

SELECT

CREATE TABLE (

Studio string

Title string

Lang

Director

Star

Rating Int

);

A	B	C	D
1	2	3	4
5	6	7	8
:	:	:	:

Movie	Studio	Title	Lang	Director	Star	Rating
-------	--------	-------	------	----------	------	--------

- "Tidy Data" (Normal forms)

Movie	Title	Lang	Runtime
	Princess Bride	English	114
	Princess Bride	Spanish	114
	Die Hard	English	128 126
	Die Hard	Polish	128
	Matrix	English	112

Movies	Title	Lang	Runtime
	Princess Bride	[Eng, Spa]	114
	Pie Herd	[Eng, Pol]	126
	Matrix	[Eng]	112

document
database

Lang	Title	Lang
PB		Eng
PB		Spa
DH		Eng
DM		Pol
M		Eng

Runtime	Title	Runtime
	PB	114
	PH	126
	M	112

Functional Dependency (FD)

$R(A B C D E \text{ ---})$

$\{A, B\} \rightarrow \{D, E\}$

If I know
A and B,

I can uniquely
obtain D, E

$\{\text{Title}\} \rightarrow \{\text{runtime}\}$

~~$\{\text{Title}\} \rightarrow \{\text{Lans}\}$~~

$$\{A\} \rightarrow \{D, E\}$$

↓ implies

$$\{AB\} \rightarrow \{D, E\}$$



$$\sim\{A\} \rightarrow \{D, E\}$$

↓ implies

$$\{A\} \rightarrow \{D\}$$

Rules

For

FDs

$$\{A\} \rightarrow \{D, E\}$$

↓

$$\{A\} \rightarrow \{A, D, E\}$$

$R(A_1, A_2, A_3, \dots, B_1, B_2, B_3, \dots)$

$\{A_1, A_2, \dots\} \rightarrow \{A_1, A_2, \dots, B_1, B_2, B_3, \dots\}$

(Super)Key

If I know these

I can uniquely ID
the entire record

Title + lang

vs

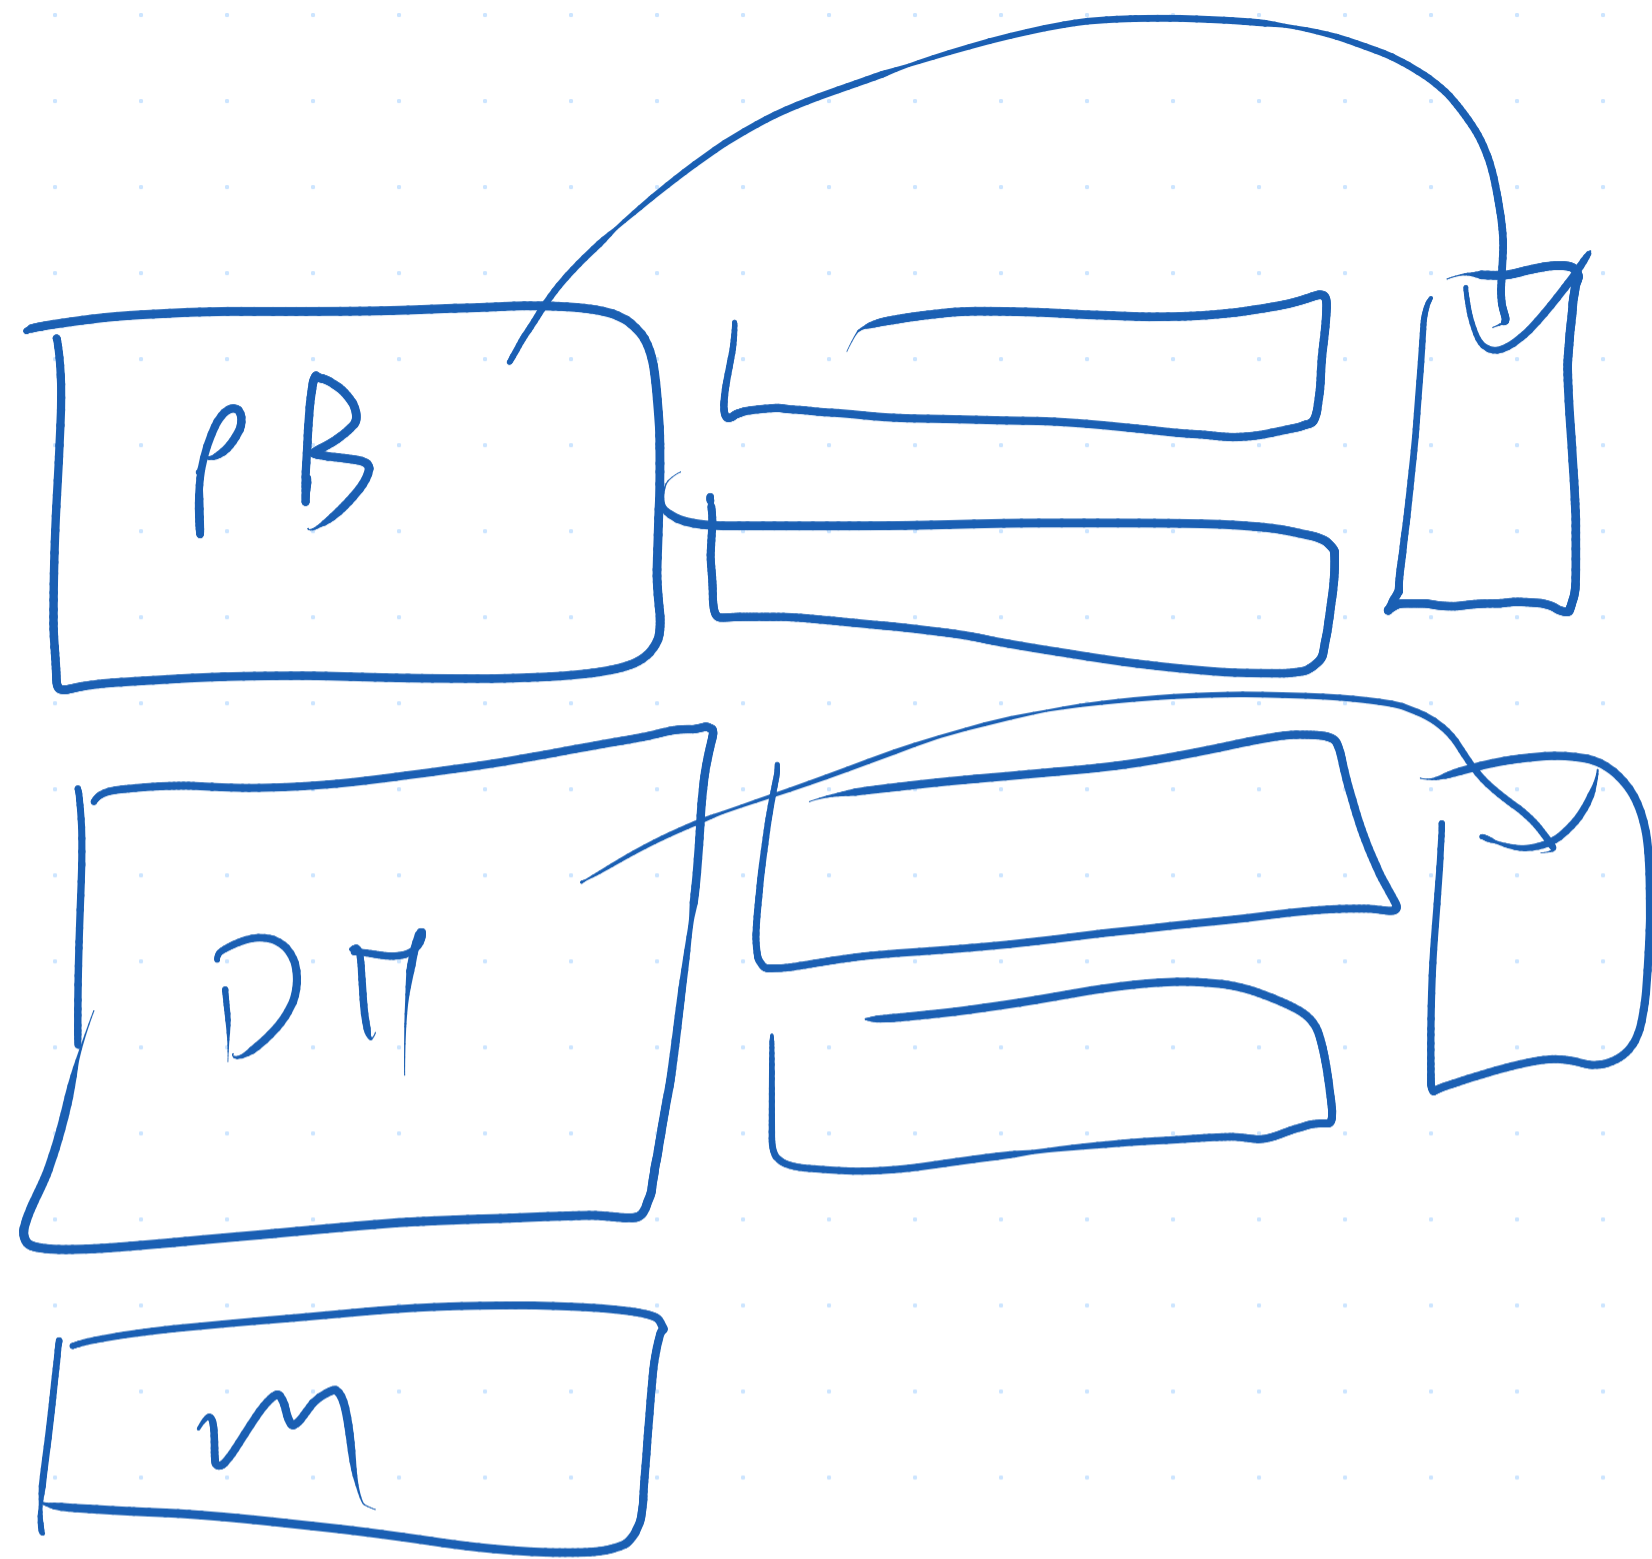
Title + lang + Runtime

] key



← superset of

] Superkeys, minimal key



FD from
 $\{Title\} \rightarrow \{Runtime\}$



Not all attr's in
 table

CHS is not a
 Key (or superkey)
 ↑

2nd Normal form
 Property

Dataset

- 1 Normal Form \rightarrow No nesting
No lists, etc. ~
- 2nd Normal form \rightarrow Any FD* must
have a superkey
on the left

