

1 B+-tree Construction

Given a collection of N data records, and three alternatives to create a B+-tree on this data set:

1. Insert each record one at a time.
2. Sort data set, then perform record wise insertion.
3. Bulk loading.

What are the I/O costs and number of page accesses for each of the alternatives?

Suppose another set of m data records arrives. What are the costs (again, I/O and page accesses) for inserting the new data set into the existing tree. Consider the following three approaches:

1. Insert each record one at a time.
2. Sort data set, then perform record wise insertion.
3. Sort data set, merge with previous data and create a new tree using bulk loading.

2 Hash indexes

Consider the Extendible Hashing index shown in Figure 1. Answer the following questions about this index:

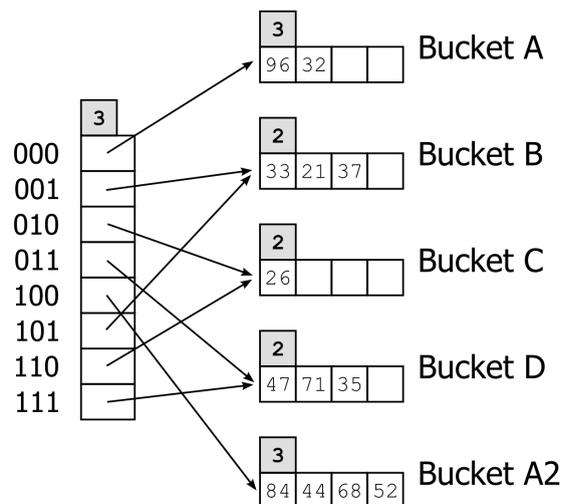


Figure 1: Extendible Hashing index.

1. Show the index after inserting an entry with hash value 20.
2. Show the index after inserting entries with hash values 49 and 85 into the original index.
3. Show the index after deleting the entry with hash value 26 into the original tree.

3 Project: Main-Memory Efficient Tree-index

Keep working on your project. Deadline was postponed to Dec 1st.