The Dictionary ADT & Binary Search Trees

Florestan Brunck

February 28, 2022

This introduction to dictionaries builds from a transcription of a lecture given by Luc Devroye in the winter semester of 2019 for the undergraduate class on Data Structures and Algorithms at McGill University (COMP 251). This lecture introduces the dictionary ADT and its implementation through binary search trees.

1   Introduction

The purpose of a **dictionary** is to store a collection of items in a way that makes them easily and quickly accessible by a search. To that effect, each object is associated a **key**, typically a value or elements of a set which has been endowed with a total order[1].

The default operations of the dictionary ADT are the SEARCH operation, together with the INSERT and DELETE operation.

Additionally, some optional operations that are often expected to be implemented in a dictionary are the sometimes called *browsing* operations. Namely the MAX/MIN operation, which return the item whose key is maximal/minimal, and the SUCCESSOR/PREDECESSOR operations, which take as input an item and return the item whose key is the next/previous value in the total order.

[1] We can for example think of an English dictionary, in which the items are the words and the key of a particular word is the word itself, seen as a number in binary (recall that each letter can be attributed a value ranging from 0 to 256 in the ASCII code.

2   Implementations

Dictionaries can be implemented with a variety of data structures, each with their own advantages and drawbacks. The particular choice of data structure must tailor to the problem at hand and make use of the nature and structure of the objects being stored.

The main data structures are:

Binary Search Trees (BST)

Balanced Trees (e.g. red-black trees)

Hash Tables

B-Trees

Tries & Suffix Trees

Because of their simplicity, binary search trees are a fundamental data structure to which we will dedicate the rest of this lecture. Some of the others data structures also have dedicated lectures in this course.

Figure 1: An implementation of a dictionary storing the numbers $\{1, 2, 3, 4, 5\}$ in a binary tree.

3 Binary Search Trees

Definition 1. A **binary search tree (BST)** is a binary tree in which each node stores exactly one key, with the added requirement that each node has the *search tree property*. For each node, this property enforces that all the nodes situated in its left (resp. right) subtree have lower (resp. higher) key values (see Fig. 8).

4 Binary Search Trees on the Entropy Scale

The problem at hand when building a dictionary is to store a list of ordered items in a way that makes them easily accessible via a search. We can therefore ask ourselves what makes BSTs good candidates for dictionaries. A useful tool in that regard is what one may call the *entropy scale*. On that scale we measure the amount of effort (i.e. the number of comparisons or use of elementary operations) needed to create the data structure from an unordered list. We can think about this amount of effort as *entropy* in the sense that it measures the amount of order in a given data structure. An unordered list realising the maximal amount of disorder with no given structure while an ordered list stands as the maximal amount of order, with the most structure. Indeed for an ordered list no effort (just printing the elements one by one) is required to output the ordered list (n comparisons).

Figure 2: The Search Tree Property.

Figure 3: The entropy scale.

We can position the various potential data structures on that scale depending on how much work it takes to construct them from an unordered list. Where do binary search trees stand on that list? Well, notice first that to output an ordered list from a binary search tree, all that is required is to traverse the tree in order. Indeed in this traversal we first visit the left subtree, then the node itself and then the right subtree, the search tree property then guarantees that the output is the ordered list. Recall then from the lecture on trees that traversing a tree takes order $O(n)$. So a binary search tree really has the same order/entropy as an ordered list!
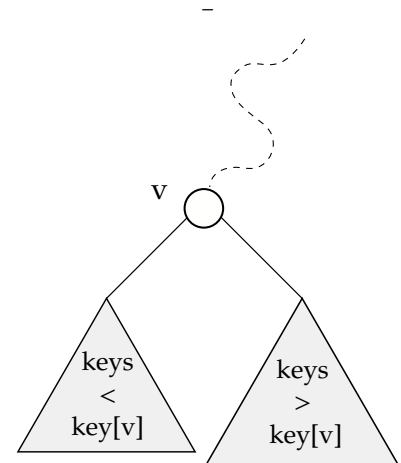
5   Basic Operations

In this section we provide an overview of the elementary operations that come with the dictionary ADT and describe their explicit implementations for binary search trees.

The first operation is the SEARCH operation, which allows us to retrieve the item associated to a given key.

SEARCH(k, x)

1   // Returns the node in a tree with root t that has the key k
2   **if** $t = $ nil or key$[t] = k$
3   return t
4   **elseif** $k < $ key$[x]$
5   SEARCH(left$[x], k$)
6   **elseif** $k > $ key$[x]$
7   SEARCH(right$[x], k$)

Exercise 2. *Re-write the SEARCH algorithm non-recursively.*

The second most elementary operation is the INSERT operation, which allows us to create and add a new item with a given key and place it at the right position in a given binary search tree.

INSERT(k, t)

1   // Inserts the node with the key k in the BST rooted at t
2   // We start by positioning ourselves at the right node in the BST
3   **while** $x \neq$ nil
4        $y \longleftarrow x$
5   **if** $k < $ key$[x]$
6            $x \longleftarrow$ left$[x]$
7   **else**
8            $x \longleftarrow$ right$[x]$
9   // We then create a new cell x with the given k
10   // at that position in the tree
11   key$[x] \longleftarrow k$
12   parent$[x] \longleftarrow y$
13   **if** $k < $ key$[y]$
14   left$[y] \longleftarrow x$
15   **else**
16   right$[y] \longleftarrow x$
17   left$[x] \longleftarrow$ nil
18   right$[x] \longleftarrow$ nil

Figure 4: The INSERT algorithm uses a travelling pair of pointers to descend the binary search tree from the root to the correct position.

