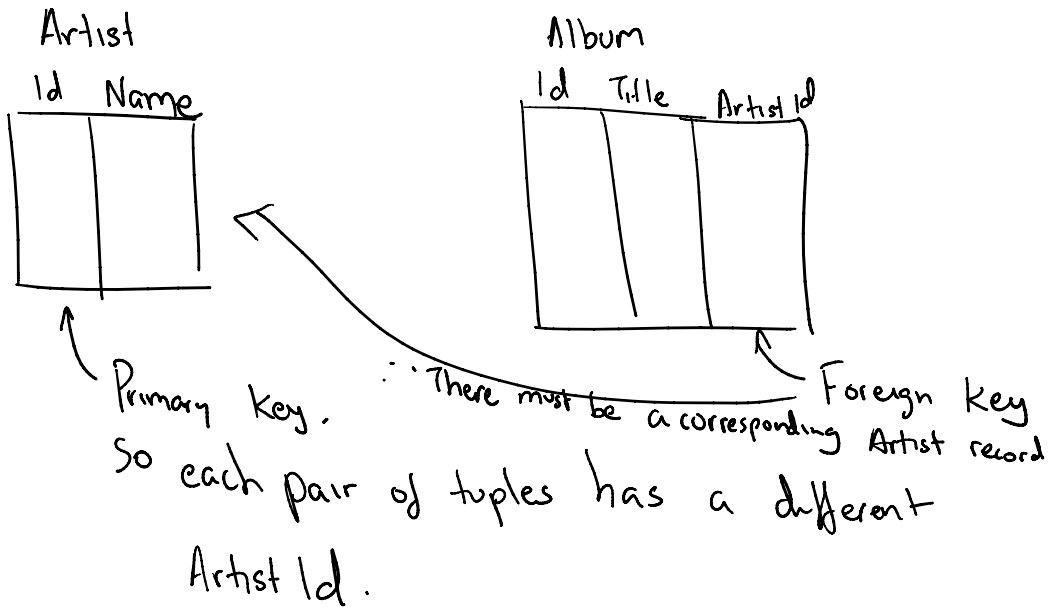
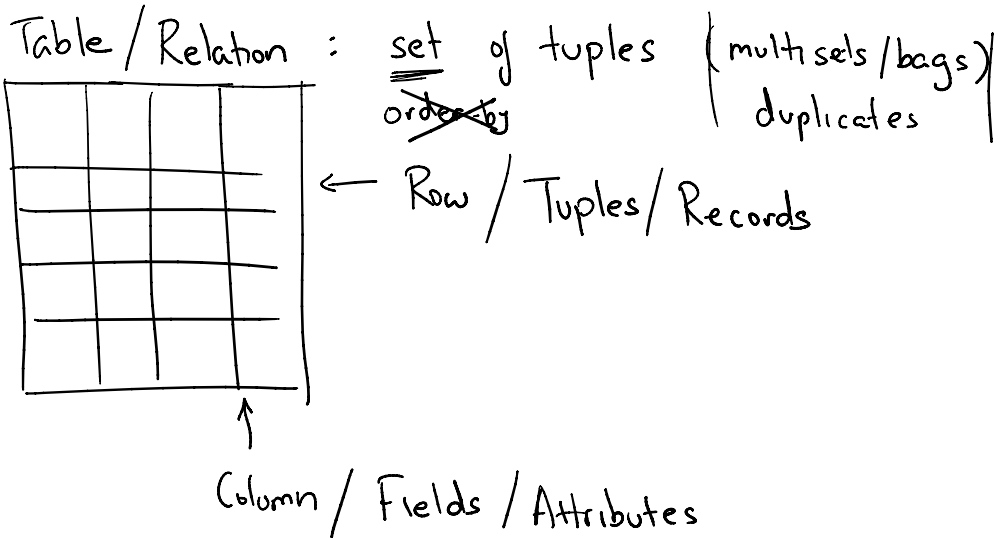


