

CSE 562: Project #2 (due November 5, 2006)

Team project. Team size: 2 students. The same teams will be selected for Project # 3. Use Java or C++ for project implementation.

Submit all work in electronic form using `submit_cse562`.

General

You will implement a file with one or more indexes that use the B+-tree data structure. They will be stored in main memory, with their persistent versions being kept on disk.

Assumptions:

1. every record type consisting of variable-length character string fields should be supported;
2. index keys consist of single fields of length at most 24 bytes;
3. key values may be repeated;
4. the indexes are dense;
5. the file is organized as a heap;
6. the B+-tree nodes have between 4 and 8 index records;
7. no block structure is assumed;
8. there is no requirement to use a specific format for the data stored on disk.

You can assume reasonable limits on the file size, record size, length of names, number of fields etc.

Interface

The following operations should be supported (“+” indicates an input argument and “-” an output argument):

- `create_file(+FileName,+Descriptor,-FileId)` where `Descriptor` is a list of fields and their lengths and `FileId` is an internal file identifier;
- `create_index(+IndexName,+Key,+FileName,-IndexId)` where `IndexId` is an internal index identifier;
- `get_record(+FileId,+Rid,-Record)`: retrieve a record given its record Id;
- `get_first_file(+FileId,-Rid,-ScanId)`: get the record Id of the first record of the file, opening a file scan;

- `get_next_file(+FileId,+ScanId,-Rid)`: get the record Id of the next record of the file in the current file scan;
- `get_rid(+FileId,+IndexRid,-Rid)`: given the Id of an index record, get the record Id from the index record;
- `get_first_index(+IndexId,+Key,-Rid,-ScanId)`: get the record Id of the first index record with the key value equal to `Key`, opening an index scan;
- `get_next_index(+IndexId,+Key,+ScanId,-Rid)`: using the index, get the next record Id of the next index record in the current index scan with the key value equal to `Key`;
- `print(+Record)`;
- `dump_file(+FileName)`: write file contents to disk;
- `dump_index(+IndexName)`: write index contents to disk;
- `restore_file(+FileName,-FileId)`: restore file contents from disk;
- `restore_index(+IndexName,-IndexId)`: restore index contents from disk;
- `view_file(+FileName)`: print the content of the file in a human-readable format;
- `view_index(+IndexName)`: print the content of the index in a human-readable format;
- `drop_file(+FileName)`: delete file;
- `drop_index(+IndexName)`: delete index;
- `exit`: exit the system closing all files and saving all the information to disk;
- `insert(+FileId,+Record)`.

The operations `get_first` and `get_next` need to keep track of the current record in the scan. There may be more than one scan on the file or index open at the same time.