



# Finding the Pitfalls in Query Performance

M.L. Kersten

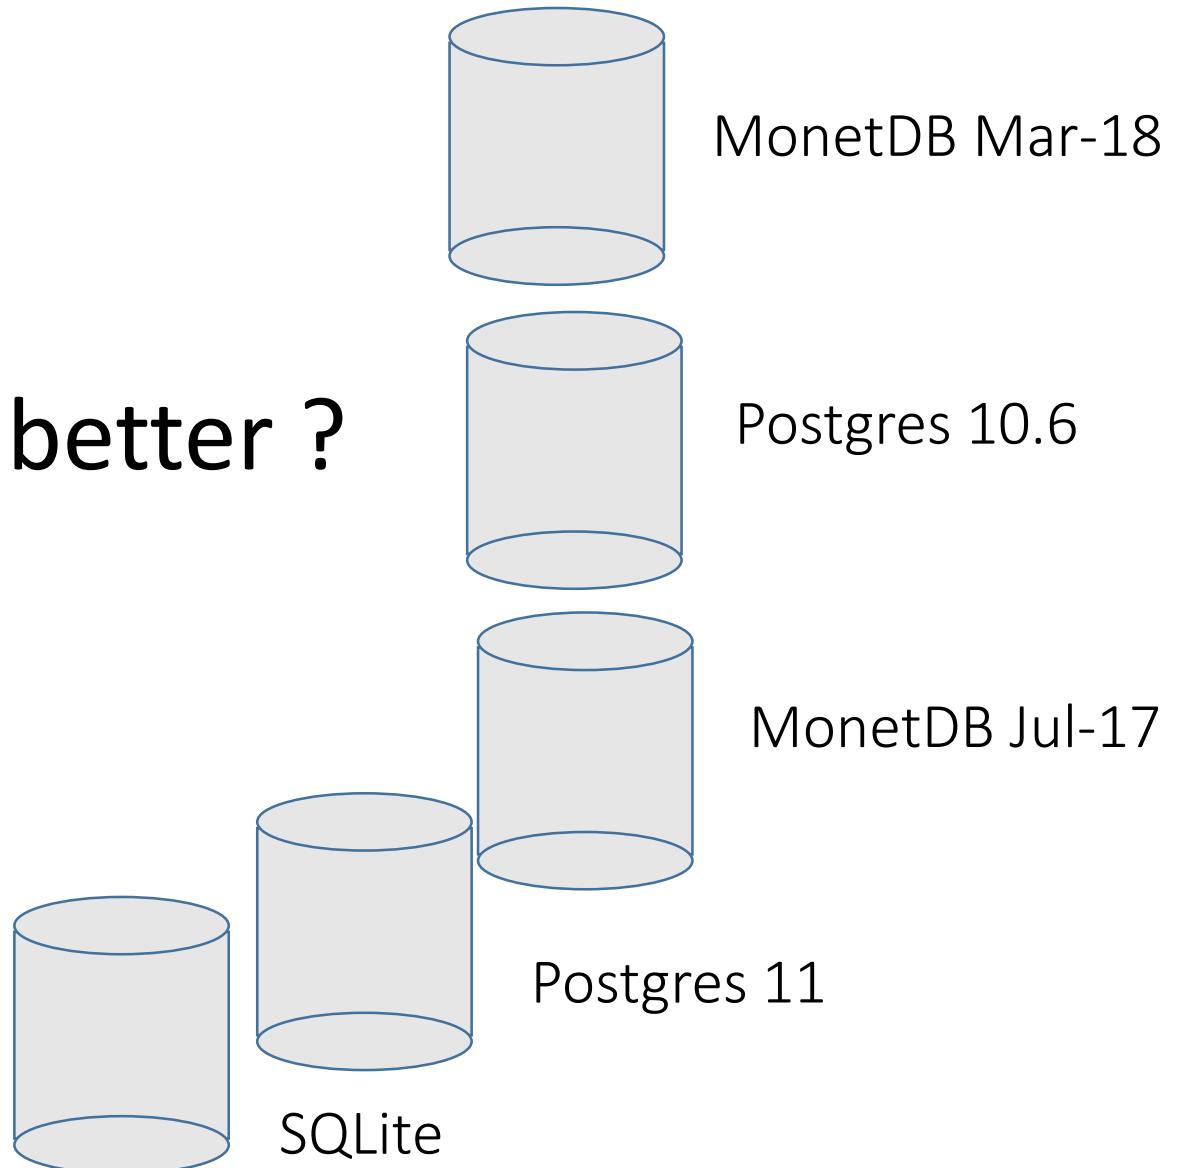
P. Koutsourakis

Y. Zhang

