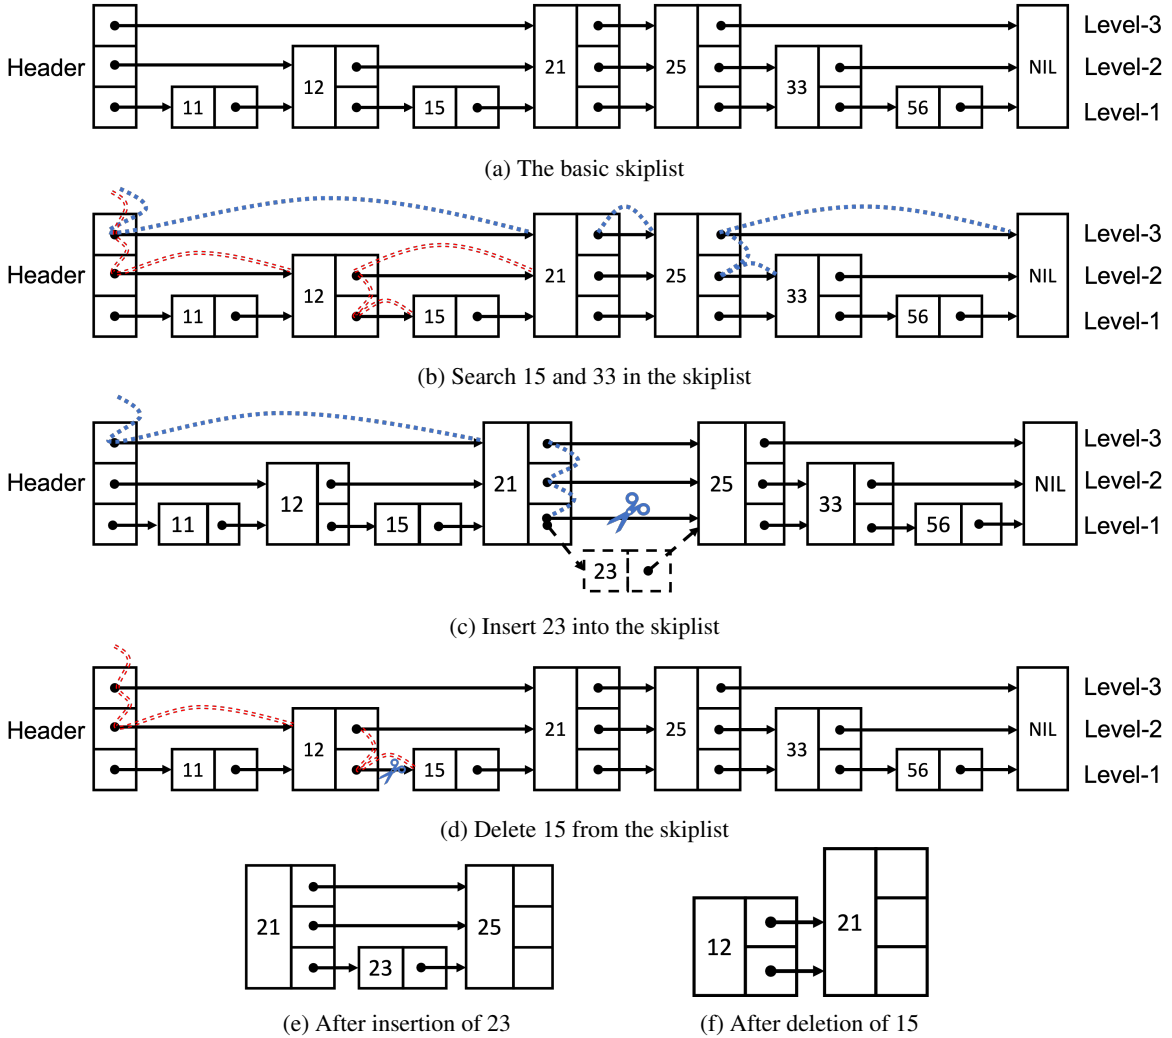


Figure 1: Skiplist and its basic operations




Figure 2: Layers in Skiplist

As in Figure 2, a skiplist can be logically divided into two parts: the index layers, and the data layer. Each query passes through the index layers, and then reaches the data layer. The role of index layers is to route each query to the data layer in $O(\log n)$ time complexity. On the other hand, the data layer stores the actual data (key and data), whereas the index layers only store key values used for routing purposes.

2.2 Operations and Cost Analysis

2.2.1 The Search Operation

Search starts from the topmost level (maximum level available) and follows the forward pointers to reach the searched item. At each node, the search key is compared with the key of the node. If the node key is the same as the search key, then the value is returned. If the node key is greater than the search key, the search process goes back to the previous node and moves one level below and follows the same process. In Figure 1b, if the search key is 33, the search process starts from the header, then to Nodes 21, 25, and finally reaches Node 33. If the search key is 15, the search checks Node 12 first, then Node 21, and then goes down of Node 12, and finally reaches Node 15.

The cost of search is proportional to the search path of the key. The expected average length of a search path is $O(\log n)$, where the list has n elements. The detailed analysis is provided in [76].

2.2.2 The Insert Operation

Algorithm 1 Determine Height of Node

procedure GETNODEHEIGHT

$level \leftarrow 1$

while $random() < p$ and $level < MaxLevel$ **do**

$level = level + 1$

return $level$

▷ $random()$ returns a random value in $[0...1)$

To insert a new data item, we first determine its height h by Algorithm 1. Then, we start to search for the key in the skiplist following the search procedure as in the previous section. However, if the height of the key is greater than one, pointers need to be modified in multiple levels. We record the pointers of the rightmost node that is to the left of the inserted node in levels- i ($i \leq h$) during the search phase, and update these pointers to reflect the insertion. In Figure 1c, suppose key 23 is inserted and its height returned by Algorithm 1 is 1, after creating Node 23, we need to update the pointer of its level-1 predecessor. Thus, Node 21 points to Node 23 as shown in Figure 1e. The expected average cost is the same as search as it takes extra constant time to update pointers after the search.

2.2.3 The Delete Operation

The delete operation is similar to the insert operation. When the search key is found, pointers are modified to reflect the deletion, as illustrated in Figure 1d. If the delete key is 15, the forward pointer of its predecessor, i.e., Node 12 needs to be modified and the resulting node is shown in Figure 1f.

2.2.4 The Range Search Operation

Range search is supported in the skiplist. Given a range $[start-key, end-key]$, we search for $start-key$ in the skiplist in the bottom level, then all the keys within the range can be found by following the bottom linked list.

2.3 Analysis of p

The value of p determines the possibility of a node to be promoted to a higher level. When $p = 1/2$, about half of the nodes in level- i are promoted to level- $(i + 1)$. If we decrease p , the search time increases while memory usage is

decreased [76]. Pugh suggests to use $p = 1/2$ if the variability in running time is not of concern, otherwise $p = 1/4$ can be used [76].

3 Contrast with B/B⁺-trees

The B-tree has been ubiquitously used in data systems [13, 24]. In this section, we study the mapping between the skiplist and the B/B⁺-tree. Given this mapping and the ubiquity of the B-tree, it is nature for the skiplist to be as popular and as ubiquitous as the B-tree in systems.

3.1 Deterministic Skiplists

One caveat about the probabilistic skiplist is that the worst case costs are unbounded. The deterministic skiplist [68] can provide a bounded cost even in the worst case. Unlike the probabilistic skiplist, the deterministic one is history-dependent.

Both the 1-2 skiplist and 1-2-3 skiplist [68] are deterministic in that both their average and worst case performance are bounded by $\Theta(\lg n)$. Example of the 1-2-3 skiplist is given in Figure 3a. For a 1-2 skiplist of N elements, there exists 1 or 2 elements at height $h - 1$ between any two nodes of height h ($h > 1$). If the forward pointers of a node are stored in an array, when a node grows, a new array has to be allocated with the pointers being copied. This can result in up to $\Theta(\lg^2 n)$ time for one insertion [68]. Munro et al. [68] propose an *array implementation of the 1-2 skiplist* to make arrays have exponentially increasing sizes. During one insertion, there is at most one array that needs to be allocated with data being copied [68]. The worst-case cost of search, insert, and delete are $\Theta(\lg n)$ [68]. These operations are explained in Section 6.

3.2 Corresponding Trees

Munro et al. [68] show that there is a one-to-one correspondence between the 1-2 skiplist and the 2-3 tree, as well as between the 1-2-3 skiplist in Figure 3a and the 2-3-4 tree in Figure 3b.

The equivalence is formally defined by Lamoureux and Nickerson [57] as *structural equivalence* and *functional equivalence*. Structural equivalence requires that data item A and B are logically adjacent to each other in both structures. In contrast, functional equivalence means that the worst-case cost functions are in the same big- O order of complexity. Two intermediate tree structures are defined in [57]: The Bd-tree and the Bd⁺-tree. To construct a Bd-tree, we use a 2-3-4 tree, and connect the tree node with its right sibling if this node has a right sibling. Figure 3c shows a Bd-tree of order 4 and is equivalent to a B-tree of order 4. The Bd⁺-tree adds the succeeding key from the parent node as the rightmost entry in each node shown in Figure 3d.

There are three pair-wise equivalences: B-tree vs. Bd-tree, Bd-tree vs. Bd⁺-tree, Bd⁺-tree vs. the horizontal array format of a deterministic skiplist. Lamoureux et al., [57] prove that a B-tree of order m is structurally and functionally equivalent to the linked list implementation of a deterministic skiplist of order m . Also, as in Figure 3d, the Bd⁺-tree is similar to the B^{link}-tree [60], where each node points to its right sibling. However, there are two differences. The Bd⁺-tree has an extra root node that contains ∞ . Moreover, each non-leaf node of the Bd⁺-tree contains an extra key. Showing the equivalence between the B-tree and the skiplist establishes the foundation that the skiplist can be used as a database index in place of a B-tree.

A detailed transformation between the skiplist and the binary search tree (BST, for short) is discussed in [27]. A skiplist can be transformed into a multiway branching search tree with edge weights (weight indicates the node height difference between parent and child nodes), and then to a BST with edge weights. To achieve a balanced BST, tree rotation is involved. The promotion (demotion) of a node resulting from an insertion or deletion in the skiplist is also performed by a rotation in the BST.

4 Concurrent Skiplists

Concurrent access and operations to the skiplist can be supported by either lock or lock-free techniques. Both implementations are possible with linear-scalability [77]. Also, there are variants of skiplists that support multi-versioning and transactional memory [93, 54, 48, 51]. In the following subsections, we discuss these techniques in further detail.

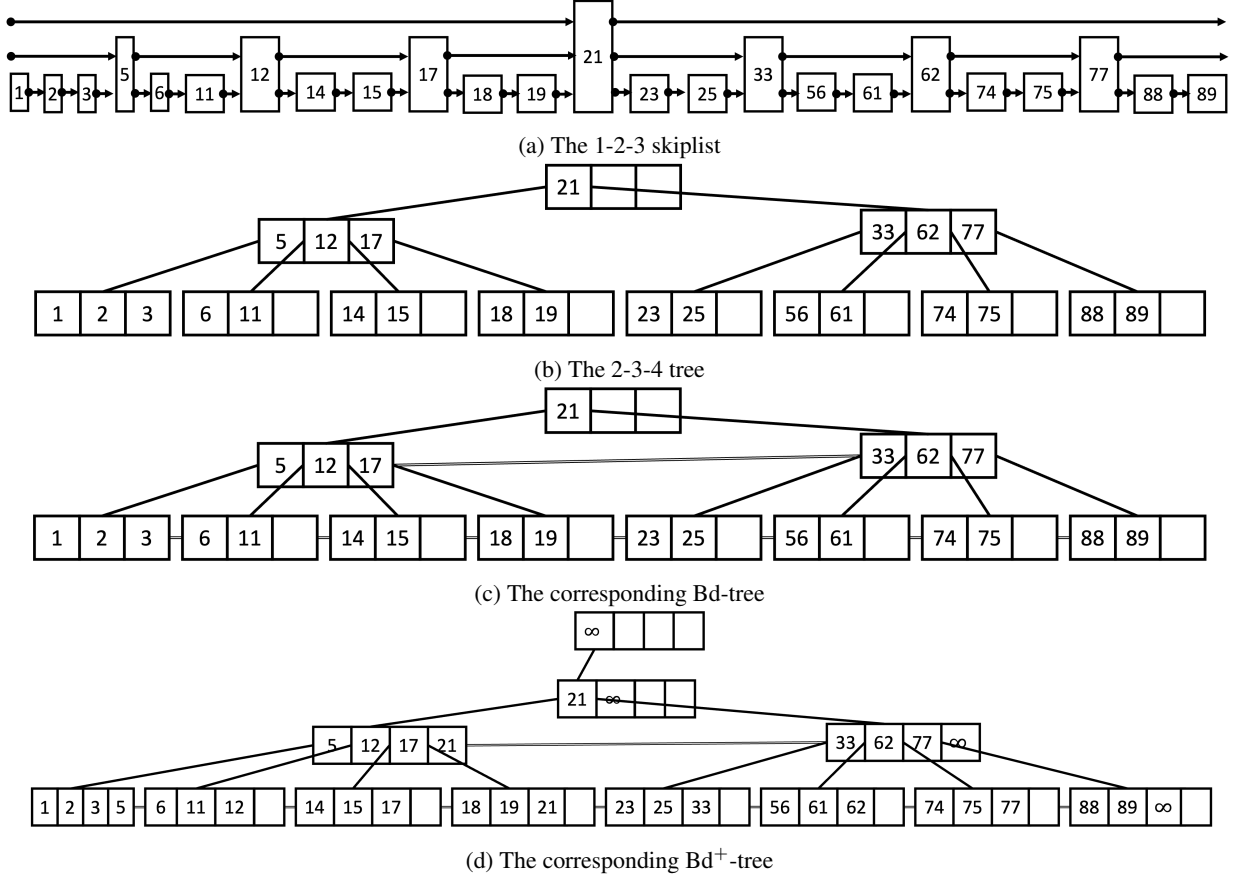


Figure 3: The 1-2-3 skiplist and its counterpart tree structures

4.1 The Lock-based Skiplist

Pugh [77] uses locks for concurrent access and manipulation of skiplists. Each node maintains an array of pointers at each level the node spans. The search procedure is the same as the basic skiplist search algorithm. Insertion starts by first searching in level-0 for a node that contains the largest key that is smaller than the inserted key. After we find the node *prev* (the node after which the new node is to be inserted), *prev* is locked and a new node is created, and is inserted after it in level-0. Next, we determine the height of the node, and increment the height of the new node one step at a time, and update the previous, new, and next node pointers at that level. We repeat the same step till we reach the desired height. Similarly, for the delete procedure, we decrement the height of the node from its current height to zero, and follow the same steps as those for insert. This approach scales linearly [77], although there is contention to hold the lock.

