DSAA 5012
Advanced Database Management for Data Science

LECTURE 5 EXERCISES
RELATIONAL ALGEBRA

EXAMPLE RELATIONAL SCHEMA

Sailor(sailorId, sName, rating, age)

Boat(boatId, bName, color)

Reserves(*sailorId, boatId*, rDate)

What is the E-R schema for this relational schema?

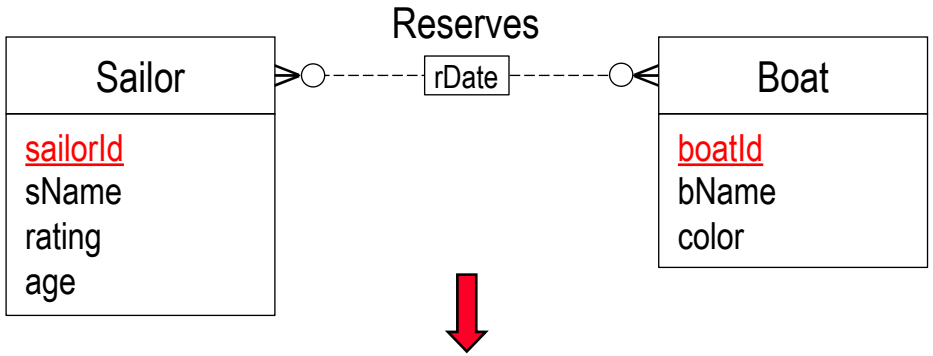

What about this schema?

Reserves(*sailorId, boatId*, rDate)

☞ **A sailor can reserve a given boat at most once!**

What do we get if we reduce Reserves?

☞ **rDate is not part of the key in the reduction!**

EXAMPLE RELATIONAL SCHEMA

Sailor(sailorId, sName, rating, age)

Boat(boatId, bName, color)

Reserves(*sailorId, boatId*, rDate)

What about this schema?

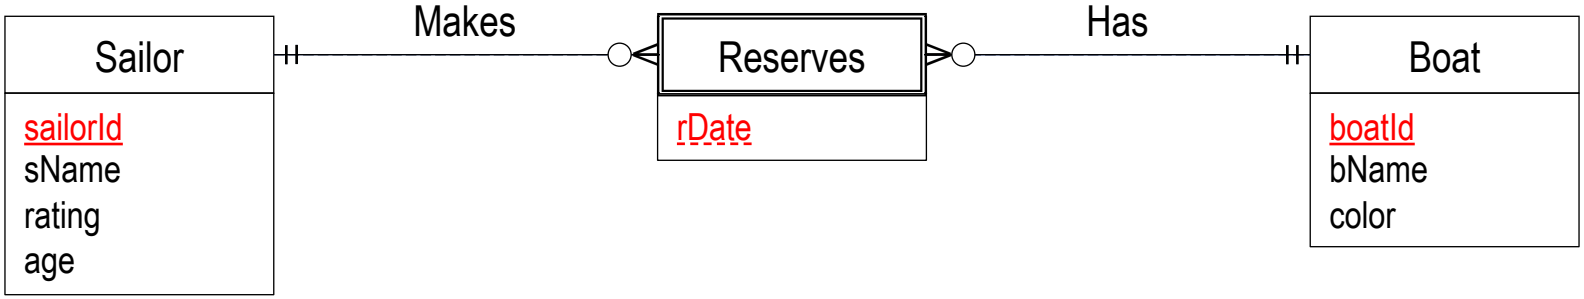

What kind of entity is Reserves? ⇒ Weak entity.

On which entity is Reserves dependent? ⇒ Both Sailor and Boat!

Is rDate a discriminator for Reserves? ⇒ Yes

What should be the cardinality constraints for Makes? ⇒ 1:N

What should be the participation constraints for Makes? ⇒ Sailor - partial; Reserves - total

What should be the cardinality constraints for Has? ⇒ 1:N

What should be the participation constraints for Has? ⇒ Boat - partial; Reserves - total

EXAMPLE RELATIONAL SCHEMA AND DATABASE

Sailor(sailorId, sName, rating, age)

Boat(boatId, bName, color)

Reserves(*sailorId*, *boatId*, rDate)

Attribute names in italics are foreign key attributes.

Sailor

sailorId	sName	rating	age
22	Dustin	7	45
29	Brutus	1	33
31	Lubber	8	55
32	Andy	8	25
58	Rusty	10	35
64	Horatio	7	35
71	Zorba	10	16
74	Horatio	9	35
85	Art	3	25
95	Bob	3	63
99	Chris	10	30

11 tuples

Reserves

sailorId	*boatId*	rDate
22	101	10/10/17
22	102	10/10/17
22	103	08/10/17
22	104	07/10/17
31	102	10/11/17
31	103	06/11/17
31	104	12/11/17
64	101	05/09/17
64	102	08/09/17
74	103	08/09/17
99	104	08/08/17

11 tuples

Boat

boatId	bName	color
101	Interlake	blue
102	Interlake	red
103	Clipper	green
104	Marine	red
105	Serenity	Cyan

5 tuples

EXERCISE 1

Find the ids of sailors who have reserved boat 103.