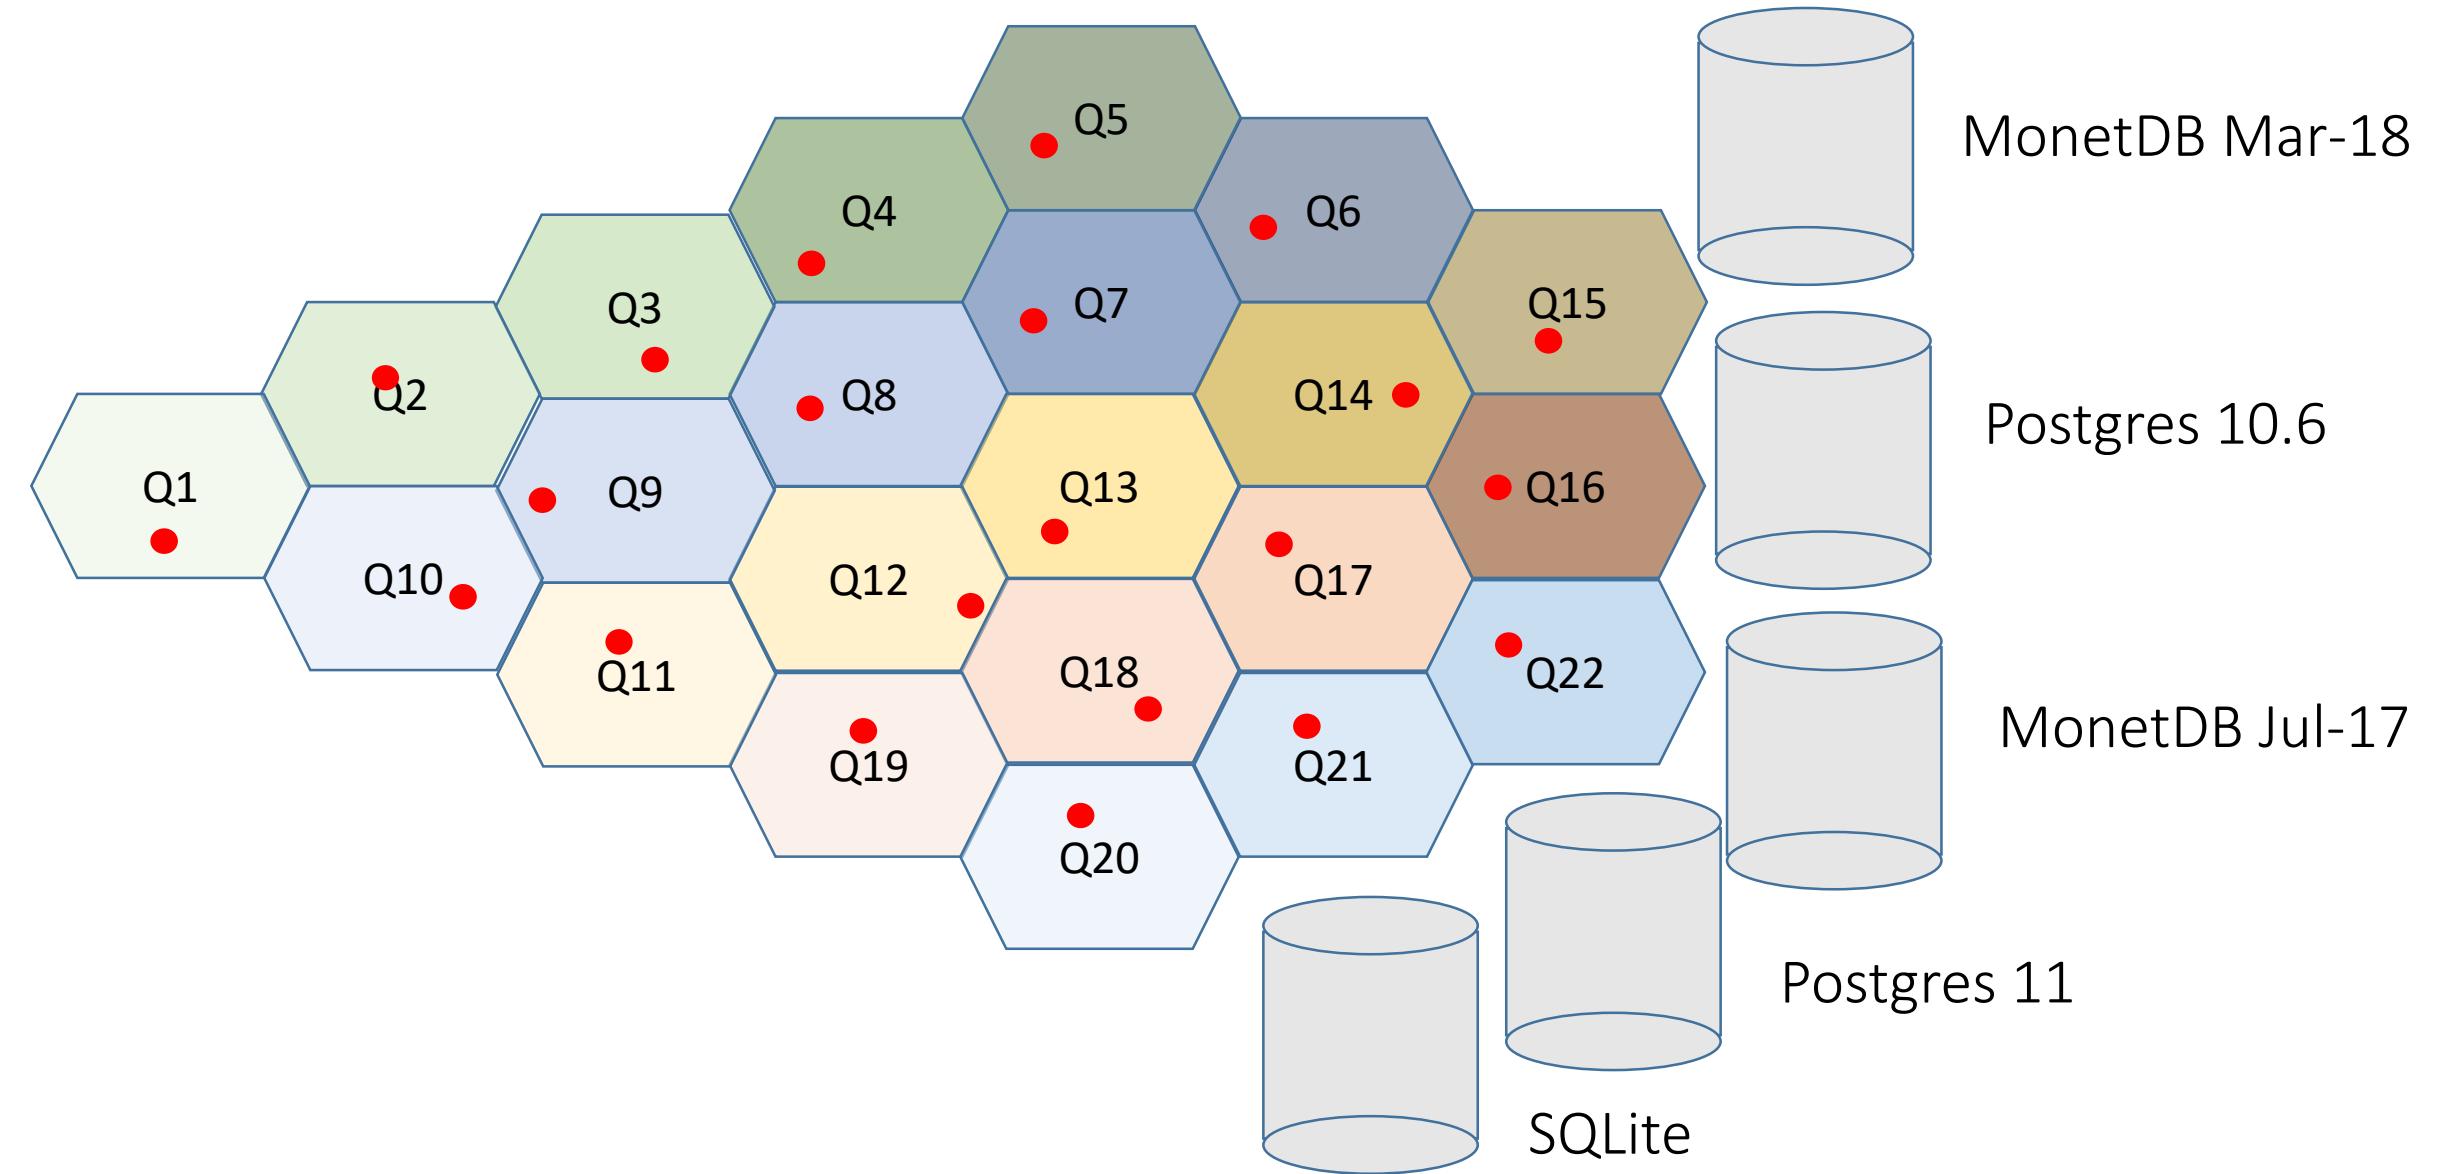
*CWI, MonetDB Solutions*

## The Challenge

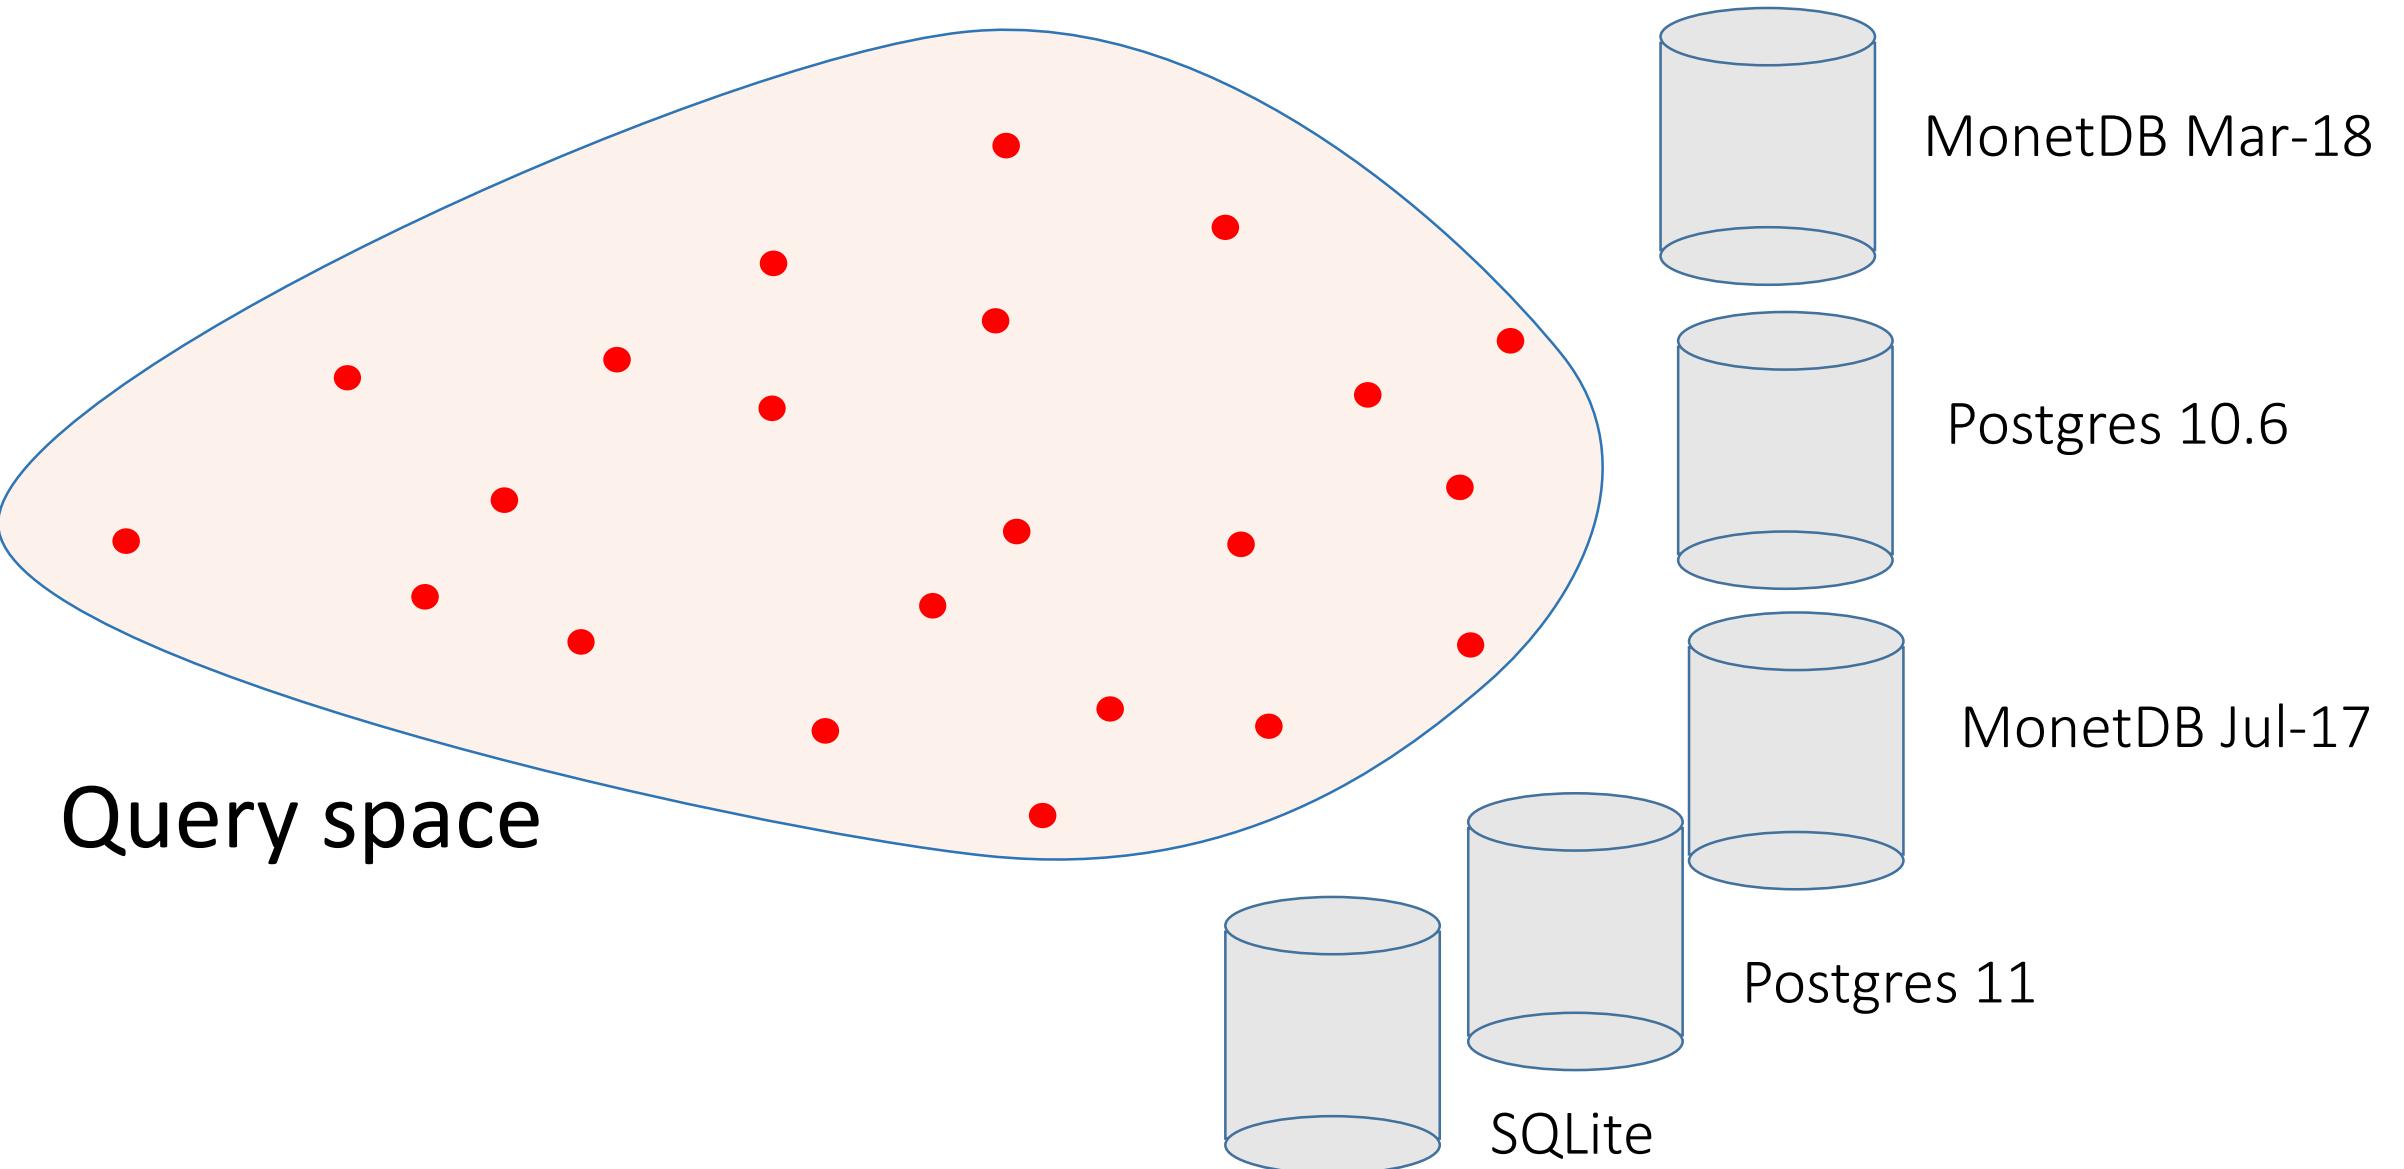
Which system is **relatively** better ?



