## SQL, the underestimated "Big Data" technology

Copyright (c) 2009-2014 by Data Geekery GmbH. Slides licensed under CC BY SA 3.0

#### No – tation

A history of databases in No-tation 1970: NoSQL = We have no SQL 1980: NoSQL = Know SQL 2000: NoSQL = No SQL! 2005: NoSQL = Not only SQL 2013: NoSQL = No, SQL! (R)DB(MS)

Seen at the 2013 O'Reilly Strata Conf: History of NoSQL by <u>Mark Madsen</u>. Picture published by <u>Edd Dumbill</u>



#### NoSQL?

# NoSQL? No, SQL!

Copyright (c) 2009-2014 by Data Geekery GmbH. Slides licensed under CC BY SA 3.0



#### Our vision at Data Geekery

- SQL dominates database systems
- SQL is very expressive
- SQL is very type safe

SQL is a device whose mystery is only exceeded by its power!



#### Me – @lukaseder



## Head of R&D at Data Geekery GmbH SQL Aficionado

- Java Aficionado

# Java developers can get back in control of SQL with jOOQ



## Big Data? NoSQL?

- You're giving up on **ACID**
- You're giving up on type safety
- You're giving up on standards
- You're giving up on **tooling**
- You're giving up on relational algebra
- You haven't asked operations
- You don't actually have «Big Data»



## Big Data? NoSQL?

- You're giving up on **ACID**
- You're giving up on type safety
- You're giving up on standards
- You're giving up on **tooling**
- You're giving up on relational algebra
- You haven't asked operations
- You don't actually have «Big Data»



#### Also Not SQL

```
@Entity @Table(name = "EVENTS")
public class Event {
    private Long id;
    private String title;
    private Date date;
    @Id @GeneratedValue(generator = "increment")
    @GenericGenerator(name = "increment", strategy = "increment")
    public Long getId() { /* ... */ }
    @Temporal(TemporalType.TIMESTAMP)
    @Column(name = "EVENT_DATE")
    public Date getDate() { /* ... */ }
```



### Also Not SQL – Annotatiomania™



Found at <a href="http://stackoverflow.com/q/17491912/521799">http://stackoverflow.com/q/17491912/521799</a>



#### Also Not SQL – JPA 3.0 Preview



Might not be true



#### Shocker! You can now write SQL in Java.





### SQL in Java 7 – JDBC





#### SQL in Java 8 – jOOQ





## Typesafe SQL in Java – jOOQ

```
DSL.using(c)
.select(s.SCHEMA_NAME, s.IS_DEFAULT)
.from(INFORMATION_SCHEMA.SCHEMATA.as("s"))
.orderBy(s.SCHEMA_NAME)
.map(rs -> new Schema(
    rs.getValue(s.SCHEMA_NAME),
    rs.getValue(s.IS_DEFAULT)
))
.forEach(System.out::println);
```



### SQL in Java 8 – Spring JDBC

```
new JdbcTemplate(
    new SingleConnectionDataSource(c, true))
.query(sql, (rs, rowNum) ->
    new Schema(
        rs.getString("SCHEMA_NAME"),
        rs.getBoolean("IS_DEFAULT")
    ))
.forEach(System.out::println);
```

dbutils...

http://commons.apache.org/

### SQL in Java 8 – Apache DbUtils

```
new QueryRunner()
    .query(c, sql, new ArrayListHandler())
    .stream()
    .map(array -> new Schema()
        (String) array[0],
        (Boolean) array[1]
    ))
    .forEach(System.out::println);
                                                commons
                                      Apache Commons
```

•

### SQL in Groovy





#### When you should use SQL – indicators

- You need JOINs, UNIONs
- You need functions, aggregations
- You need bulk reads / writes

## Calculations should be done close to the data

#### Please, run that calculation in your DB





#### SQL Trivia – NULL

-- What does this query return?
SELECT 1 AS a FROM dual
WHERE 1 IN (NULL)
UNION ALL
SELECT 2 AS a FROM dual
WHERE NOT(1 IN (NULL))