Sailor

sailorId	sName	rating	age
22	Dustin	7	45
29	Brutus	1	33
31	Lubber	8	55
32	Andy	8	25
58	Rusty	10	35
64	Horatio	7	35
71	Zorba	10	16
74	Horatio	9	35
85	Art	3	25
95	Bob	3	63
99	Chris	10	30

11 tuples

Reserves

sailorId	*boatId*	rDate
22	101	10/10/17
22	102	10/10/17
22	103	08/10/17
22	104	07/10/17
31	102	10/11/17
31	103	06/11/17
31	104	12/11/17
64	101	05/09/17
64	102	08/09/17
74	103	08/09/17
99	104	08/08/17

11 tuples

Boat

boatId	bName	color
101	Interlake	blue
102	Interlake	red
103	Clipper	green
104	Marine	red
105	Serenity	Cyan

5 tuples

EXERCISE 1

Find the ids of sailors who have reserved boat 103.

☞ **22, 31, 74**

1. Is this a solution?

$\sigma_{boatId=103}(\pi_{sailorId}Reserves)$ **X**

$\pi_{sailorId}Reserves$

sailorId
22
31
64
74

$\sigma_{boatId=103}$?

2. Is this a solution?

$\pi_{sailorId}(\sigma_{boatId=103}Reserves)$ ✓

$\sigma_{boatId=103}Reserves$

sailorId	boatId	rDate
22	103	08/10/17
31	103	06/11/17
74	103	08/09/17

$\pi_{sailorId}$

sailorId
22
31
74

EXERCISE 2

Find the names of sailors who have reserved boat 103.

Sailor

sailorId	sName	rating	age
22	Dustin	7	45
29	Brutus	1	33
31	Lubber	8	55
32	Andy	8	25
58	Rusty	10	35
64	Horatio	7	35
71	Zorba	10	16
74	Horatio	9	35
85	Art	3	25
95	Bob	3	63
99	Chris	10	30

11 tuples

Reserves

sailorId	*boatId*	rDate
22	101	10/10/17
22	102	10/10/17
22	103	08/10/17
22	104	07/10/17
31	102	10/11/17
31	103	06/11/17
31	104	12/11/17
64	101	05/09/17
64	102	08/09/17
74	103	08/09/17
99	104	08/08/17

11 tuples

Boat

boatId	bName	color
101	Interlake	blue
102	Interlake	red
103	Clipper	green
104	Marine	red
105	Serenity	Cyan

5 tuples

EXERCISE 2

Find the names of sailors who have reserved boat 103.

☞ **Dustin, Lubber, Horatio**

1. Is this a solution?

$$\pi_{sName}(\sigma_{Reserves.sailorId=Sailor.sailorId \wedge boatId=103}(Reserves \times Sailor))$$

2. Is this a solution?

$$\pi_{sName}(\sigma_{Reserves.sailorId=Sailor.sailorId}((\sigma_{boatId=103}Reserves) \times Sailor))$$

Find the names of sailors who have reserved boat 103.

EXERCISE 2: SOLUTION 1

$$\pi_{sName}(\sigma_{Reserves.sailorId=Sailor.sailorId \wedge boatId=103}(Reserves \times Sailor))$$

☞ **Dustin, Lubber, Horatio**

Reserves

sailorId	boatId	rDate
22	101	10/10/17
22	102	10/10/17
22	103	08/10/17
22	104	07/10/17
31	102	10/11/17
31	103	06/11/17
31	104	12/11/17
64	101	05/09/17
64	102	08/09/17
74	103	08/09/17
99	104	08/08/17

11 tuples

X

Sailor

sailorId	sName	rating	age
22	Dustin	7	45
29	Brutus	1	33
31	Lubber	8	55
32	Andy	8	25
58	Rusty	10	35
64	Horatio	7	35
71	Zorba	10	16
74	Horatio	9	35
85	Art	3	25
95	Bob	3	63
99	Chris	10	30

11 tuples

=

How many tuples in the result? 11 x 11 = 121 tuples!

Find the names of sailors who have reserved boat 103.

EXERCISE 2: SOLUTION 1

$$\pi_{sName}(\sigma_{Reserves.sailorId=Sailor.sailorId \wedge boatId=103}(Reserves \times Sailor))$$

☞ **Dustin, Lubber, Horatio**

Reserves X Sailor

Reserves.sailorId	boatId	rDate	Sailor.sailorId	sName	rating	age
22	101	10/10/17	22	Dustin	7	45
22	101	10/10/17	29	Brutus	1	33
22	101	10/10/17	31	Lubber	8	55
22	101	10/10/17	32	Andy	8	25
22	101	10/10/17	58	Rusty	10	35
22	101	10/10/17	64	Horatio	7	35
22	101	10/10/17	71	Zorba	10	16
22	101	10/10/17	74	Horatio	9	35
22	101	10/10/17	85	Art	3	25
22	101	10/10/17	95	Bob	3	63
22	101	10/10/17	99	Chris	10	30
22	102	10/10/17	22	Dustin	7	45
22	102	10/10/17	29	Brutus	1	33
⋮	⋮	⋮	⋮	⋮	⋮	⋮

Find the names of sailors who have reserved boat 103.