①

Lang

NATURA L JOIN

Runtime

Movies

PB Eng

PB span

DH Eng

DH Act

M Eng

PB 114

DH 126

M 112

—

—

—

—

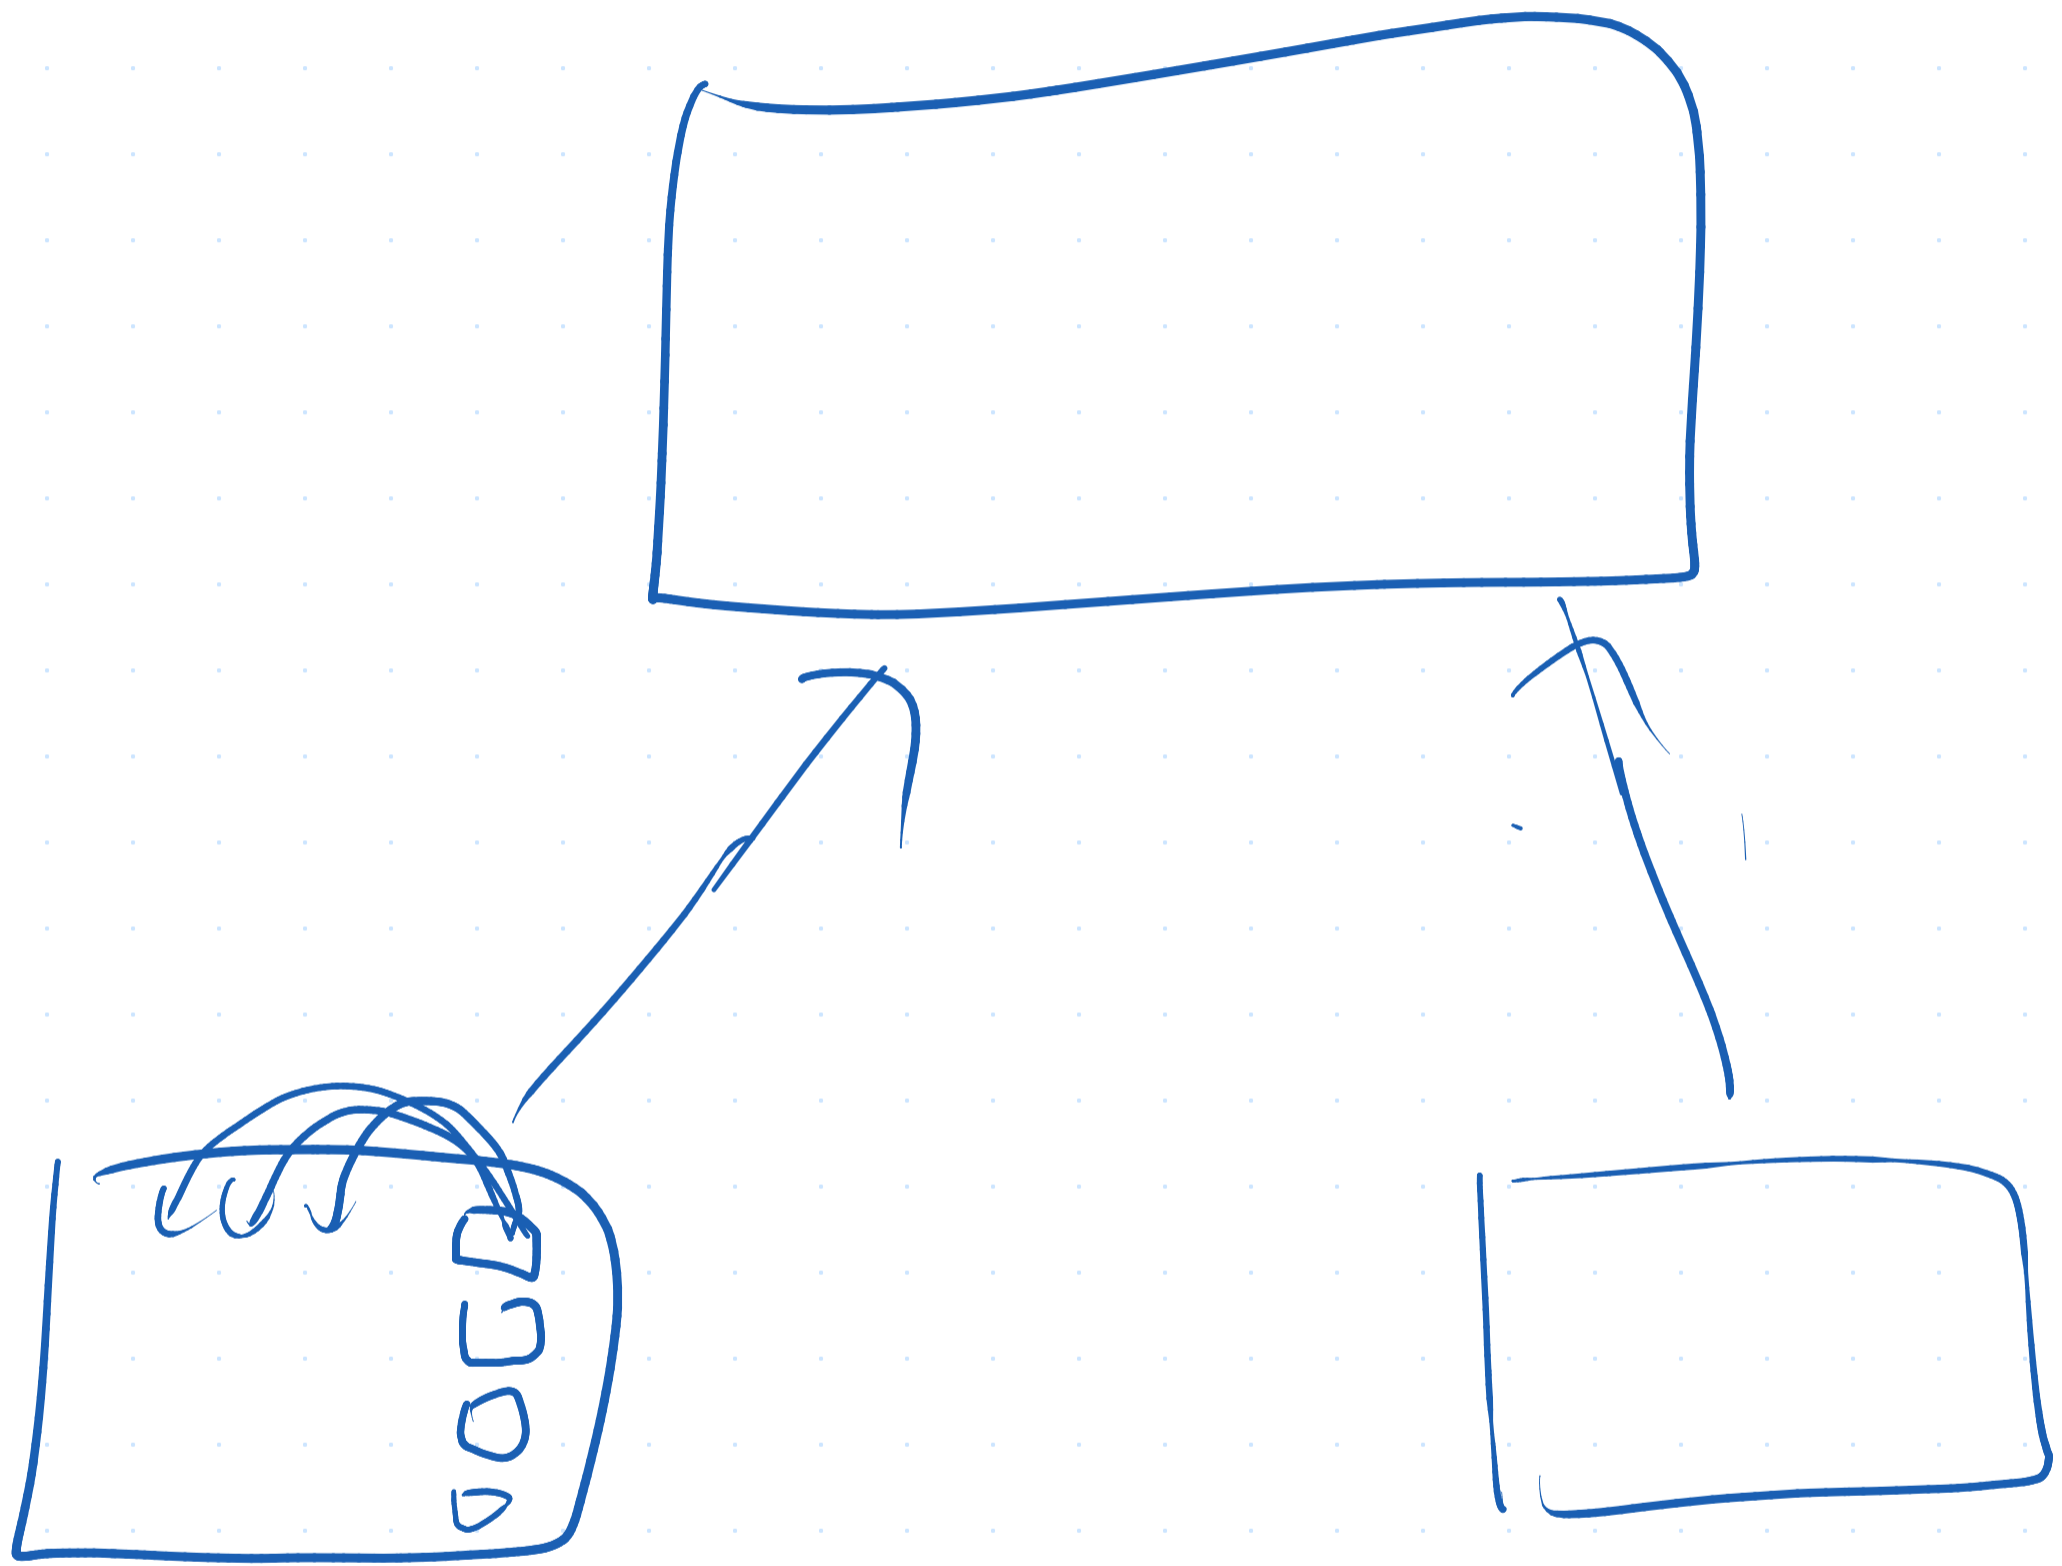
—

A Join B ON A_1, A_2, A_3

↳ For each record in A
pair it up with every
record B that has matching
values of A_1, A_2, A_3 -

NATURAL JOIN


↳ All common Attrs



Runtime | Title Runtime

Language | Title Language


Runtime
114



Language
English
Spanish

Cartesian Product

R	A	B
m	1	1
n	2	1
o	3	2




S	B	C
x	1	3
y	2	4
z	2	5

$$R.B = S.B$$

$$R.B \subseteq S.B$$

R x S	A	R.B	S.B	C
m x	1	1	1	3
m y	1	1	2	4
m z	1	1	2	5
n x	2	1	1	3
n y	2	1	2	4



← join

Equi join $R.B = S.B$

Natural join \downarrow
 $attr_s(R) \cap attr_s(S)$

~~$R.B = S.B$~~

\downarrow
 ~~\downarrow~~ ← Natural Join

TMovies = Runtime \bowtie Languages

≡ SELECT Title, Runtime, Language
FROM Runtime NATURAL JOIN
Languages

≡ SELECT Title
FROM Runtime, Languages
WHERE Runtime.Title = Lang.Title

filter (R x S)

|||
R x S

(A + B) * C
|||

C * A + C * B


$|R|$

compute $R \times S$?

$|S|$

$\hookrightarrow O(|R| \times |S|)$

runtime complexity



$R \bowtie S$

filter on θ

for r in R ;

for s in S ;

if $\theta(r, s)$;

output($r \parallel s$)

