

Implementation of LSM-Tree Key Value Store

Greeshma Yaluru

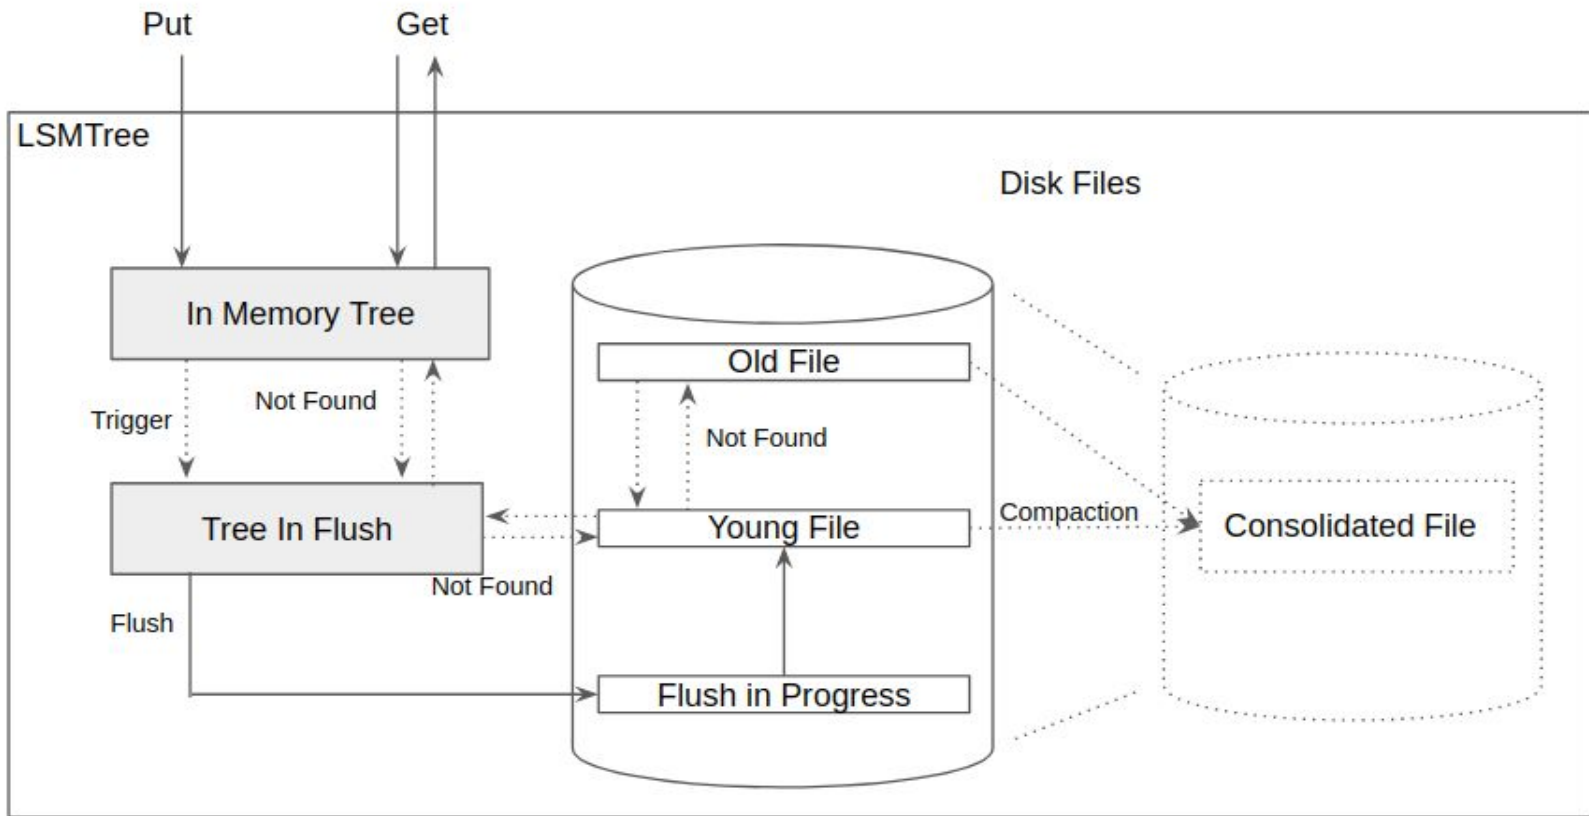
Zhenhuan Wu

What is an LSM Tree?