The Optimistic Lock-based Skiplist [45, 44]. Herlihy et al. [45, 44] have proposed an optimistic lock-based skiplist. It is optimistic because all searches are performed without locks. Only when the item is found, the node and its predecessor are locked, and a validation step ensuring the list has not been changed is performed. Each node maintains two flags, *marked* that indicates whether a logical deletion is performed or not, and *fullyLinked* that hints on whether the locked node has been added to all the levels it should reside. During insertion, if the key is not found, all the predecessors are locked up to the height of the new node H . Then, the procedure validates that the predecessors are adjacent with the successors up to H , and they are not marked for deletion. Next, the new node is created, and is inserted into the structure. Finally, *fullyLinked* is set when the insert procedure is performed, and all locks are released.

The Unrolled Skiplist [75]. The unrolled skiplist [75] uses *group mutual exclusion* as its locking technique. It groups multiple keys per node to achieve a better cache locality as well as reduce the number of pointer chases. As multiple keys share the same node, multiple threads performing the same operation on different keys in the same node can proceed concurrently. In order for a thread to join a session, [75] uses two algorithms. One algorithm uses a word to track the type of the session and the number of threads in the session. When a thread wants to join, it repeatedly reads

the word and updates the word via compare-and-swap if it joins the session successfully. The second algorithm uses an exclusive lock to protect. The threads push themselves into a queue and wait until session is compatible.

4.2 The Lock-free Skiplist




Figure 4: The lock-free linked list

Harris [40] introduces a lock-free implementation of linked lists via the atomic operation compare-and-swap (CAS). In Figure 4a, in order to insert a new node, say Node 26, the next pointer of Node 25 is atomically pointed to the new node. However, the deletion of Node 25 with one CAS as shown in Figure 4b may result in a situation in Figure 4c where the insertion of Node 26 is lost because the insert operation finds Node 25 before Node 25 is deleted. A two-step deletion can avoid this situation by marking the next pointer of Node 25 as invalid with one CAS instruction, that is referred to as a logical deletion. The second step is a physical deletion that changes the pointer atomically. The design of the lock-free linked list can be applied to the skiplist [33].

In [33], several designs of a lock-free skiplist are included. A compare-and-swap (CAS)-based design uses the idea of the lock-free linked list [40] as each level in the skiplist can be viewed as an independent linked list. As there can be inconsistencies in concurrent deletes and inserts using only one CAS as illustrated in Figure 4, a marking scheme introduced in [40] is implemented in [33]. The second lock-free skiplist [33] uses multi-word CAS (MCAS). For inserts, all writes are grouped in a batch and are installed via a single MCAS operation, while reads check that the read location is not owned by the MCAS operation. The third lock-free technique uses transactional memory [39]. Each node is treated as a separate transactional object. Concurrent operations on the same node are performed sequentially within a transaction.

Harris [40] uses a two-step process to delete a node. In the first step, we mark the node's next pointer to indicate that the node is logically deleted. A pointer can be marked by setting the lowest bit in the pointer field to 1, as the two lowest order bits are set to 0 in a 32-bit system. In the second step, we physically delete the node. Using this two-step process we can achieve a lock-free linked list. As in Figure 4d, observe that Node 25's next pointer is marked in the first step, making it logically deleted. Assume that node 26 is being concurrently inserted, it will now identify that Node 25 is logically deleted during traversal and will restart the insert process from the beginning until Node 25 is physically deleted. This will avoid the problem of disappearing nodes, as mentioned before.

The Lock-free Linked List with Three-Step Deletion [32]. The two-step deletion in [40] restarts the insert process when there is a conflict. Thus, the average cost is increased [32]. Fomitchev et al. [32] use a three-step deletion process to avoid starting a failed CAS from the start of the list. In the first step, we mark Node 25's next pointer. Next, we create a back pointer to point to the previous node, say Node 21. This is performed to prevent restarting the process from the beginning of the list. In the final step, we physically delete Node 25. By following this three-step deletion, When Node 26 is being inserted, it can start the traversal from Node 21 (with the help of the back pointer) instead of starting from the beginning of the list (Figure 5).

The Lock-free Skiplist with Three-Step Deletion [32]. Fomitchev and Ruppert [32] implement a lock-free skiplist that is composed of the above linked list. Each skiplist node that has k forward pointers is divided into k nodes vertically. Only the bottom one (*root node*) contains the data. These k nodes form a linked list vertically, and is called a *tower*. For example, in Figure 6, $B1$, $B2$ and $B3$ form a tower. Both $B2$ and $B3$ have a pointer to $B1$ where data is stored. Horizontally, a back pointer field is added to each node to facilitate deletion. Insert starts by finding the right location for the new data, and inserts the *root node* first. Other tower nodes can be added bottom up. Delete first marks




Figure 5: The 3-step deletion in a lock-free linked list

the *root node*, and utilizes the back pointers in each level as in the above lock-free linked list. When the *root node* is marked deleted, the tower nodes are superfluous. *B2* and *B3* are superfluous. Search in the skiplist helps delete superfluous nodes.




Figure 6: A lock-free skiplist

4.3 A Concurrent Priority Query Implementation Using a Skiplist

Since skiplists are ordered and provide a delete operation, they can be used to create a priority queue by deleting the first (smallest) element. Sundell and Tsigas [84] create a lock-free concurrent priority queue based on the two-step deletion and CAS primitive. They also provide an extended algorithm for lock-free priority queue using timestamps. This implementation scales well even with preemption.

4.4 Concurrent Skiplist Java Implementations

In addition to a latched implementation of the skiplist, one can implement a lock-free skiplist. `ConcurrentSkiplistMap`, implemented by Doug Lea [58], is part of the `java.util.concurrent` library. This relies on Harris's approach [40], Fomitchev and Ruppert's approach [32], or Fraser's [33]. This implementation uses Compare-And-Swap (CAS)-able node pointers. When a node needs to be deleted, it is replaced or spliced with another node that represents a deletion mark. This technique reduces space overhead by not using deletion marks, and can improve performance as well.

We represent the different states of a node by changing the node's values. If the value of a node is null or if the next node is a marker node, then the node is deleted or is in the process of being deleted. A node is said to be logically deleted when the value is null or when its next node is marked, but is physically deleted when the garbage collector collects it. A garbage collector collects a deleted node during traversal when it finds a node that is marked as logically deleted (i.e., when the node's value is null). The key idea in this approach is to use Compare And Swap (CAS) operations to detect an insert or delete in the data structure. Using CAS, we check if a previous node's next value is the same as the current node to detect modifications. If the CAS operation fails, then we have a thread that has inserted or deleted a node in the same time that we are trying to perform our operation. Thus, we re-execute the same steps till the operation succeeds.

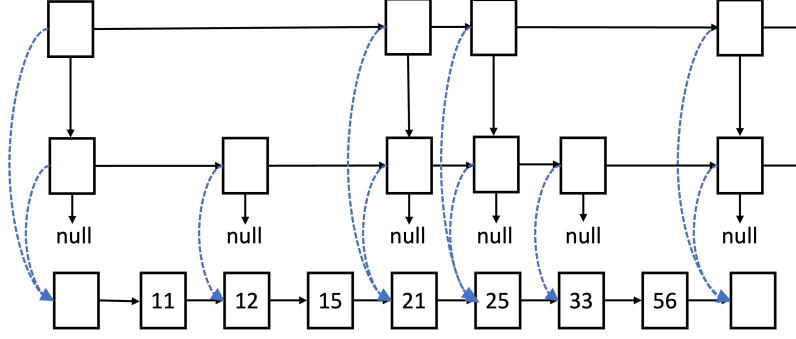


Figure 7: Java Skiplist

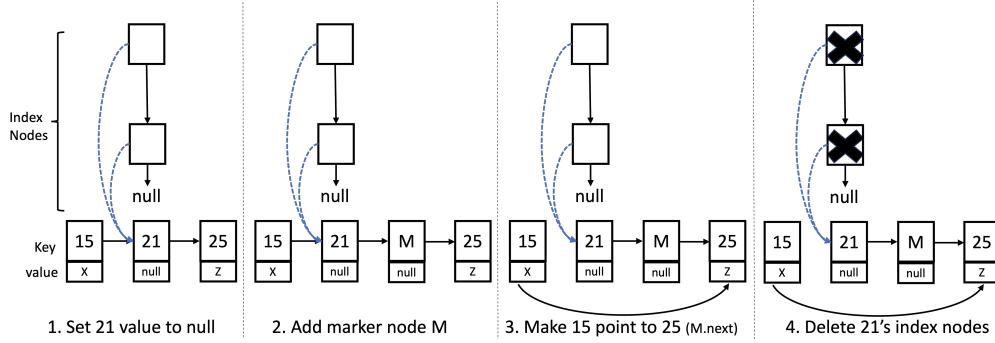


Figure 8: Java Skiplist deletion of Node 21

The search procedure is straightforward. The only difference is that if a node is marked or the next node is a marker node, then we know that the node is in the process of being deleted. Thus, the search operation should handle this case, and report `Not Found`.

In contrast, insert uses the search procedure to find the predecessor of the node we are trying to insert. Once we find the predecessor, we store it in a variable to perform CAS for future insert and delete operations. We store the previous node in a variable in order to detect a concurrent insertion or deletion. Once the CAS succeeds, the node is inserted into data layer. Afterwards, we determine the height by using the Algorithm 1, and insert the index nodes using the same CAS operation to ensure correctness. A node is said to be logically deleted if its value is set to null. To delete a node, say N , we perform the following steps. Refer to Figure 8 for illustration. In the figure, for simplicity, the index nodes for only Node 21 are displayed. Figure 8 illustrates the procedure for deleting Node 21.

1. First, using CAS, set the N 's value field from non-null to null (Figure 8-1).
2. Next, also using CAS, set the next pointer of N to point to a new marker node to avoid the modification of N 's next pointer by other concurrent threads. This could happen during inserts and deletes by other threads that are running concurrently (Figure 8-2).
3. Using CAS, check the next pointer for the node previous to N over both N and the marker, and change the previous's next pointer to point to the marker's next node. This completes the deletion of Node N in the data layer (Figure 8-3).
4. N 's index nodes are deleted as they are encountered during search traversals. Node N in the data layer will be garbage collected after all of its index nodes are deleted (Figure 8-4).

4.5 Contention-Friendly Skiplist Concurrency Techniques

The previous implementation suffers when there is a hotspot. Nodes at the higher levels of a skiplist are accessed frequently as all the queries start from the highest level. There is thread contention at the higher levels, which leads to poor performance. Thread contention primarily occurs due to inserts, when we raise the index layers, or deletes when we delete the index layers. Contention Friendly Skiplist [25] avoids this issue by decoupling inserts at the data layer from those at the index layers, similarly for deletes. It uses a deterministic skiplist to balance the nodes. Also, it uses a background thread, termed the "adaptive thread", to have the changes in the data layer propagate to the index layers.




