View metadata, citation and similar papers at core.ac.uk
brought to you by CORE
provided by Elsevier - Publisher Connector

Journal of Discrete Algorithms 14 (2012) 13–20

Contents lists available at SciVerse ScienceDirect

Journal of Discrete Algorithms

www.elsevier.com/locate/jda

# Skip lift: A probabilistic alternative to red–black trees ☆

Prosenjit Bose, Karim Douïeb *, Pat Morin

*School of Computer Science, Carleton University, Herzberg Building, 1125 Colonel By Drive, Ottawa, Ontario, K1S 5B6 Canada*

A R T I C L E   I N F O

*Article history:*
Available online 8 December 2011

*Keywords:*
Data structure
Dictionary
Skip list
Finger search

A B S T R A C T

We present the *Skip lift*, a randomized dictionary data structure inspired by the skip list [Pugh'90, Comm. of the ACM]. Similar to the skip list, the skip lift has the finger search property: given a pointer to an arbitrary element $f$, searching for an element $x$ takes expected $O(\log \delta)$ time where $\delta$ is the rank distance between the elements $x$ and $f$. The skip lift uses nodes of $O(1)$ worst-case size (for a total of $O(n)$ worst-case space usage) and it is one of the few efficient dictionary data structures that performs an $O(1)$ worst-case number of structural changes (pointers/fields modifications) during an update operation. Given a pointer to the element to be removed from the skip lift the deletion operation takes $O(1)$ worst-case time.

© 2011 Elsevier B.V. All rights reserved.

## 1. Introduction

The dictionary problem is fundamental in computer science. It asks for a data structure in the pointer machine model that stores a totally ordered set $S$ of $n$ elements and supports the operations search, insert and delete. A large number of data structures optimally solve this problem in worst-case $O(\log n)$ time per operation. Some of them guarantee an $O(1)$ worst-case number of structural changes (pointers/fields modifications) after an insertion or a deletion operation [12,19,11,13,10,6]. Note that a structural change takes $O(1)$ time.

Typically the update operations that is insert and delete, are performed in two phases: first, search for the position where the update has to take place. Second, perform the actual update and restore the balance of the structure. When the position where the new element has to be inserted or deleted is already known then the first phase of an update could be avoided. In general the first phase is considered to be part of the search operation. A dictionary that guarantees an $O(1)$ worst-case number of structural changes per update does not necessary quickly perform the second phase of the update. Much research effort has been aimed at improving the worst-case time taken by the second phase of the update: Levcopoulos and Overmars [13] presented the first search tree that takes $O(1)$ worst-case time for this second phase of the update. Later Fleischer [10] simplified this result. Brodal et al. [6] additionally guaranteed that such structures can also have the finger search property in worst-case time. These structures however are quite complicated and not really practical.

On the other hand, most randomized dictionaries are simple, practical and achieve the same performance as the result of Brodal et al. [6] in the expected sense. In the worst-case though their performance is far from optimal. Here we develop a simple randomized dictionary, called a skip lift, inspired by the skip list [18], that improves the worst-case performance of the second phase of the update operations. Namely we obtain a structure that has the finger search property in expectation and performs an $O(1)$ worst-case number of structural changes per update. Given a pointer to the element to be removed from the skip lift, the deletion operation takes $O(1)$ worst-case time.

☆ Research partially supported by NSERC and MRI.
* Corresponding author.
*E-mail addresses:* jit@cg.scs.carleton.ca (P. Bose), karim@cg.scs.carleton.ca (K. Douïeb), morin@cg.scs.carleton.ca (P. Morin).
*URL:* http://cg.scs.carleton.ca (K. Douïeb).