#### SQL Trivia – NULL





#### SQL Trivia – NULL

```
-- Nothing! It's the same as this
SELECT 1 AS a FROM dual
WHERE 1 = NULL
UNION ALL
SELECT 2 AS a FROM dual
WHERE 1 != NULL
```



YOU JUST DON'T

KNOWP

#### SQL Trivia – NULL

-- Nothing! It's the same as this SELECT 1 AS a FROM dual **SO YOU'RE TELLING** WHERE "UNKNOWN" UNION ALL SELECT 2 AS a FROM dual WHERE "UNKNOWN"

Copyright (c) 2009-2014 by Data Geekery GmbH. Slides licensed under CC BY SA 3.0



### SQL Trivia – Oracle VARCHAR2

```
-- What does this query return?
SELECT 1 AS a FROM dual
WHERE '' = ''
UNION ALL
SELECT 2 AS a FROM dual
WHERE 'a' != ''
```

Copyright (c) 2009-2014 by Data Geekery GmbH. Slides licensed under CC BY SA 3.0



### SQL Trivia – Oracle VARCHAR2

```
-- Nope! Nothing again (only in Oracle).
SELECT 1 AS a FROM dual
WHERE NULL = NULL
UNION ALL
SELECT 2 AS a FROM dual
WHERE 'a' != NULL
```



\*THAT\* EXPLAINS 1-2 THINGS

### SQL Trivia – Oracle VARCHAR2

```
-- Nope! Nothing again (only in Oracle).
SELECT 1 AS a FROM dual
WHERE NULL = NULL
UNION ALL
SELECT 2 AS a FROM dual
WHERE 'a' != NULL
```

Copyright (c) 2009-2014 by Data Geekery GmbH. Slides licensed under CC BY SA 3.0

5

## Stockholm Syndrome:

## We love JavaScript SQL

Copyright (c) 2009-2014 by Data Geekery GmbH. Slides licensed under CC BY SA 3.0

## Winston Churchill:

# SQL is the worst form of database querying, except for all the other forms.



### Let's calculate a running total

Copyright (c) 2009-2014 by Data Geekery GmbH. Slides licensed under CC BY SA 3.0



### Let's calculate a running total

ID	VALUE_DATE	AMOUNT
9997	2014-03-18	99.17
9981	2014-03-16	71.44
9979	2014-03-16	-94.60
9977	2014-03-16	-6.96
9971	2014-03-15	-65.95



### Let's calculate a running total

ID	VALUE DATE	AMOUNT	BALANCE
9997	2014-03-18	99.17	19985.81
9981	2014-03-16	71.44	19886.64
9979	2014-03-16	-94.60	19815.20
9977	2014-03-16	-6.96	19909.80
9971	2014-03-15	-65.95	19916.76



## Let's calculate a running total

ID 	VALUE_DATE	AMOUNT	BALANCE
9997   9981	2014-03-18 2014-03-16	+ <b>99.17</b>	=19985.81 +19886.64
9979 9977	2014-03-16 2014-03-16	-94.60	19815.20
9971	2014-03-15	-65.95	19916.76

Copyright (c) 2009-2014 by Data Geekery GmbH. Slides licensed under CC BY SA 3.0



## Let's calculate a running total

ID	VALUE_DATE	AMOUNT	BALANCE	
Í				
9997	2014-03-18	99.17	19985.81	
9981	2014-03-16	+71.44	=19886.64	
9979	2014-03-16	-94.60	+19815.20	
9977	2014-03-16	-6.96	19909.80	
9971	2014-03-15	-65.95	19916.76	



## Let's calculate a running total

ID	VALUE_DATE	AMOUNT	BALANCE	
9997	2014-03-18	99.17	19985.81	
9981	2014-03-16	+71.44	=19886.64	n
9979	2014-03-16	-94.60	+19815.20	n+1
9977	2014-03-16	-6.96	19909.80	
BALANCI	E(ROWn) = BALA	NCE(ROWn+1)	+ AMOUNT(RC	)Wn)
BALANCI	E(ROWn+1) = BAI	_ANCE(ROWn)	– AMOUNT(RO	)Wn)