# The Solution, TPC-H ?

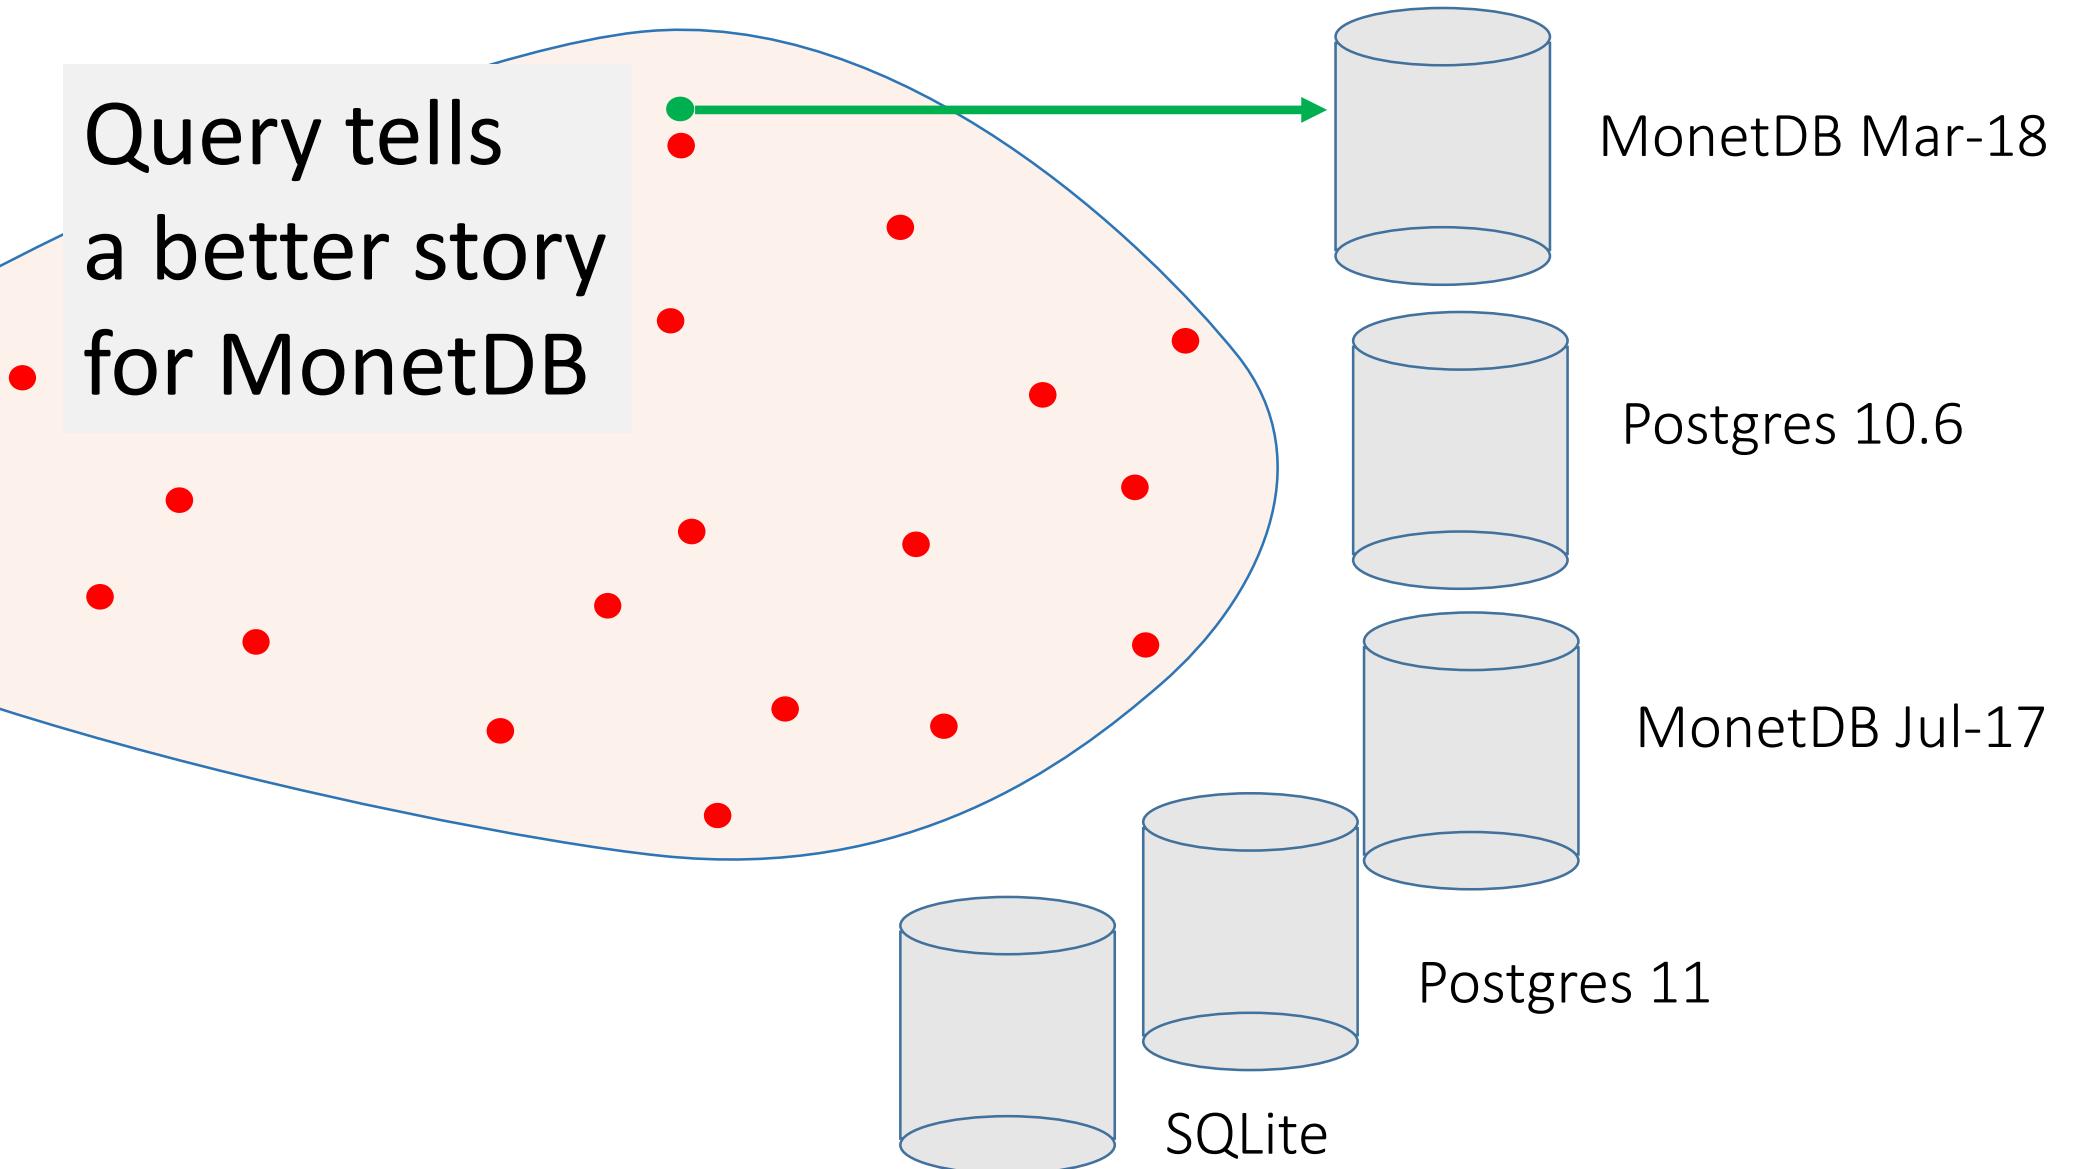


# TPC-H is a collection of ‘random points’?



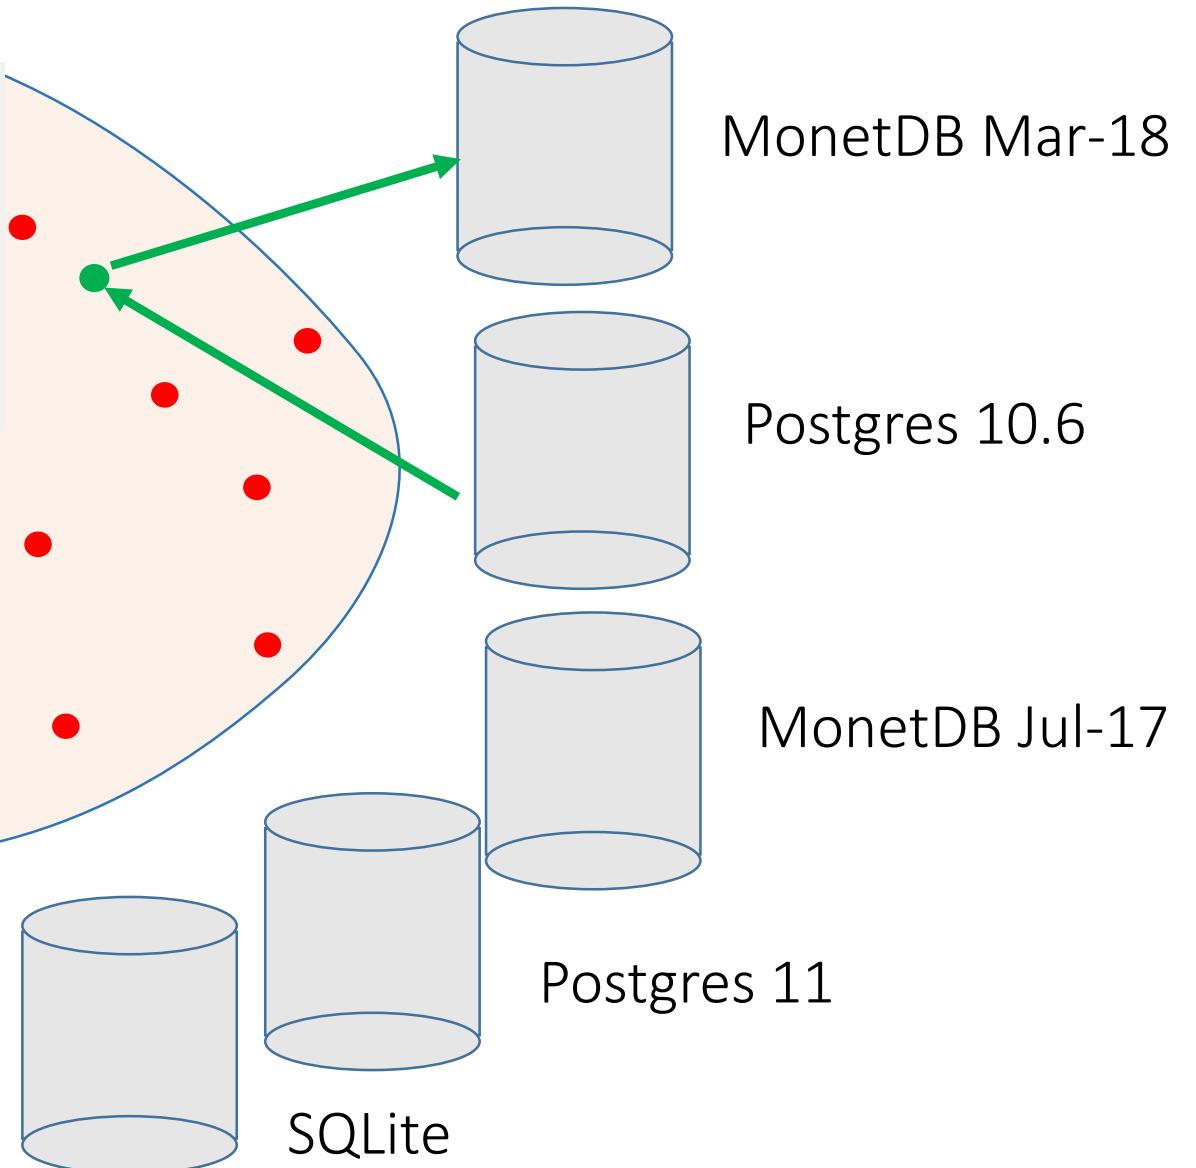
# TPC-H may miss clarifying queries

Query tells  