1570-8667/$ – see front matter © 2011 Elsevier B.V. All rights reserved.
doi:10.1016/j.jda.2011.12.017

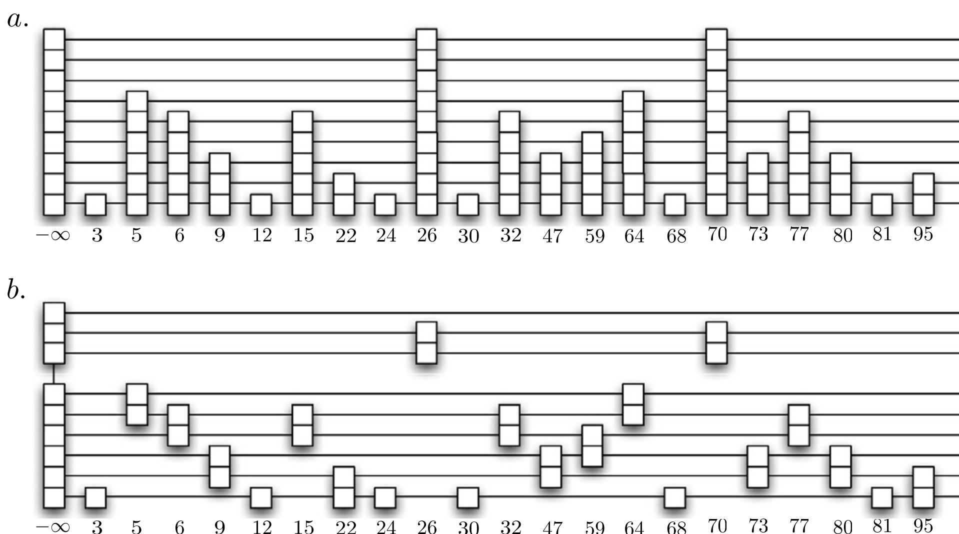


**Fig. 1.** a. Skip list, b. Skip lift.

In Section 1.1 we describe the original skip list dictionary. In Section 1.2 we mention some work related to the skip list dictionary. In Section 2 we introduce our new skip lift data structure. In Section 3 we show how to enhance the skip lift structure to allow a simple finger search. Finally, in Section 4, we give an overview of some classical randomized dictionary data structures. For each of them we briefly describe its construction and how the dictionary operations are performed. We show that, for these classical randomized dictionaries, in some situations $\Omega(n)$ structural changes are necessary to perform the update operations.

### 1.1. Skip list

The *skip list* of Pugh [18] was introduced as a probabilistic alternative to balanced trees. It is a dictionary data structure storing a totally ordered set $S$ of $n$ elements that supports insertion, deletion and search operations in $O(\log n)$ expected time. Additionally the expected number of structural changes (pointer modifications) performed on the skip list during an update is $O(1)$. A skip list is built in levels, the bottom level (level 1) is a sorted linked list of all elements in $S$. The higher levels of the skip list are build iteratively. Each level is a sublist of the previous one where each element of a level is copied to the level above with (independent) probability $p$. The copies of an element are linked between adjacent levels (see Fig. 1.a).

The *height* $h(s)$ of an element $s$ is defined as the highest level where $s$ appears. The height $H(\mathcal{L})$ of a skip list $\mathcal{L}$ is defined as $\max_{s\in\mathcal{L}} h(s)$ and the *depth* $d(s)$ of $s$ is $H(\mathcal{L}) - h(s)$. The expected height of a skip list is by definition $O(\log_{1/p} n)$. Adjacent elements on the same level are connected by their left and right pointers. The copies of the same element from two adjacent levels are connected by their up and down pointers.

#### 1.1.1. Search

To search for a given element $x$ in a skip list we start from the highest level of the sentinel element which has a key value $-\infty$. We follow the right pointers on a same level until we are about to overshoot the element $x$ that is until the element on the right has a key value strictly greater than $x$. Then we go down one level and we iterate the process until $x$ is found or until we have reached the lowest level (in this case we know that $x$ is not in $S$ and we have found its predecessor).

#### 1.1.2. Updates

To insert an element $x$ in a skip list we first determine its height in the structure. Then we start a search for $x$ in the list to find the position where $x$ has to be inserted. During the search we update the pointers of the copies of the elements that are adjacent to a newly created copy of $x$.

The deletion of an element $x$ from a skip list is straightforward given the insertion process. We first search for $x$ and we delete one by one all its copies while updating the pointers of the copies of elements that are adjacent to a copy of $x$.

