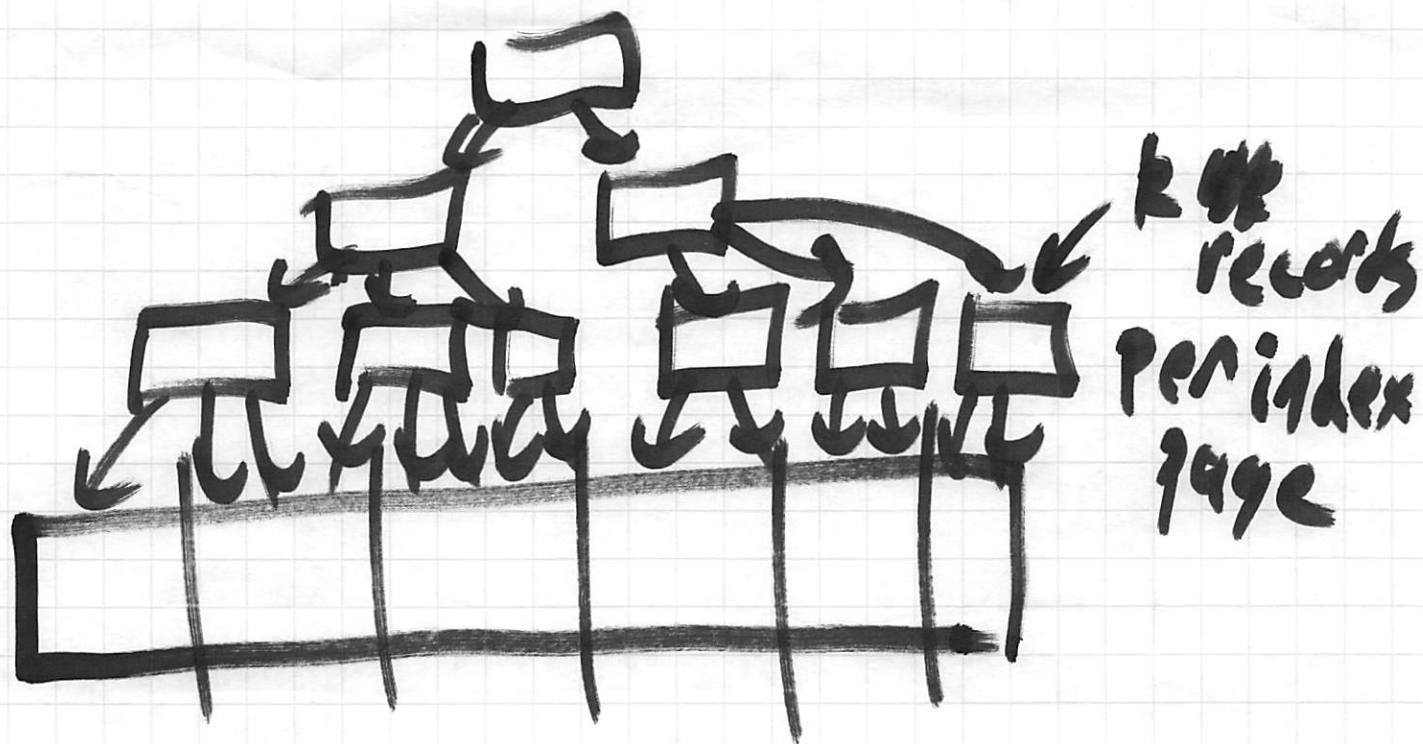


$$\log_k(N) = \frac{\log_2(N)}{\log_2(k)}$$

Index improves  
by const multiple



$N$  records  
pages

Naive Binary Search  
of  $(\log_2 N)$

w/ Index Pages

$$O(\log_k N) =$$

Idea 1: 2 (or more)

