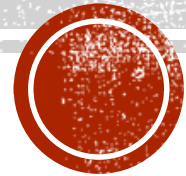


PRINCIPLES OF DATA ORGANISATION

μ -Tree




MOTIVATION

- ❧ Raw flash, no in-place update, no FTL
 - ❧ Flash translation layer
- ❧ B-Trees are good – we want to use them



B-TREE ISSUE ~ WANDERING TREE

- ❧ B-tree efficiently makes use of block-oriented storage by keeping related records together on the same page.
- ❧ Inner nodes of a B+-tree store only keys and pointers
- ❧ When a record update happens, the leaf node needs to be modified
- ❧ Since **in-place update is not supported** by the SSDs, the whole page needs to be moved into a new location
- ❧ Page move requires modification of the pointer in the parent node → iterative process ending only in the root



μ -TREE



🔗 [Kang et. al. 2007](#)

- 🔗 In μ -Tree all the nodes along the path from the root to the leaf are put together into a single flash memory page
- 🔗 μ -Tree outperforms B+-tree by up to 28% and by up to 90% with a 8KB in-memory cache.
- 🔗 Unlike B-tree, μ -Tree's nodes are of variable size
- 🔗 No significant difference between B+-Tree and μ -Tree except that the size of a node in a μ -Tree is determined by its level and the height of the tree



μ -TREE

- Leaf node always occupies half of the page.
- As the level is increased, the node size is reduced by half
- Only the root node has the same size as its children nodes



ALGORITHM

Leaves reside
in level 1

Algorithm 2 Retrieval

Input: *key* K (search predicate)

Output: *page_address* O (which points to the record corresponding to K)

```
1:  $C \leftarrow \text{GetNodeFromPage}(\text{root page address}, H)$ 
2:  $L \leftarrow H$ 
3: while  $C.\text{type} \neq \text{LEAF}$  do
4:    $K_i \leftarrow$  smallest search-key greater than  $K$ 
5:    $L \leftarrow L - 1$ 
6:   if  $K_i$  exists then
7:      $C \leftarrow \text{GetNodeFromPage}(P_i, L)$ 
8:   else
9:      $C \leftarrow \text{GetNodeFromPage}(P_m, L)$ , where  $m$  is the number of pointers in  $C$ 
10:  end if
11: end while
12: if  $K_i$  exists in  $C$ , such that  $K_i = K$  then
13:  return  $P_i$ 
14: else
15:  return  $NULL$ 
16: end if
```



ALGORITHM

Size of the
block/node to read

Offset of the
block/node to read

Algorithm 1 GetNodeFromPage

Input: *page_address* P , *level* L

Output: *node* N

1: $S \leftarrow Q/2^L$, where Q is the size of a page

2: $O \leftarrow S$

3: **if** $L = H$ **then**

4: $S \leftarrow S * 2$

5: $O \leftarrow 0$

6: **end if**

7: $N \leftarrow$ read at page P from offset O with size S

8: **return** N



ALGORITHM

Algorithm 3 Insertion

Input: *key* K , *page_address* P (which points to the record corresponding to K)

- 1: allocate a new page N
 - 2: $(R, K', P') \leftarrow \text{InsertEntry}(K, P, N, \text{root page address}, H)$
 - 3: **if** $R = \text{FULL}$ **then**
 - 4: allocate a new page N'
 - 5: $C \leftarrow \text{GetNodeFromPage}(N, H)$
 - 6: $H \leftarrow H + 1$
 - 7: $(C_l, C_r) = \text{Split}(C)$
 - 8: $C' \leftarrow \text{GetNodeFromPage}(N, H)$
 - 9: insert $(C_l.K_1, N)$ and $(C_r.K_1, N')$ into C'
 - 10: write node C_l on page N
 - 11: write node C_r on page N'
 - 12: write node C' on page N'
 - 13: **end if**
-

The root node becomes full as a result of the current insertion



ALGORITHM

Algorithm 4 InsertEntry

Input: *key* K , *page_address* P , N , B , *level* L

Output: *return_value* R , *key* K' , *page_address* P'

```
1:  $C \leftarrow \text{GetNodeFromPage}(B, L)$ 
2: if  $C.type \neq LEAF$  then
3:   find  $C.P_i$ , such that  $C.K_i \leq K < C.K_{i+1}$ 
4:   if  $C.P_i$  doesn't exist then
5:      $i \leftarrow m$ , where  $m$  is the number of pointers in  $C$ 
6:   end if
7:    $(R, K', P') \leftarrow \text{InsertEntry}(K, P, N, C.P_i, L - 1)$ 
8:    $C.P_i \leftarrow N$ 
9:   if  $R = SPLIT$  then
10:     $K \leftarrow K', P \leftarrow P', N \leftarrow P'$ 
11:   else
12:    write node  $C$  on page  $N$ 
13:    return  $R \leftarrow NULL$ 
14:   end if
15: end if
16: if  $C$  has space for  $(K, P)$  then
17:   insert  $(K, P)$  into  $C$ 
18:   write node  $C$  on page  $N$ 
19:   if  $C$  is full then
20:    return  $R \leftarrow FULL$ 
21:   else
22:    return  $R \leftarrow NULL$ 
23:   end if
24: else
25:   allocate a new page  $N'$ 
26:    $(C_l, C_r) \leftarrow \text{Split}(C)$ 
27:   insert  $(K, P)$  into  $(C_r.K_1 > K)? C_r : C_l$ 
28:   if  $C_l.type \neq LEAF$  &  $\exists C_r.P_i = N$  then
29:     swap  $C_l \leftrightarrow C_r$ 
30:   end if
31:   write node  $C_l$  on page  $N$ 
32:   write node  $C_r$  on page  $N'$ 
33:   return  $R \leftarrow SPLIT, K' \leftarrow C_r.K_1, P' \leftarrow N'$ 
34: end if
```