### 1.2. Related work

Precise analysis of the expected search cost in a skip list has been extensively studied, we refer to the thesis of Papadakis for more information [17]. Several variants of the skip list have been considered: Munro et al. [16] developed a deterministic version of the skip list, based on B-trees [3], that performs each dictionary operation in worst-case $O(\lg n)$ time. Under the

**Algorithm 1** Search($x$)

```
c ⇐ header
pred ⇐ −∞
while c ≠ x and height[c] > 1 do
    while down[c] = NIL do
        c ⇐ left[c]
    end while
    c ⇐ down[c]
    while right[c] ≠ NIL and right[c] ⩽ x do
        c ⇐ right[c]
    end while
    if pred < c then
        pred ⇐ c
    end if
end while
return pred
```

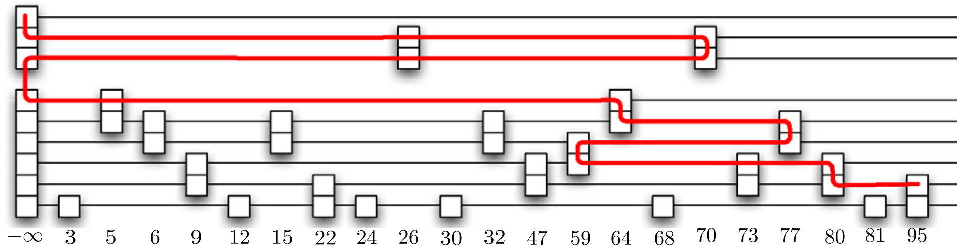


**Fig. 2.** Search path for the element 95.

assumption that the distribution of access probabilities is given, Martínez and Roura [14] developed an algorithm that minimizes the expected access time by either building an optimal static skip list in $O(n^2 \lg n)$ time or a nearly optimal one in $O(n)$ time. Bagchi et al. [2] developed the biased skip list; it manages a biased dictionary that is an ordered set $S$ of elements $x$ associated with a weight $w(x)$ and performs search, insert, delete, join, split, finger search and reweight operations in worst-case running times similar to those of biased search trees [4,9].

In the general case where access probabilities are unknown, Bose et al. [5] prove that for a class of skip lists that satisfies a weak balancing property, the working-set bound is a lower bound on the time to access any sequence. Furthermore, they develop a deterministic self-adjusting skip list whose running time matches the working-set bound, thereby achieving dynamic optimality in this class (both in internal and external memory).

## 2. Skip lift

The expected amount of extra information per element (number of copies) in a standard skip list [18] is constant. In the worst-case this number can reach $\Omega(\log n)$. Hence the number of structural changes in a skip list during an update is $\Omega(\log n)$ in the worst-case. Here we present a slight modification of the skip list data structure (as the title of the paper suggests) which guarantees, in the worst-case, a constant amount of extra information per element and a constant number of structural changes per update.

A skip lift is a light version of the skip list where copies of elements have been removed from specific levels. A skip lift only keeps the copies of an element in the two highest levels where it would appear in the skip list. Every other copy of an element is removed. The copies of the elements at the same level are connected with their left and right pointers. Additionally the two copies of an element are connected with their up and down pointers (see Fig. 1.b). Each copy stores its height in an extra height field.

A level of the skip lift is empty if no element of the set $S$ appears in it. The copies of the sentinel element appearing in an empty level are deleted. The remaining copies of the sentinel element are connected with their up and down pointers. A copy of the sentinel element at height $+\infty$ is explicitly maintained, this copy is called the *header* of the skip lift.

### *2.1. Search*

To search for a given element $x$ in a skip lift we start at the header of the list. We follow the right pointers on the same level until we see that we are about to overshoot the element $x$ that is until the element on the right has a key value strictly greater than $x$. If it is possible we go down to the next non-empty level. Otherwise we follow the left pointers until we find an element which allows us to go down to the next non-empty level. Then we iterate the process until $x$ is found or when we have reached the lowest level (in this case $x$ is not in $S$ and we know its predecessor). This procedure is described in detail in Algorithm 1 and illustrated in Fig. 2.

**Lemma 1.** *A skip lift supports a search operation in $O(\frac{1}{p}\log_{1/p} n)$ expected time, where n is the number of elements in the skip lift and p is the probability for an element in level i to appear in level $i+1$.*

**Proof.** The expected length of the search path in a skip lift $\mathcal{L}$ corresponds to the expected number of vertical steps plus the expected number of horizontal steps. The number of vertical steps performed during a search is upper bounded by the height $H(\mathcal{L})$ of the skip lift which has an expected value of $\log_{1/p} n + \frac{1}{1-p}$ (cf. Section named "Probabilistic Analysis of Search Cost" of the paper [18]). The expected height of a skip lift corresponds exactly to the expected height of a skip list.