EXERCISE 2: SOLUTION 1

$$\pi_{sName}(\sigma_{Reserves.sailorId=Sailor.sailorId \wedge boatId=103}(Reserves \times Sailor))$$

☞ **Dustin, Lubber, Horatio**

$\sigma_{Reserves.sailorId=Sailor.sailorId \wedge boatId=103}(Reserves \times Sailor)$

Reserves.sailorId	boatId	rDate	Sailor.sailorId	sName	rating	age
22	103	08/10/17	22	Dustin	7	45
31	103	06/11/17	31	Lubber	8	55
74	103	08/09/17	74	Horatio	9	35

Apply π_{sName} to above result:

sName
Dustin
Lubber
Horatio

✓

Find the names of sailors who have reserved boat 103.

EXERCISE 2: SOLUTION 2

$\pi_{sName}(\sigma_{Reserves.sailorId=Sailor.sailorId}((\sigma_{boatId=103}Reserves) \times Sailor))$

☞ **Dustin, Lubber, Horatio**

$\sigma_{boatId=103}$Reserves

sailorId	boatId	rDate
22	103	08/10/17
31	103	06/11/17
74	103	08/09/17

X

Sailor

sailorId	name	rating	age
22	Dustin	7	45
29	Brutus	1	33
31	Lubber	8	55
32	Andy	8	25
58	Rusty	10	35
64	Horatio	7	35
71	Zorba	10	16
74	Horatio	9	35
85	Art	3	25
95	Bob	3	63
99	Chris	10	30

11 tuples

=

How many tuples in the result? 3 x 11 = 33 tuples!

Find the names of sailors who have reserved boat 103.

EXERCISE 2: SOLUTION 2

$$\pi_{sName}(\sigma_{Reserves.sailorId=Sailor.sailorId}((\sigma_{boatId=103}Reserves) \times Sailor))$$

☞ **Dustin, Lubber, Horatio**

$(\sigma_{boatId=103}Reserves) \times Sailor$

Reserves.sailorId	boatId	rDate	Sailor.sailorId	sName	rating	age
22	103	08/10/17	22	Dustin	7	45
22	103	08/10/17	29	Brutus	1	33
22	103	08/10/17	31	Lubber	8	55
22	103	08/10/17	32	Andy	8	25
22	103	08/10/17	58	Rusty	10	35
22	103	08/10/17	64	Horatio	7	35
22	103	08/10/17	71	Zorba	10	16
22	103	08/10/17	74	Horatio	9	35
22	103	08/10/17	85	Art	3	25
22	103	08/10/17	95	Bob	3	63
22	103	08/10/17	99	Chris	10	30
31	103	06/11/17	22	Dustin	7	45
31	103	06/11/17	29	Brutus	1	33
31	103	06/11/17	31	Lubber	8	55
⋮	⋮	⋮	⋮	⋮	⋮	⋮

Find the names of sailors who have reserved boat 103.

EXERCISE 2: SOLUTION 2

$$\pi_{sName}(\sigma_{Reserves.sailorId=Sailor.sailorId}((\sigma_{boatId=103}Reserves)\ X\ Sailor))$$

☞ **Dustin, Lubber, Horatio**

$\sigma_{Reserves.sailorId=Sailor.sailorId}((\sigma_{boatId=103}Reserves)\ X\ Sailor)$

Reserves.sailorId	boatId	rDate	Sailor.sailorId	sName	rating	age
22	103	08/10/17	22	Dustin	7	45
31	103	06/11/17	31	Lubber	8	55
74	103	08/09/17	74	Horatio	9	35

Apply π_{sName} to above result:

sName
Dustin
Lubber
Horatio

✓

EXERCISE 2

Find the names of sailors who have reserved boat 103.

☞ **Dustin, Lubber, Horatio**

1. Is this a solution? ✓

$$\pi_{sName}(\sigma_{Reserves.sailorId=Sailor.sailorId \wedge boatId=103}(Reserves \times Sailor))$$

Initial result: 121 tuples

2. Is this a solution? ✓

$$\pi_{sName}(\sigma_{Reserves.sailorId=Sailor.sailorId}((\sigma_{boatId=103}Reserves) \times Sailor))$$

Initial result: 33 tuples

To be continued …

Find the names of sailors who have reserved boat 103.

EXERCISE 2: SOLUTION 3

$\pi_{sName}((\sigma_{boatId=103}Reserves)$ JOIN Sailor)

☞ **Dustin, Lubber, Horatio**

$\sigma_{boatId=103}$Reserves

sailorId	boatId	rDate
22	103	08/10/17
31	103	06/11/17
74	103	08/09/17

JOIN

Sailor

sailorId	sName	rating	age
22	Dustin	7	45
29	Brutus	1	33
31	Lubber	8	55
32	Andy	8	25
58	Rusty	10	35
64	Horatio	7	35
71	Zorba	10	16
74	Horatio	9	35
85	Art	3	25
95	Bob	3	63
99	Chris	10	30

11 tuples

=

How many tuples in the result? 3 tuples!