## **L**How can we do it?

Copyright (c) 2009-2014 by Data Geekery GmbH. Slides licensed under CC BY SA 3.0



"

#### How can we do it?

- In Java
- Calculate on UPDATE
- Nested SELECT
- Recursive SQL
- Window functions
- MODEL clause (Oracle)
- Stored procedures



### How can we do it? – With SQL!



- Window functions
- MODEL clause (Oracle)

- Stored procedures



# **G**Using nested SELECTs

Copyright (c) 2009-2014 by Data Geekery GmbH. Slides licensed under CC BY SA 3.0



#### SELECT t1.\*, t1.current balance - ( SELECT NVL(SUM(amount), 0) FROM v transactions t2 WHERE t2.account\_id = t1.account\_id (t2.value date, t2.id) > AND (t1.value\_date, t1.id) ) AS balance FROM v transactions t1 WHERE t1.account id = 1 ORDER BY t1.value date DESC, t1.id DESC











### Using nested SELECTs

ID 	VALUE_DATE	AMOUNT	BALANCE
9997	2014-03-18	-(99.17)	+19985.81
9981	2014-03-16	-(71.44)	19886.64
9979	2014-03-16	-(-94.60)	19815.20
9977	2014-03-16	-6.96	=19909.80
9971	2014-03-15	-65.95	19916.76



1	d	Operation	Name	A-Rows	A-Time
	0	SELECT STATEMENT		50	00:00:00.77
	1	SORT AGGREGATE		1101	00:00:00.76
*	2	TABLE ACCESS BY INDEX ROWID	T_TRANSACTIONS	605K	00:00:00.69
*	3	INDEX RANGE SCAN	I_TRX_ACCO_ID	1212K	00:00:00.21
	4	SORT ORDER BY		50	00:00:00.77
	5	NESTED LOOPS		1101	00:00:00.01
	6	TABLE ACCESS BY INDEX ROWID	T_ACCOUNTS	1	00:00:00.01
*	7	INDEX UNIQUE SCAN	SYS_C006991	1	00:00:00.01
	8	TABLE ACCESS BY INDEX ROWID	T_TRANSACTIONS	1101	00:00:00.01
*	9	INDEX RANGE SCAN	I_TRX_ACCO_ID	1101	00:00:00.01
	100				

# **G**Using recursive SQL

Copyright (c) 2009-2014 by Data Geekery GmbH. Slides licensed under CC BY SA 3.0



#### We need to number transactions

ID	VALUE_DATE	AMOUNT	TRANSACTION_NR
9997	2014-03-18	99.17	1
9981	2014-03-16	71.44	2
9979	2014-03-16	-94.60	3
9977	2014-03-16	-6.96	4
9971	2014-03-15	-65.95	5



```
CREATE OR REPLACE VIEW v_transactions_by_time
AS
SELECT
 t.*,
  ROW NUMBER() OVER (
    PARTITION BY account id
    ORDER BY t.value date DESC,
                 t.id DESC
  ) AS transaction number
FROM
  v transactions t;
```



```
EX
2
WITH ordered with balance (
  account id, value date, amount, balance, transaction number
AS (
  SELECT t1.account id, t1.value date, t1.amount, t1.current balance,
        t1.transaction number
       v transactions by time t1
  FROM
  WHERE t1.transaction number = 1
  UNION ALL
  SELECT t1.account_id, t1.value date, t1.amount, t2.balance - t2.amount,
        t1.transaction number
        ordered with balance t2
  FROM
  JOIN v transactions by time t1
        t1.transaction number = t2.transaction_number + 1
  ON
       t1.account id = t2.account id
  AND
SELECT *
        ordered with balance
FROM
        account id = 1
WHERE
ORDER BY transaction number ASC
```



```
WITH ordered_with_balance (
  account id, value date, amount, balance, transaction number
AS (
  SELECT t1.account id, t1.value date, t1.amount, t1.current balance,
         t1.transaction number
  FROM
       v transactions by time t1
  WHERE t1.transaction number = 1
  UNION ALL
  SELECT t1.account id, t1.value date, t1.amount, t2.balance - t2.amount,
         t1.transaction number
        ordered with balance t2
  FROM
        v transactions by time t1
  JOIN
        t1.transaction number = t2.transaction number + 1
  ON
       t1.account id = t2.account id
SELECT *
FROM
        ordered with balance
         account id = 1
WHERE
ORDER BY transaction number ASC
```



```
WITH ordered with balance (
  account id, value date, amount, balance, transaction number
AS (
  SELECT t1.account id, t1.value date, t1.amount, t1.current balance,
         t1.transaction number
        v transactions by time t1
  FROM
  WHERE t1.transaction_number = 1
  UNION ALL
  SELECT t1.account id, t1.value date, t1.amount, t2.balance - t2.amount,
         t1.transaction number
        ordered with balance t2
  FROM
        v transactions by time t1
  JOIN
        t1.transaction_number = t2.transaction_number + 1
  ON
        t1.account id = t2.account id
SELECT *
FROM
        ordered with balance
         account id = 1
WHERE
ORDER BY transaction number ASC
```