What are some main components in the Tree

Name a few partitioning strategies?






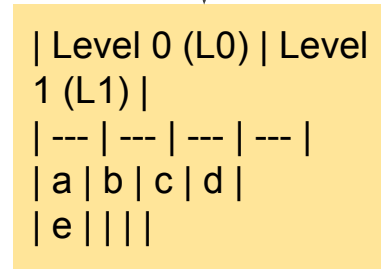
Buffer

1. Insert the key-value pair ("a", 1) into the write buffer.
2. Insert the key-value pair ("b", 2) into the write buffer.
3. Insert the key-value pair ("c", 3) into the write buffer.
4. Insert the key-value pair ("d", 4) into the write buffer.


Flush into
disk



Insert "e"



Insert "f"



Leveling Strategy

<p>Time: t_1 New sstable in Level 0</p>	<p>Level 0 10 210</p> <p>Level 1 1 100 200 400</p>
<p>Time: t_2 After compacting Level 0 into Level 1</p>	<p>Level 0</p> <p>Level 1 1 10 100 200 210 400</p>
<p>Time: t_3 New sstable in Level 0</p>	<p>Level 0 20 220</p> <p>Level 1 1 10 100 200 210 400</p>
<p>Time: t_4 After compacting Level 0 into Level 1</p>	<p>Level 0</p> <p>Level 1 1 10 20 100 200 210 220 400</p>
<p>Time: t_5 New sstable in Level 0</p>	<p>Level 0 30 330</p> <p>Level 1 1 10 20 100 200 210 220 400</p>
<p>Time: t_6 After compacting Level 0 into Level 1</p>	<p>Level 0</p> <p>Level 1 1 10 20 30 100 200 210 220 330 400</p>