Records per page $\rightarrow \left[\frac{|R|}{P_r} \right]$ records from R

$\rightarrow \left[\frac{|S|}{P_s} |R| \right]$ records from S

$$\text{Runtime} = O(|R| + |S|)$$

$$I/O = O\left(\frac{|R|}{p} + \frac{|S| \cdot |R|}{p}\right) = O\left(\frac{|S| \cdot |R|}{p}\right)$$

$$\text{Memory} = O(1)$$

← file ← file
 $R' \in \text{sort}(R)$

$S' \in \text{sort}(S)$

↖ e.g. 2-pass Merge Sort

on join key e.g. B in $R.B = S.B$




Table 1



①

①



②

②

②

②

Table 2




$O(N)$
rec

Merge-Join

$O(N)$
rec

U

U



$\curvearrowright O(N^2)$

Merge Join

Runtime:

$$O(N^2)$$
$$\sim (N+N)$$

Assumes sorted
tables

+ $O(N \log N)$
otherwise

$$O(|T_1| + |T_2| + |T_1 \cap T_2|)$$

$$\text{if } |T_1 \cap T_2| = O(N)$$

$$\text{and } |T_1| = O(N)$$

$$|T_2| = O(N)$$

$$\hookrightarrow O(N)$$

IO complexity

$$O(\underbrace{|T1| + |T2|}_{\text{reads}} + \underbrace{|T1 \cap T2|}_{\text{writes}})$$

Memory complexity

$\uparrow + N \log(N)$
if unsorted



$$O(\uparrow)$$

Sum of fanout bounds
of the two tables

B




↪



$x+1 \% 2 \text{ } 2$

Hash Join (2 Pass)