Separate copies of  
the index

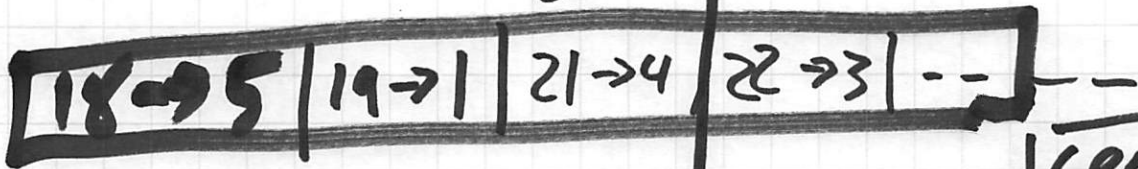
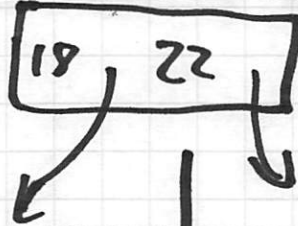
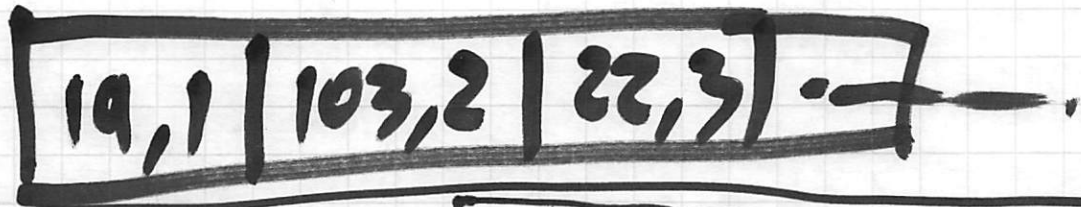
pro: easy

con: 2 copies of all data

Idea 2: Hierarchical sort

pro: Saves space

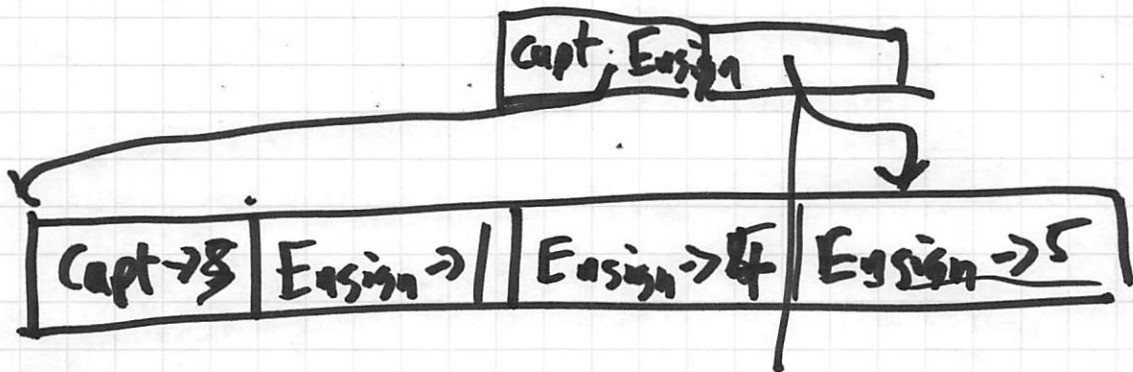
con: limited query support



Secondary Index

Idea: 2nd & subsequent indexes  
Store only unique Id of records

Pro: Less Space (so maybe faster)



Con:

#4

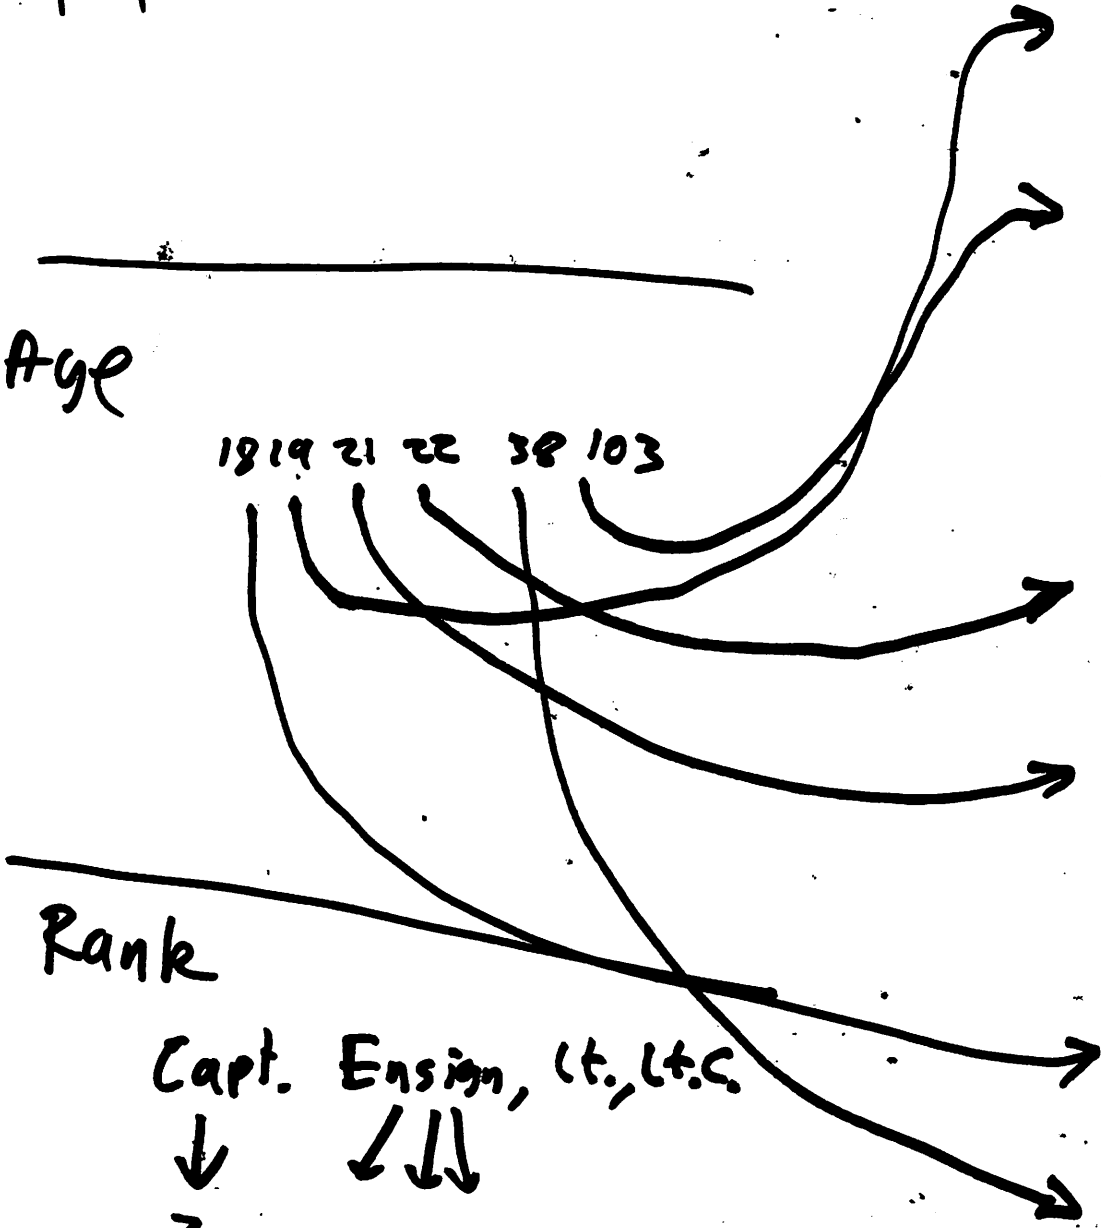
Age

18 19 21 22 38 103

Rank

Capt. Ensign, Lt., Lt.C.

