


# CSE 350

## Advanced Data Structures

Topic 9: Storage in Trees □□

# File Layout





## Goals

- Fast record lookups
  - ↳ CPU
  - ↳ IO
- Low wasted space → Fast Scans
- Data mutability

# Flat File

Table  $\rightarrow$  Bag  
 $\rightarrow$  Set  
 $\rightarrow$  List  
 } of records (Tuples)



Idea for name

$(P_0, R_1)$   
 $(P_0, R_2)$   
 $\vdots$   
 $R_3$

Permitted  $\rightarrow$  Key

$\rightarrow$  1  
 2  
 3



| City    | Type       | Address |
|---------|------------|---------|
| Buffalo | Electrical | -       |
| Deppw   | Electrical | -       |
| Buffalo | -          | -       |

Find  $P_5 R_4$   $\rightarrow$  Load Page 5  
 $\rightarrow$  Find  $R_4$





Header



Optimization  
Use a Stack

## Finding Records

How is a record identified?

↳ Pointer (Page, Record index)

↳ Key

(Any attribute to id a collection of records)

What are the tradeoffs of each?

Key

⊖ Table may not have one

⊕

⊖ Need to map key to location  
(mapping function)

Pointer

⊕

⊖ SQL only talks in attributes

⊕ Fast lookups

How can we change the file layout to make key lookups faster?

Problem


↳ No information relating key to location by default

↳ Hash Table

↳ Sort it

# Idea 1: Sort the Records

How do we find a record?



How Much IO?

$N$  records

$P$  records / page

IO

0

1

2

$\vdots$

$\log_{\frac{N}{P}}$

$= \log_2 N - \log_2 P$

Search Space


$\frac{N}{P}$  pages

$\frac{N}{P} \cdot \frac{1}{2}$  pages

$\frac{N}{P} \cdot \frac{1}{4}$

$\frac{N}{P} \cdot \frac{1}{2^i}$

1 page



What is the IO complexity of binary search?

On a sorted file

$\log_2 N - \log_2 P$  Pages read

$O(\log_2(N))$  IO

⤴ Not great

What is the Memory complexity of binary search?

on a sorted file

$P + C$  ← arguments  
etc

$O(1)$  ← fantastic!  
Pase being  
analyzed



## Relative Sizes

128B record  
8B key

— Permit Record

— Permit ID (int)

$$1000 \approx 1024 = 2^{10}$$

How big is 20 million records?


How big is 20 million keys?

$$2^{31} \approx 2GB$$

$$2^{27} \approx 128MB$$

## Idea 2: Track key information for each page

Sorted File



$K_1, K_2, K_3, K_4, \dots, K_{100} < K_{101}, \dots, K_{200}$

$K_1 < K_2 < K_3 \dots$

Idea: Look-up table

|     |           |       |
|-----|-----------|-------|
| 1   | $K_1$     | $P_0$ |
| 3   | $K_2$     | $P_0$ |
| 4   | $K_3$     | $P_0$ |
| ... | ...       | ...   |
| 182 | $K_{100}$ | $P_0$ |
| 195 | $K_{101}$ | $P_1$ |
| ... | ...       | ...   |
|     | $K_{200}$ | $P_1$ |

$N$  Keys + page ids  $\rightarrow$  much smaller than  $N$  records  $\frac{1}{2^{7/24}}$   
 using our example  $\approx \frac{1}{2^3} = \frac{1}{8}$  the size

Idea: Collapse Look-up Table

array index  
 0:  $K_1 - K_{100} : P_0$   
 1:  $K_{101} - K_{200} : P_1$   
 2:  $K_{201} - K_{300} : P_2$   
 ...