$O(N)$

\approx IO and Runtime

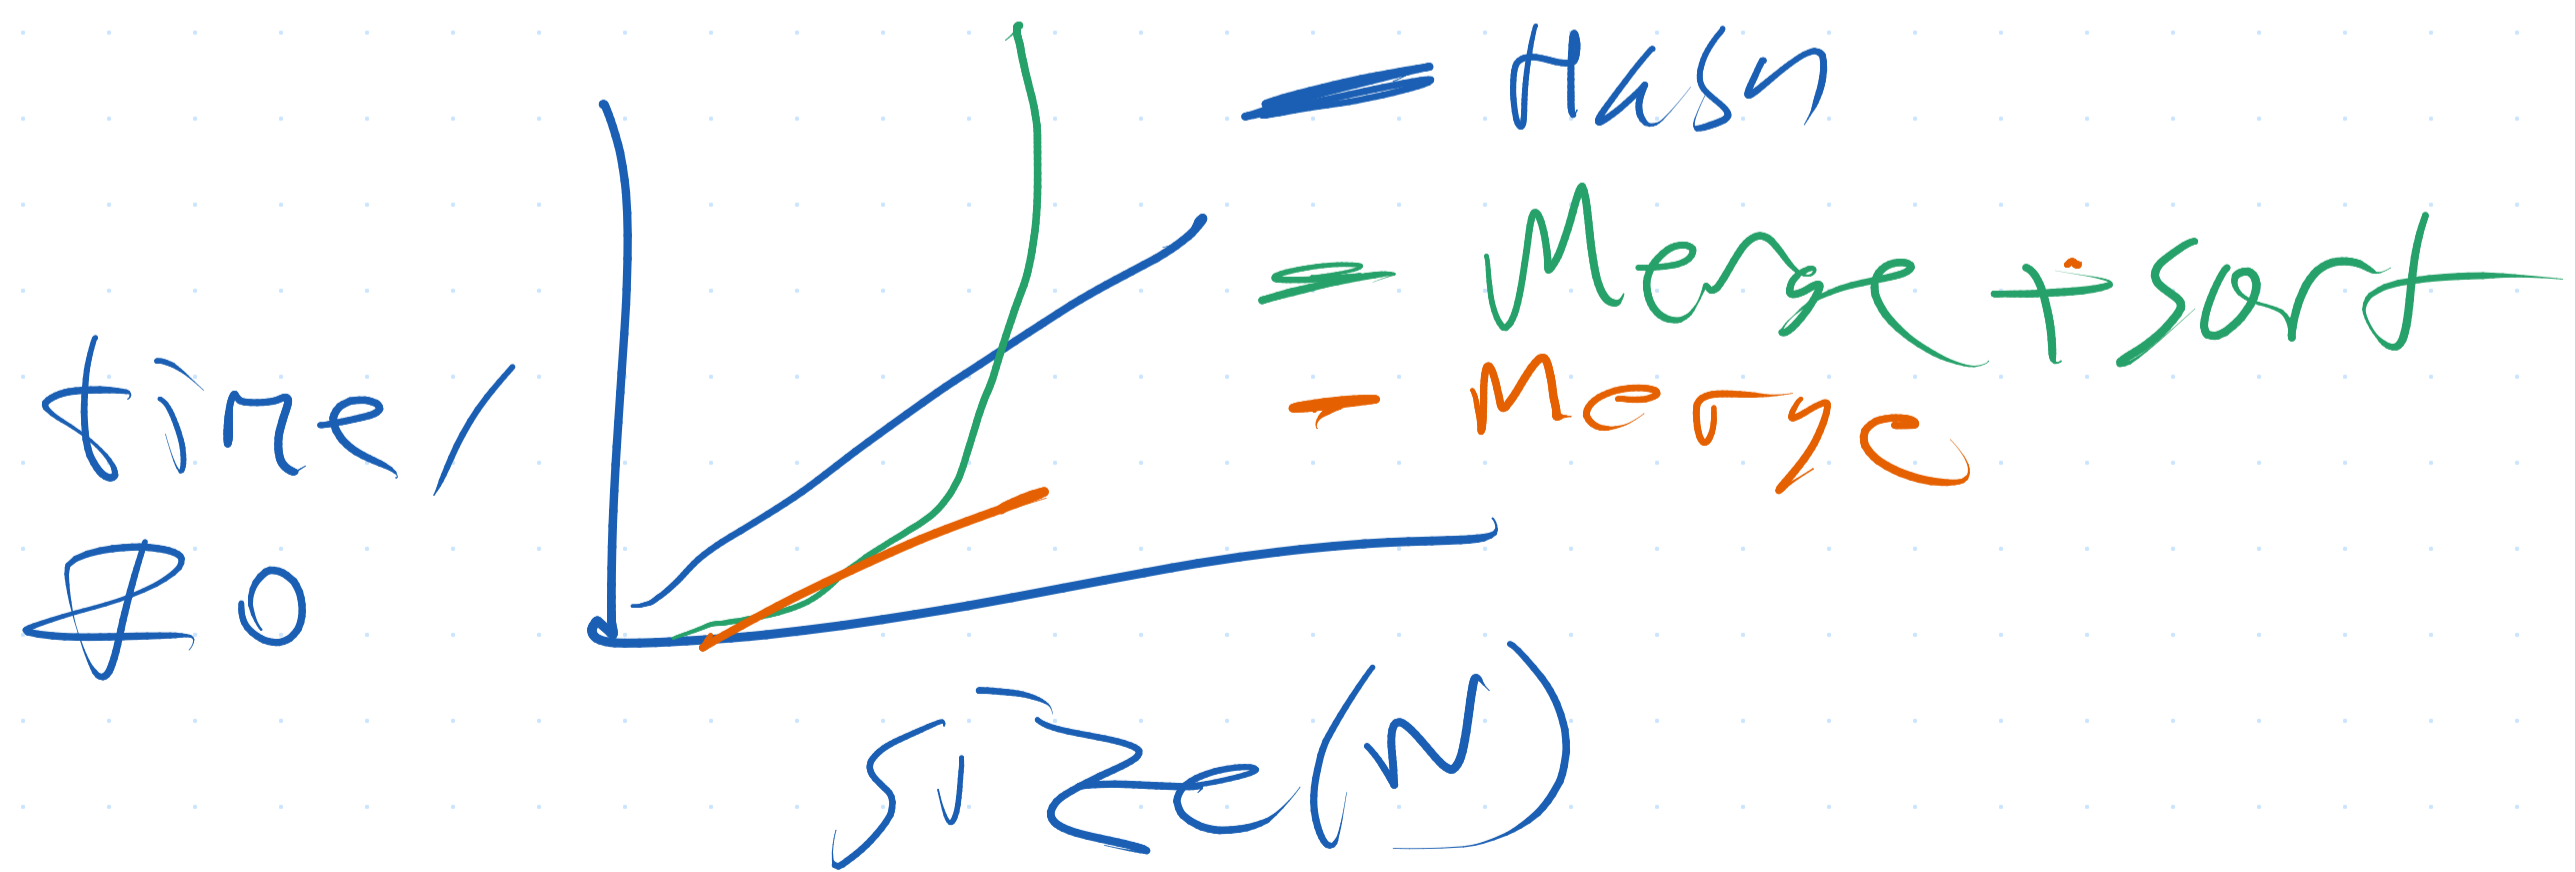
k bins \rightarrow

$\frac{N}{k}$ records per pass


2) m m


Total I/O = $O(|T_1| + |T_2| + |T_1 \cap T_2|)$ $O(\frac{N}{k})$ expected records per page
(and runtime)

one extra round of write + read (higher constant)