Figure 9: Contention Friendly Skiplist

Another key change is the use of a doubly linked node in the data layer, as shown in Figure 9. Doubly linked nodes are needed to prevent the repeated find predecessor function calls that can be performed in only $O(1)$ time when using doubly-linked nodes compared to $O(\log n)$ time for the earlier approach. With these advantages, the contention-friendly skiplist, also termed the no-hotspot skiplist, outperforms the previous concurrent skiplists in the following ways:

- **Search:** Search is performed identically to the Java implementation (Section 4.4), except that we do not use the `findPredecessor` function whenever CAS fails. We just use the previous pointer instead. Using the previous pointer of the doubly-linked list speeds up the skiplist as `findPredecessor` is used quite extensively in the Java Implementation, and with the doubly-linked list, one can avoid invoking it completely.
- **Insert:** Inserts are first performed at the data layer. The adaptive thread raises the level of the node to maintain balance.
- **Delete:** Deletes take place at the data layer in a way similar to how the deletion is performed in Java’s implementation. Because we have decoupled deletes in the index level, this responsibility is handled by the adaptive thread.
- **Adaptive Thread:** The adaptive thread is a background thread that performs the raising towers, lowering towers, and clean up operations.
- **Raising Towers:** Because we use a deterministic skiplist, whenever we encounter m nodes with the same height we raise the middle node. Thus, balance is maintained in the skiplist.
- **Lowering Towers:** Whenever the number of deleted nodes hits a threshold number, the adaptive thread reduces the entire skiplist level by deleting the bottommost index layer. We choose to delete the bottommost index layer as it has many items, and is less frequently accessed compared to the higher layers.
- **Clean-up:** During traversals, if there is any logically deleted node, it is garbage collected either by any thread that performs operations or by an adaptive thread.

4.6 The Rotating Skiplist

Although the no-hotspot skiplist avoids contention in the upper levels and remains lock-free, its performance does not scale well with the increase in the size of the number of elements due to high cache misses [28]. The no-hotspot skiplist uses linked nodes, and has poor spatial locality. It also requires linear time to lower the level of the skiplist during rebalancing. The *Rotating Skiplist* [28] addresses these issues by using a *wheel* and a global variable named “ZERO”. ZERO is used to indicate the lowest logical level of all the arrays, as shown in Figure 10 where ZERO initially points to level 0. A wheel is an array node, where by using an array, we improve spatial locality, and increase cache hits. The Rotating Skiplist uses modular arithmetic to wrap around the array node ($\text{level} = (\text{current level} + \text{ZERO}) \bmod \text{max-levels}$). The Rotating skiplist has wheels that “rotate” during the lowering of the skiplist to balance the structure, and hence providing good time complexity for all the operations. The Rotating skiplist contains a doubly-linked list in the bottom-most level, similar to [25], where this level contains the actual data. All the operations except Lowering Towers are performed in a similar manner to that of a no-contention skiplist.

Lowering Towers is performed by one atomic instruction by incrementing the ZERO variable by one. As in Figure 11, ZERO is changed from 0 to 1. All operations that are operating concurrently to lowering towers will continue to operate normally. However, any new operation will now ignore Level 0, and operate on Levels 1 and 2 only. After all operations operating on Level 0 are completed, the garbage collector physically deletes Level 0.




Figure 10: Rotating Skiplist




Figure 11: Lowering Towers

The Rotating Skiplist overcomes several problems, e.g., having low-cache miss ratio, scales well with the increase in size of the skiplist, and also lowers the level of the skiplist in one atomic instruction. However, the main disadvantage is that the space complexity of the Rotating Skiplist is $O(n \log n)$ [28]. This space is caused due to the fixed wheel capacity that is predetermined at the start time for each wheel [28]. We use a rotating skiplist to mitigate contention hotspots, and reduce low cache misses, and thus enables the rotating skiplist to achieve peak performance of 200 MOPS/sec, and scales well with the increase in size of the skiplist [28].

4.7 Multi-Versioned Skiplist

Multiversion Concurrency Control (MVCC) is an optimistic concurrency control technique used in database management systems. MVCC keeps multiple copies of each data item that is being modified. Each user can see a *snapshot* of the database at a particular time corresponding to the time when the user's transaction has started. Snapshot isolation [15] guarantees that a transaction, say T , observes a database that is produced by all the transactions that are committed before T starts. Studies have been conducted to compare multiversion concurrency control algorithms, e.g., [16, 79].

JellyFish. Both the skiplist of X-Engine [48] and JellyFish [93] are designed with the observation that multiple versions of the same key are linked according to their timestamps at the bottom layer of the skiplist. Searching for the key within certain timestamps may require to traverse the list in reverse order if the same key at the upper layer has an obsolete timestamp. For example, searching for Key 12 in Figure 12a lands at Value E with Timestamp 4 that is obsolete because of a more recent value at Timestamp 6. Instead, both X-Engine's skiplist and JellyFish organize a vertical list to store multiple versions of the same key (Figure 12b). The insertion of a new key follows the same process as that in the ordinary skiplist. However, for updates, JellyFish adds the newer version into the vertical list with the most recent version at the second place of the list (The very first head of the vertical list is part of the horizontal bottom layer and is not moved). The nodes for Key 12 are arranged in ordered timestamp as in Figure 12b. The *Get* operation starts at the upper layer of the skiplist, and searches the target data until it hits the vertical list at the bottom layer. The *Scan* operation traverses the bottom layer, and does not have to scan through the obsolete key-value pairs. The read operations, both the get and range searches, benefit most from this design.

Jiffy. Jiffy [54] is another multi-versioned skiplist. Similar to JellyFish [93], Jiffy also keeps a vertical list of revisions at each entry in the bottom layer. Each revision is associated with a version number and contains multiple key-value pairs in the range from the rooted skiplist entry and the next skiplist entry. Revision uses a copy-on-write technique and is installed atomically via compare-and-swap to the list. The obsolete revisions are garbage collected. Jiffy also supports batch updates. Several update operations that need to be performed atomically are put in one batch. All update operations are stored in one batch descriptor. The revisions created by the batch updates reference the batch descriptor. A node covering a smaller range has smaller revisions, and smaller revisions are better for write-heavy workloads. On the other hand, larger ones are more suited for read-heavy workloads. Thus adaptation to workloads is performed by adjusting the node size. Jiffy uses an autoscaling policy for the dynamic adjustments. By monitoring the ratio of time spent on reading and updating, Jiffy decides an optimal size of a node and performs node merges (so that a node covers larger ranges) or splits (so that nodes cover smaller ranges) accordingly.

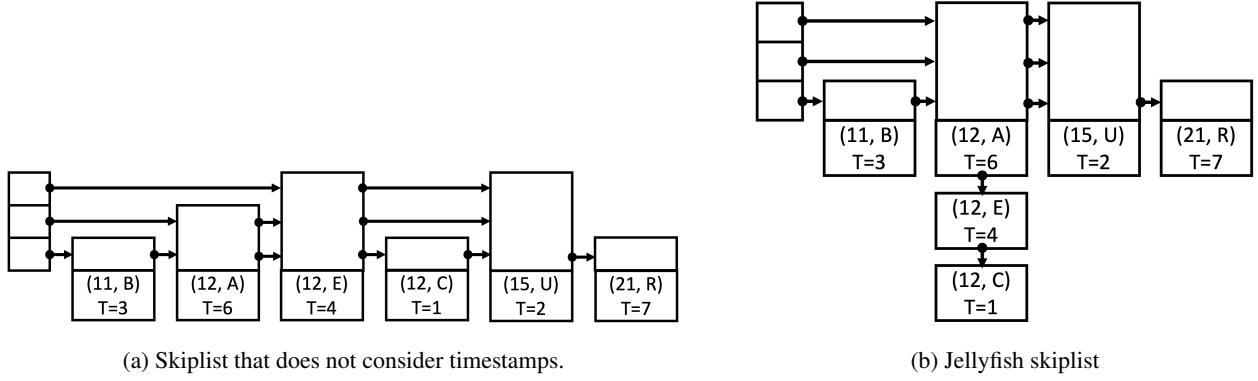


Figure 12: Multi-versioned skiplists with timestamp.

5 Variants of the Skiplist: Taxonomy and Related Systems

Figure 13 illustrates the evolution of the skiplist variants over time, and the various systems that make use of these variants. All skiplist variants originate from the first skiplist [76]. Thus, we omit the connection from the first skiplist paper.

5.1 Skiplists as Database Indexes

The equivalence between the skiplist and the B/B^+ -tree has been demonstrated in Section 3. Thus, skiplists are used often in database systems. *SingleStore*, previously known as MemSQL [4], uses the skiplist as an index for its row-stores. Nitro uses a lock-free skiplist as its core index [55]. ListDB [53] is a LSMT-based key-value store for Non-Volatile Memory (NVM), or Persistent Memory (PM). ListDB has a NUMA-aware skiplist on persistent memory to realize a skiplist-based key-value store. Hash-based key-value stores over RDMA mostly support point queries. A skiplist, supporting range searches and being highly concurrent, is a good fit for an RDMA-based key-value store. RS-store [47] is a skiplist-based key-value store using RDMA.

5.2 Skiplists as Part of an Index

Many Key-Value stores use a skiplist to store the in-memory data for fast concurrent insertion while maintaining data in sorted order, e.g., LevelDB [34], RocksDB [6], X-engine [48], Redis [7], HBase [5] etc. Skiplists have also been used in combination with other data structures to improve the scalability of in-memory component of key-value stores, e.g., S3 [94], FloDB [12], TeksDB [35].

5.3 Skiplists In Operating Systems

The skiplist has been investigated for use in the Linux kernel since 2001. Its proposed use is in replacement of red-black trees. One skiplist implementation outperforms the red-black tree in a multi-threaded testing [3]. MuQSS, a process scheduler, is also based on the skiplist [2]. A more detailed kernel skiplist investigation can be found in [1].

5.4 Skiplist Versatility

As the height of each inserted node is determined via randomness, the skiplist does not require rebalancing, which is different from the B^+ -tree. *PebblesDB* [78] combines the skiplist with Log-Structured Merge Trees (LSMT) in the fragmented LSMT (FLSM). The skiplist is a hierarchical structure with the most coarse-grained level on the top. This structure, together with the probabilistic nature, has inspired algorithms, e.g., Hierarchical Navigable Small World (HNSW) [64] with proximity graphs replacing linked lists. We will discuss more in Section 13.

6 Enforcing Determinism

We overview deterministic skiplists in Section 3. In this section, we discuss the search, insert, and delete operations in 1-2 deterministic skiplists [68]. For insert and delete, the key point is to ensure the gap invariant that between any two consecutive nodes at Height h , there are at most 2 nodes at Height $h - 1$ [68].

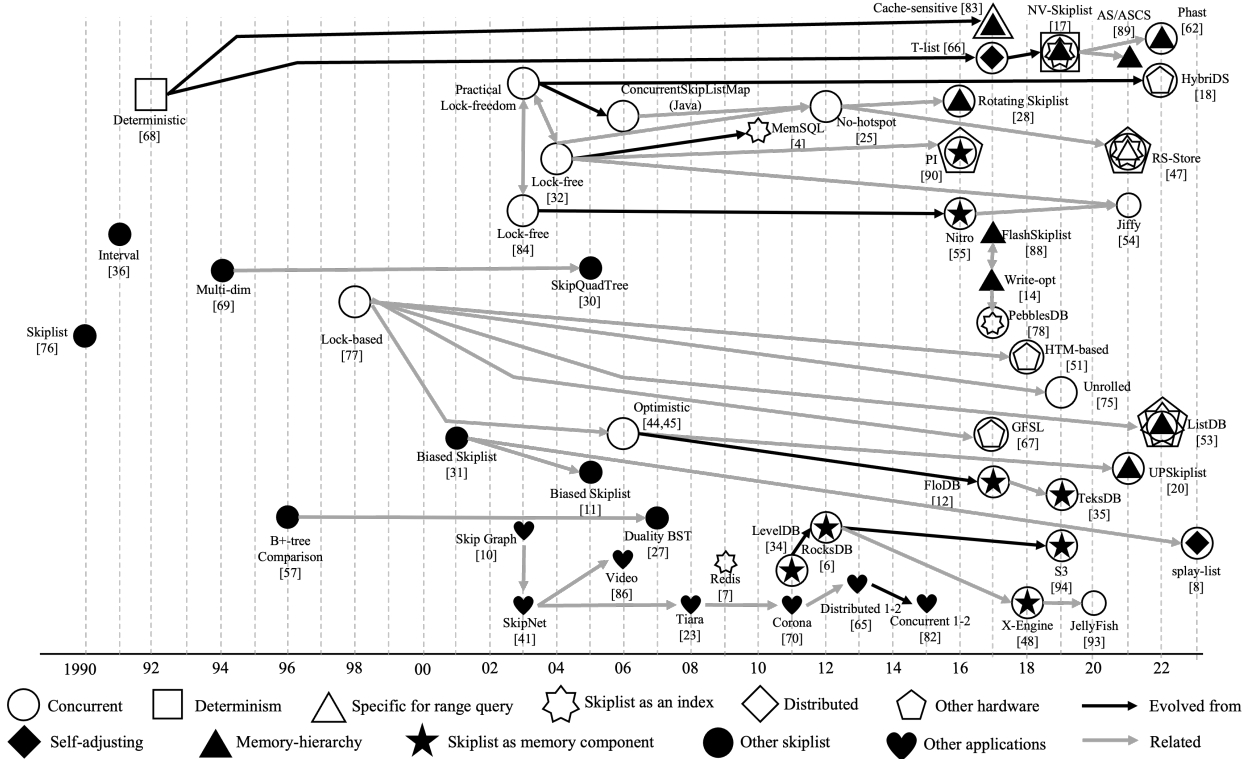


Figure 13: The taxonomy of skiplist variants and their applications.