Id	Operation	Name	A-Rows	A-Time
0	SELECT STATEMENT		50	00:00:35.29
1	SORT ORDER BY		50	00:00:35.29
* 2	VIEW		1101	00:00:35.29
3	UNION ALL (RECURSIVE WITH) BREADTH FIRST		9999	00:00:35.28
* 4	VIEW	V_TRANSACTIONS_BY_TIME	9	00:00:00.03
* 5	WINDOW SORT PUSHED RANK		18	00:00:00.03
6	NESTED LOOPS		9999	00:00:00.01
7	NESTED LOOPS		9999	00:00:00.01
8	TABLE ACCESS FULL	T_ACCOUNTS	10	00:00:00.01
* 9	INDEX RANGE SCAN	I_TRX_ACCO_ID	9999	00:00:00.01
10	TABLE ACCESS BY INDEX ROWID	T_TRANSACTIONS	9999	00:00:00.01
* 11	HASH JOIN		9990	00:00:35.08
12	VIEW	V_TRANSACTIONS_BY_TIME	<b>11</b> M	00:00:29.13
13	WINDOW SORT		<b>11</b> M	00:00:27.19
14_	NESTED LOOPS		11M	00:00:13.62
15	NESTED LOOPS		<b>1</b> 1M	00:00:03.89
16	INDEX FAST FULL SCAN	SYS_C006991	11450	00:00:00.06
* 17	INDEX RANGE SCAN	I_TRX_ACCO_ID	11M	
18	TABLE ACCESS BY INDEX ROWID	T_TRANSACTIONS	11M	
	PUMP		9999	00:00:00.01

# Using window functions

"

Copyright (c) 2009-2014 by Data Geekery GmbH. Slides licensed under CC BY SA 3.0

```
SELECT
 t.*,
 t.current_balance - NVL(
   SUM(t.amount) OVER (
     PARTITION BY t.account_id
                  t.value_date DESC,
     ORDER BY
                  t.id DESC
     ROWS BETWEEN UNBOUNDED PRECEDING
          AND
               1
                            PRECEDING
  0) AS balance
FROM
       v_transactions t
WHERE t.account id = 1
ORDER BY t.value_date DESC,
        t.id
                     DESC
```







## Using window functions

ID	VALUE_DATE	AMOUNT	BALANCE	
9997	2014-03-18	-(99.17)	+19985.81	
9981	2014-03-16	-(71.44)	19886.64	
9979	2014-03-16	-(-94.60)	19815.20	
9977	2014-03-16	-6.96	=19909.80	
9971	2014-03-15	-65.95	19916.76	



I	d	Operation	Name	A-Rows	A-Time
	0	SELECT STATEMENT		50	00:00:00.01
	1	WINDOW SORT		50	00:00:00.01
	2	NESTED LOOPS		1101	00:00:00.01
	3	TABLE ACCESS BY INDEX ROWID	T_ACCOUNTS	1	00:00:00.01
*	4	INDEX UNIQUE SCAN	SYS_C006991	1	00:00:00.01
	5	TABLE ACCESS BY INDEX ROWID	T_TRANSACTIONS	1101	00:00:00.01
*	6	INDEX RANGE SCAN	I_TRX_ACCO_ID	1101	00:00:00.01