Join






$R \cdot key$
 $= S \cdot key$

Nested Loop
Join

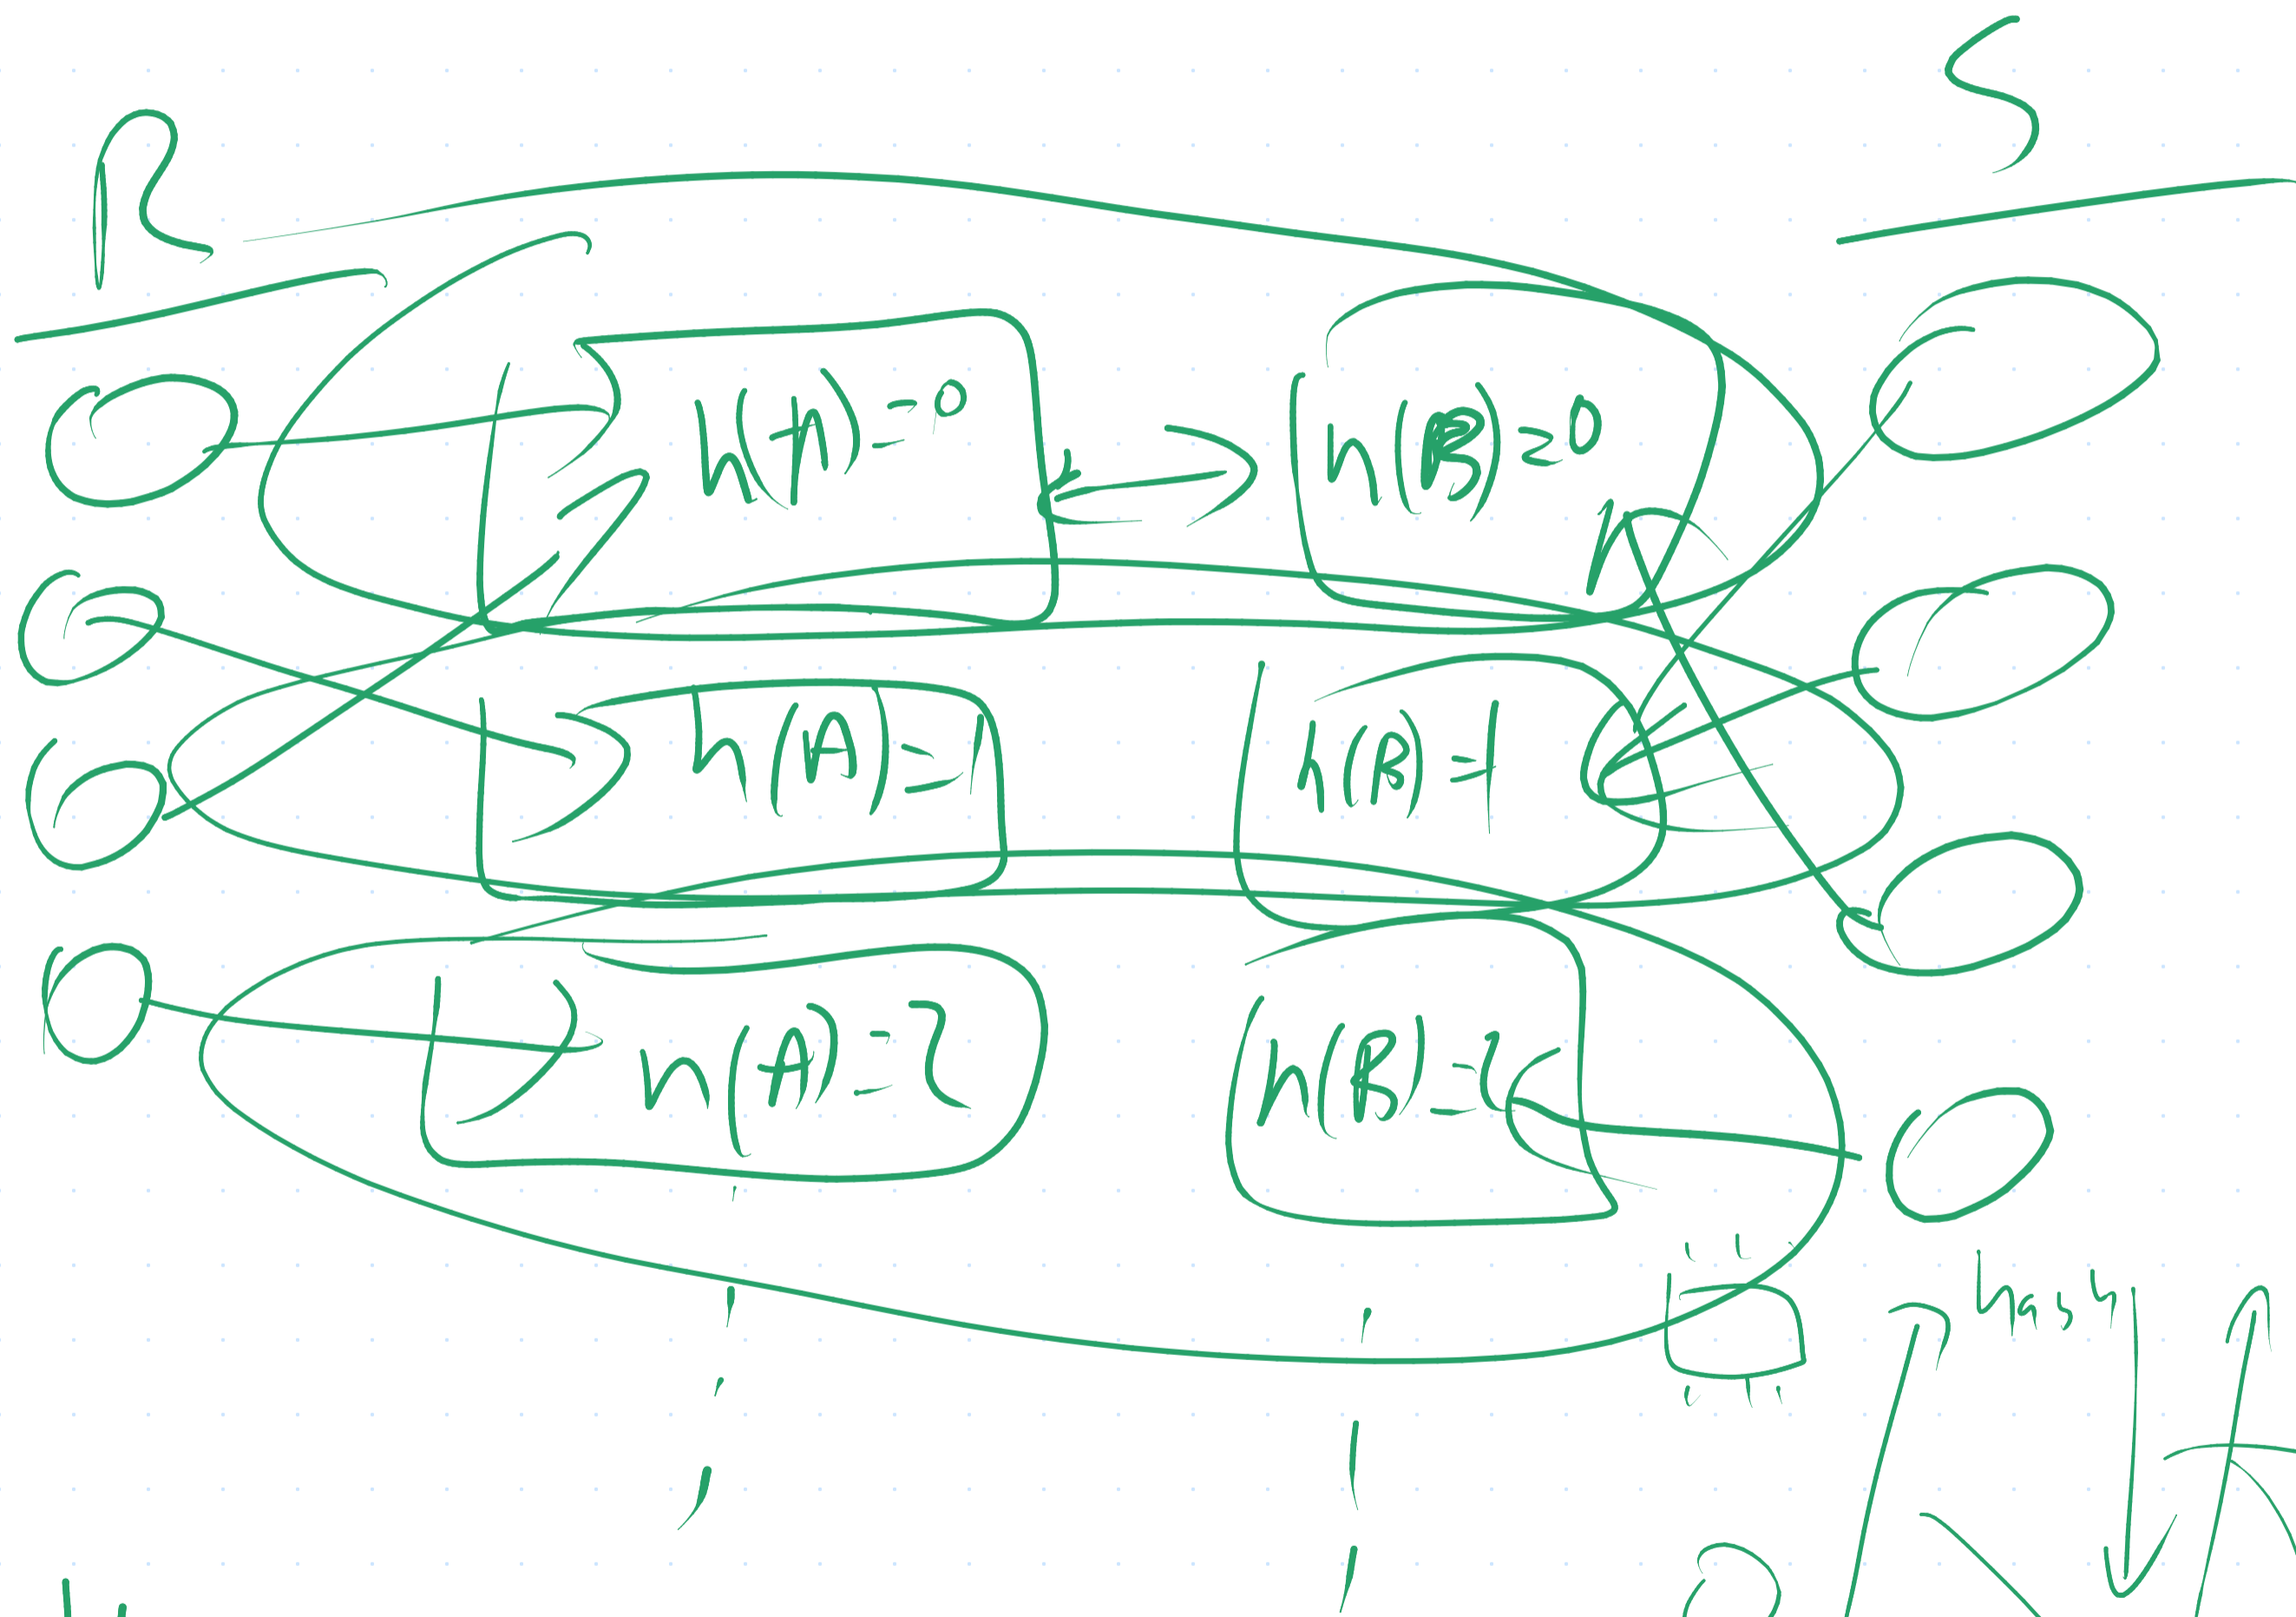
for r in R
for s in S
if $test(r, s)$
emit r, s