Implementation

```
#include "Buffer.h"

Value Buffer::get(int key) {
    if ((key < this->min) || (key > this->max)) return Value(false);
    for (auto &pair: this->pairs) {
        if ((pair.first == key) && (pair.second.visible)) {
            return pair.second;
        }
    }
    return Value(false);
}

void Buffer::put(int key, Value val) {
    if (this->size() == 0) {
        this->max = key;
        this->min = key;
    } else if (this->max < key) { this->max = key; }
    else if (this->min > key) { this->min = key; }
```

```
61_templatedb > src > templatedb > diskrun.hpp > DiskRun<K, V> > run_No
#include <string>

using namespace std;

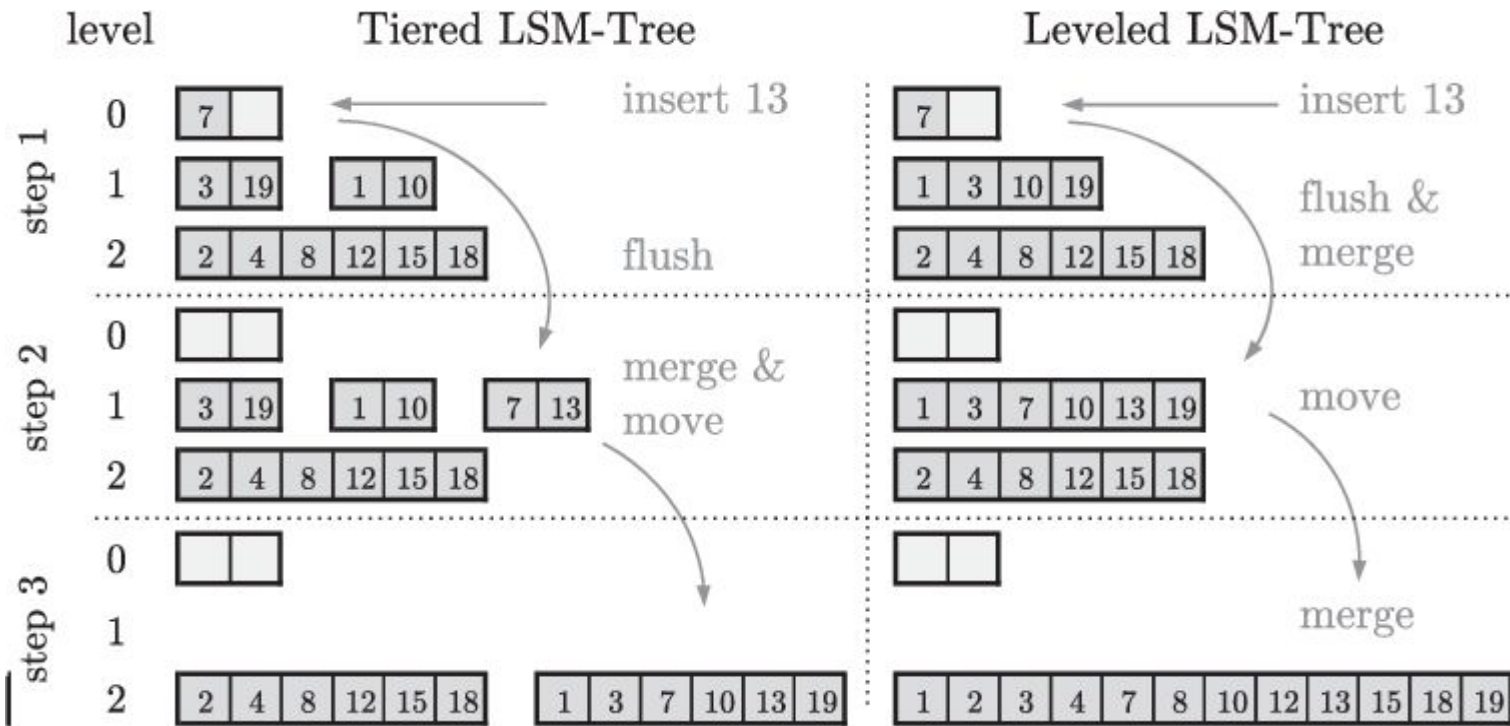
template<typename K, typename V>
class DiskRun : public Run<K, V> {
    typedef Pair<K, V> KV_pair;
    int capacity;
    int entries_in_page;
    int level;
    int run_No;
    int entries_num;
    int page_num;
    int page_size;
    bool doExist;
    K* fence_pointer;
    K MIN;
    K MAX;
    string dir;

public:
    DiskRun(int capacity, int pagesize, int level, int run_No) {
        MIN = 0;
        MAX = 0;
        dir = "./data/LSM_L" + to_string(level) + "_R" + to_string(run_No) + ".run";
        this->capacity = capacity;
        this->page_size = pagesize;
        this->level = level;
        this->run_No = run_No;
        doExist = false;
    }
};
```