Search in the 1-2 deterministic skiplist is the same as that in the probabilistic skiplist. There is one advantage in the 1-2 skiplist that, during search, it is guaranteed that after two horizontal steps there is a descentance of one level [68]. An insert in the 1-2 skiplist is performed by first adding the item at Height 1 (bottom level). This may invalidate the gap invariant by introducing a third item of Height 1 in a row. Thus, the middle item is promoted to Height 2. This is repeated until reaching the height of the skiplist [68]. To maintain the gap invariant during deletion, demoting, or promoting a node may be needed [68]. For example, in Figure 3a, deleting Node 5 promotes Node 3 to a higher level, and the gap invariant is still maintained.

In contrast to a probabilistic skiplist, the height of a node in a deterministic skiplist is not decided by consulting a random number generator, rather by maintaining the gap invariant. In the 1-2 deterministic skiplist of n different items, the number of horizontal pointers in the worst can never exceed $2n$ [68]. Since the insertion process mentioned above involves items to grow, this can cause copying of horizontal pointers if these pointers are stored in an array. Then an insertion can take up to $\Theta(\log^2 n)$ time [68]. If horizontal pointers are stored in a linked list, this reduces the time complexity but increases the storage cost up to $6n$. Munro et al. [68] use an array of exponentially increasing heights and the allocated space is smaller than $2.282n$ pointers [68]. Also, the time complexity of insertion is $\Theta(\log n)$ [68].

Both the probabilistic and the deterministic skiplists use active inserts, and deletes, meaning that updates are immediately propagated to the index layers after the operations, which makes these structures balanced.

Network overlay algorithms, e.g., [23, 70, 41, 42, 82, 65], use the deterministic skiplist. Deterministic skiplists are widely used by network maintenance algorithms. Dynamic overlay networks are applied in peer-to-peer systems. As the skiplist supports logarithmic search and update, and can also be made deterministic, the skiplist has become a valid candidate for network overlay algorithms [23, 82, 65, 70].

7 Accommodating for Skewed Data Access Patterns in Skiplists

Data structures often work efficiently when access patterns are uniform, but access patterns are often biased or skewed in many situations, e.g. locality in reference results in memory [85], disk [71] and buffer management [26]. Many data structures have been proposed to deal with these access patterns.

The Biased Skiplist (BSL) [31]. Ergun et al. [31] introduces a skiplist for biased access patterns. The skiplist is constructed similarly to the ordinary skiplist except that the keys are partitioned into *classes*, and data is copied either automatically or randomly into the upper levels. The motivation for replicating some keys is that the frequently accessed keys can be replicated and stored in the higher levels in the skiplist and hence are likely to be found faster. To assign keys to different classes, keys are ordered in ascending order by their rank $r(k)$ ($r(k)$ is the number of distinct keys accessed since the last access to k). This ordered list is partitioned into classes $C_1, C_2, \dots, C_{\log n}$ contiguously where class size is $|C_i| = 2^{i-1}$. Class 1 is the smallest rank, which suggests that the data should be accessed early in the upper levels of the skiplist. During construction, the height of a key is determined by its class as well as a random number (details in [31]). Searching a key is similar to searching in an ordinary skiplist. After the key is found, the rank of the key becomes 1, and it is moved to the front of the rank ordered linked list. Keys in other classes need to be moved to different classes due to the class size constraint. In BSL, the expected search time for key k is $O(\log r(k))$ where $r(k)$ is the rank of key k . In a sense, this mimics Shannon’s coding theory [80], where frequently accessed items get assigned a shorter code length. Similarly, in BSL, a frequently accessed item, say i , will require less probes to the skiplist before i is found, the reason being that i will be replicated into the upper levels in the skiplist as i gets accessed more frequently, and hence is found faster.

The (a, b) -biased Skiplist And Randomized biased Skiplist [11]. Bagchi, Buchsbaum, and Goodrich [11] introduce another notion of a biased skiplist, where each key is assigned a weight w_i . The goal is to achieve faster search time than $O(\log n)$ for highly weighted items. Bagchi et al. devise two biased skiplists: a deterministic one and a randomized one. The (a, b) -biased skiplist is a deterministic skiplist where the height of the item $h_i \geq r_i$ and r_i is the rank defined as $\lfloor \log_a w_i \rfloor$. Two invariants need to be enforced: (1) There are no more than b consecutive items of Height i , where i is within the height of the skiplist; and (2) For Node x and all i that are between x ’s rank and height, there are at least a nodes of height $i - 1$ between x and any consecutive node of height at least i . The (a, b) -biased skiplist can access an item in $O(1 + \log(W/w_i))$ time in the worst case [11]. The randomized version assigns the height of an item to be the sum of the rank-determined number as well as a random number. It can achieve the above bound in the expected case [11].

The Self-adjusting Skiplist (SASL) [21, 22]. SASL is a skiplist that can promote or demote its node’s height [21, 22]. Assume that there are n data elements in the skiplist and the height of Node s is bounded by $H = \Theta(\log n)$. SASL also uses the access pattern and a random number to determine the height of a node, similar to [31]. The horizontal linked lists of SASL are grouped into multiple bands, which are similar to classes in [31]. The size of the band grows exponentially. Refer to Figure 14 for illustration. Assume that there are 20 data elements in the skiplist and 3 bands. As in the figure, B_1 consists of list L_7 while B_2 consists of lists L_5 and L_6 , etc. Each item in the skiplist is associated with two integers, a random value r and a deterministic value d . r is used as the random height within a band and d equals to the height of some band B_i . A node, say s , resides in B_i if $d(s) = H(i)$. For example, in Figure 14, if $r(15) = 1$ and $d(15) = 4$ so Key 15 is in B_2 as $H(2) = 4$. The shadowed part of Key 15 in Figure 14 is the deterministic part, which is height of 4. The height of Key 15 inside B_2 is $r(15) = 1$. Search proceeds as in the ordinary skiplist. Adjustment is done after search. If the target element x resides in B_i , x is promoted to B_1 . $r(x)$ is not discarded for future demotion. For each B_j where $j \in [1, i - 1]$, a random element is chosen to be demoted from B_j to B_{j+1} . This search and restructuring takes $O(2^i)$ expected time. The detailed proof is in [21, 22].




Figure 14: Self-adjusting skiplist (SASL)

The T-List [66]. The T-list [66] is a skiplist that is constructed and modified during search process. The intuition is similar to the 1-2 deterministic skiplist [68] where the search process can only visit no more than a fixed number of nodes in the same level. Mei et al. [66] define a fixed *span* to limit the number of horizontal steps. Unlike the deterministic skiplist, T-List does not enforce a strong constraint. Since insertion always starts by searching, the T-list is a byproduct of the search process. During search, when the number of consecutive horizontal steps exceeds a predefined threshold *span*, that node is promoted to a higher level. Refer to Figure 15 for illustration. Let *span* be equal to 2. From the figure, searching for Key 25 requires 3 horizontal steps in Level-0, which exceeds the threshold 2. Thus, Key 25 is

promoted to a higher level and next time Key 25 can be found in fewer steps. To promote a node, only two nodes need to be locked for update: the current modifying node and the first node that is on this level before current node. This can increase concurrency. Delete is performed logically by inserting an invalidating entry.






Figure 15: The T-list

The Splay-list [8]. The Splay-list [8] is a skiplist that is constructed based on data access frequency. The idea is to promote the frequently accessed nodes to higher levels in the skiplist, and demote the less frequently accessed nodes. The maximum height of a node, say u , is denoted by h_u . The children nodes of u at Height h are denoted by C_u^h which is also the subtree of u at h . The subtree of u includes the successor nodes of u at Height $h < h_u$, e.g., the subtree for Node 15 contains itself, Node 21 and Node 25 (shaded nodes in Figure 16a). In Figure 16a, both the head and tail store the number of accesses of the subtree. Let H be the head. For each node, the left box stores the key, and the right box stores the number of accesses to the subtree. Note that the height starts at 0. In the example, $C_{15}^1 = \{15, 21, 25\}$; $C_{12}^1 = \{12\}$, $C_H^1 = \{H, 11\}$, $C_H^2 = \{H, 11, 12, 15, 21, 25\}$ including the subtree of Node 15. Assume that the total number of accesses to the skiplist is $m = 10$. Thus, the height of the tree is $\lfloor \log(m) \rfloor = 3$. For the descent condition, compute the sum (C_1) of accesses of the subtrees of the neighboring nodes u and v . If the sum is less than a descent threshold, the node is demoted, and is merged into the subtree of its left predecessor. Similarly, for the ascent condition, compute the sum (C_2) of accesses of the subtree. If it is greater than a threshold, the node is promoted. In the example, after *contains*(25), Key 25 is found, and the skiplist needs to be adjusted in a backward pass. The first step is to increment the access number for Node 25. To decide whether to demote Node 21, $C_1 = C_{15}^0 + C_{21}^0 = 2$, which is below the threshold. For the promotion computation of node 21, $C_2 = C_{21}^1 + C_{25}^0 = 3$ exceeds the threshold. Thus, the skiplist becomes as in Figure 16b, where Node 21 is promoted to Level-1 and the access number of Node 15 at Level-1 is changed to 1. Similar adjustments are performed to Node 15 as in Figure 16c. More detail on how to set the threshold can be found in [8].



(a) The splaylist



(b) First three steps in backward pass after *contains*(25). (c) Next three steps in backward pass after *contains*(25).

Figure 16: The Splaylist and the promotion/demotion process

8 Skiplists with Unrolling nodes

In the initial skiplist [76], visiting a node requires multiple pointer traversals. Pointer chasing may not be optimal for caches and on-disk structures. Several skiplist variants group multiple data items into one node. This grouping is referred to as *unrolling*. Next, we introduce conditions where unrolling can be applied to improve certain aspects of skiplists.

The Cache-sensitive Skiplist [83]. The cache-sensitive skiplist [83] has been introduced as the linked list offers poor spatial locality. The nodes of a linked list are not necessarily located contiguously next to each other in memory. This results in a cache miss every time a pointer is chased when we move from one element of the next in the linked list. In order to address this issue, Sprenger, Zeuch, and Leser [83] use a linearized fast lane that stitches the entire index layer into a single array. The cache-sensitive Skiplist eliminates the use of pointers, and calculates the positions of the child nodes by using arithmetic. This saves space, and reduces cache misses. Moreover, SIMD instructions [95] can be used to further speed up operations. These benefits have the disadvantage that the index layers become rigid, i.e., whenever an insert occurs, the node is only inserted in the data layer, and after the number of inserts reaches a threshold, the cache-sensitive skiplist reconstructs the new index layers. Similarly, for a delete operation, we delete the node in the data layer, but we update all the corresponding positions in the index layers with the smallest value greater than the deleted node. This is performed to avoid false positives during search in the cache-sensitive skiplist. The cache-sensitive skiplist cannot use null value, as it cannot be compared against the search value during traversals, so the only logically valid option is to copy the next greater value. Due to all the design choices, the cache-sensitive skiplist achieves performance comparable to that of the Adaptive Radix Tree [61], and offers fast range search performance. Another in-memory skiplist [91] separates the data layer from the index layer, and is SIMD-friendly as multiple entries are stored per index node, and can be loaded into a single SIMD register.

The Write-optimized Skiplist [14]. Bender et al. [14] present a write-optimized skiplist that works on external-memory of block size B to achieve asymptotically better insertion performance than the B-tree and offer similar range search performance. Each node contains pivots and a buffer, but the amount of space assigned to each is variable. Based on the promotion probability [14], on Level $L_{i \geq 1}$, each node has $\Theta(B^\epsilon)$ pivots in expectation, and on Level L_0 , there are expected $\Theta(B^{1-\epsilon})$ pivots. The smallest pivot per node is termed *leader*. In Figure 17b, the bold elements are pivots and the number of pivots per node is different. All the data items in the buffer are greater than the leader of the same node, but are smaller than the leader of the successor node in the same level. Figure 17a gives the same skiplist as the one in Figure 3a but in a different representation. Insertions are first buffered in the buffer area in the node until the buffer overflows. A flush operation is triggered to distribute each data item in the buffer to the corresponding child node. If the height, say h_e , of data item e is greater than the current level, this suggests that e needs to become a pivot in the lower levels. This involves node splitting and establishing new parent-child node link. Insertions take $O((\log_{B^\epsilon} N / (B^{1-\epsilon}))$ amortized I/Os in expectation and with high probability. Its range search returning K elements achieves $O(\log_{B^\epsilon} N + K/B)$ I/Os w.h.p for $0 < \epsilon < 1$ and block size B [14].




Figure 17: The write-optimized skiplist.

