

PRINCIPLES OF DATA ORGANISATION

Fagin



MOTIVATION

- Key, pointer pairs ~ index
- Hashed file organization
- Dynamic hashing
 - We can add new records (without performance penalty)
 - We do not need to specify the amount of data beforehand
- Collapsing a trie (prefix tree)
 - We use a growing part of the tree



FAGIN'S EXTENDIBLE HASHING

- 🔗 Fagin 1979
- 🔗 In general: hash function $h(k)$ returns a string of bits
 - ✂ But we do not need all of them all the time
- 🔗 **Directory based**
 - ✂ Level of indirection = we do not need a continuous space in the secondary memory
 - ✂ Global depth d_G
 - Bits needed to tell any pair of records from **different buckets** apart
 - ✂ Local depth d_L (own for each bucket/page)
 - Number of bits common to all records in a bucket
 - $2^{(d_G-d_L)}$ = how many directory records point to a page
- 🔗 Hash function (uniform, fast,...) provides d_G -long address of the directory entry with a pointer to the bucket/page
- 🔗 Overflowing causes a change in the structure of the directory (d_g, d_L) and the primary file
 - 🔗 Adding new blocks or modifying (doublong) the directory



FAGIN – EXAMPLE

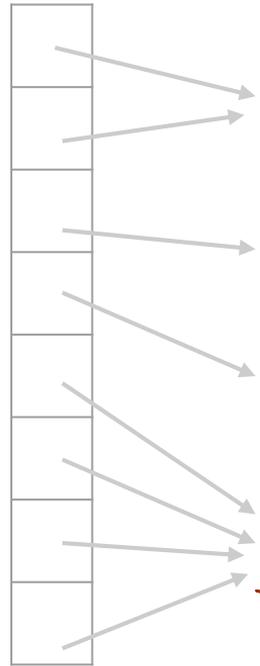
Bits we use

$d_G = 3$

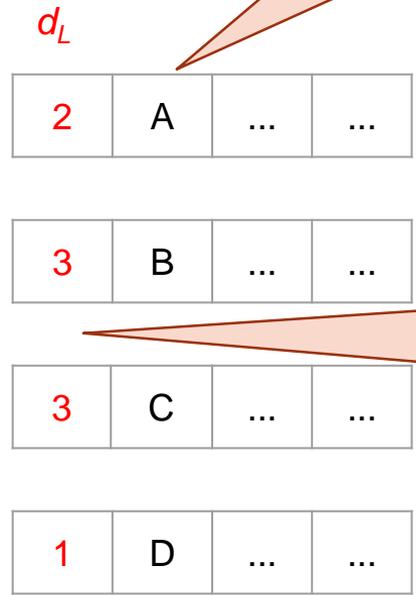
000
001
010
011
100
101
110
111

In the main memory

DIRECTORY



Buckets/pages in the secondary memory



$d_L == d_G \rightarrow$
doubling the
directory and
splitting the
page

$d_L < d_G \rightarrow$ more than one pointer =
the page can be split, d_L++

FAGIN — FIND

Finding a record with a key k

- ④ Compute $k' = H(k)$
- ④ Compute $k'' = h_{d_G}(k')$
- ④ Access page pointed to by the directory record with key k''
- ④ Scan the accessed page for record with key k
 - ④ If the record is not found, it is not present in the file



FAGIN — INSERT

Inserting a record R with a key k

- ④ Find a page where the record R should be inserted
- ④ If the page is not full, insert R into the page and return
- ④ If the page is full, split the page
 - ④ Locally
 - ④ Globally



FAGIN — SPLIT LOCAL

Split page P if $d_L(P) < d_G$

- Allocate new page Q
- Modify the directory pointers originally pointing to P so that, e.g., half of them having common first $d_L(P)$ bits followed by 0 point to P and rest of them point to Q
- Set $d_L(P) = d_L(Q) = d_L(P) + 1$
- Reinsert all the data from P



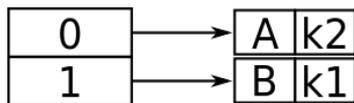
FAGIN – SPLIT GLOBAL

Split page P if $d_L(P) == d_G$

🔗 Double the directory size

🔗 $d_G = d_G + 1$

🔗 For each page Q , set the pointers so that if Q was pointed to by an entry with a bit key x , now it is pointed to by entries with keys starting with $x0$ and $x1$

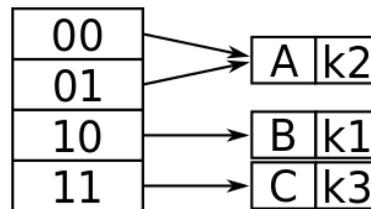


$H(k1) = 100100$

$H(k2) = 010110$



$H(k3) = 110110$



FAGIN

- ❧ The performance stays more or less constant with increasing number of stored records
- ❧ The directory might not fit in the main memory
- ❧ If the block factor is low, many splits can occur leaving many pages empty

