



Sequential scan



$$\text{COST} = \text{PAGES}(R)$$

Index Scan




$M=2$
total

$\text{COST} = \text{COST OF SCANNING INDEX}$
 $+ \text{COST OF \# DATA PAGES READ}$

DISTINCT (DUPLICATE REMOVAL)

SELECT DISTINCT A, B
FROM R;



COST = PAGES(R)

If unique A, B values
fit in memory

GROUP BY

SELECT A, SUM(B)
FROM R
GROUP BY A




$M=1$

Is A new?

No, Update
Sum(count
of B)

YES $\rightarrow$ A, SUM(B), count=1



$M-1$

COST = PAGES(R)

Assuming $M-1$

is enough to

keep the
result in
memory)