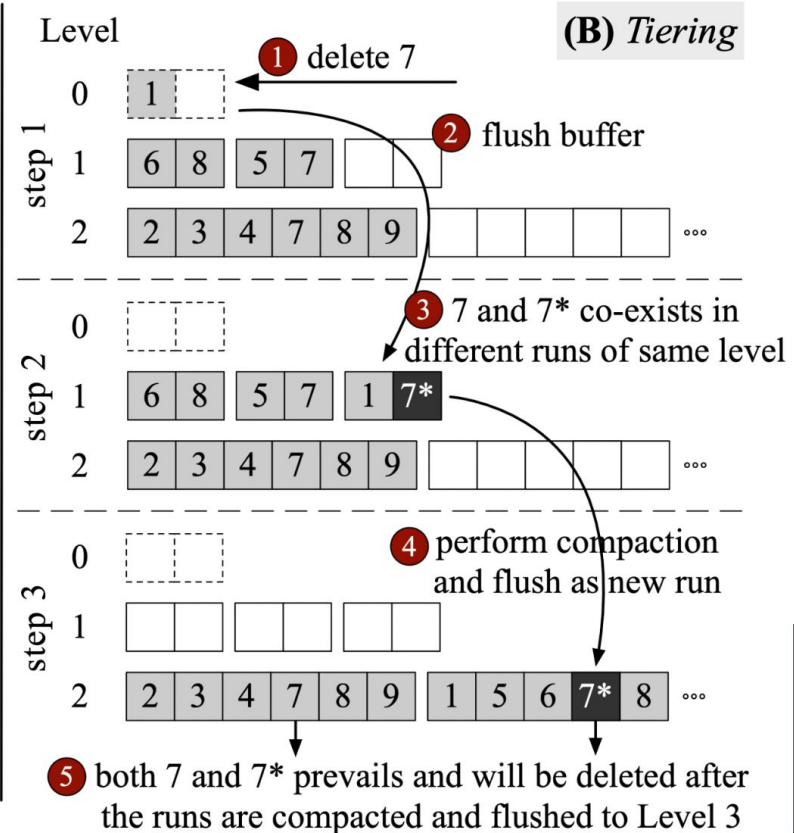
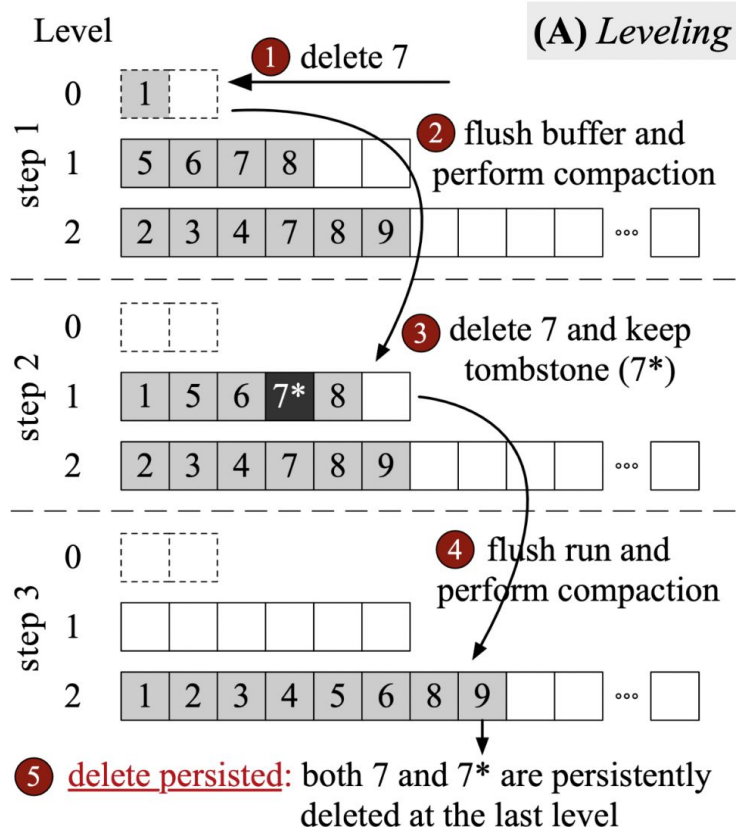
```
#include <vector>
#include <cmath>
#include "buffer.hpp"
#include "diskrun.hpp"
#include "bloomfilter.hpp"

class LSM {
private:
    std::vector<Buffer*> buffers;
    std::vector<DiskRun*> runs;
    int maxBufferSize;
    int runSize;
    int level;
    int numLevels;
    double mergeThreshold;
    BloomFilter bloomFilter;

public:
    LSM(int bufferSize, int runSize, int numLevels, double mergeThreshold, int bloomFilterSize) :
        maxBufferSize(bufferSize),
        runSize(runSize),
        level(0),
        numLevels(numLevels),
        mergeThreshold(mergeThreshold),
        bloomFilter(bloomFilterSize)
    {
        // Initialize the buffers and runs
        for (int i = 0; i < numLevels; i++) {
            buffers.push_back(new Buffer(maxBufferSize));
        }
    }
};
```