Unit 3: Programming with Relations

SQL / Relational algebra



SQL: select, insert, update, create table

↑

TD. 1.11


... , update , create table

Drop table


(select _____
from Table₁, Table₂
where _____)

Union
Intersects


Sets : Union



Intersection





Selection (where)



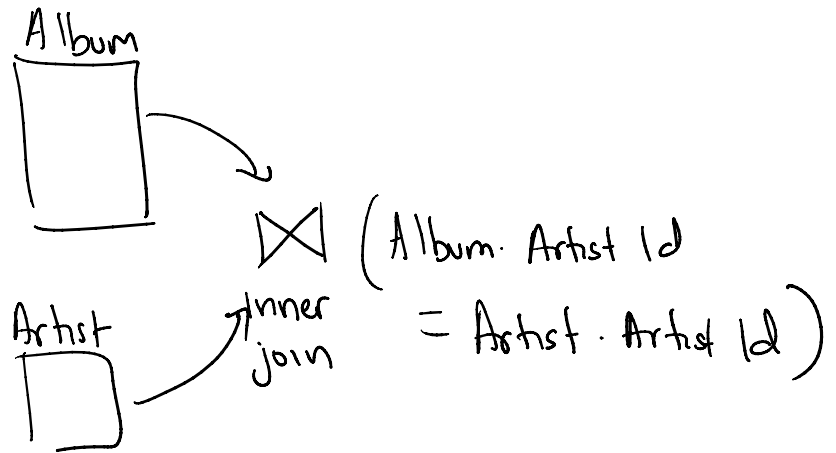
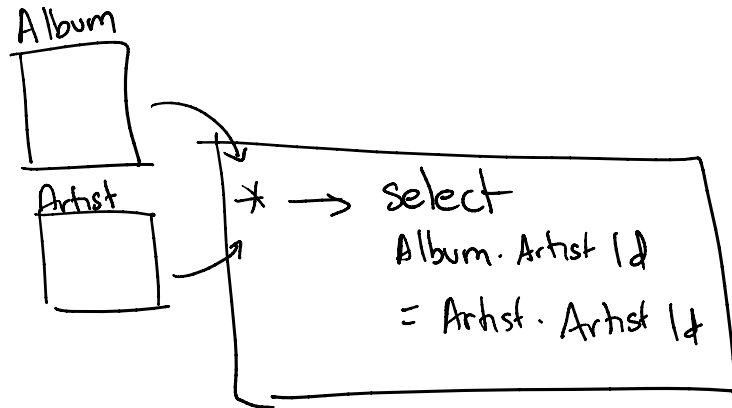

pred

Projection (Select)