Conversely, we also need to be able to delete an item with a given key and update the tree so that it still retains the search tree property. If node x has no children (case 1), then one can just set the appropriate child pointer of parent$[x]$ to nil. If x has only one child (as in the figure on the right), then one can simply bypass the node by adjusting two pointers (see case 2 in the figure). Case 3 occurs when x has at most one child and is also the root. The trickier case 4 covers the situation when x has two children. In that instance, one could let ℓ denote the rightmost node in the left subtree of x, and transfer key$[\ell]$ to x. Since ℓ has no right child, we can simply delete ℓ as described above.

Figure 5: Case 2 of the DELETE algorithm.

DELETE(x)

1   // Delete the node x when it has two children
2   // Let $L = $ left$[x]$ denote the left subtree of x
3   Let $\ell \longleftarrow $ maximum(L)
4   key$[x] \longleftarrow $ key$[\ell]$
5   DELETE(ℓ) (as ℓ has at most one child)
6   // Note that equivalently we could
7   // replace ℓ with $r \longleftarrow $ minimum(R)
8   // where $R = $ right$[x]$ denotes the right subtree of x
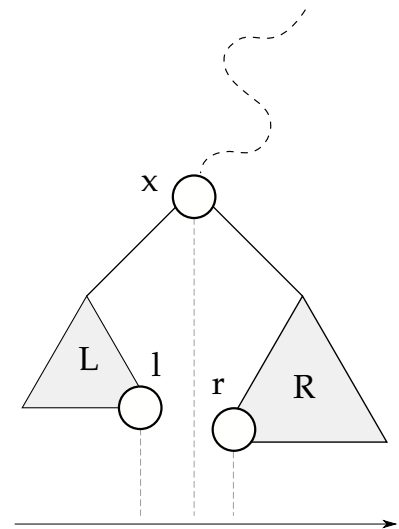

Figure 6: Case 4 of the DELETE algorithm.

SUCCESSOR(x)

1   **//** Returns the node with the next key in the total order.
2   **if** right[x] $\neq$ nil
3   return MINIMUM(right[x])
4   **else**
5       $y \longleftarrow$ parent[x]
6   **while** $y \neq$ nil & $x =$ right[y]
7       $x \longleftarrow y$
8       $y \longleftarrow$ parent[x]
9   return y

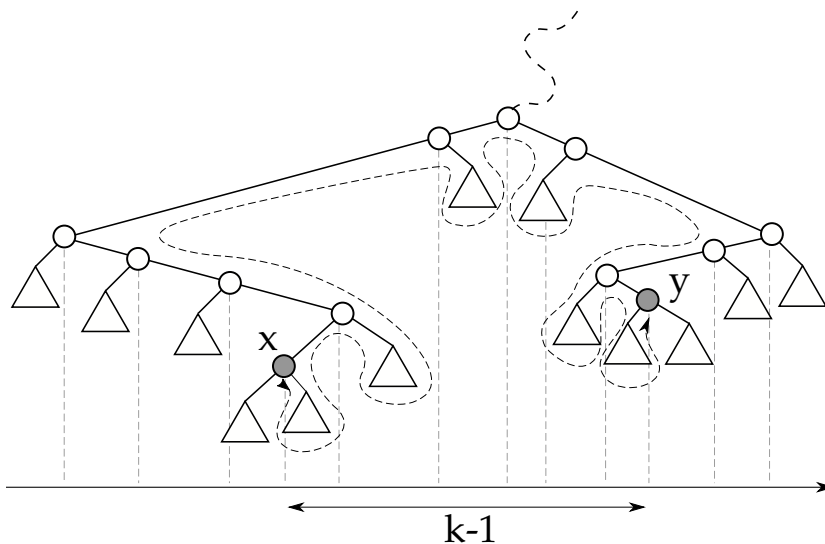Figure 7: An example of an ordered rooted tree (or plane tree).

Since we go down a tree level at each call, the running time is at most the height of the binary search tree.

6   Browsing in BSTs

The BROWSING operation can be simply implemented by repeated application of the SUCCESSOR operation. Namely, browsing k-steps can be realised by k successive calls of the SUCCESSOR operation started at a given node.

What is the complexity of this operation? Well, the SUCCESSOR operation takes $O(h)$ so a naive first answer would be to multiply this by the number of steps to obtain a running time of $O(h \cdot k)$. However we can show that the running time is $O(h + k)$.

Figure 8: The BROWSING operation, obtained by iterating the SUCCESSOR operation.

To see this notice that the repeated application of the SUCCESSOR operation amounts in the path indicated in dashed lines which consists namely of two components:

1) Going down and back up each subtree between the start and end

node (all the subtrees visited in between in a preorder traversal, think "along the coastline" as in our lectures on trees). In each subtree an amount of time equal to twice the amount of nodes in these subtree (going down and back up). But there are $k - 1$ nodes in between the start and end node by construction, which means that the total time spent is $2 \cdot (k - 1)$

2) Going up the tree from the start node to the least common ancestor of the start and end node, and then climbing back down to the end node. This clearly takes at most time $2 \cdot h$ where h is the height of the tree.

The overall cost of the BROWSING operation is thus $O(h + k)$.

Figure 9: The complexity of the BROWS-ING operation can be broken down in two components.