Now we are going to bound the number of horizontal steps. By construction an element has a probability $p^i$ to be of height $i$. At any level $i$ of $\mathcal{L}$ only elements of height $i$ and $i+1$ can appear with probability $1/(1+p)$ and $p/(1+p)$, respectively the probability of being of height $i$ or $i+1$ under the condition to appear in level $i$. This means that from any position in level $i$ the expected number of horizontal steps required to reach an element of height $i$ is at most

$$E\left[NB\left(1, \frac{1}{p+1}\right)\right] + 1 = 1 + p. \tag{1}$$

Where $NB(s,q)$ denotes a random variable (negative binomial distribution) equal to the number of failures seen before the $s$th success in a series of random independent trials where the probability of a success in a trial is $q$ and $E[NB(s,q)] = s(1-q)/q$. Similarly the expected number of horizontal steps required to reach an element of height $i+1$ in level $i$ is at most

$$E\left[NB\left(1, \frac{p}{p+1}\right)\right] + 1 = \frac{1+p}{p}. \tag{2}$$

Consider $e(i,x)$ the element of height $i$ that has the greatest key value smaller than $x$. The search path to an element $x$ in $\mathcal{L}$ traverses all elements $e(i,x)$ with $h(x) \leqslant i \leqslant H(\mathcal{L})$. These are the only elements where the search path performs a down step. Between each of these $e(i,x)$ elements, the search path traverses horizontally a certain number of other elements. On expectation this number differs depending on whether the path goes from left to right or right to left. If the path goes from right to left this expected number corresponds to Eq. (1) otherwise it corresponds to Eq. (2). The probability that the search path goes from left to right on level $i$ is $1/(p+1)$. This corresponds to the probability of seeing $e(i,x)$ before $e(i+1,x)$ from the position of $x$ on level $i$ which also corresponds to the probability that an element of height $i$ appears on level $i$. Respectively the probability that the search path goes from right to left on level $i$ is $p/(p+1)$. Hence the expected number of horizontal steps performed between each element $e(i,x)$ is

$$(p+1)\frac{p}{p+1} + \frac{1+p}{p}\frac{1}{1+p} = p + \frac{1}{p}.$$

The expected cost to access the first element $e(H(\mathcal{L}),x)$ is smaller than the expected number of elements of height greater or equal to $\log_{1/p} n$ which is $1/p$. Thus total expected number of horizontal steps is upper bounded by

$$\sum_{i=1}^{H(\mathcal{L})}\left(p + \frac{1}{p}\right) + \frac{1}{p} = H(\mathcal{L})\left(p + \frac{1}{p}\right) + \frac{1}{p}.$$

Therefore the expected length of a search path in a skip lift is

$$H(\mathcal{L}) + H(\mathcal{L})\left(p + \frac{1}{p}\right) + \frac{1}{p} = \frac{H(\mathcal{L})+1}{p} + (p+1)H(\mathcal{L}) = O\left(\frac{\log_{1/p} n}{p}\right). \qquad \square$$

*2.2. Updates*

To insert an element $x$ in a skip lift we first determine its height $h(x)$ in the structure. Then we start to search for $x$ in the list to find the position where $x$ has to be inserted that is its position in levels $h(x)$ and $h(x)-1$. Once we find these positions, the copies of the element $x$ (an item storing the value $x$ linked with a down pointer to a copy of itself) are inserted in the corresponding level. This is performed similarly to the insertion of an element in a standard doubly-linked list. If the level where the copy of $x$ has to be inserted is empty then we create a new copy of the sentinel element and insert it in the skip lift (seeing all copies of the sentinel element as a doubly-linked list). This process is described in detail in Algorithm 2. We assume that $x$ is not in the set $S$ (otherwise we could simply search for $x$ before performing the actual insert operation).