for block in \mathcal{A} need
 for s in S need
 for r in block \mathcal{M}
 if test r, s
 emit r, s

Many tests


$$\rightarrow R.A = S.B$$



2 Pass Hash

hash join
2N reads N writes

$$R.A = S.B$$



Sort increasing on A


Sort increasing on B

Similar to sort-merge
[sort] Merge Join

R N records

$R \neq S.B$

S M records



One-pass (Grace) Hash Join

Nested Loop

Block Nested Loop

Sort Merge (w/o sort)

2 Pass Hash

1 Pass Hash

W

$N \cdot M$

$N \cdot M$

$O(1)$

$O(N+M)$

$O(W+M)$

L

$N \cdot M$

$\frac{N \cdot M}{B}$

$O(1)$

$O(W+M)$

$O(N+M)$

$M \cdot N$

$O(1)$

B


$O(1)$

B

$|R|$ or $|S|$

smaller constant than 2 pass

Movies	Title	ID	Runtime	Hours	Genre
	A1	A2	A3	A4	A5



Movies	Title	ID	Runtime	gross	Genre
	r1				
	r2				
	r3				



1 billion records @ 50 Bytes/record

↳ 50GB

Q: Average Runtime of Action Maps

Genre + Runtime: 12GB

меню

A7

→ [(rid, value), ...]

Movies =

Title (RIID, value)
ID (RIID, value)

~~Title~~ ~~ID~~ ~~Runtime~~
~~Gross~~ ~~Genre~~

Runtime (RIID, value)

Gross (RIID, value)

Genre (RIID, value)

Runtime ~~Gross~~

Idea 1

No RID

↳ Use position in Array instead

Pros - Less overhead (No RIDs)

Cons - Fixed sort order, Expensive to find
RIDs for values
- Harder to use w/ variable-sized data

Example: Genre

Categorical Data

↳ "Sci-Fi" → 0

"Action" → 1

"Drama" → 2


"Thriller" → 3

8 bytes


↓

1 byte

Dictionary




↳ Dictionary encoded data



'0' for the next 50 entries $\langle 0, 50 \rangle$
'1' for the next 100 entries $\langle 1, 100 \rangle$
...

Run-length encoding

193?



(0, 50)	(1, 100)	(2, 80)	(3, 75)	(4, 120)
---------	----------	---------	---------	----------

0 → 0 49 → 0 50 → 1 149 → 1, 150 → 2

↳ Pro/Small

Can/Requires sorted runs in the data

0 → { 0, 10, 13, 23, 27, }
value record ID
45 45 45 45
 $10^7 781 0^8 10^9$

1 → { }

2 → { }



0, 10, 13, 23, 27

└──┬──┬──┬──┘

10 3 10 4