Find the names of sailors who have reserved boat 103.

EXERCISE 2: SOLUTION 3

$\pi_{sName}((\sigma_{boatId=103}Reserves)\ JOIN\ Sailor)$

☞ **Dustin, Lubber, Horatio**

$(\sigma_{boatId=103}Reserves)\ JOIN\ Sailor$

Reserves.sailorId	boatId	rDate	Sailor.sailorId	sName	rating	age
22	103	08/10/17	22	Dustin	7	45
31	103	06/11/17	31	Lubber	8	55
74	103	08/09/17	74	Horatio	9	35

Apply π_{sName} to above result:

sName
Dustin
Lubber
Horatio

✓

EXERCISE 2: SUMMARY

Find the names of sailors who have reserved boat 103.

☞ **All three queries get the correct answer, BUT ...**

1. Is this a solution? ✓

$\pi_{sName}(\sigma_{Reserves.sailorId=Sailor.sailorId \wedge boatId=103}(Reserves \times Sailor))$

Initial result: 121 tuples

2. Is this a solution? ✓

$\pi_{sName}(\sigma_{Reserves.sailorId=Sailor.sailorId}((\sigma_{boatId=103}Reserves) \times Sailor))$

Initial result: 33 tuples

3. Is this a solution? ✓

$\pi_{sName}((\sigma_{boatId=103}Reserves) \text{ JOIN } Sailor)$

Initial result: 3 tuples

Query Optimization

Relational DBMSs do such optimizations based on relational algebra.

EXERCISE 3

Find the names of sailors who have reserved a red boat.

Sailor

sailorId	sName	rating	age
22	Dustin	7	45
29	Brutus	1	33
31	Lubber	8	55
32	Andy	8	25
58	Rusty	10	35
64	Horatio	7	35
71	Zorba	10	16
74	Horatio	9	35
85	Art	3	25
95	Bob	3	63
99	Chris	10	30

11 tuples

Reserves

sailorId	*boatId*	rDate
22	101	10/10/17
22	102	10/10/17
22	103	08/10/17
22	104	07/10/17
31	102	10/11/17
31	103	06/11/17
31	104	12/11/17
64	101	05/09/17
64	102	08/09/17
74	103	08/09/17
99	104	08/08/17

11 tuples

Boat

boatId	bName	color
101	Interlake	blue
102	Interlake	red
103	Clipper	green
104	Marine	red
105	Serenity	Cyan

5 tuples

EXERCISE 3: SOLUTION 1

Find the names of sailors who have reserved a red boat.

☞ **Dustin, Lubber, Horatio, Chris**

Is this a solution?

$\pi_{sName}((\sigma_{color='red'}Boat)$ JOIN Reserves JOIN Sailor)

$\sigma_{color='red'}$Boat

boatId	bName	color
102	Interlake	red
104	Marine	red

JOIN

Reserves

sailorId	boatId	rDate
22	101	10/10/17
22	102	10/10/17
22	103	08/10/17
22	104	07/10/17
31	102	10/11/17
31	103	06/11/17
31	104	12/11/17
64	101	05/09/17
64	102	08/09/17
74	103	08/09/17
99	104	08/08/17

=

How many tuples in the result?
6 tuples!

How many columns in the result?
5 columns!

Find the names of sailors who have reserved a red boat.

EXERCISE 3: SOLUTION 1

$\pi_{sName}((\sigma_{color='red'}Boat)$ JOIN Reserves JOIN Sailor)

☞ **Dustin, Lubber, Horatio, Chris**

$(\sigma_{color='red'}Boat)$ JOIN Reserves

bName	color	sailorId	boatId	rDate
Interlake	red	22	102	10/10/17
Marine	red	22	104	07/10/17
Interlake	red	31	102	10/11/17
Marine	red	31	104	12/11/17
Interlake	red	64	102	08/09/17
Marine	red	99	104	08/08/17

JOIN

Sailor

sailorId	sName	rating	age
22	Dustin	7	45
29	Brutus	1	33
31	Lubber	8	55
32	Andy	8	25
58	Rusty	10	35
64	Horatio	7	35
71	Zorba	10	16
74	Horatio	9	35
85	Art	3	25
95	Bob	3	63
99	Chris	10	30

=

How many tuples in the result?
6 tuples!

How many columns in the result?
8 columns!

Find the names of sailors who have reserved a red boat.

EXERCISE 3: SOLUTION 1

$\pi_{sName}((\sigma_{color='red'}Boat)\ JOIN\ Reserves\ JOIN\ Sailor)$

☞ **Dustin, Lubber, Horatio, Chris**

$(\sigma_{color='red'}Boat)\ JOIN\ Reserves\ JOIN\ Sailor$

bName	color	sailorId	boatId	rDate	sName	rating	age
Interlake	red	22	102	10/10/17	Dustin	7	45
Marine	red	22	104	07/10/17	Dustin	7	45
Interlake	red	31	102	10/11/17	Lubber	8	55
Marine	red	31	104	12/11/17	Lubber	8	55
Interlake	red	64	102	08/09/17	Horatio	7	35
Marine	red	99	104	08/08/17	Chris	10	30

Apply π_{sName} to above result:

sName
Dustin
Lubber
Horatio
Chris

✓

EXERCISE 3: SOLUTION 2

Find the names of sailors who have reserved a red boat.

