

CSCIX370 Database Management

Spring 2014

Project 3: Implement your own Index Structures.

Due: Feb 25 (Monday), 8 AM.

In this project you will implement your own index structures. Same as in Project 2 you need to store your tuples in a FileList. You have to choose implementation of B+tree AND (Linear Hash OR Extendable Hash) for index.

Reuse your implementation and source code in Projects 1 & 2. Three new files are BpTree.java, LinHash.java and ExtHash.java (at the course web page). Use the test cases in the main functions of BpTree.java, LinHash.java or ExtHash.java. More test cases will be posted later.

BpTree.java provides B+Tree maps that are multi-level index structures and provide efficient access for both point and range queries. LinHash.java and ExtHash.java provide hash maps that use the Linear Hashing and Extendable Hashing algorithms. A hash table is created that is an array of buckets. Since indexing will take more time to process than the last time you created (in Project 2), it will be more convenient to build the index using main memory rather than the disk. For this stage of indexing, you'll also need key and value pairs created as maps. Last time (in Project 2), you've created TreeMap objects and this time you need to swap those to the maps that will be used for B+Tree and (Linear Hashing or Extendable Hashing) algorithms. Those maps are given in the starter code and you need to implement indexing using those map objects.

Your program must be thoroughly documented (generate javadoc). Use the @author tag for each class and method. Each method should have a single author. The coding workload should be split roughly in half. We will check this by examining the @author tags. Please make sure that the output of your program is easy to understand. Provide a flag for turning on/off your tracing/debugging messages in your program's output – if necessary.

Programming language: Java 7 is required for the project.

What to submit: Please submit

- all source code
- all the javadoc files
- a readme file

The readme file should contain: your names, how to compile and run your code and other specifications you want to make. Please pack all your files in a zip package with the file name: "project3" + last names of group members. For example: project3_chen_kim_wong_allen.zip.

How to submit: Mail your ".zip" file to the TA (Sara Vahid, savahid@gmail.com) with subject line "[X370] Project 3 Submission". If you submit more than once the latest copy will be considered (we encourage only a single submission). See course page or syllabus for the policy on late submissions.

An electronic copy of this project and links to source code templates to be used can be found at the course web page:

http://www.cs.uga.edu/~budak/db_sp14.html

Also every group member needs to submit a peer evaluation form at the time of project submission. This form can be e-mailed to TA in a separate e-mail.