To delete an element $x$ from a skip lift we first search the two copies of $x$ using the search operation described above. Once we found the copies of $x$ we delete them from their corresponding level. This is performed similar to the deletion of an element in a standard doubly-linked list. If the deletion of the copies of $x$ creates an empty level, we remove the corresponding copy of the sentinel element. This process is described in detail in Algorithm 3.

**Theorem 2.** *The skip lift supports search, insert and delete operations in $O(\frac{1}{p}\log_{1/p} n)$ expected time and requires $O(n)$ worst-case space. The total number of structural changes performed during an update is $O(1)$ in worst-case.*

**Algorithm 2** Insert($x$)

```
c ⇐ header
h ⇐ randomLevel()
while height[c] ⩾ h do
    while down[c] = NIL do
        c ⇐ left[c]
    end while
    if c = −∞ and height[down[c]] < h and h < height[c] then
        e ⇐ new element(−∞, h)
        down[e] ⇐ down[c]
        down[c] ⇐ e
    end if
    c ⇐ down[c]
    while right[c] ≠ NIL and right[c] ⩽ x do
        c ⇐ right[c]
    end while
    if height[c] = h then
        right[x] ⇐ right[c]
        left[x] ⇐ c
        left[right[c]] ⇐ x
        right[c] ⇐ x
        x ⇐ down[x]
        if x ≠ NIL then
            h ⇐ h − 1
        end if
    end if
end while
```

**Algorithm 3** Delete($x$)

```
c ⇐ header
while height[c] > 1 do
    while down[c] = NIL do
        c ⇐ left[c]
    end while
    c ⇐ down[c]
    while right[c] ≠ NIL and right[c] ⩽ x do
        c ⇐ right[c]
    end while
    while c = x do
        right[left[x]] ⇐ right[x]
        if right[x] ≠ NIL then
            left[right[x]] ⇐ left[x]
        else if left[x] = −∞ then
            down[up[left[x]]] ⇐ down[left[x]]
            if down[left[x]] ≠ NIL then
                up[down[left[x]]] ⇐ up[left[x]]
            end if
            delete left[x]
        end if
        c′ ⇐ c
        c ⇐ down[c]
        delete(c′)
    end while
end while
```

## 3. Finger search

A data structure satisfies the finger search property if searching for an element $x$ given a pointer, called *finger*, to an arbitrary element $f$ requires logarithmic time in the rank distance between $x$ and $f$ in the set of ordered elements. It is possible to describe a finger search operation on the skip lift (as described in the previous section) but it is a bit complicated. Instead we show how to enhance the skip lift structure in order to simplify the description of the finger search. The *enhanced skip lift* maintains an extra copy of each element at the bottom level. This copy is linked to the lowest copy of the corresponding element above the bottom level with the up and down pointers.

We can search for an element $x$ in an enhanced skip lift starting at the bottom copy of any element $f$ to which we are given an initial pointer. Assume without loss of generality that the key value of the element $x$ is greater than that of $f$ (the opposite case is symmetric). The finger search can be decomposed into an *up phase* and a *down phase*. The up phase behaves as the inverse of the search operation described in Algorithm 1 and the down phase is similar to Algorithm 1.

The search path described by the following algorithm traverses only elements that are between $f$ and $x$ in the skip lift. We start the search from the bottom copy of $f$ then from any current position we follow the left pointers on the same

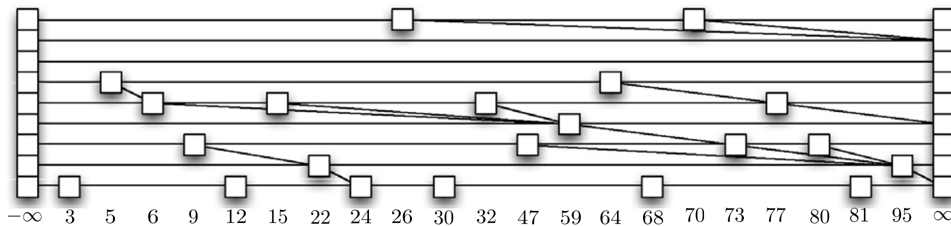


**Fig. 3.** Modified skip list.

level until the element on the left has a key value strictly smaller than $f$. If it is possible we go one level up (if the up pointer jumps over more than one level then we do not take it). Otherwise we follow the right pointers until we find an element which allows us to go one level up or when the element on the right has a key value greater than $x$ (this last case corresponds to the end of the up phase). From the current position, the down phase consists of following the right pointers on the same level until the element on the right has a key value strictly greater than $x$. If it is possible we go down by one level (if the down pointer jumps over more than one level then we do not take it). Otherwise we follow the left pointers until we find an element which allows us to go down by one level. Then we iterate the process until $x$ is found or until we have reached the lowest level (in this case $x$ is not in $S$ and we know its predecessor).