↓ ↓↓  
3



One Index for each attr

Easy

Space Hog

Hierarchical Indexes

Saves space

Still Easy

Limited Query Support

Secondary Indexes

→ Saves space

→ can be faster

→

often  
Need to  
load more  
pages

Given a query for rank =       $\wedge$  age =       
PRO CON

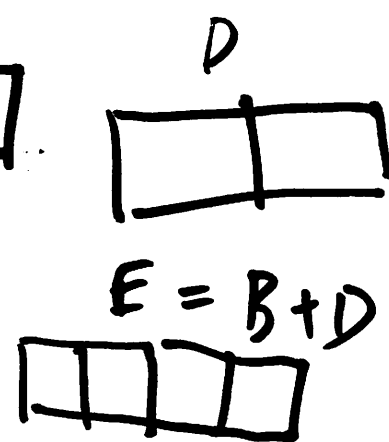
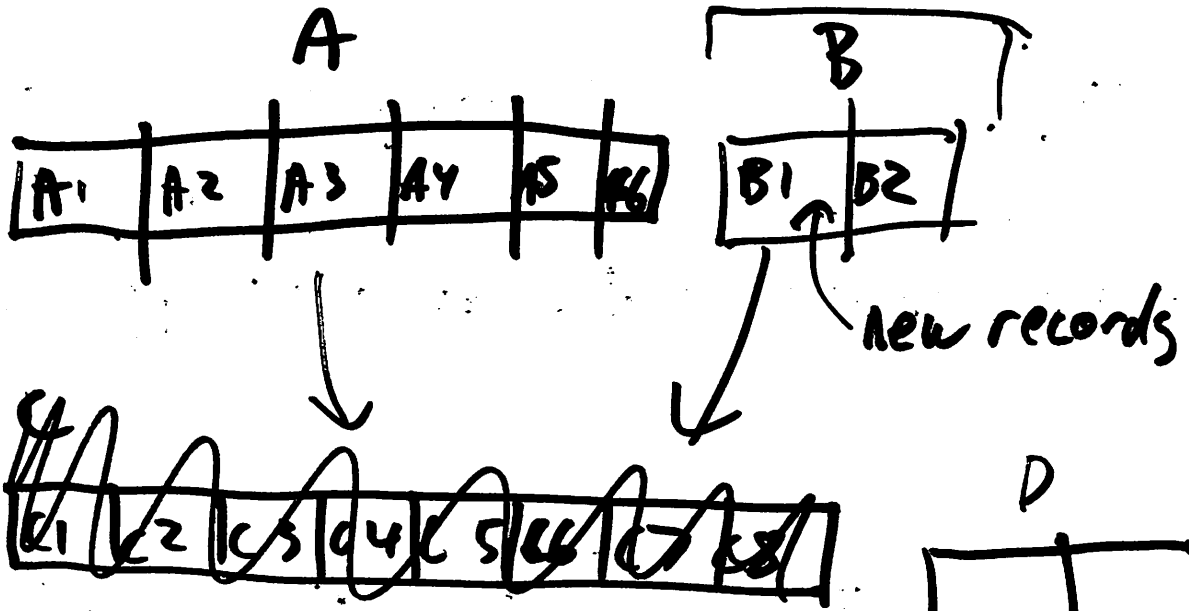
Option 1: Pick an index (either rank or age)

Option 2: Scan the full data

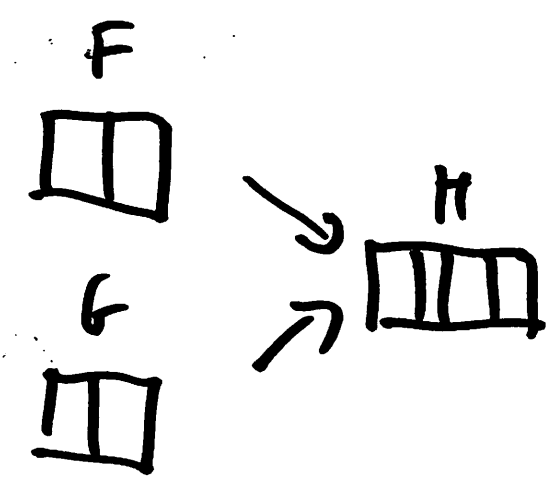
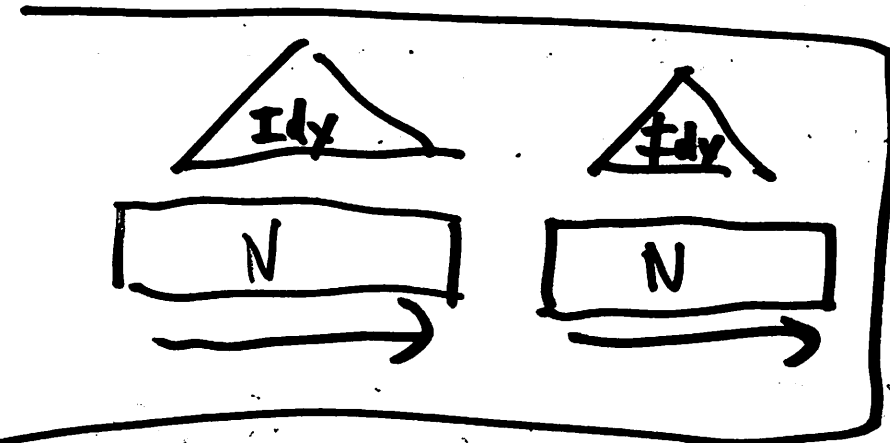
Option 3: Find all <sup>identifiers</sup> keys for rank, for age  
+ compute intersection

Option 4: Hierarchical Index

~~Question~~ Question for Piazza ~~29~~



~~W~~



# LSM-Tree