☞ **Dustin, Lubber, Horatio, Chris**

$$\pi_{sName}((\sigma_{color='red'}Boat) \text{ JOIN } Reserves \text{ JOIN } Sailor)$$

Can you give a more efficient solution in terms of result size?

$$\pi_{sName}((\pi_{boatId}(\sigma_{color='red'}Boat)) \text{ JOIN } Reserves \text{ JOIN } Sailor)$$

$\sigma_{color='red'}Boat$

boatId	bName	color
102	Interlake	red
104	Marine	red

After selecting red boats, first project onto boatId before doing the join since the name and color of the boat is not needed for the query. Thus, only the boatId is "carried" when evaluating the rest of the query.

Find the names of sailors who have reserved a red boat.

EXERCISE 3: SOLUTION 2

$\pi_{sName}((\pi_{boatId}(\sigma_{color='red'}Boat))$ JOIN Reserves JOIN Sailor)

☞ **Dustin, Lubber, Horatio, Chris**

$\pi_{boatId}(\sigma_{color='red'}Boat)$

boatId
102
104

JOIN

Reserves

sailorId	boatId	rDate
22	101	10/10/17
22	102	10/10/17
22	103	08/10/17
22	104	07/10/17
31	102	10/11/17
31	103	06/11/17
31	104	12/11/17
64	101	05/09/17
64	102	08/09/17
74	103	08/09/17
99	104	08/08/17

=

How many tuples in the result? 6 tuples!

How many columns in the result? 3 columns!

Find the names of sailors who have reserved a red boat.

EXERCISE 3: SOLUTION 2

$\pi_{sName}((\pi_{boatId}(\sigma_{color='red'}Boat))$ JOIN Reserves JOIN Sailor)

☞ **Dustin, Lubber, Horatio, Chris**

$(\pi_{boatId}(\sigma_{color='red'}Boat))$ JOIN Reserves

sailorId	boatId	rDate
22	102	10/10/17
22	104	07/10/17
31	102	10/11/17
31	104	12/11/17
64	102	08/09/17
99	104	08/08/17

JOIN

Sailor

sailorId	sName	rating	age
22	Dustin	7	45
29	Brutus	1	33
31	Lubber	8	55
32	Andy	8	25
58	Rusty	10	35
64	Horatio	7	35
71	Zorba	10	16
74	Horatio	9	35
85	Art	3	25
95	Bob	3	63
99	Chris	10	30

=

How many tuples in the result? 6 tuples!

How many columns in the result? 6 columns!

Find the names of sailors who have reserved a red boat.

EXERCISE 3: SOLUTION 2

$\pi_{sName}((\pi_{boatId}(\sigma_{color='red'}Boat))$ JOIN Reserves JOIN Sailor)

☞ **Dustin, Lubber, Horatio, Chris**

$(\sigma_{color='red'}Boat)$ JOIN Reserves JOIN Sailor

sailorId	boatId	rDate	sName	rating	age
22	102	10/10/17	Dustin	7	45
22	104	07/10/17	Dustin	7	45
31	102	10/11/17	Lubber	8	55
31	104	12/11/17	Lubber	8	55
64	102	08/09/17	Horatio	7	35
99	104	08/08/17	Chris	10	30

Apply π_{sName} to above result:

sName
Dustin
Lubber
Horatio
Chris

✓

Find the names of sailors who have reserved a red boat.

EXERCISE 3: SUMMARY

Solution 1

$\pi_{sName}((\sigma_{color='red'}Boat)$ JOIN Reserves JOIN Sailor)

(6 tuples, 5 columns) + (6 tuples, 8 columns)

Solution 2

$\pi_{sName}((\pi_{boatId}(\sigma_{color='red'}Boat))$ JOIN Reserves JOIN Sailor)

(6 tuples, 3 columns) + (6 tuples, 6 columns)

☞ **Solution 2 is more efficient in terms of tuple size.**

Query Optimization

Relational DBMSs do such optimizations based on relational algebra.

Upload your completed exercise worksheet to Canvas by **11 p.m.**

EXERCISES 4, 5, 6

Sailor(sailorId, sName, rating, age)

Boat(boatId, bName, color)

Reserves(*sailorId*, *boatId*, rDate)

Exercise 4: Find the names of sailors who have reserved either a red or a green boat.

Exercise 5: Find the names of sailors who have reserved both a red and a green boat.

Exercise 6: Find the ids of sailors who have made at least two reservations on the same date. [Hint: You need to join Reserves with itself.]

DO NOT try to optimize the queries.

Just try to get a solution.

EXERCISE 4

Find the names of sailors who have reserved either a red or a green boat.