a better story  
for MonetDB

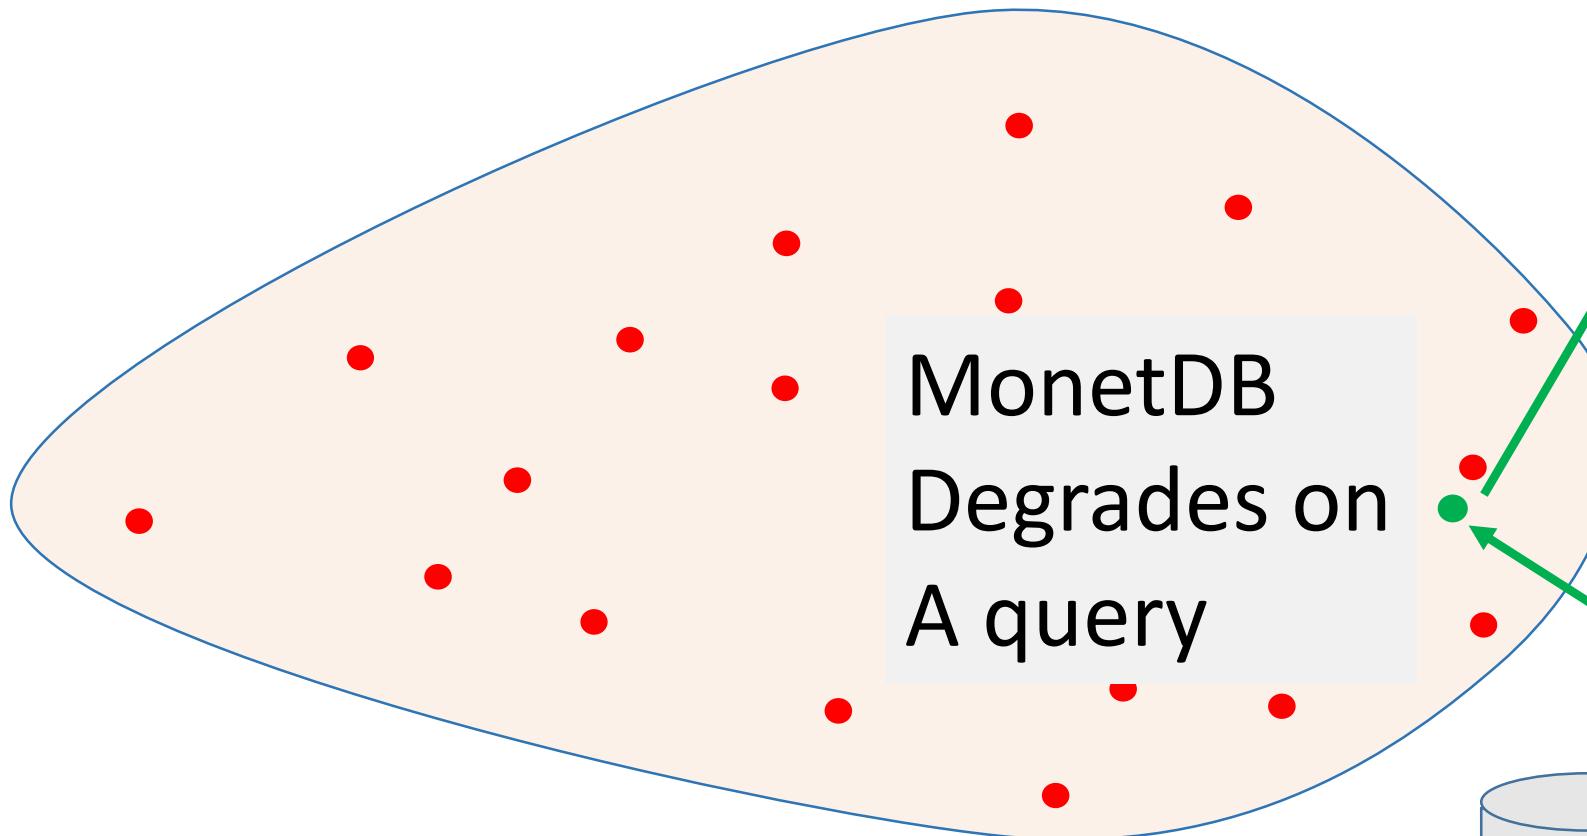


# TPC-H may miss discriminative queries

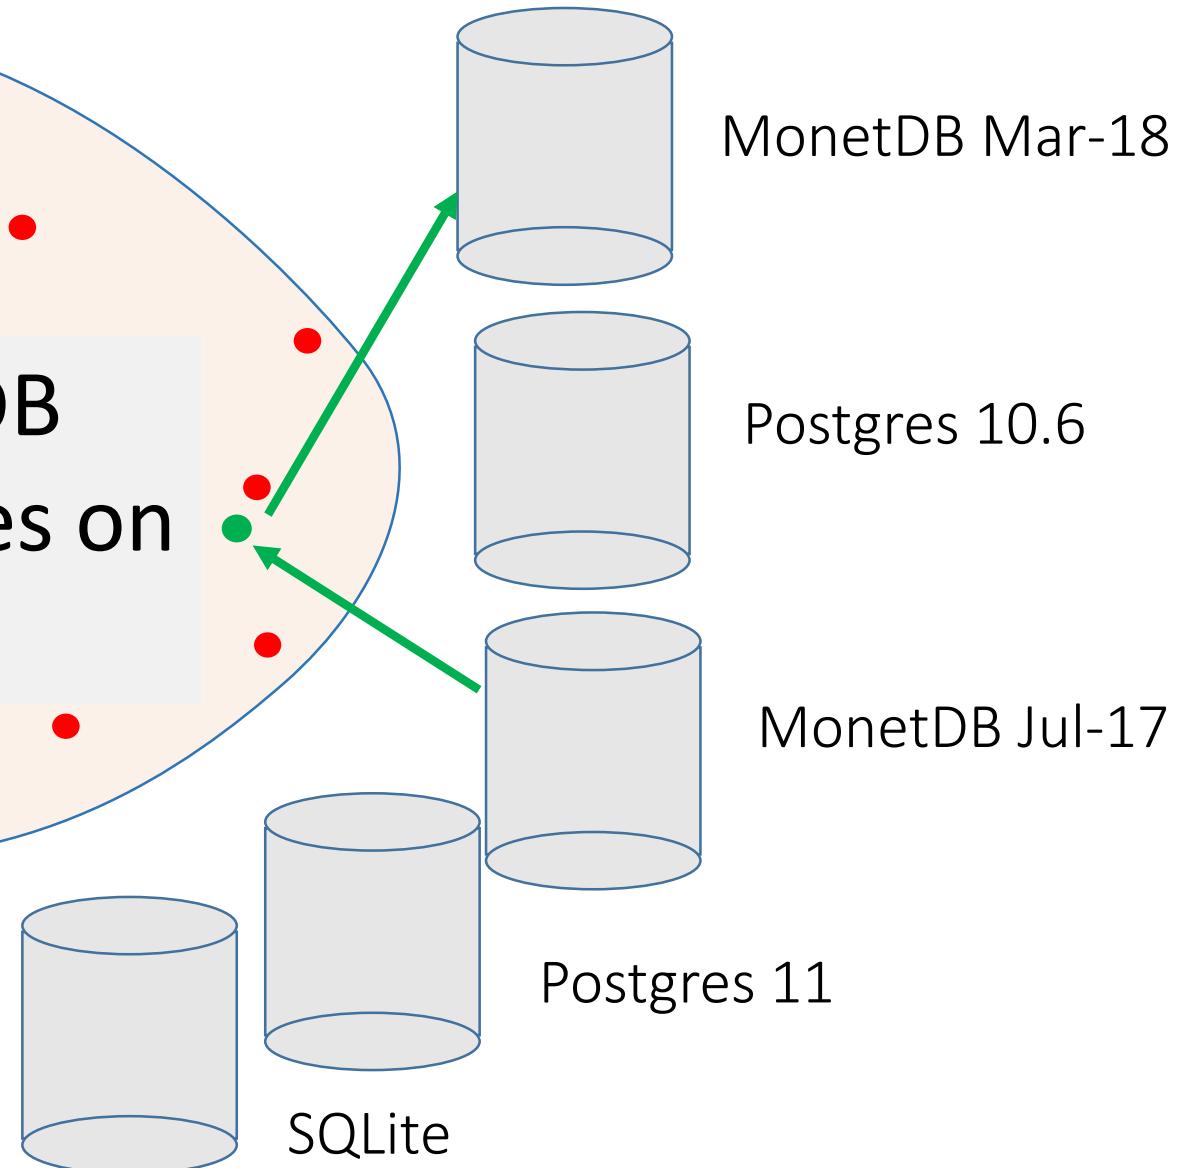
Postgres  
is relatively better  
than MonetDB