The FlashSkipList [88]. The FlashSkipList [88] is designed for external flash memory. Entries are grouped in chunks, where the chunks can be accessed through skiplist nodes. The skiplist is constructed based on the largest entry per chunk. This design saves the space for pointers and is I/O efficient as one I/O brings in multiple entries. Also, batched updates can be applied to the same chunk at once [88]. More detail about the *FlashSkipList* is covered in Section 9.1.1.

The PhaST [62]. Persistent Memory (PM) is another example of external memory. PhaST [62] also adopts the unrolling design to have multiple key-value pairs in the same node. We will explain the whole design in details in Section 9.1.2.

Concurrent Unrolled Skiplist [75]. Unrolling has also been applied to skiplists that are designed for better concurrency [75]. Each node is associated with one anchor key and stores multiple keys. Each level of linked lists is ordered by the anchor keys per node. The keys that are within the node lie in the range of the co-localized anchor key and the anchor key in the successor node on the same level. Unrolling helps reduce the number of pointers in the skiplist, and thus improves cache performance. Since the number of nodes is reduced, the number of promotions/demotions during

skiplist modification is reduced. In most cases, insertions and deletions do not cause structural changes, then locking a single node is enough to perform that update.

9 Hardware Implications

In this section, we explore the implications of new hardware on the designs of the skiplist. This includes the impact of new types of memory including flash and persistent non-volatile memories, and CPU architectures including NUMA awareness and GPUs.

9.1 External Memory Storage

9.1.1 Flash Memory

Flash Memory is asymmetric. Its read and write latencies are not equal. It has low read latency but high write latency. Moreover, flash memory has the erase-before-write constraint that increases the write cost. The FlashSkiplist data structure [88] tackles the constraints posed in flash memory by using a write-optimized component, a read-optimized component, and a dynamic rearrangement strategy to achieve overall good performance. Since writes are costly, the FlashSkiplist buffers the writes into a write-optimized component that uses an append-only strategy to achieve high write throughput. As shown in Figure 18, lists $L1 - L3$ comprise the write optimized component. Any new updates will be appended to the lists in the topmost level of FlashSkiplist.

The read-optimized component is comprised of chunks of pages. As shown in Figure 18 the chunks $C1 - C6$ form the read-optimized component. Each node in the flash skiplist points to a chunk, and the value of the node corresponds to the maximum value allowed in the particular chunk.

One important feature of the FlashSkiplist is the use of dynamic rearrangement that restricts the length of the write-optimized list so that the read cost is manageable, and also the writes can be batched and transferred to the read-optimized skiplist component to decrease the overall write costs. Insert, delete, search, and rearrangement operations are described below:

- **Insert:** An insert operation is simply an append to lists in the topmost level of FlashSkiplist. This takes $O(1)$ time. In order to reduce the cost further, the FlashSkiplist uses a linked list where the newly inserted node points back to the old node, instead of the old node pointing to the newly inserted node. This way, one can modify the newly inserted node, and keep the write-optimized structure intact free of any erases. As shown in fig 18 when a new command insert 68 arrives, we append it to the list $L1$ in the topmost level of the FlashSkiplist.
- **Delete:** A delete operation is inserting a ghost entry into the top level list. Ghost entry is an entry that indicates that the element has been deleted. For example, as in Figure 18, in List $L2$ Element 34 is a ghost element, meaning that 31 has been deleted.
- **Search:** A search takes place in two steps. The first step is a Skiplist search that finds the corresponding list or chunk in which the key is present. The second step is to search within a linked list or a chunk. If the key is in a list, then we perform linked list search. However, if the key is found in a chunk, we perform a binary search.
- **Rearrangement:** Rearrangement is performed on a single list. Rearrangement of a list L_i will cause each element of the list to be pushed to a lower level. As in Figure 19, rearranging the list $L1$ causes Element 10 to be appended to List $L2$ and Element 68 to be pushed to Level 2. The rearrangement of a list in Level 1 will cause each element of the list to be written to the chunk. For example, if list $L3$ is pushed element 27 will be written to chunk $C3$. Rearrangement is performed to ensure that the length of the lists falls within a threshold, and is also to perform efficient batch write operations to the chunks (read-optimized components).

9.1.2 Persistent Memory

Intel®Optane™ DC Persistent Memory (PM, for short), also called Non-Volatile Memory (NVM), is a class of storage that is both byte-addressable as well as persistent at the same time [49]. To guarantee data durability in PM, persistence barrier including cacheline flush plus store fence, is needed. In order to fit the skiplist into PM, the traditional skiplist needs to be modified such that any updates to the skiplist are persisted immediately to the PM storage [92]. Oukid et al. [72] compare a persistent skiplist with a traditional DRAM-based skiplist. They find that both read and write performances degrade in the persistent skiplist because of the high latency of PM and the random access pattern of skiplist. Efforts have been put into designing performant persistent skiplists in recent years [17, 59, 20, 89, 62]. We give an overview of persistent skiplists below.

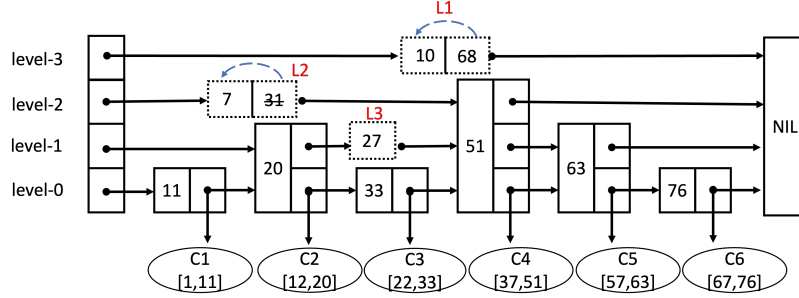


Figure 18: The FlashSkiplist




Figure 19: FlashSkiplist Rearrangement of list L1

NV-skiplist [17]. The NV-skiplist [17] is a persistent skiplist that has been tested on a PM emulator. The NV-skiplist adopts designs that reduce the PM writes: (1) Multiple entries are grouped into one node, and remain unsorted within the node. (2) A bitmap is used to identify empty slots. (3) Selective persistence where only the last level is persisted in PM while the internal levels are stored in DRAM to keep only last level consistent and rebuild the internal levels upon system failure. The NV-skiplist has introduced two optimizations. The first optimization uses the deterministic design [66] to optimize search performance. The second optimization requires the whole key space to be partitioned into several groups to increase scalability.

UPSkiplist [20]. Before diving into the details of UPSkiplist [20], it is essential to introduce RECIPE [59]. RECIPE [59] is an approach to convert DRAM-resident indexes to crash-consistent indexes for PM. These indexes need to follow three conditions: (1) Write operations are visible to other threads with a single hardware-atomic store. (2) Reads and writes are non-blocking; writes fix the inconsistencies. (3) Reads are non-blocking, and writes are blocking as writes do not fix the inconsistencies. However, the lock-free skiplist proposed by Herlihy et al. [46] does not satisfy these conditions as one write operation involves multiple hardware-atomic stores, and writes are non-blocking as well as non-repairing.

The UPSkiplist [20] uses an epoch mechanism to solve this issue. Each node is given an epochID. When this epoch is different from the current epoch, the threads will compete to claim by a hardware-atomic store and fix this inconsistency. Thus, non-blocking-non-repairing operations in the lock-free skiplist proposed by Herlihy et al. [46] are transformed into repairing ones.

Atomic Skiplist (AS) and Atomic and Selective Consistency Skiplist (ASCS) [89]. Write latency is higher than read latency in PM. Moreover, excessive PM writes can result in hardware failure. In order to reduce PM writes, Xiao et al. [89] uses a log-free design and all PM writes are guaranteed to be atomic via the instructions CLFLUSH and MFENCE. During inserts, AS orders the pointer updates such that the skiplist is recoverable without a log: (1) A list L of skiplist pointers that need to be updated is created; (2) A data node is created and flushed to PM; (3) Update the pointer in this data node to point to its successor; (4) Update the pointer of its predecessor; (5) Update the other pointers in L . Since each PM write is atomic, AS is recoverable even if the system crashes during insertion. ASCS relaxes the consistency criterion in that only the bottommost list is guaranteed to be consistent while the other levels are persistent but are not necessarily consistent upon system failure. Upon recovery from a failure, internal levels are rebuilt. This results in speedy recovery compared to [17] where the internal levels are volatile.

PhaST [62]. The (Partitioned Hierarchical Skiplist) [62] (PhaST, for short) is another skiplist that is designed for PM. PhaST lowers the overall height of the skiplist to reduce the random access in between layers. This is achieved by partitioning the data based on their range. Each partition keeps its own skiplist index nodes. PhaST caches frequent keys in DRAM, and groups multiple data entries in one node similar to the other write-optimized skiplists [14]. But

unlike traditional skiplists, PhaST keeps entries within one node unsorted. A node is augmented with a bitmap to indicate the unused slots and a fingerprint array. Writes are concurrent but node splitting is exclusive. In order to protect reads from reading intermediate state, reads need to validate the maximum key of each node after the read operation as the maximum key is the last to be updated during a split process. Experimental results show that PhaST has higher throughput than the NV-skiplist [17].

9.2 NUMA environment

Parallel In-memory Skiplist-based Index (PI) and Parallel In-memory Skip List (PSL) [90, 91]. PI and PSL are highly efficient skiplists that work in NUMA environments. Skiplists have a sorted data layer. PI and PSL partition the data layer into disjoint key ranges, and assign each key range to a NUMA node along with the corresponding index layers. In PI and PSL, each node is an array. Using an array node allows for good spatial locality, and enables the use of SIMD that speeds up comparisons during the search procedure. In addition to using array nodes, queries are also batch-processed. Incoming queries are grouped into a fixed batch size, where each batch is sorted based on the key. After a batch is formed, each query in the batch is routed to the corresponding NUMA node that operates over the key value. PI and PSL are made concurrent and lock-free.

A background thread is dedicated to guarantee load balance within the NUMA environment. In PI and PSL we use a background thread to reconstruct the index layers when the number of deletes and updates reaches a threshold. This maintains the balance in skiplist. PI and PSL perform queries and reconstruction in the following way:

- **Search:** A search operation will be routed to the corresponding NUMA node, and then will perform search in the usual way. However, comparisons are faster with the help of SIMD operations (since both PI and PSL use array nodes).
- **Insert:** An insert operation will be routed to the corresponding NUMA node. Next, the node is only inserted in the data layer with its height attribute determined using Algorithm 1. After the number of inserts exceeds a particular threshold, the background thread will propagate changes to the index layers.
- **Delete:** A delete operation will be routed to the corresponding NUMA node, and will only mark the node as deleted. After the number of logically deleted nodes exceeds a threshold, the background thread will delete the marked node in the data layer, and propagate the changes into the index layers.
- **Range Search:** If the range overlaps multiple NUMA nodes, it is split into different disjoint sub-queries (start point, end point) for the corresponding NUMA nodes, and each subquery is run there. Then results are recombined later.
- **Reconstruction:** After the number of updates and deletes exceeds a threshold, the index layers are discarded. First, all logically deleted nodes are removed from the data layer. Then, each NUMA node shares its data layer with all the other NUMA nodes. After sharing, each NUMA node will reconstruct the index layer using the height attribute of each item in its data layer.
- **Self-adjusted threading:** It is not always the case that the skiplist operations are uniformly spread among the NUMA nodes. Some NUMA nodes can have more queries routed to them. In this case, PI uses self-adjusted threading by spawning threads proportional to the number of queries it has. In the scenario that the number of threads uses all the hardware threads, then the NUMA node will offload some skiplist operations to a remote NUMA node.

Similar to the Cache Sensitive Skiplist, a designated background thread is dedicated to reconstructing the index layer of the skiplist.

The Braided Skiplist [53]. Since NUMA effect is more severe in PM compared to DRAM, ListDB [53] proposes a NUMA-aware braided skiplist that effectively reduces the NUMA effects of PM. The NUMA-oblivious skiplist may have data items scattered across NUMA nodes, and thus searching involves both local and remote NUMA access. The braided skiplist is built on the skiplist invariant that each upper layer is a sub-list of the bottom layer, and the upper layer does not have to be a sub-list of its next layer. The upper layer pointers can only point to data within the same NUMA node; while the bottom layer is connected via local and remote pointers such that any data is guaranteed to be found starting from any skip list.

9.3 Other Hardware

9.3.1 GFSL [67]

GPU performance is impaired by random memory access and execution divergence. Thus, GFSL [67] modifies the skiplist to make it GPU-friendly. Multiple keys are grouped into a chunk that is a cache-aligned. Each chunk or skiplist node is equipped with a lock and a next pointer. Threads are divided into teams, same or smaller than a warp-size. A warp of threads execute the same instruction. The number of entries per chunk is the same as the number of threads in a team. During search, a team of threads reads the entire chunk with each reading one entry. Next, the team of threads uses intra-warp operations to decide to continue the same level or descend to the next level.