☞ **Dustin (22), Lubber (31), Horatio (64), Horatio (74), Chris (99)**

$\pi_{sName}(\pi_{boatId}(\sigma_{color='red' \vee color='green'} Boat))$ JOIN Reserves JOIN Sailor)

Identify all red or green boats ($\sigma_{color='red' \vee color='green'} Boat$), then find sailors who have reserved one of these boats (… JOIN Reserves JOIN Sailor).

π_{boatId} is a nice optimization but is not strictly needed to answer the query.

EXERCISE 5: SOLUTION 1

Is this a solution?

Find the names of sailors who have reserved both a red and a green boat.

☞ **Dustin (22), Lubber (31)**

$\pi_{sName}((\sigma_{color='red' \wedge color='green'}Boat)$ JOIN Reserves JOIN Sailor)

Boat

boatId	bName	color
101	Interlake	blue
102	Interlake	red
103	Clipper	green
104	Marine	red
105	Serenity	cyan

No! Why?

☞ **Nothing is selected! Why?**

The condition color='red' ∧ color='green' can *never* be satisfied!

Find the names of sailors who have reserved both a red and a green boat.

EXERCISE 5: SOLUTION 2

Is this a solution?

$\pi_{sName}((\sigma_{color='red' \vee color='green'}Boat)$ JOIN Reserves JOIN Sailor)

☞ **Dustin (22), Lubber (31)**

$(\sigma_{color='red' \vee color='green'}Boat)$ JOIN Reserves JOIN Sailor

bName	color	sailorId	boatId	rDate	sName	rating	age
Interlake	red	22	102	10/10/17	Dustin	7	45
Marine	red	22	104	07/10/17	Dustin	7	45
Interlake	red	31	102	10/11/17	Lubber	8	55
Marine	red	31	104	12/11/17	Lubber	8	55
Interlake	red	64	102	08/09/17	Horatio	7	35
Marine	red	99	104	08/08/17	Chris	10	30
Clipper	green	22	103	08/10/17	Dustin	7	45
Clipper	green	31	103	06/11/17	Lubber	8	55
Clipper	green	74	103	08/09/17	Horatio	7	35

What's the problem?

π_{sName} ⇒

sName
Dustin
Lubber
Horatio
Chris

X

The condition color='red' $\vee$ color='green' includes sailors who have reserved only a red or only a green boat, as well as both a red and a green boat!

Must identify sailors who have reserved red boats, sailors who have reserved green boats, then find the intersection.

Find the names of sailors who have reserved both a red and a green boat.

EXERCISE 5: SOLUTION 3

Is this a solution? (intersect join result)

$$\pi_{sName}((\sigma_{color='red'}Boat)\ JOIN\ Reserves\ JOIN\ Sailor) \cap (\sigma_{color='green'}Boat)\ JOIN\ Reserves\ JOIN\ Sailor)$$

☞ **Dustin (22), Lubber (31)**

$(\sigma_{color='red'}Boat)$ JOIN Reserves JOIN Sailor

bName	color	sailorId	boatId	rDate	sName	rating	age
Interlake	red	22	102	10/10/17	Dustin	7	45
Marine	red	22	104	07/10/17	Dustin	7	45
Interlake	red	31	102	10/11/17	Lubber	8	55
Marine	red	31	104	12/11/17	Lubber	8	55
Interlake	red	64	102	08/09/17	Horatio	7	35
Marine	red	99	104	08/08/17	Chris	10	30

∩

$(\sigma_{color='green'}Boat)$ JOIN Reserves JOIN Sailor

bName	color	sailorId	boatId	rDate	sName	rating	age
Clipper	green	22	103	08/10/17	Dustin	7	45
Clipper	green	31	103	06/11/17	Lubber	8	55
Clipper	green	74	103	08/09/17	Horatio	7	35

→ The result is empty!

Find the names of sailors who have reserved both a red and a green boat.

EXERCISE 5: SOLUTION 4

Is this a solution? (intersect on sName)

$$\pi_{sName}((\sigma_{color='red'}Boat)\ JOIN\ Reserves\ JOIN\ Sailor) \cap \pi_{sName}((\sigma_{color='green'}Boat)\ JOIN\ Reserves\ JOIN\ Sailor)$$

☞ **Dustin (22), Lubber (31)**

Since sName is not unique, there may be incorrect tuples in the intersection (i.e., Horatio is not unique).

$(\sigma_{color='red'}Boat)$ JOIN Reserves JOIN Sailor

bName	color	sailorId	boatId	rDate	sName	rating	age
Interlake	red	22	102	10/10/17	Dustin	7	45
Marine	red	22	104	07/10/17	Dustin	7	45
Interlake	red	31	102	10/11/17	Lubber	8	55
Marine	red	31	104	12/11/17	Lubber	8	55
Interlake	red	64	102	08/09/17	Horatio	7	35
Marine	red	99	104	08/08/17	Chris	10	30

π_{sName} →

sName
Dustin
Lubber
Horatio
~~Chris~~

$(\sigma_{color='green'}Boat)$ JOIN Reserves JOIN Sailor

bName	color	sailorId	boatId	rDate	sName	rating	age
Clipper	green	22	103	08/10/17	Dustin	7	45
Clipper	green	31	103	06/11/17	Lubber	8	55
Clipper	green	74	103	08/09/17	Horatio	7	35

π_{sName} →

sName
Dustin
Lubber
Horatio

$\cap$ π_{sName} →

sName
Dustin
Lubber
Horatio

X

Find the names of sailors who have reserved both a red and a green boat.