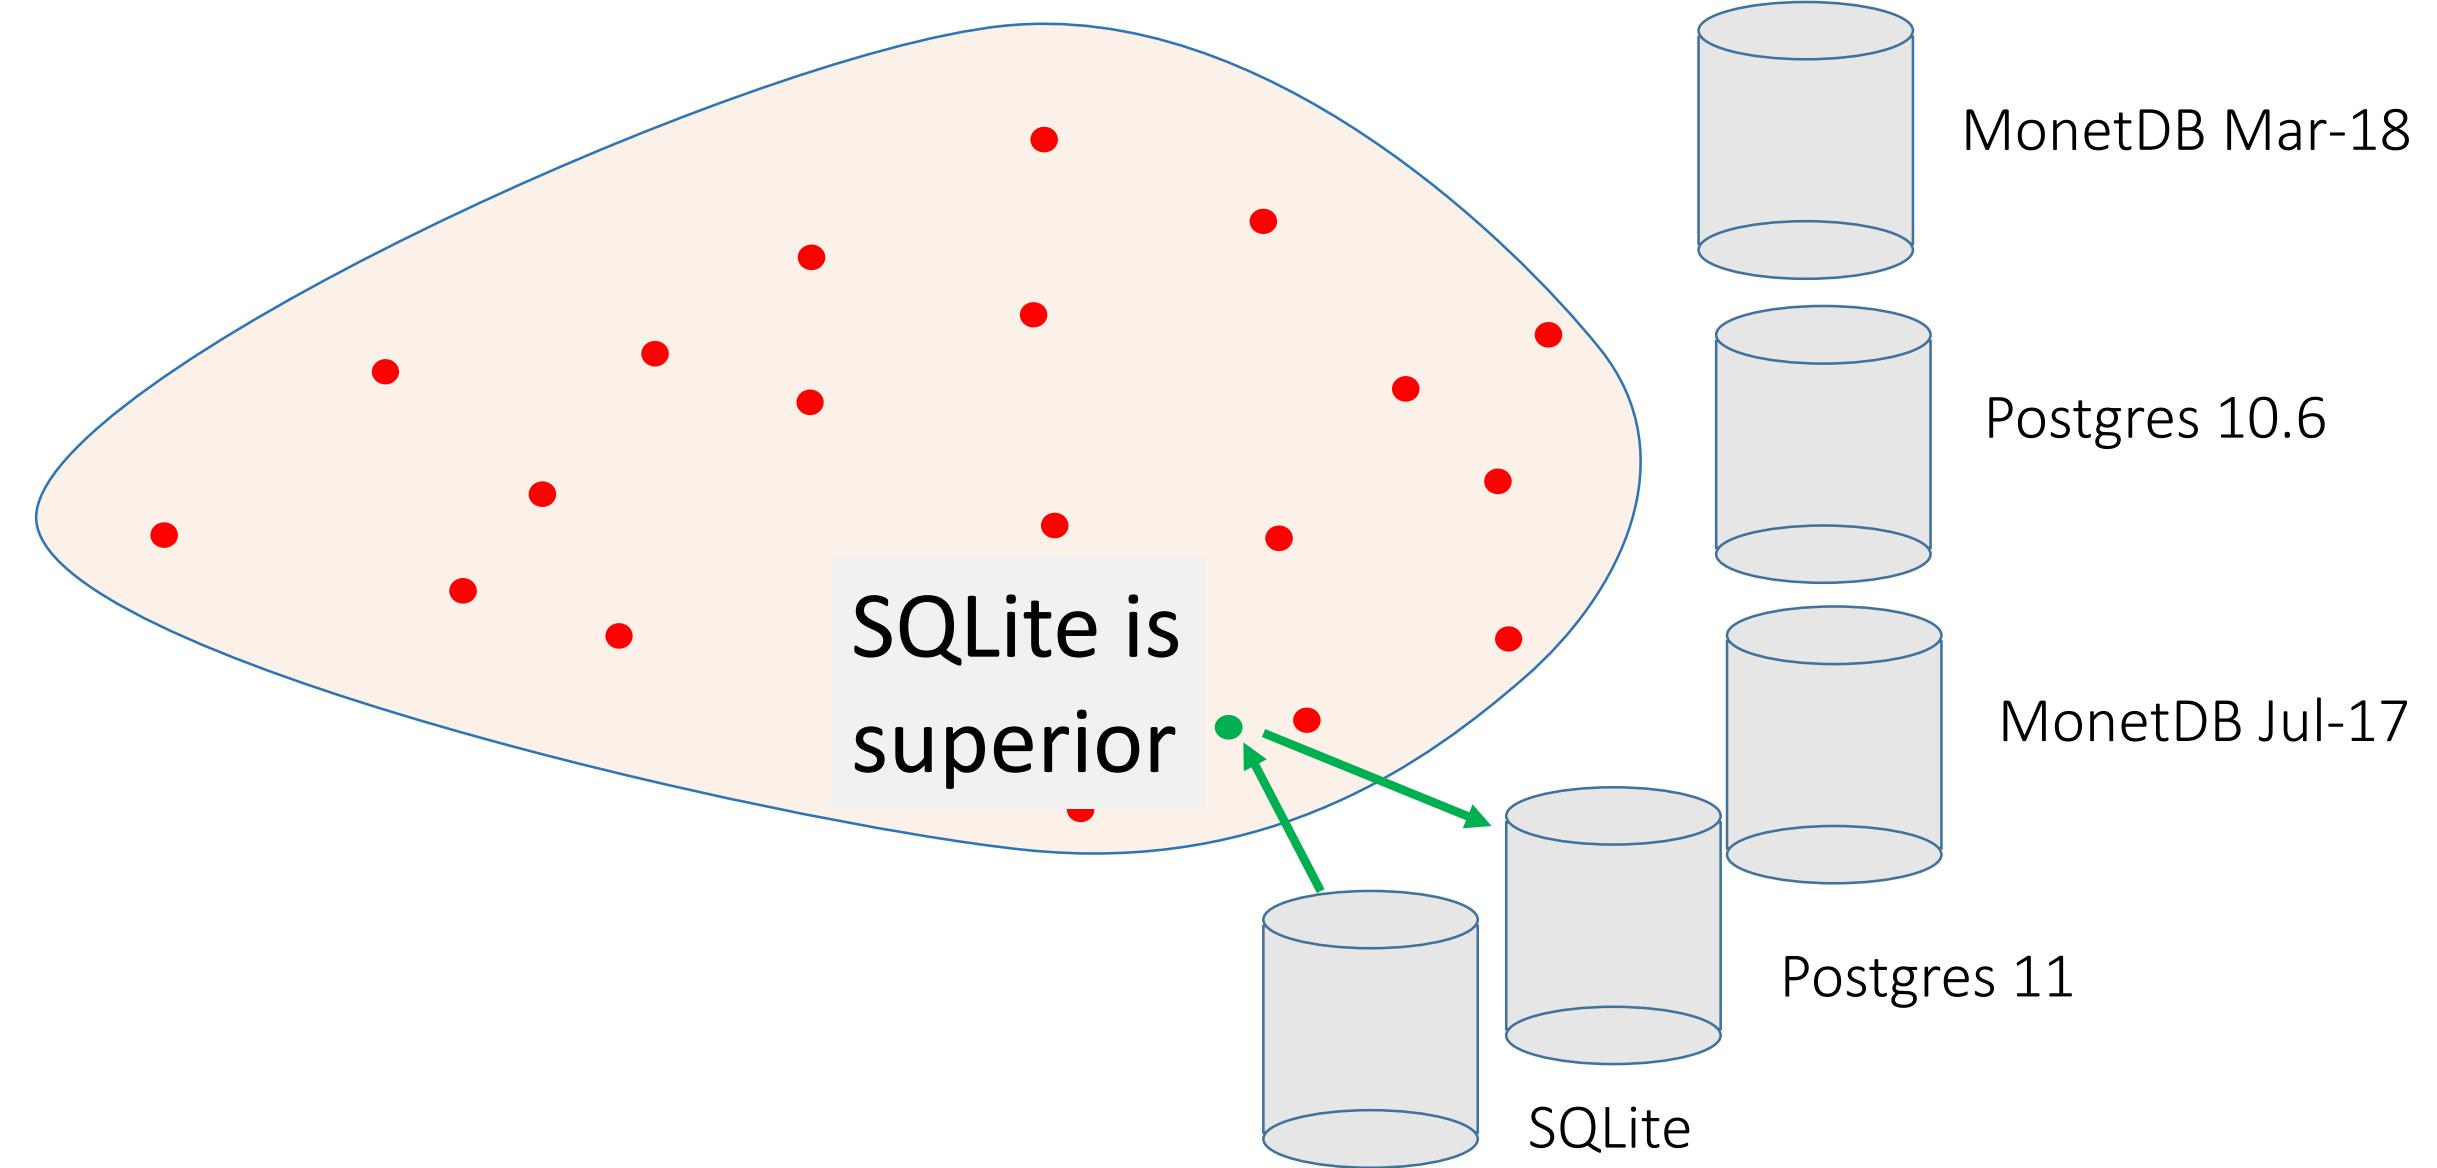
# TPC-H may miss discriminative queries



MonetDB  
Degrades on  
A query



# TPC-H may miss discriminative queries

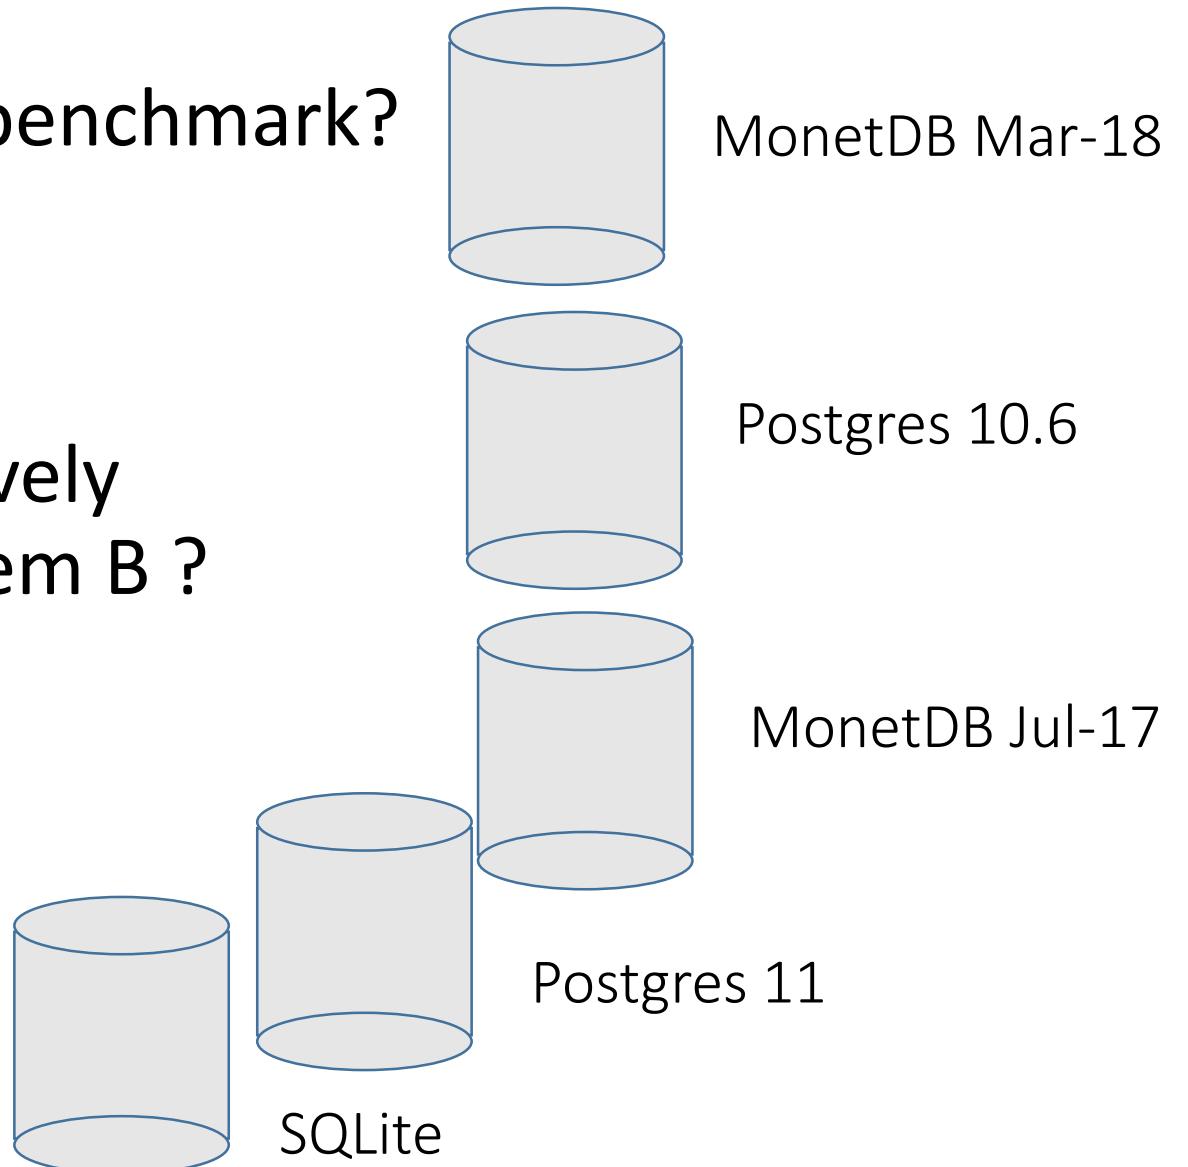


## The Challenge

- Which system is better on a benchmark?