**Theorem 3.** *Finger searching for an element $x$ given a finger pointing to an arbitrary element $f$ in an enhanced skip lift takes $O(\frac{1}{p}\log_{1/p}\delta)$ time where $\delta$ is the rank distance between the finger and the search element $x$.*

**Proof.** The search path traverses only elements that are between $f$ and $x$ in the skip lift. The sublist between $f$ and $x$ contains $\delta$ elements by definition. Thus the expected height of this sublist is $O(\log_{1/p}\delta)$ [18]. In each level we perform $O(1/p)$ expected steps since this corresponds to the expected number of steps needed to find an element of height $i$ or $i+1$ from any position on level $i$. Therefore the total length of the search path is $O(\frac{1}{p}\log_{1/p}\delta)$. □

## 4. Overview of randomized dictionaries

We present an overview of some classical randomized dictionary data structures. For each of them we briefly describe its construction and how search, insertion and deletion operations are performed. It is easy to realize that the structural changes performed during an update operation can in some situations involve $\Omega(n)$ elements of the structure. Of course those situations are very unlikely to happen but are not impossible. The skip lift is the first efficient randomized dictionary that guarantees an $O(1)$ number of structural changes per update.

### *4.1. Modified skip list*

A *modified skip list*, introduced by Cho and Sahni [8], is a variant of the skip list that uses nodes of constant worst-case size (containing $O(1)$ pointers). The modified skip list structure is a skip list where all copies of an element are deleted except for its highest copy. Thus an element $x$ only appears on the level $h(x)$. Each element $x$ has three pointers: right[$x$], left[$x$] and down[$x$]. The pointers right[$x$] and left[$x$] point to the elements on level $h(x)$ to the right and the left of $x$, respectively. The pointer down[$x$] points to the element on level $h(x)-1$ that has the smallest key value greater than $x$. Two sentinel elements with key value $-\infty$ and $\infty$ are maintained, a copy of these elements appear in every level. The down pointer of a copy of a sentinel element points to the copy of itself on the level below (see Fig. 3).

#### *4.1.1. Search*

To search for a given element $x$ in a modified skip list we start from the highest level of the sentinel element with key value $-\infty$. We follow the right pointers on a same level until the element on the right has a key value strictly greater than $x$. From this point we follow the left pointer then we immediately go down one level by following the down pointer from this left element. The process is iterated until $x$ is found or until we have reached the lowest level (in which case we know that $x$ is not in $S$ and we have found its predecessor).

#### *4.1.2. Updates*

The insert and delete operations require to search the position of $x$ in the list. When inserting an element $x$ in a modified skip list only one copy is created in the level $h(x)$ and the down pointer of $x$ is set to the element in level $h(x)-1$ that has the smallest key value greater than $x$. When deleting an element $x$ from a modified skip list we have to update the down pointers of all the elements from level $h(x)+1$ that are pointing to $x$ by setting them to the element on the right of $x$.

A degenerate situation is when all elements in the structure have height 2 except for the very last one (with the greatest key value). Deleting the last element would force the modification of the down pointer of all elements, implying an $\Omega(n)$ number of structural changes in the structure. A similar situation occurs if we insert an element just before the last one.

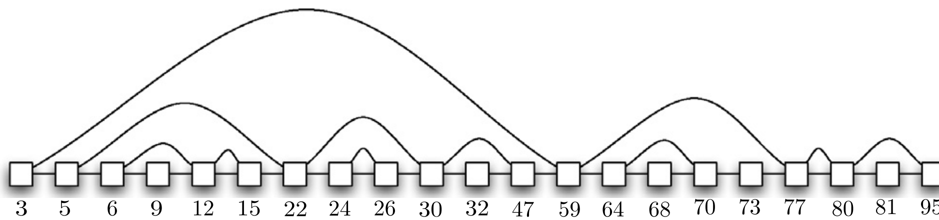


**Fig. 4.** Jumplist.

### 4.2. Treap

A *treap* is a randomized data structure introduced by Aragon and Seidel [1]. It is structured as a binary search tree structure, so the left and the right subtrees of any node only contain elements of smaller or greater key value, respectively. Each element of $S$ is given a random priority. The treap is built such that the root is the minimum-priority node and the priority of any non-root node must be greater than or equal to the priority of its parent (heap-ordering property).

