

CSE 350

Advanced Data Structures


Topic 13: Indexing

Bulk loading trees

P_0	P_1	P_2	P_3	...
<29, ...>	<35, ...>	<100, ...>	<17, ...>	...
<28, ...>	<43, ...>	<42, ...>	<61, ...>	
<12, ...>	<30, ...>	<56, ...>	<32, ...>	
<95, ...>	<92, ...>	<33, ...>	<34, ...>	

$O(N)$ reads (each page once) = $O(N)$ I/Os
 $O(N)$ writes (each B+Tree page once)

8np
 part
 per
 layer = $\log N$
 Pages



<12, ...>	<30, ...>	<35, ...>	<61, ...>	
<17, ...>	<32, ...>	<42, ...>	<92, ...>	
<28, ...>	<33, ...>	<43, ...>	<95, ...>	
<29, ...>	<34, ...>	<56, ...>	<100, ...>	

Clustered



12	→ P ₀
17	→ P ₀
34	→ P ₃

0	12	17
R	28	29

30	33
32	34

directory pages


leaf pages

data pages

B+ Tree

~~R → tuple~~

R → address



Unclustered

City


SELECT Count (*)

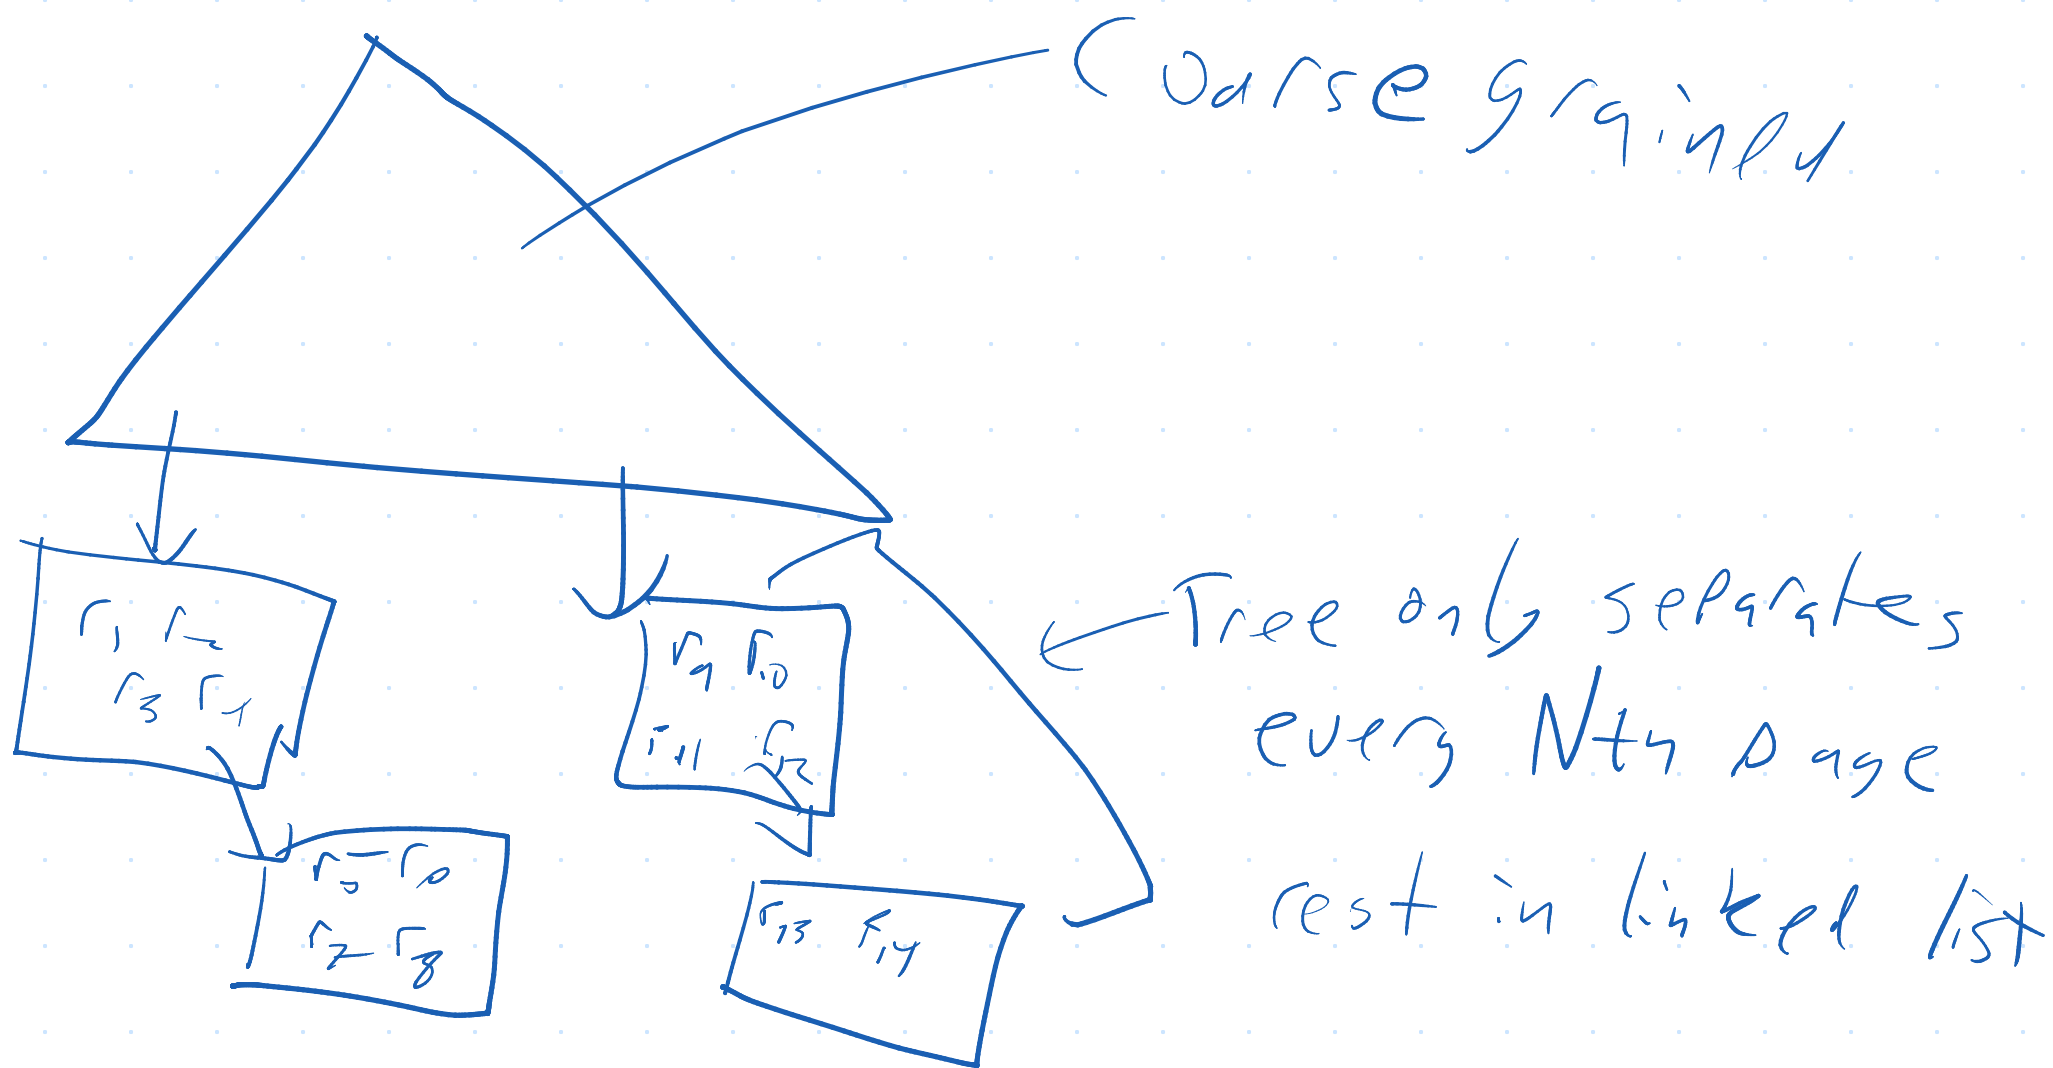
WHERE City = Towanda

↳ Index has one entry per record already
Data not needed!