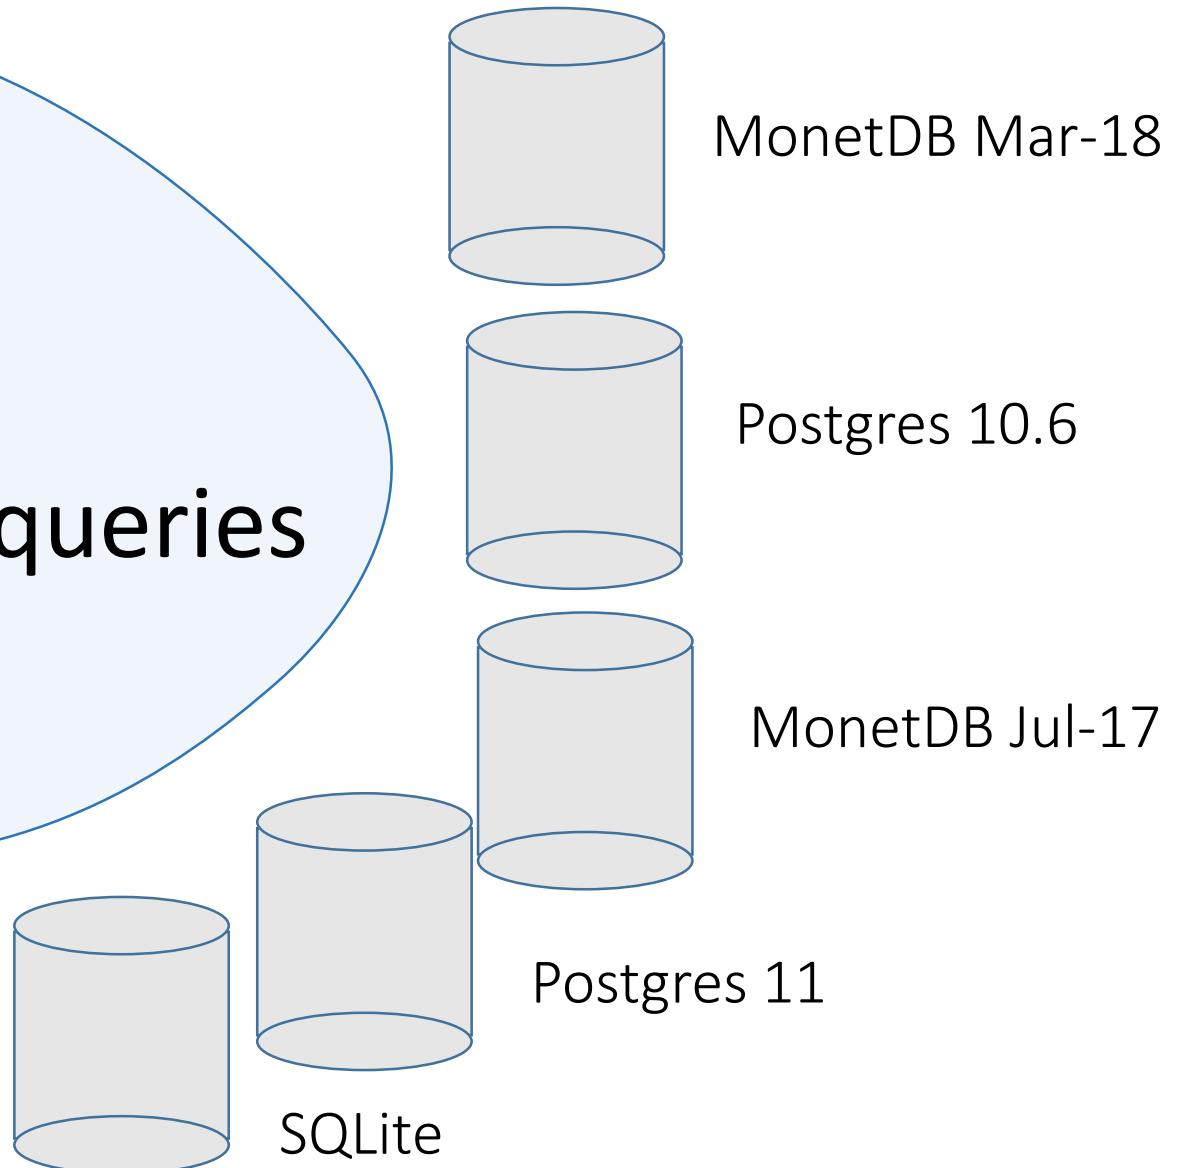
- What queries perform relatively better on system A than system B ?



