

UPC code

Products | Name | ~~Id~~ | Description | Price

[string]

- Most expensive product?
- Are there missing descriptions?
- List the products in Alphabetical order
- What is the price for a given UPC code

let records: Vec<Product> = _____

for (record in records) The thing I want ~

{ if (record.upc == " ") }

↳ return record.price

}
}

SELECT price] ← what I want (outputs)

FROM products] ← where it comes from

WHERE upc = "41"] ← what I want (records)

[ORDER BY] ←

$(10^9)^4 \approx 2^{40}$
1 Trillion ints (unique)

↳ If, an int part of that set


64bit


$$2^{40} \text{ ints} = 2^{48} \text{ bytes} \quad \frac{1}{4} \text{ petabyte}$$


Assumptions

— We can sort the data


binary search
over on-disk
array








A hand-drawn sketch of a zigzag line, consisting of several connected horizontal and vertical segments, resembling a sawtooth pattern.





$$\# \text{ of } I_{O_s} = \text{Log}(N) - B$$

F

S > V C M





$$y = mx + b$$



$$y = \frac{a + bx + cx^2}{}$$

$$y = f(x) \quad \leftarrow \text{CDF}$$

← Page →




a b | pet



← 10 24 4 bytes in 1/2

← v?

Binary search: V less than M




← 1PB sorted data

↑ filesystem read

1024 samples at 1TB intervals
(1 page)




← 1PB sorted data





if V is between b and c , then V must be in the 1TB - 2TB range

1024 samples at 1TB intervals
(1 page)



← 1TB sorted data





ISAM
 e + n o d
 t r u c t u r e
 a d e x

- 1
- 10^8
- 10^{18}
- 10^{28} 1K Pages
- 10^{38} 1K Pages
- 1 TB

$O(\log_{1024}(N))$ vs $O(\log_2(N))$ for B+ search
 6 steps 2^{38}

No space in between



4 byte int

Stored at addresses
divisible by 4




size of

Stored in sorted order


record is known & fixed

Problems


- No variable length records
- Expensive inserts/deletes



but need ← Can alloc at end
 a way to by making
 free pages in file bigger
 the middle



get from free list
 O(1) compute ID [(> alloc page (/ → page id) → make file bigger
 (> free page (page id) → (/] → prepend to free list
 (> get page (page id) → start byte




Option 1: N -element array of records
 where $N \cdot \text{sizeof}(\text{record}) < 4\text{kb}$

↳ assumption: $\text{sizeof}(\text{record})$ is fixed and known

- 2 general assumptions
- $\text{sizeof}(\text{record}) < 4\text{kb}$
 - each page is dedicated to one specific format

Option 2: Delimited records
 (N in CSV)



Option 3: Array of pointers to first byte of variable-size record

Locating Records


— < Page #, Record # > → Array → Which pos in array?
→ Delimited → How many preceding delims?
→ pointer → Which pointer to follow?

→ find the 4kb chunk in the file

Get

<3, 2>

② find pointer



① Find page (3)


← record →



Predictable sized fields → Option 1

Variable sized fields → Option 2

Option 3 ← mix of both



Variable sized component starts at byte 93