9.3.2 Near-memory Processing [18]

Near-memory processing (NMP) is becoming popular to alleviate memory bottleneck. Choe et al. [18] present a hybrid skiplist for near-memory processing. The hybrid structure is managed by two sets of threads: host CPU and near-memory compute unit and the total accessible main memory includes host-accessible memory and NMP-capable memory. The data structure is split into two parts managed by those two sets of threads. With the hierarchical nature of the skiplist, the upper levels are placed in the host-managed portion to have an approximate size of the last-level cache, and the remaining levels are kept in NMP cores. A technique, termed *flat-combining* [43, 19, 63], is used for supporting concurrency in the NMP portion. Flat-combining data structures have a single combiner thread to apply operations on the skiplist. Each concurrent thread posts its operation requests in the publication list. The publication list is used for sending and retrieving requests. The NMP core is single-threaded, and the combiner thread iterates over the list to check for the unserved requests. Insertions are applied to the NMP-managed portion first while deletions are applied to the host-managed portion first.

10 The Multi-dimensional Skiplist

Similar to its tree-based counterparts, skiplists can also be used to index multi-dimensional data. In this section, we discuss skiplists that can be used to index multi-dimensional data. We assume that there are n total points, and each point is in the k -dimensional space.

10.1 Multi-dimensional Skiplist

Skiplists have been extended to support multi-dimensional data [69]. Three versions of the skiplists have been introduced.

- Version 1: This version uses projection to represent all data points. We create k deterministic skiplists, one for each dimension. To insert a point, we simply insert the corresponding value of the i^{th} dimension in the i^{th} skiplist. Each node points to the memory location of the point p .
- Version 2: In the second version, we have k skiplists, each ordered by a dimension, but in addition, each node has k additional dimensional pointers that point to the same node's corresponding location in the other $k - 1$ skiplists, and one-dimensional pointer points to itself.
- Version 3: In this version, we only store a point in only one node instead of storing the same point in k different nodes (once in each skiplist). This considerably reduces memory usage as we do not repeatedly store the points. One disadvantage of this approach is that if the points are not uniformly distributed in all the dimensions, then the multi-dimensional skiplist degenerates into a single skiplist.

10.2 The Skip QuadTree

QuadTrees and OctTrees are famous for indexing multi-dimensional data as they partition the space into quadrants and hypercubes, respectively. However, quadtrees and octTrees can have high depth $O(n)$, and also provide poor update and search times in the worst case. Eppstein et al. [30] combine a skiplist and a compressed quadtree to achieve a balanced structure that can be used to answer point queries, approximate range count, and approximate nearest-neighbor searching queries.

11 The Interval Skiplist

One important variant of the skiplist is the interval skiplist [36]. Given intervals $(a_1, b_1), (a_2, b_2), \dots, (a_n, b_n)$ and a number x , we can find all the intervals that overlap the number x using the interval skiplist in $O(\log n + L)$ time complexity, where n is the number of intervals and L is the number of matching intervals.

Given an interval $I = (a, b)$, the interval skiplist places Nodes a and b into the skiplist. The interval skiplist places marker ‘I’ to represent the interval on forward edges using containment and maximality rules [36]. A containment rule ensures that a marker ‘I’ will be placed on a forward edge going from X to Y , if I contains interval (X, Y) . A maximality rule ensures that there is no forward edge corresponding to an interval (X', Y') that lies within I and contains (X, Y) . Moreover, a node X will have markers for all intervals I , if $X \in I$.

Figure 20 shows the interval skiplist for the three intervals $A = [11, 15]$, $B = [12, 33]$, $C = [21, 56]$. To find all intervals that contain the number 27, we perform a search operation for 27 on the interval skiplist, and compute the union of all the markers in edges along the search path and also the marker of the last node we stopped search at. The search path of 27 is shown in Figure 21. Let S be the set containing all the intervals that overlap the Number 27. Initially, S is empty, when we traverse the forward edge from Node 21 to Node 33, we add B to S , and finally, when we end our search, we add all markers of Node 21 to the set S . The resulting set is $\{B, C\}$, so Number 27 overlaps Intervals $B = [12, 33]$ and $C = [21, 56]$.

The Interval skiplist [36] has been realized in the Ariel Active Database system [37]. Later, the interval skiplist has been shown to perform similarly to tree-based alternatives [38].




Figure 20: Interval Skiplist corresponding to intervals $A = [11, 15]$, $B = [12, 33]$, $C = [21, 56]$




Figure 21: Search Path for finding all intervals that overlap 27

12 The Skiplist as a Priority Queue

A priority queue offers two operations: the Insert, and Deletemin operations. The skiplist is already ordered, and hence can be modified easily to behave as a priority queue. Two variants of the skiplist serve this purpose.

The Lock-Based Priority Queue [81]. In the lock-based priority query [81], the skiplist is extended using lock-based concurrent skiplists [77]. Each node has a flag that is set to false while inserting and is set to true when logically deleted. Inserts are performed in exactly the same way as in the concurrent skiplist. Deletemin is performed by traversing the base linked list and finding the first node that is not marked as deleted. Then, the process locks and marks the node as deleted, and then is physically deleted. This allows extending the skiplist into a priority queue. Because the skiplist is a distributed structure and does not require rebalancing, it is a good candidate to act as a concurrent priority queue with multiple processors.

The SprayList [9]. The sprayList [9] is another priority queue implementation using a lock-free skiplist and a custom “spray” function to achieve a concurrent relaxed priority queue that scales well in high-contention scenarios. The DeleteMin operation in the spraylist returns an element among the first $O(q \log^3 q)$ elements in the list with high probability, where q is the number of threads. This relaxed concurrent priority queue reduces the contention among

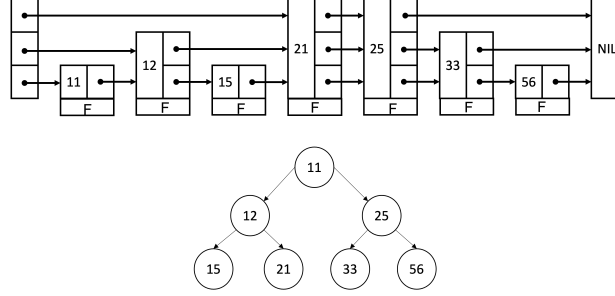


Figure 22: Skiplist and the equivalent priority queue

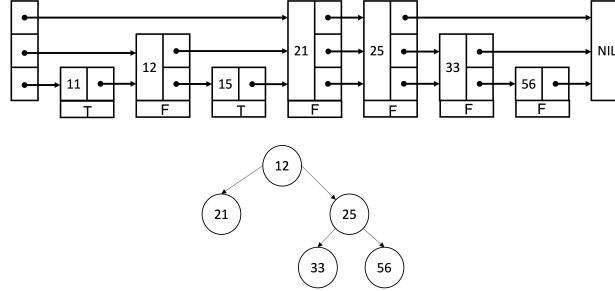


Figure 23: Skiplist with 11 and 15 deleted and the equivalent priority queue

threads to compete for the minimal node. The `DeleteMin` operation uses Operation spray that starts at a starting level, performs a random walk that goes right horizontally and down vertically to the bottom level, to find the returned node. Dummy nodes are padded in the front of the spraylist such that the very first few nodes can also be returned by the spray operation. Detailed description of the spraylist and discussion of its parameters can be found in [9].

13 Skiplist Applications

Because of their appealing features, skiplists have been applied in various systems. They serve a variety of purposes. For example, a skiplist can serve as an in-memory component in big data systems [34, 6, 48, 7, 5, 94, 12, 35]. Also, skiplists can serve as indexes for entire data sets within one machine or across a cluster of machines [4, 55, 53, 47]. Finally, skiplists can be used in network overlay algorithms, e.g., [23, 70, 41, 42, 82, 65], use the deterministic skiplist.

13.1 LSMT-based Systems

Log Structure Merge Trees are used in NoSQL applications and databases, e.g., in the database systems LevelDB [34], RocksDB [29], and FloDB [12]. LSM has multiple levels, each with exponentially increasing memory capacity. All the arriving data, e.g., key-value pairs, are first sent to the LSM's in-memory component. When the in-memory component gets filled, the key-value pairs are compacted into a block format, and then is flushed into permanent storage afterwards. The choice of data structure for the memory component of many LSM realizations is the skiplist, as it is a lightweight data structure that maintains the data in sorted order. Operations, e.g., insert, delete, find, and range search can be performed in $O(\log n)$ time. Once the main-memory component that stores the skiplist becomes full, we compact this main-memory component that is already sorted into a block to be persisted in stable storage. LevelDB and RocksDB use Probabilistic Skiplists. In RocksDB, skiplists support concurrent inserts and lookups.

13.1.1 FloDB [12]

An LSM data store has an in-memory component and a disk component. All the operations, e.g., inserts, updates, and scans will first operate on the in-memory component and will then proceed into the subsequent levels on stable storage, e.g., on disk. Popular LSM implementations, e.g., LevelDB [34] and RocksDB [6] use a skiplist as their in-memory component. FloDB [12] creates a custom two-level in-memory structure that improves performance with the increase in main memory size.

The in-memory component can either be sorted, e.g., as in the case of a skiplist, or unsorted, e.g., as in the case of a hashtable. Experiments [12] show that the hashtable is very good at dealing with write heavy workloads, whereas skiplists are very good at dealing with read- and scan-heavy workloads. FloDB’s main contribution is the use of a concurrent hash-table, termed the membuffer, and a skiplist, termed the memtable. This two-level structure in main memory gets the best of both data structures.

In FloDB, the membuffer (hash-table) is smaller in size compared to the memtable (skiplist). The hash-table takes the l -most significant bits as a hash function. Whenever the membuffer is full, it is compacted into the memtable. And when the memtable is full, it is flushed into disk. The point lookup, insert, and range search operations are described below:

- Point lookup (key): Point lookup first checks the membuffer, and if it is not found then we search the memtable and then in the disk.
- Insert(Key, Value): Inserts take place first in membuffer, and if the membuffer is full, then inserts are directed into the memtable.
- Range search (lo, hi): Since range search operations are faster in a skiplist, we by default flush the data from the membuffer into the skiplist and then perform the scan there.

In addition to the above operations, FloDB allows bulk inserts that insert a batch of key-value pairs in ascending order into the skiplist. FloDB [12] stores the predecessor node’s path in order to speed up insertion time for a node. It takes advantage of neighborhood proximity in order to improve speed. Moreover, it supports concurrency.

13.1.2 S3 [94]

Similar to FloDB, S3 proposes a novel in-memory data structure that improves the efficiency of the overall database. S3 can seamlessly be integrated into LevelDB and RocksDB, and is comparable in speed to Adaptive Radix Tree (ART) [46]. S3 uses Fast Architecture Sensitive Tree (FAST) [52]; an index tree that fits in the cache on top of a semi-ordered skiplist. In addition, it uses neural networks to predict the guard entries and achieves comparable throughput to that of ART. Search, insert, and compaction operations are described below:

- Search (key): Search first proceeds to the root of FAST, and finally reaches a guard entry in a skiplist. Then, it proceeds to find the guard range in which the searched key is located using the regular search procedure of the skiplist.
- Insert(Key, Value): After searching for the proper place to insert an item, it is placed within its guard range.
- Compaction: During compaction, we get a data distribution and query distribution. These are fed into the neural network to predict good guard entries for FAST and the skiplist data structures.

13.1.3 TeksDB [35]

Han et al. [35] observe that no single data structure can achieve both fast point access and fast range search at the same time. Also, simple combination of multiple data structures that are each optimal for one operation cannot work [35]. Thus, TeksDB proposes a hybrid data structure by “judiciously weaving” two complementary structures: a CuckooHash [74] and a skiplist. CuckooHash is responsible for point lookups and updates while the skiplist serves range searching. To save space and avoid duplicating data among the two data structures, the key-value pairs are kept in a pool, and the two data structures only keep pointers to this pool.

13.2 Skiplist-based Key-Value Stores

In this section, we highlight several already deployed key-value store systems that heavily utilize skiplists to map the keys to their corresponding values. The advantages of the skiplist are that they keep the keys sorted, and hence facilitate fast search, as well as allow for fast inserts and deletes as well as range searches. Moreover, they allow for high concurrency. Specifics of how certain key-value store systems make use of the skiplist data structure is given below.

13.2.1 PebblesDB [78]

As the height of each inserted node is determined via randomness, the skiplist does not require rebalancing, which is different from the B^+ tree. *PebblesDB* [78] combines the skiplist with Log-Structured Merge Trees (LSMT) in the context of the fragmented LSMT (FLSM). Ordinary LSMT flushes memory buffer to disk files, termed *sstables* and compacts *sstables* in a hierarchical way [73]. Since *sstable* is immutable, the same data can be written multiple times

through its lifetime during compactions, thus data write amplification is high [78]. FLSM adopts the probabilistic design of the skiplist in choosing which keys from the inserted keys in each level to split that level’s key space. These chosen keys are termed *guards* as no files can span across guards. Each guard has its associated sstables. Guards are chosen from the inserted keys in a similar way as to that of the skiplist. This design alleviates the issue of write amplification in the LSMT [78].