~~$5 \cdot 2^8$  bytes~~  
 $2 \cdot 2^8$  bytes

$P = 100$

~~$2 \frac{N}{P}$  keys~~ vs  $N$  records  
 much much smaller  
 $\frac{1}{1000}$  the size

Find 312

I/O

Fence Pointer Table

|     |     |     |     |     |      |      |      |      |      |
|-----|-----|-----|-----|-----|------|------|------|------|------|
| 192 | 274 | 458 | 703 | 904 | 1089 | 1285 | 1358 | 1584 | 1821 |
|-----|-----|-----|-----|-----|------|------|------|------|------|

$P_1$   $P_2$   $P_3$   $P_4$   $P_5$   $P_6$   $P_7$  ...

← Array of keys

How big?

$O(\frac{N}{P})$

One key per page  $\approx O(N)$

|       |       |       |       |       |       |       |       |       |       |          |          |
|-------|-------|-------|-------|-------|-------|-------|-------|-------|-------|----------|----------|
| $P_0$ | $P_1$ | $P_2$ | $P_3$ | $P_4$ | $P_5$ | $P_6$ | $P_7$ | $P_8$ | $P_9$ | $P_{10}$ | $P_{11}$ |
|       | 0     | 192   | 274   | 458   | 703   | 904   | 1089  | 1285  | 1358  | 1584     | 1821     |
|       | 178   | 273   | 412   | 611   | 891   | 1063  | 1273  | 1326  | 1566  | 1723     | 2001     |

↑  
I/O 1 \*

↑  
I/O 2

$O(\frac{N}{P})$  pages of fence pointer data


## Variation: Unordered Data Pages

| $P_0$ | $P_1$    | $P_2$      | $P_3$      | $P_4$        | $P_5$      | $P_6$      | $P_7$        | $P_8$       | $P_9$        | $P_{10}$     | $P_{11}$     |
|-------|----------|------------|------------|--------------|------------|------------|--------------|-------------|--------------|--------------|--------------|
|       | 0<br>178 | 192<br>273 | 274<br>412 | 1285<br>1326 | 703<br>891 | 458<br>611 | 1089<br>1273 | 904<br>1063 | 1358<br>1566 | 1584<br>1773 | 1821<br>2001 |

Linked list for order

[ 192  
274  
1285  
703  
458  
1089

[ 192  $P_1$   
274  $P_2$   
458  $P_6$   
703  $P_5$   
904  $P_8$   
1089  $P_7$



What is the IO complexity of binary search?

w/FP Table

"2"

But really  $O(N)$

$\nwarrow$   $N$  divided by a BIG constant

---

Variation: Keep FP Table In memory

$\hookrightarrow$  IO Exactly

What is the Memory complexity of binary search?

w/ FP Table

1 Page + small constant

---


Variation: Keep FP Table In Memory

$O(\frac{N}{p})$  memory



# Idea 3: Layers

| $P_0$   | $P_1$   | $P_2$ | $P_3$    | $P_4$      | $P_5$      | $P_6$      | $P_7$      | $P_8$       | $P_9$        | $P_{10}$     | $P_{11}$     |
|---------|---------|-------|----------|------------|------------|------------|------------|-------------|--------------|--------------|--------------|
| Layer 2 | Layer 1 |       | 0<br>178 | 192<br>273 | 274<br>412 | 458<br>611 | 703<br>891 | 904<br>1063 | 1089<br>1273 | 1285<br>1326 | 1358<br>1566 |





Directory  
Pages



Layer 2



ISAM Tree

index  
+  
clusters  
e  
+  
n  
o  
d  
e

Layer 1



$P_3$   $P_4$   $P_5$   $P_6$   $P_7$

$P_8$   $P_9$   $P_{10}$   $P_{11}$

|     |     |     |     |     |
|-----|-----|-----|-----|-----|
| 0   | 192 | 274 | 458 | 703 |
| 178 | 273 | 412 | 611 | 891 |

|      |      |      |      |  |
|------|------|------|------|--|
| 904  | 1089 | 1285 | 1358 |  |
| 1063 | 1273 | 1326 | 1566 |  |

Data  
Pages

How deep is the tree?

What is the IO complexity of binary search?

What is the Memory complexity of binary search?

What needs to change to allow changes?

# Idea 4: B+ Trees