EXERCISE 5: SOLUTION 5

Is this a solution?
(intersect on sailorId, sName)

$$\pi_{sName}(\pi_{sailorId,\ sName}((\sigma_{color='red'}Boat)\ JOIN\ Reserves\ JOIN\ Sailor) \cap \pi_{sailorId,\ sName}((\sigma_{color='green'}Boat)\ JOIN\ Reserves\ JOIN\ Sailor))$$

☞ **Dustin (22), Lubber (31)**

$(\sigma_{color='red'}Boat)$ JOIN Reserves JOIN Sailor

bName	color	sailorId	boatId	rDate	sName	rating	age
Interlake	red	22	102	10/10/17	Dustin	7	45
Marine	red	22	104	07/10/17	Dustin	7	45
Interlake	red	31	102	10/11/17	Lubber	8	55
Marine	red	31	104	12/11/17	Lubber	8	55
Interlake	red	64	102	08/09/17	Horatio	7	35
Marine	red	99	104	08/08/17	Chris	10	30

$\pi_{sailorId,\ sName}$ →

sailorId	sName
22	Dustin
31	Lubber
~~64~~	~~Horatio~~
~~99~~	~~Chris~~

$(\sigma_{color='green'}Boat)$ JOIN Reserves JOIN Sailor

bName	color	sailorId	boatId	rDate	sName	rating	age
Clipper	green	22	103	08/10/17	Dustin	7	45
Clipper	green	31	103	06/11/17	Lubber	8	55
Clipper	green	74	103	08/09/17	Horatio	7	35

$\pi_{sailorId,\ sName}$ →

sailorId	sName
22	Dustin
31	Lubber
~~74~~	~~Horatio~~

$\cap$ π_{sName} →

sName
Dustin
Lubber

✓

Find the names of sailors who have reserved both a red and a green boat.

EXERCISE 5: SOLUTION 6

Is this a solution?
(join on sName)

$$\pi_{sName}(\pi_{sailorId,\ sName}((\sigma_{color='red'}Boat)\ JOIN\ Reserves\ JOIN\ Sailor)\ JOIN_{sName}\ \pi_{sailorId,\ sName}((\sigma_{color='green'}Boat)\ JOIN\ Reserves\ JOIN\ Sailor))$$

☞ **Dustin (22), Lubber (31)**

$\pi_{sailorId,\ sName}((\sigma_{color='red'}Boat)\ JOIN\ Reserves\ JOIN\ Sailor)$

sailorId	sName
22	Dustin
31	Lubber
64	Horatio
99	Chris

$JOIN_{sName}$

$\pi_{sailorId,\ sName}((\sigma_{color='green'}Boat)\ JOIN\ Reserves\ JOIN\ Sailor)$

sailorId	sName
22	Dustin
31	Lubber
74	Horatio

Since sName is not unique, there may be incorrect tuples in the join (i.e., there are two *different* sailors with the same name, Horatio).

=

R1.sailorId	sName	R2.sailorId
22	Dustin	22
31	Lubber	31
64	Horatio	74

π_{sName}

sName
Dustin
Lubber
Horatio

X

Find the names of sailors who have reserved both a red and a green boat.

EXERCISE 5: SOLUTION 7

Is this a solution?
(join on sailorId)

$$\pi_{sName}(\pi_{sailorId,\ sName}((\sigma_{color='red'}Boat)\ JOIN\ Reserves\ JOIN\ Sailor)\ JOIN_{sailorId}\ \pi_{sailorId,\ sName}((\sigma_{color='green'}Boat)\ JOIN\ Reserves\ JOIN\ Sailor))$$

☞ **Dustin (22), Lubber (31)**

$\pi_{sailorId,\ sName}((\sigma_{color='red'}Boat)\ JOIN\ Reserves\ JOIN\ Sailor)$

sailorId	sName
22	Dustin
31	Lubber
64	Horatio
99	Chris

$JOIN_{sailorId}$

$\pi_{sailorId,\ sName}((\sigma_{color='green'}Boat)\ JOIN\ Reserves\ JOIN\ Sailor)$

sailorId	sName
22	Dustin
31	Lubber
74	Horatio

=

sailorId	R1.sName	R2.sName
22	Dustin	Dustin
31	Lubber	Lubber

π_{sName}

sName
Dustin
Lubber

✓

EXERCISE 6: SOLUTION 1

Find the ids of sailors who have made at least two reservations on the same date.

☞ 22

We need to use rename: ρ_{R1}(Reserves), ρ_{R2}(Reserves)

$$\pi_{R1.sailorId}(\sigma_{R1.sailorId=R2.sailorId \wedge R1.rDate=R2.rDate \wedge R1.boatId<>R2.boatId}(\rho_{R1}(Reserves) \times \rho_{R2}(Reserves)))$$

Or equivalently:

$$\pi_{R1.sailorId}(\rho_{R1}(Reserves) \text{ JOIN}_{R1.sailorId=R2.sailorId \wedge R1.rDate=R2.rDate \wedge R1.boatId<>R2.boatId} \rho_{R2}(Reserves))$$

Find the ids of sailors who have made at least two reservations on the same date.