#### 4.2.1. Search

To search for a given element $x$, we use the standard binary search algorithm in a binary search tree independently of the priorities.

#### 4.2.2. Updates

To insert a new element $x$ into the treap, we first generate a random priority for $x$. We perform a search for $x$ in the treap. If $x \in S$ we do nothing otherwise we make $x$ a child of the last element visited during the search. Then $x$ is rotated up as long as its priority is smaller than the priority of its parent or when $x$ becomes the new root.

To delete a node $x$ from the treap three cases are considered. If $x$ is a leaf, we simply remove it. If $x$ has a single child, we remove $x$ from the treap and make the child of $x$ the new child of its the parent (or make the child of $x$ the root if $x$ had no parent). Finally, if $x$ has two children, swap its position in the treap with its predecessor, resulting in one of the previously discussed cases. In this final case, the swap may violate the heap-ordering property, so additional rotations may need to be performed to restore it.

A degenerate situation is when the tree is a path of $n$ elements. Inserting an element at the end of the path with a given priority that is smaller than any priority in the tree would bring the new inserted element to the root. This is performed by a sequence of $\Omega(n)$ rotations that is an $\Omega(n)$ number of structural changes in the tree. A similar situation can occur when deleting an element.

### 4.3. Randomized binary search tree

A *randomized binary search tree* is another dictionary data structure developed by Martínez and Roura [15]. Each subtree of a random search tree is itself a random search tree. The root of such a tree is chosen uniformly at random among the elements of $S$ that is with probability $1/n$. The remainder of the tree is defined iteratively.

#### 4.3.1. Search

To search for a given element $x$, we use the standard binary search algorithm in a binary search tree.

#### 4.3.2. Updates

To insert a new element $x$ into a random search tree $T$ we proceed as follows: with probability $1/(|T|+1)$ the element $x$ has to be the root of the new tree. In this case the tree $T$ is split at $x$ and the two obtained subtrees are attached as the children of $x$. Otherwise we iterate the process on the left (right) subtree if $x$ is smaller (greater) than the key value of the root.

To delete an element $x$ from a random search tree $T$, we search for it in $T$. Once it is found we remove it and we replace the subtree rooted at $x$ by a newly created subtree obtained by joining the left and right subtree of $x$ (this joining procedure is fully described in [15]).

A degenerate situation would be when the tree is a path of $n$ elements so that the key of the elements from the root to the leaf are alternatively greater and smaller than $x$. Assume we insert a new element with key value $x$. It could be that $x$ has to be inserted has the root of the tree. In this case we split the tree at $x$ which requires an $\Omega(n)$ number of structural changes in the tree. The inverse situation can occur when deleting an element.

### 4.4. Jumplist

A *jumplist* of Brönnimann et al. [7] is a randomized data structure inspired by the randomized tree. It is a linked list data structure ordered by key value whose nodes are endowed with an additional pointer, the *jump pointer* (see Fig. 4). An

element $x$ of a jumplist has a next[$x$] pointer which points to the immediate successor of $x$ in $S$. Additionally an element has a jump[$x$] pointer which points to an element further on the list to the right of $x$. The jumplist is constructed as follows: the element $j$ pointed to by the jump pointer of the head of the list is chosen uniformly at random among the elements in the list. This assignment divides the list into two independent sublists that are built recursively using the same random process. This construction ensures that the jump pointers do not cross.

#### *4.4.1. Search*

The jumplist is based on the *jump-and-walk* strategy: whenever possible use the jump pointer to speed up the search, and walk along the list otherwise. So to search for an element $x$ we use the jump pointer until we are about to overshoot $x$ in which case we follow the next pointer. We iterate this process until we find the element $x$ or until the next pointer leads us to an element with greater key value than $x$ (in this case we know that $x$ is not in $S$ and we have found its predecessor).