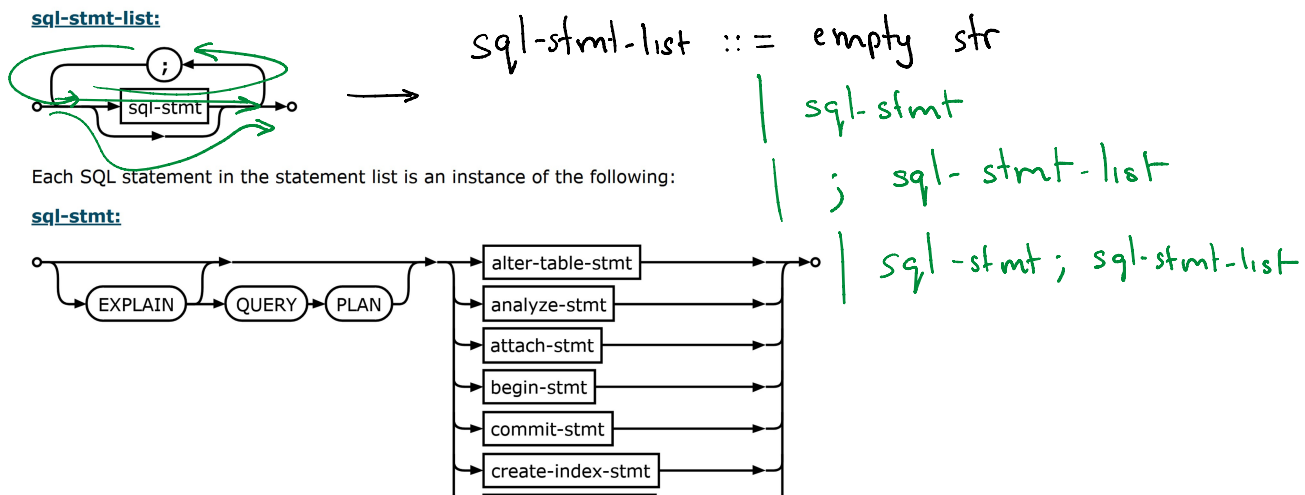


Cartesian product



RA	Select	Project	Joins, Cartesian product	Union	Intersectm
SQL	where	select	Join(s)	Union	Intersect where exists
	σ	Π	\bowtie *	\cup	\cap
			\Join		difference
					minus

Railroad Diagram of SQLite's SQL dialect



Defining Relational Algebra

Table = set of tuples

$Q ::= T$ (Just a table)

| $Q_1 \cup Q_2$ (Union)

| $Q_1 \cap Q_2$ (Intersection)

| $Q_1 \setminus Q_2$ (Set difference)

| $\sigma_{\text{filter}} Q$ (Select)

| $\pi_{\text{cols}} Q$ (Project)

| $Q_1 * Q_2$ (Cartesian product)

| $Q_1 \bowtie Q_2$ (join)

inner outer left outer natural

Ex: All employees celebrating their birthday today

$\sigma_{\text{dob} = \text{today}}$ Employee

Ex: Funding a travel subsidy

σ salary < 90K Employee
 \wedge daily commute > 40min

Ex: Interstate commuters

Employee		
Name	Home Zip	Work Zip

Zip	
Zip	state

$$\sigma_{\substack{z_1.s \\ \neq z_2.s}} \left(\Pi_{E.N, z_1.s, z_2.s} \left(\left((E \bowtie_{E.H2=z_1.z} Z_1) \bowtie_{E.W2=z_2.z} Z_2 \right) \right) \right)$$

Evaluating RA Queries (at 20000ft)

$\sigma_f Q$: result := \emptyset
 for t in Q :
 if $f t$:
 result := result $\cup \{t\}$

$\pi_{cs} Q$: ...

$Q_1 \bowtie_{Q_1 \cdot c_1 = Q_2 \cdot c_2} Q_2$: for t_1 in Q_1 :
 for all matchable t_2 in Q_2 :
 emit ($t_1 \cdot t_2$)


Loop Nest Join Accelerated
 using indices

Hash join

Question: How do we evaluate joins quickly?

Triangle Query

select n_1 n_2 n_3
 from edge e_1 edge e_2 edge e_3
 ...



from edge e_1 edge e_2 edge e_3 ! /
 where $e_1 \cdot v_1 = e_2 \cdot v_2$ and $\begin{matrix} 6 \\ 7 \end{matrix}$ Edge
 $e_2 \cdot v_1 = e_3 \cdot v_2$ and v_1 v_2
 $e_3 \cdot v_1 = e_1 \cdot v_2$ 1 5
 $n_1 = e_1 \cdot v_1$ $n_2 = e_2 \cdot v_1$ $n_3 = e_3 \cdot v_1$ 1 2
5 2
2 3
4 6
⋮

for $e_1 \in \text{Edge}$

for $e_2 \in \text{Edge}$

if $e_1 \cdot v_2 = e_2 \cdot v_1$ and

if $(e_2 \cdot v_2, e_1 \cdot v_1) \in \text{Edge}$,

then emit $(e_1 \cdot v_1, e_2 \cdot v_1, e_2 \cdot v_2)$

Triangle Query :

— There is an algorithm which takes $O(n^2)$ time

— There exist Edge tables where result contains $O(n^{1.76})$ tuples

result contains $O(n^{1.76})$ tuples