# SQL *scalpel*

## Find the discriminative queries

The Challenge

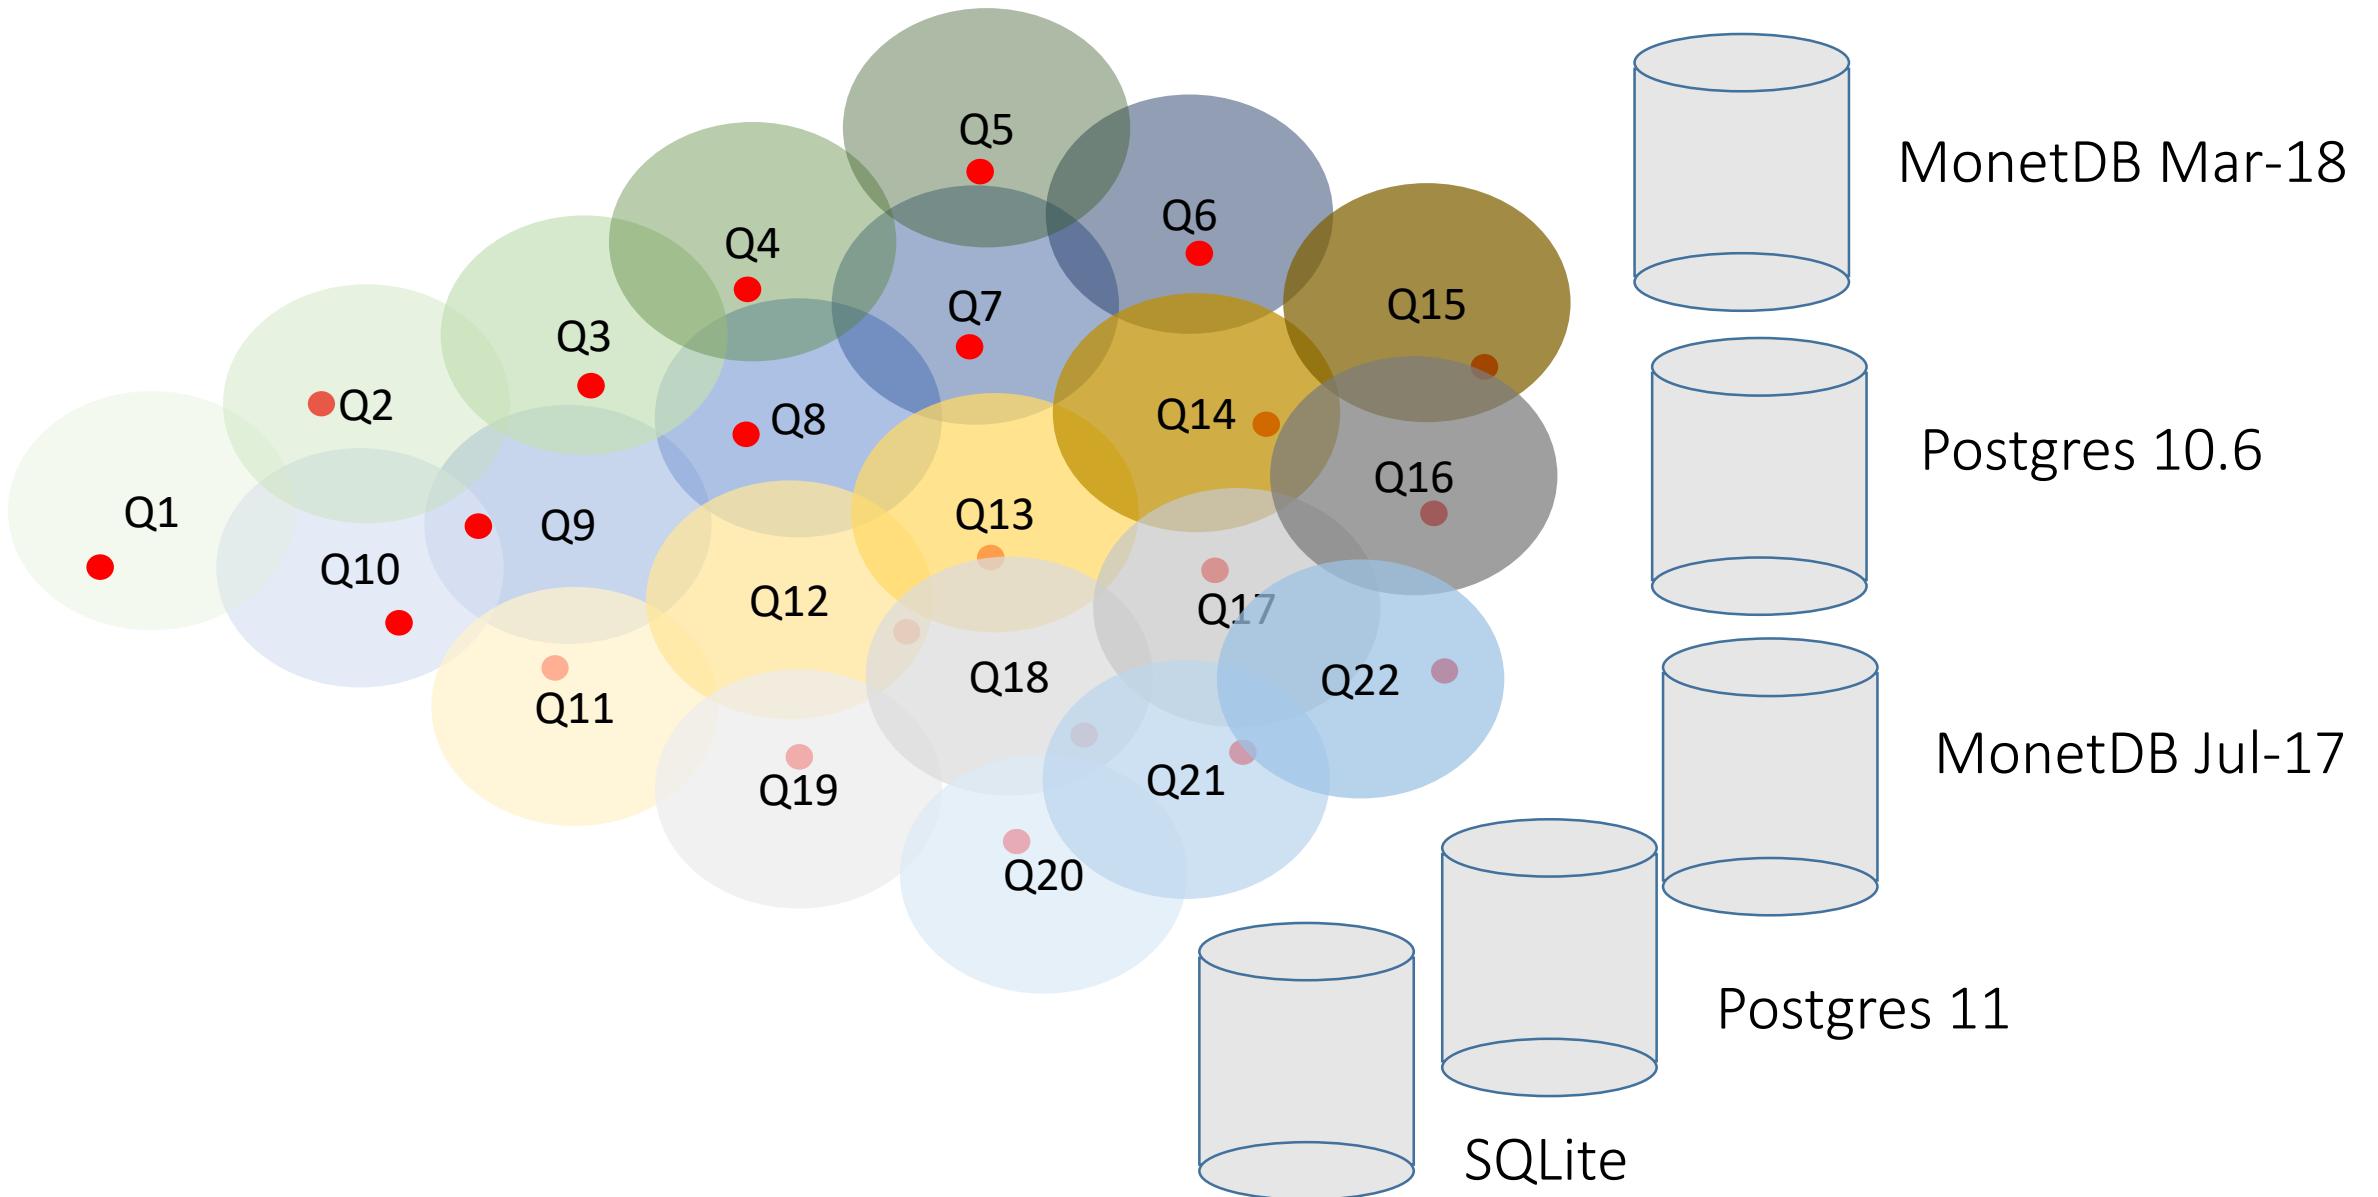


# SQL *scalpel*

Find the discriminative queries

- The database schema and data distribution are given
- A collection of business inspired queries are available
- **No direct access** to the Database/DBMS/Platform

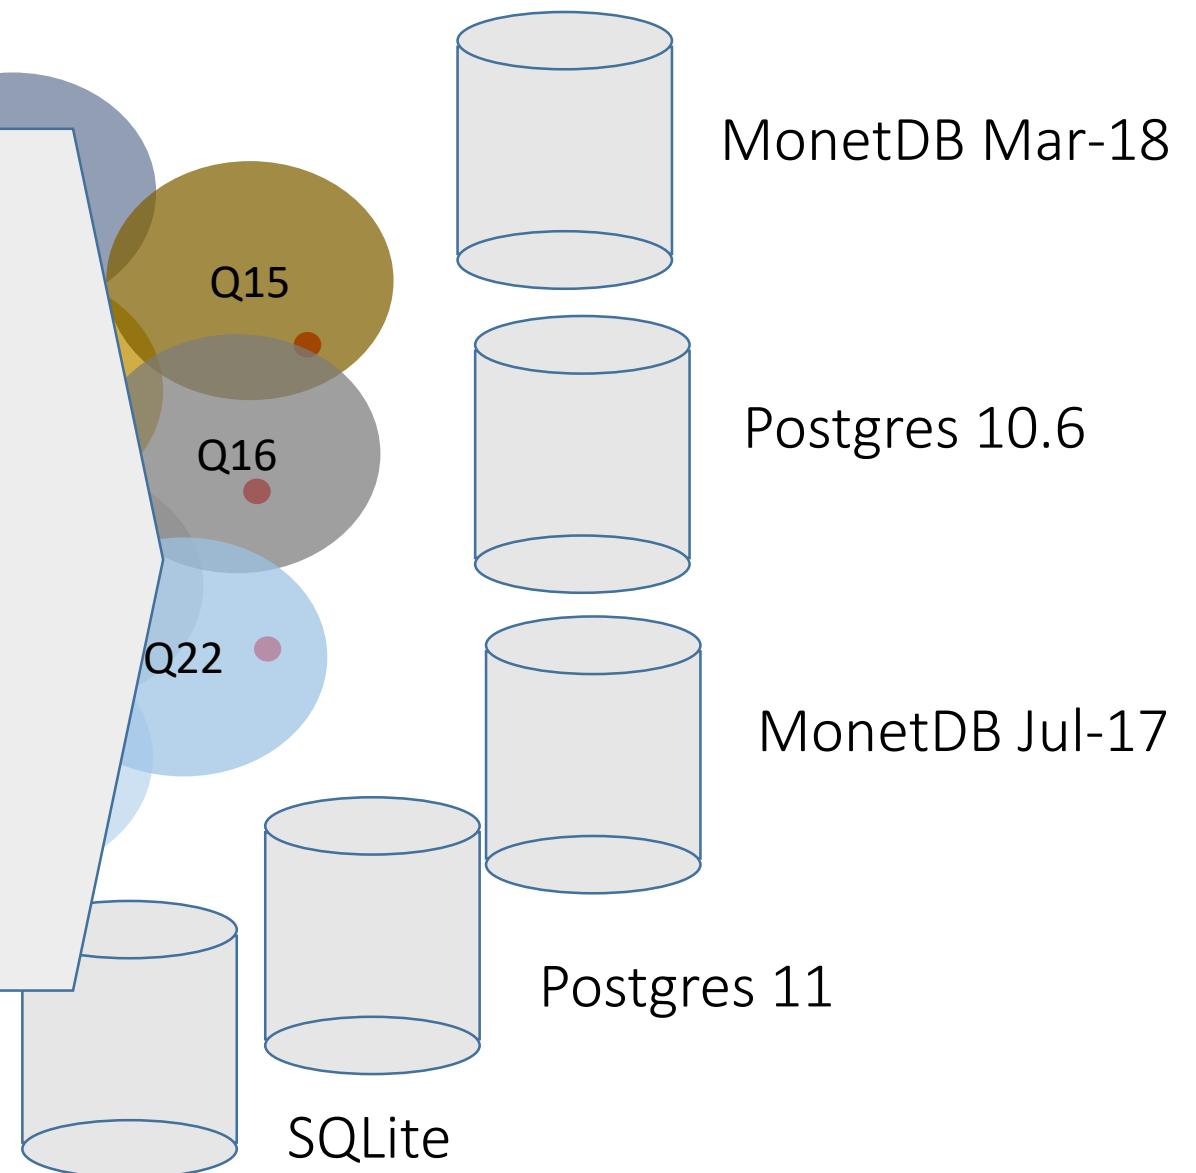
# The Solution, TPC-H as a start



# The Solution, TPC-H as a start

```
-- Query 6

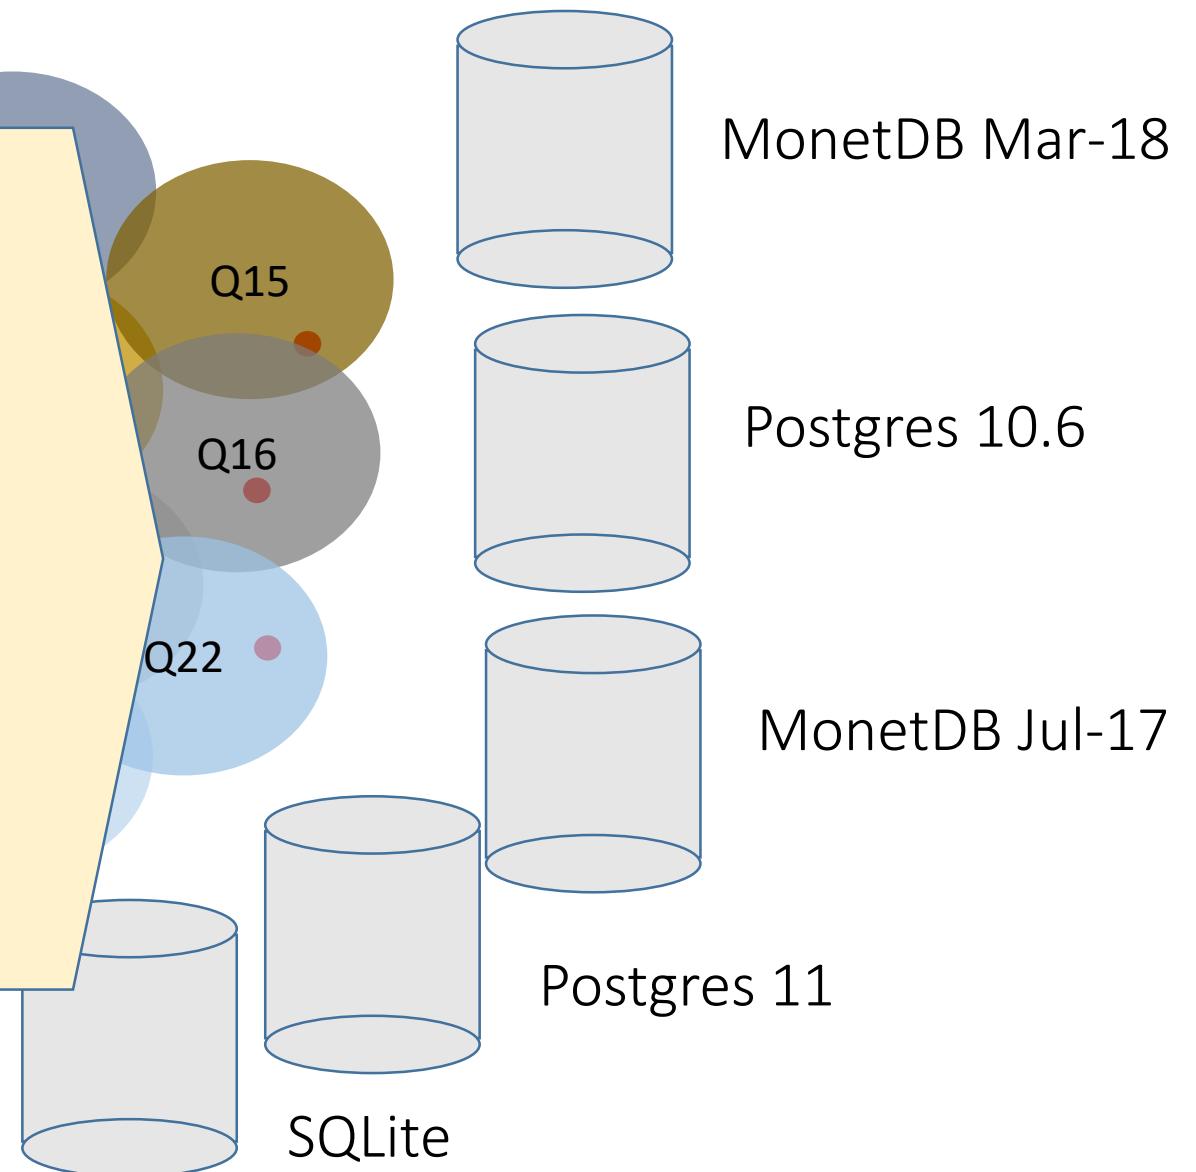
select
    sum(l_extendedprice * l_discount) as revenue
from
    lineitem
where
    l_shipdate >= date '1994-01-01' and
    l_shipdate < date '1994-01-01' + interval '1' year and
    l_discount between .06 - 0.01 and .06 + 0.01 and
    l_quantity < 24;
```