13.2.2 RS-store [47]

Remote Direct Memory Access (RDMA, for short) is a high-speed network communication mechanism with low latency that allows direct access to memory in remote nodes [50]. RS-store [47] is a skiplist-based key-value store using RDMA. RS-store’s skiplist, termed the R-skiplist, follows a block-based approach, where it groups multiple skiplist entries into one block. This saves communication cost as one communication can fetch a block instead of one node of the skiplist. For increased scalability, concurrency control is required on the server-side as well as between the client and the server in the R-skiplist. On the server side, an exclusive access strategy is adopted to omit latches. The R-skiplist and tasks are partitioned into several partitions, where each partition is accessed exclusively by one thread. Write operations that are only performed on the server side are serialized within one partition. Since partitions are generated based on the top level nodes, more than one thread can visit the same top level block causing conflicts. It is resolved by a latch-free reference counting technique (Refer to [47] for further detail). For controlling concurrency between the server and the client, remote read atomicity is guaranteed by a verification-based technique. A flag on the block is set before every write operation, and a remote read verifies the flag before reading the fetched block. The R-skiplist supports range search without copying and sending data from the server to the client. After the client issues a range search request, the server returns once with the starting address of the block, and the client can issue remote RDMA reads to retrieve the next blocks.

13.2.3 ListDB [53]

ListDB [53] proposes a skiplist that is aware of the Non-Uniform Memory Access architecture (NUMA, for short) [56], i.e., a NUMA-aware skiplist on persistent non-volatile memory (PM, for short) [49], and uses a skiplist to build a key-value store. ListDB is a three-level LSMT with one volatile level in DRAM and L0 and L1 levels on PM. During data flushing from DRAM to L0, ListDB restructures the log entries in PM as a skiplist. Thus, write amplification is reduced. Once the pointers between skiplist nodes are persisted via `clflush` (details in Section 9.1.2), this skiplist is now in L0. Later, this L0 skiplist can be merge-sorted with the existing skiplist in L1 in-place.

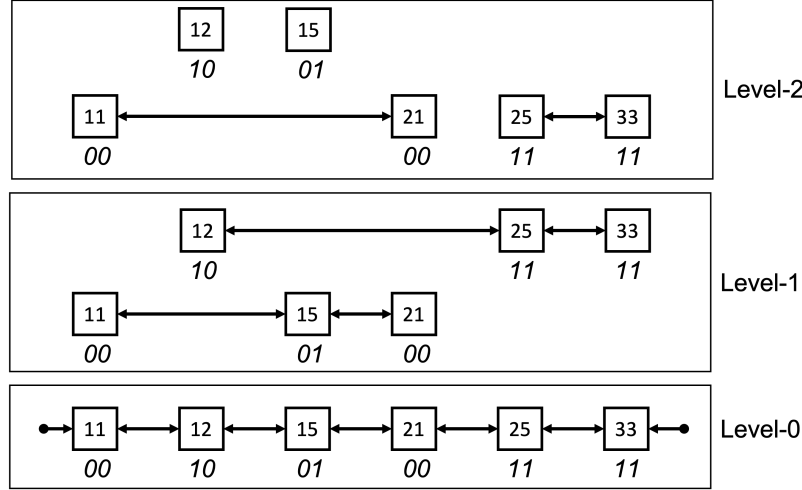
13.3 Network Overlay algorithms

Due to the balanced nature and the ability to efficiently search and update the skiplist, it is widely used in network overlay algorithms. Tiara [23] uses a sparse 0-1 distributed skiplist, and applies a self-stabilizing algorithm for the skiplist. Nor et al. propose the Corona algorithm [70] that constructs a self-stabilizing deterministic skiplist in message-passing systems. A distributed deterministic 1-2 skiplist has been proposed [65] and has been later extended [82] to support concurrency. Skiplists have also been used in asynchronous video streaming [86] and VCR interactions [87] because of their distributed and balanced nature.

The skip graph [10] is a distributed data structure based on the skiplist. The list in a skip graph is doubly-linked. With N data items, the number of levels is $O(\log N)$. However, in a skip graph, there can be many lists at one level. The bottom level (Level-0) contains all data items in a doubly linked list. The upper level contains the same amount of data items but links fewer data items in one linked list. Each node participates in one of the lists until the node becomes a singleton node that is not linked with any other nodes after $O(\log N)$ level elevation on average. Each data item x is given a random word of membership vector $m(x)$. For example, in Figure 24, $m(15) = 01$. At Level- i , all items that share the same prefix of Length i are in the same list, e.g., 11, 15, and 21 have Prefix 0 at Level-1 while 11 and 21 have Prefix 00 at Level-2. Search and insert are similar as in the skiplist. The skip graph is often used in peer-to-peer networks. It supports range searches while distributed hashtables cannot. The skip graph exhibits logarithmic insert and delete performance, and can tolerate node failures without becoming disconnected. More detailed analysis and discussion can be found in [10].

13.4 Hierarchical Navigable Small World Graphs

Hierarchical Navigable Small World (HNSW) is used for K-Approximate Nearest Neighbour Search (K-ANNS) [64]. This algorithm can be seen as an extension to the skiplist with proximity graphs replacing the linked list component in each level [64].

Figure 24: A skip graph with $\lceil O(\log N) \rceil = 3$ levels.

14 Concluding Remarks

The skiplist has become ubiquitous, and is being used extensively in big data systems. This paper presents a survey of many variants of skiplists, and addresses the motivation behind each variant and its tradeoffs and its potential applications. We have surveyed new skiplist variants, their motivation, and highlighted how they accommodate new functions, and how they apply in various domains. We have presented a taxonomy of skiplist variants and the related big data systems that make use of these variants. We have reviewed the existing systems that use skiplists. We hope that this survey serve as a guide to researchers, practitioners, and end-users.

15 Acknowledgements

Walid G. Aref acknowledges the support of the National Science Foundation under Grant Number IIS-1910216.

References

- [1] A kernel skiplist implementation, Accessed: 2023-10-09.
- [2] The muqss cpu scheduler, Accessed: 2023-10-09.
- [3] skiplists for range indexes, Accessed: 2023-10-09.
- [4] Singlestore, Accessed: 2023-10-10.
- [5] Apache hbase, Accessed: 2024-01-09.
- [6] Meta, Accessed: 2024-01-09.
- [7] Redis, Accessed: 2024-01-09.
- [8] AKSENOV, V., ALISTARH, D., DROZDOVA, A., AND MOHTASHAMI, A. The splay-list: A distribution-adaptive concurrent skip-list. *Distributed Computing* (2023), 1–24.
- [9] ALISTARH, D., KOPINSKY, J., LI, J., AND SHAVIT, N. The spraylist: A scalable relaxed priority queue. In *Proceedings of the 20th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming* (2015), pp. 11–20.
- [10] ASPNES, J., AND SHAH, G. Skip graphs. *Acm transactions on algorithms (talg)* 3, 4 (2007), 37–es.
- [11] BAGCHI, A., BUCHSBAUM, A. L., AND GOODRICH, M. T. Biased skip lists. *Algorithmica* 42 (2005), 31–48.
- [12] BALMAU, O., GUERRAUI, R., TRIGONAKIS, V., AND ZABLOTCHI, I. Flodb: Unlocking memory in persistent key-value stores. In *Proceedings of the Twelfth European Conference on Computer Systems* (2017), pp. 80–94.
- [13] BAYER, R., AND MCCREIGHT, E. Organization and maintenance of large ordered indices. In *Proceedings of the 1970 ACM SIGFIDET (Now SIGMOD) Workshop on Data Description, Access and Control* (1970), pp. 107–141.

- [14] BENDER, M. A., FARACH-COLTON, M., JOHNSON, R., MAURAS, S., MAYER, T., PHILLIPS, C. A., AND XU, H. Write-optimized skip lists. In *Proceedings of the 36th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems* (2017), pp. 69–78.
- [15] BERENSON, H., BERNSTEIN, P., GRAY, J., MELTON, J., O’NEIL, E., AND O’NEIL, P. A critique of ansi sql isolation levels. *ACM SIGMOD Record* 24, 2 (1995), 1–10.
- [16] BERNSTEIN, P. A., AND GOODMAN, N. Multiversion concurrency control—theory and algorithms. *ACM Transactions on Database Systems (TODS)* 8, 4 (1983), 465–483.
- [17] CHEN, Q., LEE, H., KIM, Y., YEOM, H. Y., AND SON, Y. Design and implementation of skiplist-based key-value store on non-volatile memory. *Cluster Computing* 22 (2019), 361–371.
- [18] CHOE, J., CROTTY, A., MORESHET, T., HERLIHY, M., AND BAHAR, R. I. Hybrids: Cache-conscious concurrent data structures for near-memory processing architectures. In *Proceedings of the 34th ACM Symposium on Parallelism in Algorithms and Architectures* (2022), pp. 321–332.
- [19] CHOE, J., HUANG, A., MORESHET, T., HERLIHY, M., AND BAHAR, R. I. Concurrent data structures with near-data-processing: An architecture-aware implementation. In *The 31st ACM Symposium on Parallelism in Algorithms and Architectures* (2019), pp. 297–308.
- [20] CHOWDHURY, S., AND GOLAB, W. A scalable recoverable skip list for persistent memory. In *Proceedings of the 33rd ACM Symposium on Parallelism in Algorithms and Architectures* (2021), pp. 426–428.
- [21] CIRIANI, V., FERRAGINA, P., LUCCIO, F., AND MUTHUKRISHNAN, S. Static optimality theorem for external memory string access. In *The 43rd Annual IEEE Symposium on Foundations of Computer Science, 2002. Proceedings.* (2002), IEEE, pp. 219–227.
- [22] CIRIANI, V., FERRAGINA, P., LUCCIO, F., AND MUTHUKRISHNAN, S. A data structure for a sequence of string accesses in external memory. *ACM Transactions on Algorithms (TALG)* 3, 1 (2007), 1–23.
- [23] CLOUSER, T., NESTERENKO, M., AND SCHEIDELER, C. Tiara: A self-stabilizing deterministic skip list. In *Stabilization, Safety, and Security of Distributed Systems: 10th International Symposium, SSS 2008, Detroit, MI, USA, November 21-23, 2008. Proceedings 10* (2008), Springer, pp. 124–140.
- [24] COMER, D. Ubiquitous b-tree. *ACM Computing Surveys (CSUR)* 11, 2 (1979), 121–137.
- [25] CRAIN, T., GRAMOLI, V., AND RAYNAL, M. No hot spot non-blocking skip list. In *2013 IEEE 33rd International Conference on Distributed Computing Systems* (2013), IEEE, pp. 196–205.
- [26] DAN, A., AND TOWSLEY, D. An approximate analysis of the lru and fifo buffer replacement schemes. In *Proceedings of the 1990 ACM SIGMETRICS conference on Measurement and modeling of computer systems* (1990), pp. 143–152.
- [27] DEAN, B. C., AND JONES, Z. H. Exploring the duality between skip lists and binary search trees. In *Proceedings of the 45th annual southeast regional conference* (2007), pp. 395–399.
- [28] DICK, I., FEKETE, A., AND GRAMOLI, V. A skip list for multicore. *Concurrency and Computation: Practice and Experience* 29, 4 (2017), e3876. e3876 cpe.3876.
- [29] DONG, S., KRYCZKA, A., JIN, Y., AND STUMM, M. Evolution of development priorities in key-value stores serving large-scale applications: The {RocksDB} experience. In *19th USENIX Conference on File and Storage Technologies (FAST 21)* (2021), pp. 33–49.
- [30] EPPSTEIN, D., GOODRICH, M. T., AND SUN, J. Z. The skip quadtree: a simple dynamic data structure for multidimensional data. In *Proceedings of the twenty-first annual symposium on Computational geometry* (2005), pp. 296–305.
- [31] ERGUN, F., CENK ŞAHINALP, S., SHARP, J., AND SINHA, R. K. Biased skip lists for highly skewed access patterns. In *Algorithm Engineering and Experimentation: Third International Workshop, ALENEX 2001 Washington, DC, USA, January 5–6, 2001 Revised Papers 3* (2001), Springer, pp. 216–229.
- [32] FOMITCHEV, M., AND RUPPERT, E. Lock-free linked lists and skip lists. In *Proceedings of the twenty-third annual ACM symposium on Principles of distributed computing* (2004), pp. 50–59.
- [33] FRASER, K. Practical lock-freedom. Tech. rep., University of Cambridge, Computer Laboratory, 2004.
- [34] GHEMAWAT, S., AND DEAN, J. Leveldb, 2011.
- [35] HAN, Y., KIM, B. S., YEON, J., LEE, S., AND LEE, E. Tekstdb: weaving data structures for a high-performance key-value store. *Proceedings of the ACM on Measurement and Analysis of Computing Systems* 3, 1 (2019), 1–23.
- [36] HANSON, E. N. The interval skip list: A data structure for finding all intervals that overlap a point. In *Workshop on Algorithms and Data Structures* (1991), Springer, pp. 153–164.