EXERCISE 6: SOLUTION 1 (cont'd)

$$\pi_{R1.sailorId}(\sigma_{R1.sailorId=R2.sailorId \wedge R1.rDate=R2.rDate \wedge R1.boatId<>R2.boatId}(R1 \times R2))$$

R1

sailorId	boatId	rDate
22	101	10/10/17
22	102	10/10/17
22	103	08/10/17
22	104	07/10/17
31	102	10/11/17
31	103	06/11/17
31	104	12/11/17
64	101	05/09/17
64	102	08/09/17
74	103	08/09/17
99	104	08/08/17

X

R2

sailorId	boatId	rDate
22	101	10/10/17
22	102	10/10/17
22	103	08/10/17
22	104	07/10/17
31	102	10/11/17
31	103	06/11/17
31	104	12/11/17
64	101	05/09/17
64	102	08/09/17
74	103	08/09/17
99	104	08/08/17

=

Find the ids of sailors who have made at least two reservations on the same date.

EXERCISE 6: SOLUTION 1 (cont'd)

$$\pi_{R1.sailorId}(\sigma_{R1.sailorId=R2.sailorId \wedge R1.rDate=R2.rDate \wedge R1.boatId<>R2.boatId}(R1 \times R2))$$

$\sigma_{R1.sailorId=R2.sailorId \wedge R1.rDate=R2.rDate \wedge R1.boatId<>R2.boatId}$

R1.sailorId	R1.boatId	R1.rDate	R2.sailorId	R2.boatId	R2.rDate
22	101	10/10/17	22	101	10/10/17
22	101	10/10/17	22	102	10/10/17
22	101	10/10/17	22	103	08/10/17
22	101	10/10/17	22	104	07/10/17
22	101	10/10/17	31	102	10/11/17
22	101	10/10/17	31	103	06/11/17
22	101	10/10/17	31	104	12/11/17
22	101	10/10/17	64	101	05/09/17
22	101	10/10/17	64	102	08/09/17
22	101	10/10/17	74	103	08/09/17
22	101	10/10/17	99	104	08/08/17
22	102	10/10/17	22	101	10/10/17
22	102	10/10/17	22	102	10/10/17
22	102	10/10/17	22	103	08/10/17
22	102	10/10/17	22	104	07/10/17
⋮	⋮	⋮	⋮	⋮	⋮

$\pi_{R1.sailorId}$ =

sailorId
22

EXERCISE 6: SOLUTION 1 (cont'd)

What do we get if we omit R1.rDate=R2.rDate?

$\sigma_{R1.sailorId=R2.sailorId \wedge R1.boatId<>R2.boatId}$

R1.sailorId	R1.boatId	R1.rDate	R2.sailorId	R2.boatId	R2.rDate
22	101	10/10/17	22	102	10/10/17
22	101	10/10/17	22	103	08/10/17
22	101	10/10/17	22	104	07/10/17
22	102	10/10/17	22	101	10/10/17
22	102	10/10/17	22	103	08/10/17
22	102	10/10/17	22	104	07/10/17
22	103	08/10/17	22	101	10/10/17
22	103	08/10/17	22	102	10/10/17
22	103	08/10/17	22	104	07/10/17
22	104	07/10/17	22	101	10/10/17
22	104	07/10/17	22	102	10/10/17
22	104	07/10/17	22	103	08/10/17
31	102	10/11/17	31	103	06/11/17
31	102	10/11/17	31	104	12/11/17
31	103	06/11/17	31	102	10/11/17
31	103	06/11/17	31	104	12/11/17
⋮	⋮	⋮	⋮	⋮	⋮

Sailors who have made more than one reservation.

$\pi_{R1.sailorId}$ =

sailorId
22
31
64

EXERCISE 6: SOLUTION 1 (cont'd)

What do we get if we omit R1.boatId<>R2.boatId?

$\sigma_{R1.sailorId=R2.sailorId \wedge R1.rDate=R2.rDate}$

R1.sailorId	R1.boatId	R1.rDate	R2.sailorId	R2.boatId	R2.rDate
22	101	10/10/17	22	101	10/10/17
22	101	10/10/17	22	102	10/10/17
22	102	10/10/17	22	101	10/10/17
22	102	10/10/17	22	102	10/10/17
22	103	08/10/17	22	103	08/10/17
22	104	07/10/17	22	104	07/10/17
31	102	10/11/17	31	102	10/11/17
31	103	06/11/17	31	103	06/11/17
31	104	12/11/17	31	104	12/11/17
64	101	05/09/17	64	101	05/09/17
64	102	08/09/17	64	102	08/09/17
74	103	08/09/17	74	103	08/09/17
99	104	08/08/17	99	104	08/08/17

Sailors who have made at least one reservation.

$\pi_{R1.sailorId}$ =

sailorId
22
31
64
74
99