Clustered vs Unclustered

Dense, Sparse, Coarse-grained Indexes

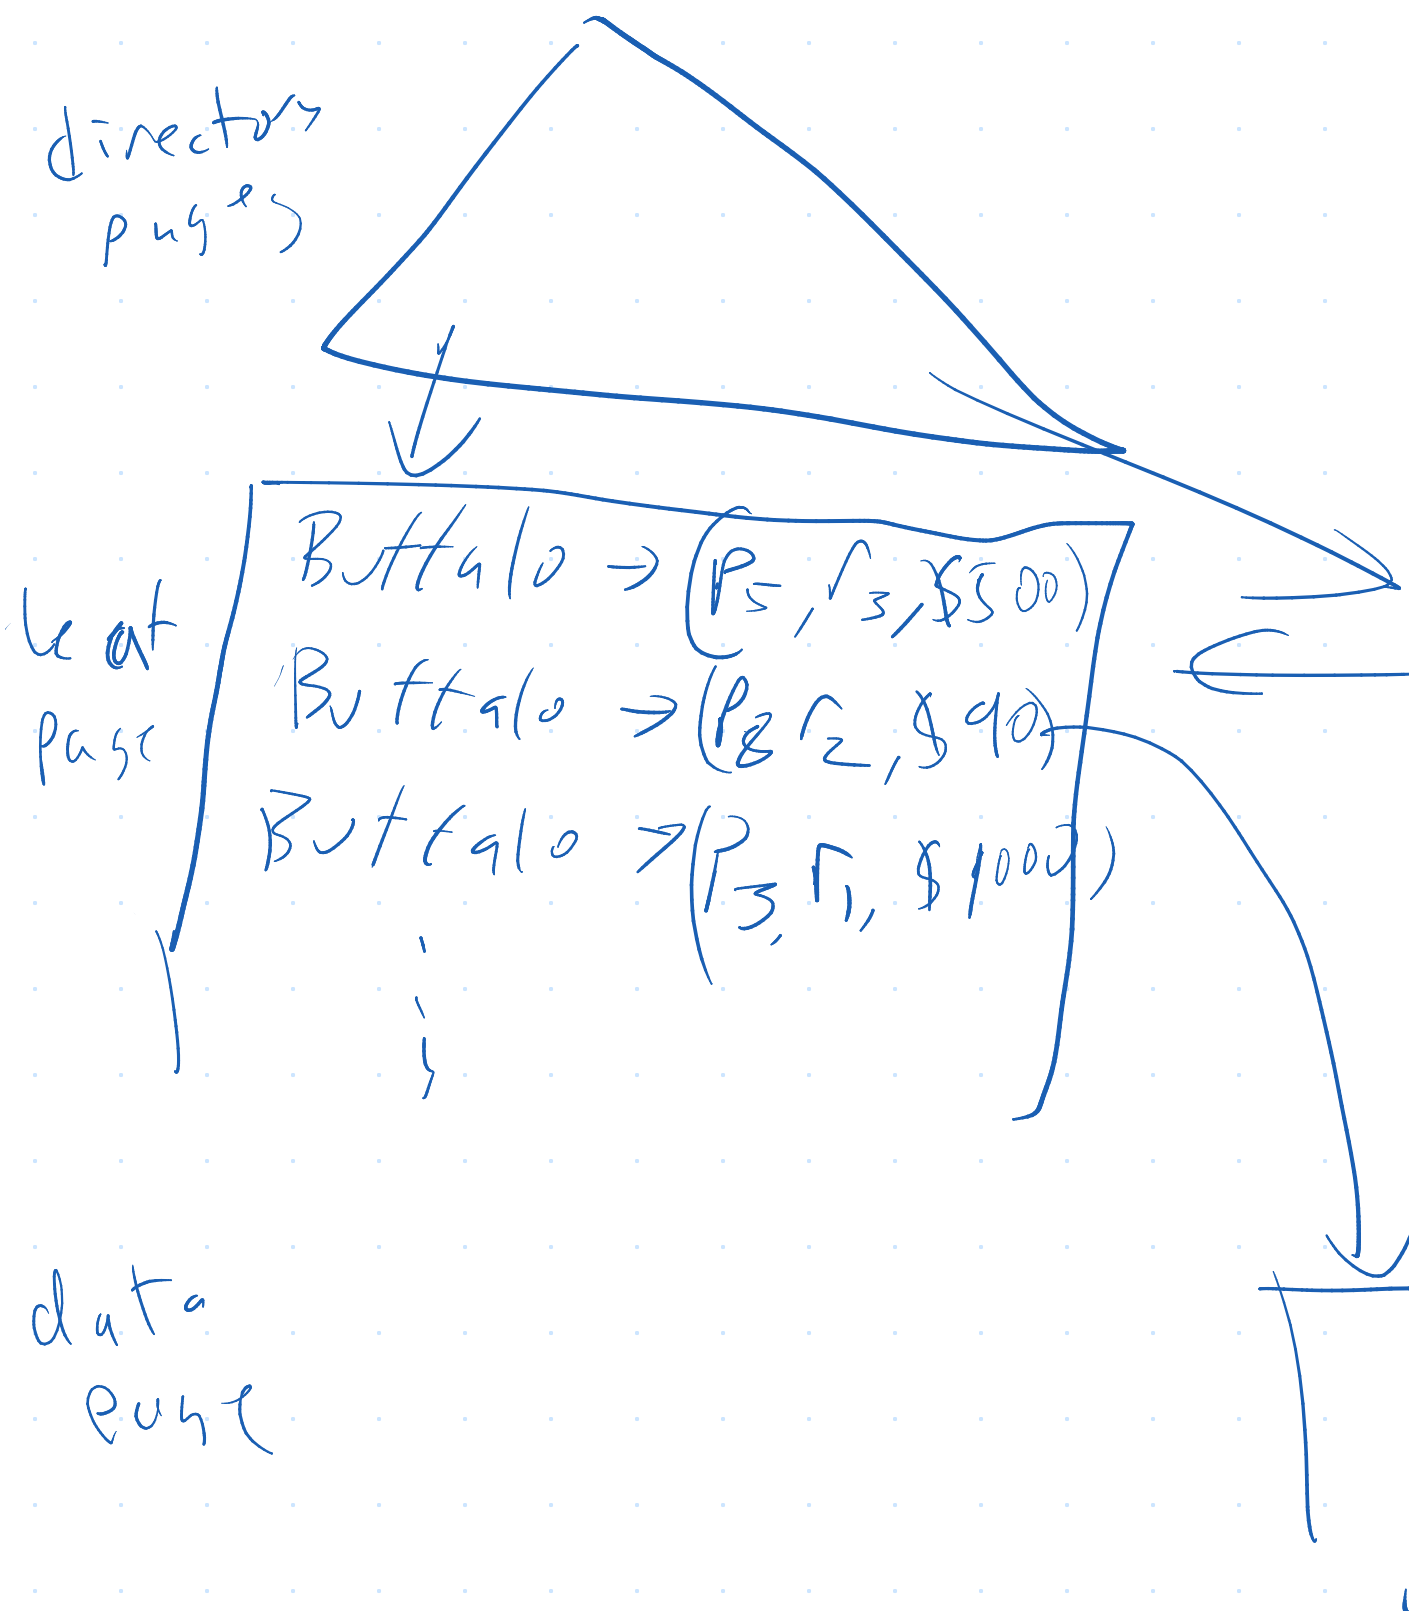




Useful

- Much duplication
- queries looking for large ranges

Covering Indexes



SELECT SUM(PRICE)
FROM Permits
WHERE CITY = Buffalo

City, price are covered
covered attrs in query
can be answered only
w/ index
w/o data pages

Composite Indexes

(City, Price)

Index w/ multiple Attrs as a key

↳ by lexical sort on City, Price in a B+Tree

useful for filtering on

- = City, \geq city

- = City = Price

- = city \geq Price

But not as

just price