Tiering Strategy



Tiering Strategy

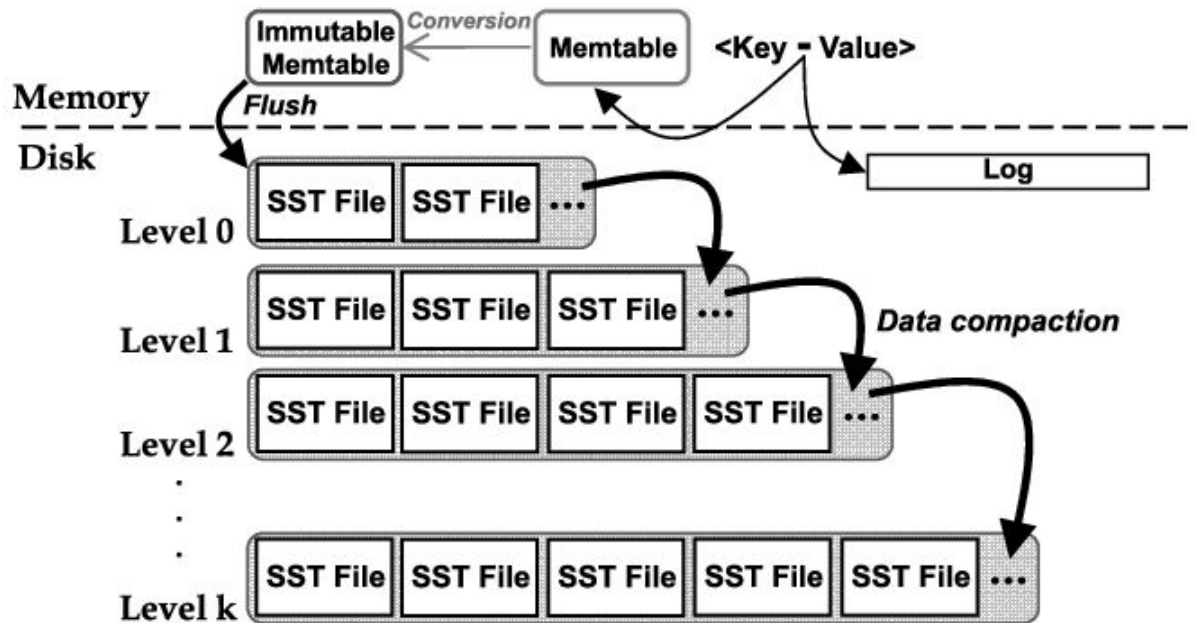


FIGURE 1. An overview of LSM-tree in LevelDB.

Solution Approach

- Understanding the framework
- Buffer & Memory
- Disk & Leveling
- Tiering (in-progress)



Solution Approach

- Understanding the framework

```
class Value
{
public:
    std::vector<int> items;
    bool visible = true;
    int range = 0;
    Value() {}
    Value(bool _visible) {visible = _visible;}
    Value(bool _visible, int _range) {visible = _visible; range = _range;}
    Value(std::vector<int> _items) { items = _items;}

    bool operator==(Value const & other) const
    {
        return (visible == other.visible) && (items == other.items);
    }
}
```

Solution Approach

- Understanding the framework

```
class Value
{
public:
    std::vector<int> items;
    bool visible = true;
    int range = 0;
    Value() {}
    Value(bool _visible) {visible = _visible;}
    Value(bool _visible, int _range) {visible = _visible; range = _range;}
    Value(std::vector<int> _items) { items = _items;}

    bool operator ==(Value const & other) const
    {
        return (visible == other.visible) && (items == other.items);
    }
}
```

Preliminary Results



Thank You

Questions?