- [37] HANSON, E. N. The design and implementation of the ariel active database rule system. *IEEE Transactions on Knowledge and Data Engineering* 8, 1 (1996), 157–172.
- [38] HANSON, E. N., AND JOHNSON, T. Selection predicate indexing for active databases using interval skip lists. *Information Systems* 21, 3 (1996), 269–298.
- [39] HARRIS, T., LARUS, J., AND RAJWAR, R. *Transactional memory*. Springer Nature, 2022.
- [40] HARRIS, T. L. A pragmatic implementation of non-blocking linked-lists. In *International Symposium on Distributed Computing* (2001), Springer, pp. 300–314.
- [41] HARVEY, N. J., DUNAGAN, J., JONES, M., SAROIU, S., THEIMER, M., AND WOLMAN, A. Skipnet: A scalable overlay network with practical locality properties.
- [42] HARVEY, N. J., AND MUNRO, J. I. Brief announcement: Deterministic skipnet. In *Proceedings of the twenty-second annual symposium on Principles of distributed computing* (2003), pp. 152–152.
- [43] HENDLER, D., INCZE, I., SHAVIT, N., AND TZAFRIR, M. Flat combining and the synchronization-parallelism tradeoff. In *Proceedings of the twenty-second annual ACM symposium on Parallelism in algorithms and architectures* (2010), pp. 355–364.
- [44] HERLIHY, M., LEV, Y., LUCHANGCO, V., AND SHAVIT, N. A provably correct scalable concurrent skip list. In *Conference On Principles of Distributed Systems (OPODIS)*. Citeseer (2006), p. 103.
- [45] HERLIHY, M., LEV, Y., LUCHANGCO, V., AND SHAVIT, N. A simple optimistic skiplist algorithm. In *Structural Information and Communication Complexity: 14th International Colloquium, SIROCCO 2007, Castiglione, Italy, June 5-8, 2007. Proceedings 14* (2007), Springer, pp. 124–138.
- [46] HERLIHY, M., SHAVIT, N., LUCHANGCO, V., AND SPEAR, M. *The art of multiprocessor programming*. Newnes, 2020.
- [47] HUANG, C., HU, H., QI, X., ZHOU, X., AND ZHOU, A. Rs-store: Rdma-enabled skiplist-based key-value store for efficient range query. *Frontiers of Computer Science* 15, 6 (2021), 156617.
- [48] HUANG, G., CHENG, X., WANG, J., WANG, Y., HE, D., ZHANG, T., LI, F., WANG, S., CAO, W., AND LI, Q. X-engine: An optimized storage engine for large-scale e-commerce transaction processing. In *Proceedings of the 2019 International Conference on Management of Data* (2019), pp. 651–665.
- [49] IZRAELEVITZ, J., YANG, J., ZHANG, L., KIM, J., LIU, X., MEMARIPOUR, A., SOH, Y. J., WANG, Z., XU, Y., DULLOOR, S. R., ET AL. Basic performance measurements of the intel optane dc persistent memory module. *arXiv preprint arXiv:1903.05714* (2019).
- [50] KALIA, A., KAMINSKY, M., AND ANDERSEN, D. G. Design guidelines for high performance {RDMA} systems. In *2016 USENIX Annual Technical Conference (USENIX ATC 16)* (2016), pp. 437–450.
- [51] KARDARAS, M., SIAKAVARAS, D., NIKAS, K., GOUMAS, G., AND KOZIRIS, N. Fast concurrent skip lists with htm, 2018.
- [52] KIM, C., CHHUGANI, J., SATISH, N., SEDLAR, E., NGUYEN, A. D., KALDEWEY, T., LEE, V. W., BRANDT, S. A., AND DUBEY, P. Fast: fast architecture sensitive tree search on modern cpus and gpus. In *Proceedings of the 2010 ACM SIGMOD International Conference on Management of Data* (New York, NY, USA, 2010), SIGMOD ’10, Association for Computing Machinery, p. 339–350.
- [53] KIM, W., PARK, C., KIM, D., PARK, H., CHOI, Y.-R., SUSSMAN, A., AND NAM, B. {ListDB}: Union of {Write-Ahead} logs and persistent {SkipLists} for incremental checkpointing on persistent memory. In *16th USENIX Symposium on Operating Systems Design and Implementation (OSDI 22)* (2022), pp. 161–177.
- [54] KOBUS, T., KOKOCIŃSKI, M., AND WOJCIECHOWSKI, P. T. Jiffy: a lock-free skip list with batch updates and snapshots. In *Proceedings of the 27th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming* (2022), pp. 400–415.
- [55] LAKSHMAN, S., MELKOTE, S., LIANG, J., AND MAYURAM, R. Nitro: a fast, scalable in-memory storage engine for nosql global secondary index. *Proceedings of the VLDB Endowment* 9, 13 (2016), 1413–1424.
- [56] LAMETER, C. Numa (non-uniform memory access): An overview: Numa becomes more common because memory controllers get close to execution units on microprocessors. *Queue* 11, 7 (2013), 40–51.
- [57] LAMOUREUX, M., NICKERSON, B., ET AL. On the equivalence of b-trees and deterministic skip list.
- [58] LEA, D. Concurrentskiplistmap, Accessed on Jan 7, 2024.
- [59] LEE, S. K., MOHAN, J., KASHYAP, S., KIM, T., AND CHIDAMBARAM, V. Recipe: Converting concurrent dram indexes to persistent-memory indexes. In *Proceedings of the 27th ACM Symposium on Operating Systems Principles* (2019), pp. 462–477.

- [60] LEHMAN, P. L., AND YAO, S. B. Efficient locking for concurrent operations on b-trees. *ACM Transactions on Database Systems (TODS)* 6, 4 (1981), 650–670.
- [61] LEIS, V., KEMPER, A., AND NEUMANN, T. The adaptive radix tree: Artful indexing for main-memory databases. In *2013 IEEE 29th International Conference on Data Engineering (ICDE)* (2013), IEEE, pp. 38–49.
- [62] LI, Z., JIAO, B., HE, S., AND YU, W. Phast: Hierarchical concurrent log-free skip list for persistent memory. *IEEE Transactions on Parallel and Distributed Systems* 33, 12 (2022), 3929–3941.
- [63] LIU, Z., CALCIU, I., HERLIHY, M., AND MUTLU, O. Concurrent data structures for near-memory computing. In *Proceedings of the 29th ACM Symposium on Parallelism in Algorithms and Architectures* (2017), pp. 235–245.
- [64] MALKOV, Y. A., AND YASHUNIN, D. A. Efficient and robust approximate nearest neighbor search using hierarchical navigable small world graphs. *IEEE transactions on pattern analysis and machine intelligence* 42, 4 (2018), 824–836.
- [65] MANDAL, S., CHAKRABORTY, S., AND KARMAKAR, S. Deterministic 1–2 skip list in distributed system. In *2012 2nd IEEE International Conference on Parallel, Distributed and Grid Computing* (2012), IEEE, pp. 296–301.
- [66] MEI, F., CAO, Q., WU, F., AND LI, H. A concurrent skip list balanced on search. In *Advanced Parallel Processing Technologies: 12th International Symposium, APPT 2017, Santiago de Compostela, Spain, August 29, 2017, Proceedings 12* (2017), Springer, pp. 117–128.
- [67] MOSCOVICI, N., COHEN, N., AND PETRANK, E. A gpu-friendly skiplist algorithm. In *2017 26th International Conference on Parallel Architectures and Compilation Techniques (PACT)* (2017), Ieee, pp. 246–259.
- [68] MUNRO, J. I., PAPADAKIS, T., AND SEDGEWICK, R. Deterministic skip lists. In *Proceedings of the third annual ACM-SIAM symposium on Discrete algorithms* (1992), pp. 367–375.
- [69] NICKERSON, B. G. Skip list data structures for multidimensional data. Tech. rep., 1998.
- [70] NOR, R. M., NESTERENKO, M., AND SCHEIDELER, C. Corona: A stabilizing deterministic message-passing skip list. *Theoretical Computer Science* 512 (2013), 119–129.
- [71] O’NEIL, E. J., O’NEIL, P. E., AND WEIKUM, G. The lru-k page replacement algorithm for database disk buffering. *Acm Sigmod Record* 22, 2 (1993), 297–306.
- [72] OUKID, I., LEHNER, W., KISSINGER, T., WILLHALM, T., AND BUMBULIS, P. Instant recovery for main memory databases. In *CIDR* (2015).
- [73] O’NEIL, P., CHENG, E., GAWLICK, D., AND O’NEIL, E. The log-structured merge-tree (lsm-tree). *Acta Informatica* 33, 4 (1996), 351–385.
- [74] PAGH, R., AND RODLER, F. F. Cuckoo hashing. *Journal of Algorithms* 51, 2 (2004), 122–144.
- [75] PLATZ, K., MITTAL, N., AND VENKATESAN, S. Concurrent unrolled skiplist. In *2019 IEEE 39th International Conference on Distributed Computing Systems (ICDCS)* (2019), IEEE, pp. 1579–1589.
- [76] PUGH, W. Skip lists: a probabilistic alternative to balanced trees. *Communications of the ACM* 33, 6 (1990), 668–676.
- [77] PUGH, W. Concurrent maintenance of skip lists. Tech. rep., 1998.
- [78] RAJU, P., KADEKODI, R., CHIDAMBARAM, V., AND ABRAHAM, I. Pebblesdb: Building key-value stores using fragmented log-structured merge trees. In *Proceedings of the 26th Symposium on Operating Systems Principles* (2017), pp. 497–514.
- [79] REED, D. P. Naming and synchronization in a decentralized computer system.
- [80] SHANNON, C. E. A mathematical theory of communication. *The Bell system technical journal* 27, 3 (1948), 379–423.
- [81] SHAVIT, N., AND LOTAN, I. Skiplist-based concurrent priority queues. In *Proceedings 14th International Parallel and Distributed Processing Symposium. IPDPS 2000* (2000), IEEE, pp. 263–268.
- [82] SINGH, R., CHAKRABORTY, S., AND KARMAKAR, S. Concurrent deterministic 1–2 skip list in distributed message passing systems. *International Journal of Parallel, Emergent and Distributed Systems* 30, 2 (2015), 135–174.
- [83] SPRENGER, S., ZEUCH, S., AND LESER, U. Cache-sensitive skip list: Efficient range queries on modern cpus. In *International Workshop on Accelerating Data Analysis and Data Management Systems, International Workshop on In-Memory Data Management and Analytics* (2017), Springer, pp. 1–17.
- [84] SUNDELL, H., AND TSIGAS, P. Fast and lock-free concurrent priority queues for multi-thread systems. *Journal of Parallel and Distributed Computing* 65, 5 (2005), 609–627.

- [85] TANENBAUM, A. S., AND WOODHULL, A. S. *Operating systems: design and implementation*, vol. 68. Prentice Hall Englewood Cliffs, 1997.
- [86] WANG, D., AND LIU, J. Peer-to-peer asynchronous video streaming using skip list. In *2006 IEEE International Conference on Multimedia and Expo* (2006), IEEE, pp. 1397–1400.
- [87] WANG, D., AND LIU, J. A dynamic skip list-based overlay for on-demand media streaming with vcr interactions. *IEEE Transactions on Parallel and Distributed Systems* 19, 4 (2008), 503–514.
- [88] WANG, H., AND FENG, J. Flashskiplist: indexing on flash devices. In *Proceedings of the ACM Turing 50th Celebration Conference-China* (2017), pp. 1–10.
- [89] XIAO, R., FENG, D., HU, Y., WANG, F., WEI, X., ZOU, X., AND LEI, M. Write-optimized and consistent skiplists for non-volatile memory. *IEEE Access* 9 (2021), 69850–69859.
- [90] XIE, Z., CAI, Q., JAGADISH, H., OOI, B. C., AND WONG, W.-F. Pi: a parallel in-memory skip list based index. *arXiv preprint arXiv:1601.00159* (2016).
- [91] XIE, Z., CAI, Q., JAGADISH, H., OOI, B. C., AND WONG, W.-F. Parallelizing skip lists for in-memory multi-core database systems. In *2017 IEEE 33rd International Conference on Data Engineering (ICDE)* (2017), IEEE, pp. 119–122.
- [92] YANG, J., KIM, J., HOSEINZADEH, M., IZRAELEVITZ, J., AND SWANSON, S. An empirical guide to the behavior and use of scalable persistent memory. In *18th USENIX Conference on File and Storage Technologies (FAST 20)* (2020), pp. 169–182.
- [93] YEON, J., KIM, L., HAN, Y., LEE, H. G., LEE, E., AND KIM, B. S. Jellyfish: A fast skip list with mvcc. In *Proceedings of the 21st International Middleware Conference* (2020), pp. 134–148.
- [94] ZHANG, J., WU, S., TAN, Z., CHEN, G., CHENG, Z., CAO, W., GAO, Y., AND FENG, X. S3: a scalable in-memory skip-list index for key-value store. *Proceedings of the VLDB Endowment* 12, 12 (2019), 2183–2194.
- [95] ZHOU, J., AND ROSS, K. A. Implementing database operations using simd instructions. In *Proceedings of the 2002 ACM SIGMOD International Conference on Management of Data* (New York, NY, USA, 2002), SIGMOD ’02, Association for Computing Machinery, p. 145–156.