#### *4.4.2. Updates*

To insert an element $x$ in a jumplist $J$ we proceed as follows: with probability $1/|J|$ the element $x$ has to be the element pointed by the jump pointer of the head of the list. In this case the whole list is rebuilt from scratch. Otherwise $x$ is inserted in one of its sublists. In the case where $x$ has to be inserted as the new head of a sublist, a process that does not rebuild the sublist from scratch is called to maintain the randomness of the structure.

Since an insertion could cause the reconstruction of the entire jumplist, this operation requires an $\Omega(n)$ number of structural changes in the list.

## References


[1] C.R. Aragon, R. Seidel, Randomized search trees, Algorithmica 16 (1996) 464–497.
[2] A. Bagchi, A.L. Buchsbaum, M.T. Goodrich, Biased skip lists, Algorithmica 42 (1) (2005) 31–48.
[3] R. Bayer, E. McCreight, Organization and maintenance of large ordered indexes, Acta Informatica 1 (1972) 173–189.
[4] S.W. Bent, D. Sleator, R. Tarjan, Biased search trees, SIAM Journal on Computing 14 (3) (1985) 545–568.
[5] P. Bose, K. Douïeb, S. Langerman, Dynamic optimality for skip lists and B-trees, in: Proceedings of the Nineteenth Annual ACM–SIAM Symposium on Discrete Algorithms (SODA'08), 2008, pp. 1106–1114.
[6] G.S. Brodal, G. Lagogiannis, C. Makris, A.K. Tsakalidis, K. Tsichlas, Optimal finger search trees in the pointer machine, Journal of Computer and System Sciences 67 (2) (2003) 381–418.
[7] H. Brönnimann, F. Cazals, M. Durand, Randomized jumplists: A jump-and-walk dictionary data structure, in: Proceedings of the 12th Annual Symposium on Theoretical Aspects of Computer Science (STACS'03), in: LNCS, vol. 2607, 2003, pp. 283–294.
[8] S. Cho, S. Sahni, Weight-biased leftist trees and modified skip lists, ACM Journal on Experimental Algorithmics 3 (1998) 2.
[9] J. Feigenbaum, R. Tarjan, Two new kinds of biased search trees, Bell System Technical Journal 62 (10) (1983) 3139–3158.
[10] R. Fleischer, A simple balanced search tree with $O(1)$ worst-case update time, International Journal of Foundations of Computer Science 7 (1996) 137–149.
[11] B. Haeupler, S. Sen, R.E. Tarjan, Rank-balanced trees, in: 11th International Symposium on Algorithms and Data Structures (WADS'09), 2009, pp. 351–362.
[12] R.S. Leonidas, J. Guibas, A dichromatic framework for balanced trees, in: Proc. 19th IEEE Symp. on Foundations of Computer Science (FOCS'78), 1978, pp. 8–21.
[13] C. Levcopoulos, M. Overmars, A balanced search tree with $O(1)$ worst-case update time, Acta Informatica 26 (3) (1988) 269–277.
[14] C. Martinez, S. Roura, Optimal, nearly optimal static weighted skip lists, Technical report, LSI-95-34-R, Dept. Llenguatges i Sistemes Informàtics, Universitat Politèchnica de Catalunya, 1995.
[15] C. Martínez, S. Roura, Randomized binary search trees, Journal of the ACM 45 (2) (1998) 288–323.
[16] I. Munro, T. Papadakis, R. Sedgewick, Deterministic skip lists, in: Proceedings of the Third Annual ACM–SIAM Symposium on Discrete Algorithms (SODA'92), 1992, pp. 367–375.
[17] T. Papadakis, Skip lists and probabilistic analysis of algorithms, PhD thesis, University of Waterloo, Department of Computer Science and Faculty of Mathematics, 1993. Available as Tech. Report CS-93-28.
[18] W. Pugh, Skip lists: A probabilistic alternative to balanced trees, Communications of the ACM 33 (6) (1990) 668–676.
[19] R.E. Tarjan, Updating a balanced search tree in $O(1)$ rotations, Information Processing Letters 16 (5) (1983) 253–257.