To Balance or to Rebuild? – An Experimental Study of Randomly Built Binary Search Trees

István Fekete[a], Roberto Giachetta[b], Péter Kovács[a]

[a]Department of Algorithms and Their Applications
Eötvös Loránd University
e-mail: {fekete, kpeter}@inf.elte.hu

[b]Department ot Software Technology and Methodology
Eötvös Loránd University
e-mail: groberto@inf.elte.hu

Abstract

Our poster proposes a simple method for maintaining the height of binary search trees and presents an experimental performance analysis by simulating the building process of these trees.

The binary search tree is a classical data structure for storing records with specified keys. The running time of its main operations (lookup, insertion and deletion of records) are determined by the height of the tree. Several methods are known to prevent linear extension of the tree height. One of them is applied in the well-known AVL-tree structure, in which the tree height is balanced by preserving a local condition for each node with the so-called tree rotations. The overall time of the main operations (along with the occasional maintenance of the tree) are bounded by $1.44 \log_2 n$, where n denotes the number of records stored in the tree.

We propose an entirely different approach, which we call RBS-tree. After each insertion or deletion, the overall height of the tree is checked. If it exceeds $c \log_2 n$ for a specified $c \geq 1$ constant, the tree is rebuilt from scratch to be fully balanced. This operation can only be achieved with a linear algorithm, thus it should not be applied frequently. By simulating the building of these two types of search trees, it was checked if a proper $c \geq 1$ constant can be selected for which the amortized performance of RBS-tree is comparable to AVL-tree. Supposing and generating uniform distribution of the data keys, we were able to confirm this conjecture in our experimental study.

Keywords: binary search tree, AVL-tree, height limitation, experimental analysis, simulation

MSC: 68P05