#### FASGINATING memegenerator.net der CC BY SA 3.0



"

# Using the Oracle MODEL clause



```
SELECT account id, value date, amount, balance
FROM (
  SELECT id, account id, value date, amount,
         current balance AS balance
 FROM v_transactions
) t
WHERE account id = 1
MODEL
  PARTITION BY (account_id)
  DIMENSION BY (
    ROW NUMBER() OVER (ORDER BY value date DESC, id DESC) AS rn
  MEASURES (value date, amount, balance)
  RULES (
    balance[rn > 1] = balance[cv(rn) - 1] - amount[cv(rn) - 1]
ORDER BY rn ASC
```



```
SELECT account id, value date, amount, balance
FROM (
  SELECT id, account id, value date, amount,
         current balance AS balance
 FROM v transactions
) t
WHERE account id = 1
MODEL
 PARTITION BY (account_id)
 DIMENSION BY (
    ROW_NUMBER() OVER (ORDER BY value_date DESC, id DESC) AS rn
 MEASURES (value date, amount, balance)
  RULES (
    balance[rn > 1] = balance[cv(rn) - 1] - amount[cv(rn) - 1]
ORDER BY rn ASC
```



#### 

SU	IMME 🔹	:	×	✓ f <sub>×</sub>	-	=C3-B3
	А		В	С		D
1	value_date		amount	balar	ice	
2	17.03.2014		15.87	13222	.45	
3	16.03.2014		-33.14	13206	.58	
4	16.03.2014		-93.77	=C3-B3		
5	13.03.2014		10.65	13333	.49	
6	11.03.2014		19.16	13322	.84	
7	11.03.2014		-59.25	13303	.68	
8	11.03.2014		94.86	13362	.93	
9	10.03.2014		80.42	13268	.07	
10	10.03.2014		38.43			
11	09.03.2014		-4.41			
12	08.03.2014		80.45	X		
13	07.03.2014		-56.45			
-						

#### -- does it look familiar?

Id	Operation	Name	A-Rows	A-Time
0   1   2   3   4   5  * 6  * 7	SELECT STATEMENT SORT ORDER BY SQL MODEL ORDERED WINDOW SORT NESTED LOOPS TABLE ACCESS BY INDEX ROWID INDEX UNIQUE SCAN	T_ACCOUNTS SYS_C006991 T_TPANSACTTONS	50 50 1101 1101 1101 1 1 1 1101	00:00:00.02         00:00:00.02         00:00:00.02         00:00:00.02         00:00:00.01         00:00:00.01         00:00:00.01         00:00:00.01         00:00:00.01         00:00:00.01



WHEN A PLAN CONTAINS A MODEL CLAUSE memegenerator.net)H. Slides licensed under CC BY SA 3.0-



## The MODEL clause is Oracle's most powerful and underused feature



#### Our vision at Data Geekery - Revisited

- SQL dominates database systems
- SQL is expressive
- SQL is type safe

## SQL is a device whose mystery is only exceeded by its power!

#### Our vision at Data Geekery - Revisited

## ٰ کا نے LiooQ is the best way to write SQL in Java



```
SELECT
 t.*,
  t.current_balance - NVL(
   SUM(t.amount) OVER (
      PARTITION BY t.account_id
                  t.value_date DESC,
      ORDER BY
                   t.id
                               DESC
      ROWS BETWEEN UNBOUNDED PRECEDING
           AND
                1
                             PRECEDING
    ),
  0) AS balance
FROM
       v_transactions t
WHERE t.account id = 1
ORDER BY t.value_date DESC,
         t.id
                      DESC
```







## Thank you

## More free Java / SQL knowledge on:

- Blog: <u>http://blog.jooq.org</u>
- Twitter: <u>@JavaOOQ</u>