# The Solution, TPC-H as a start

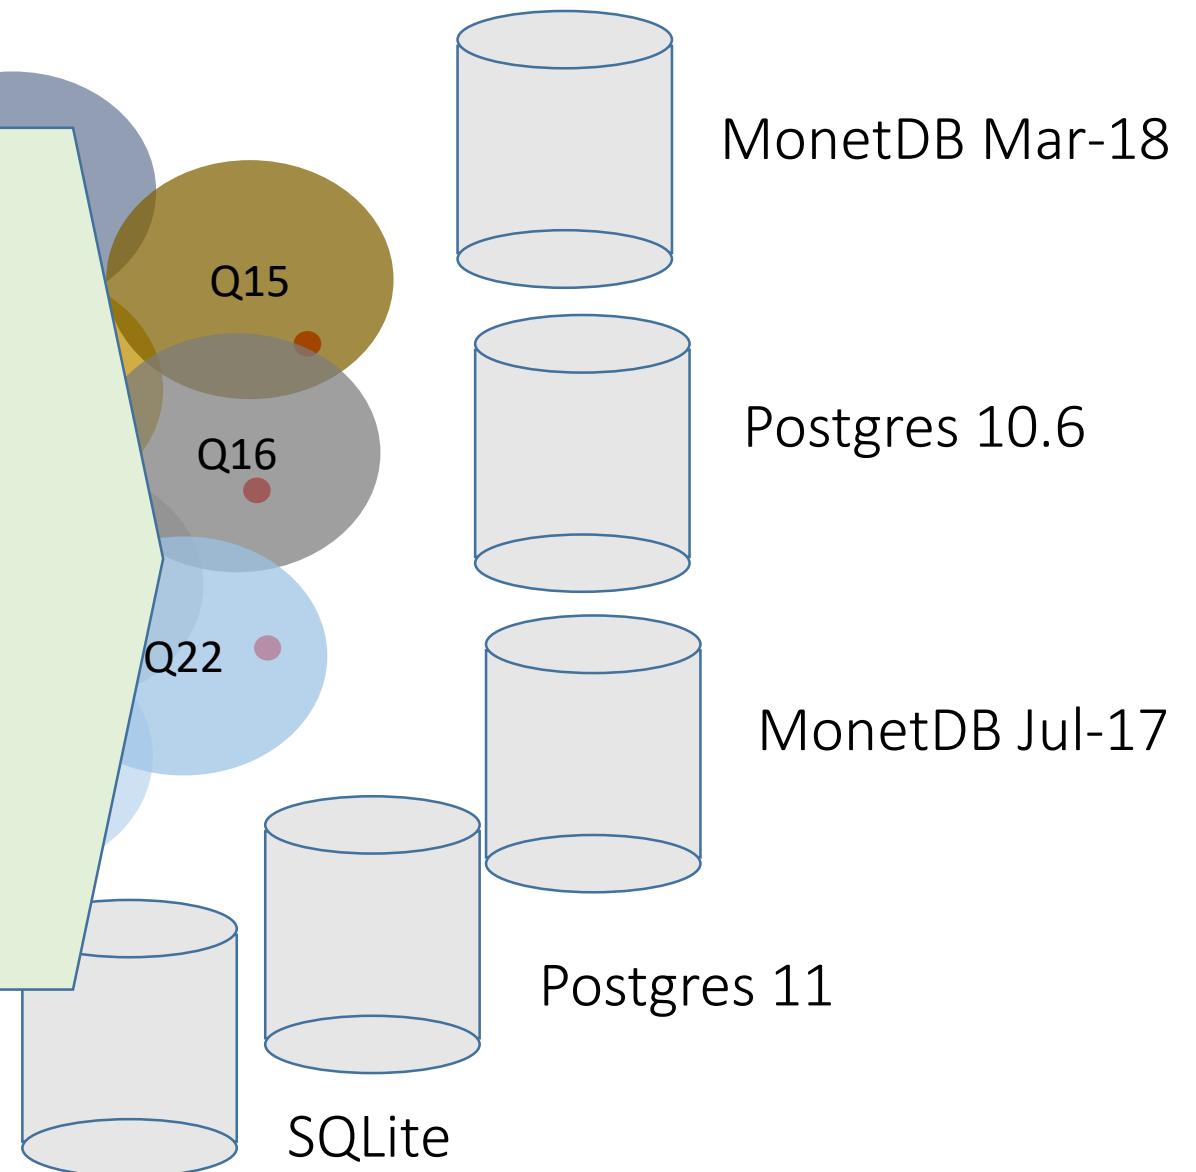
```
-- Query 6

select
    sum(l_extendedprice * l_discount) as revenue
from
    lineitem
where
    l_shipdate >= date '1994-01-01'
```



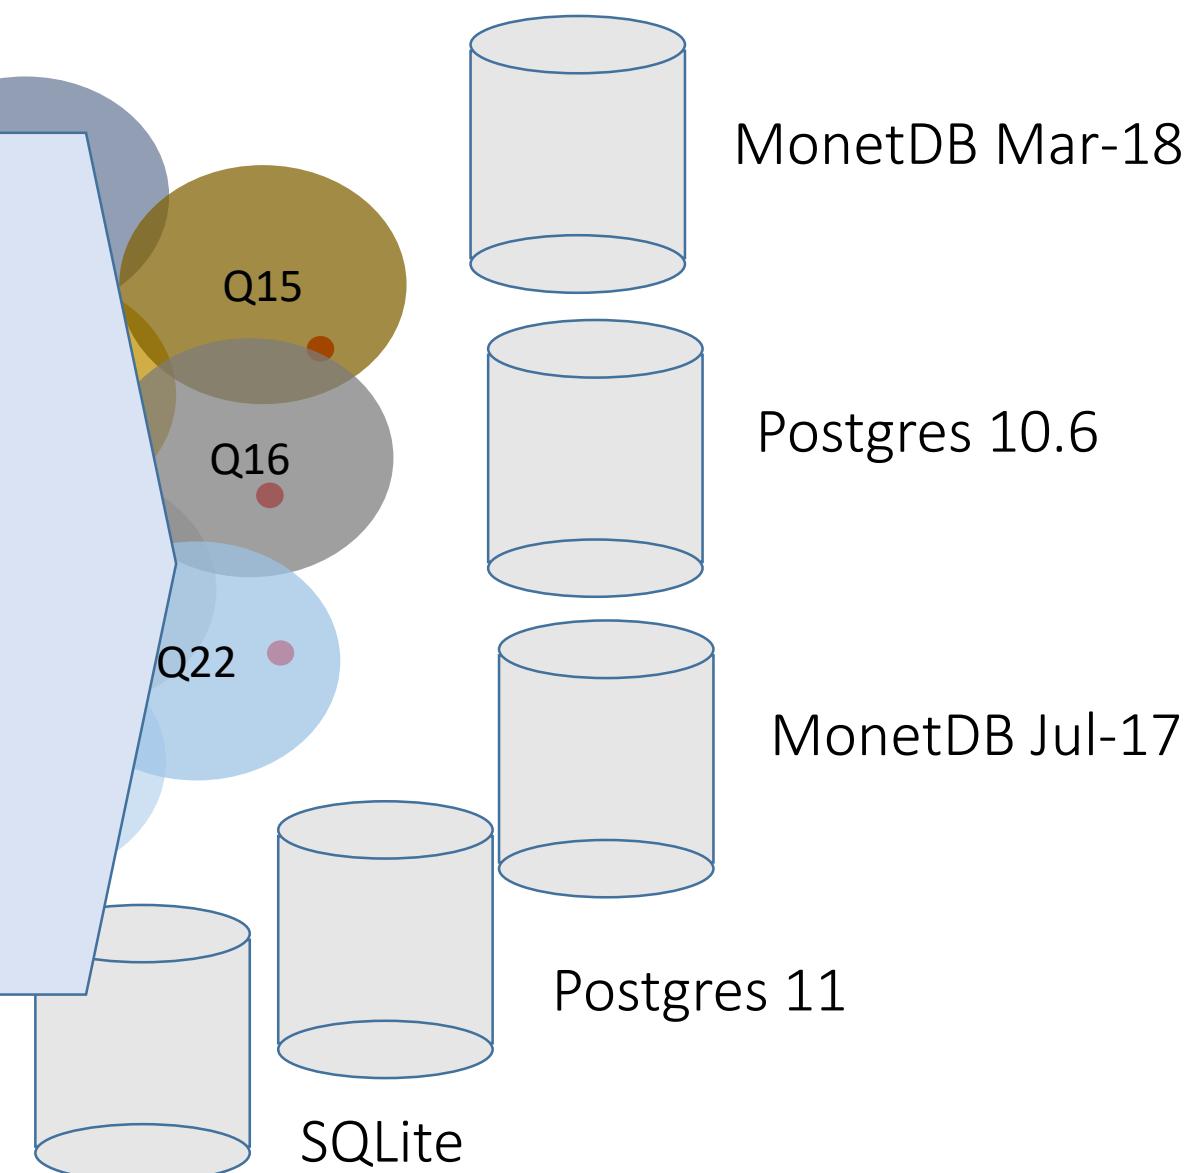
# The Solution, TPC-H as a start

```
-- Query 6  
  
select  
    sum(l_extendedprice * l_discount) as revenue  
from  
    lineitem  
where  
    l_discount between .06 - 0.01 and .06 + 0.01 and  
    l_quantity < 24;
```



# The Solution, TPC-H as a start

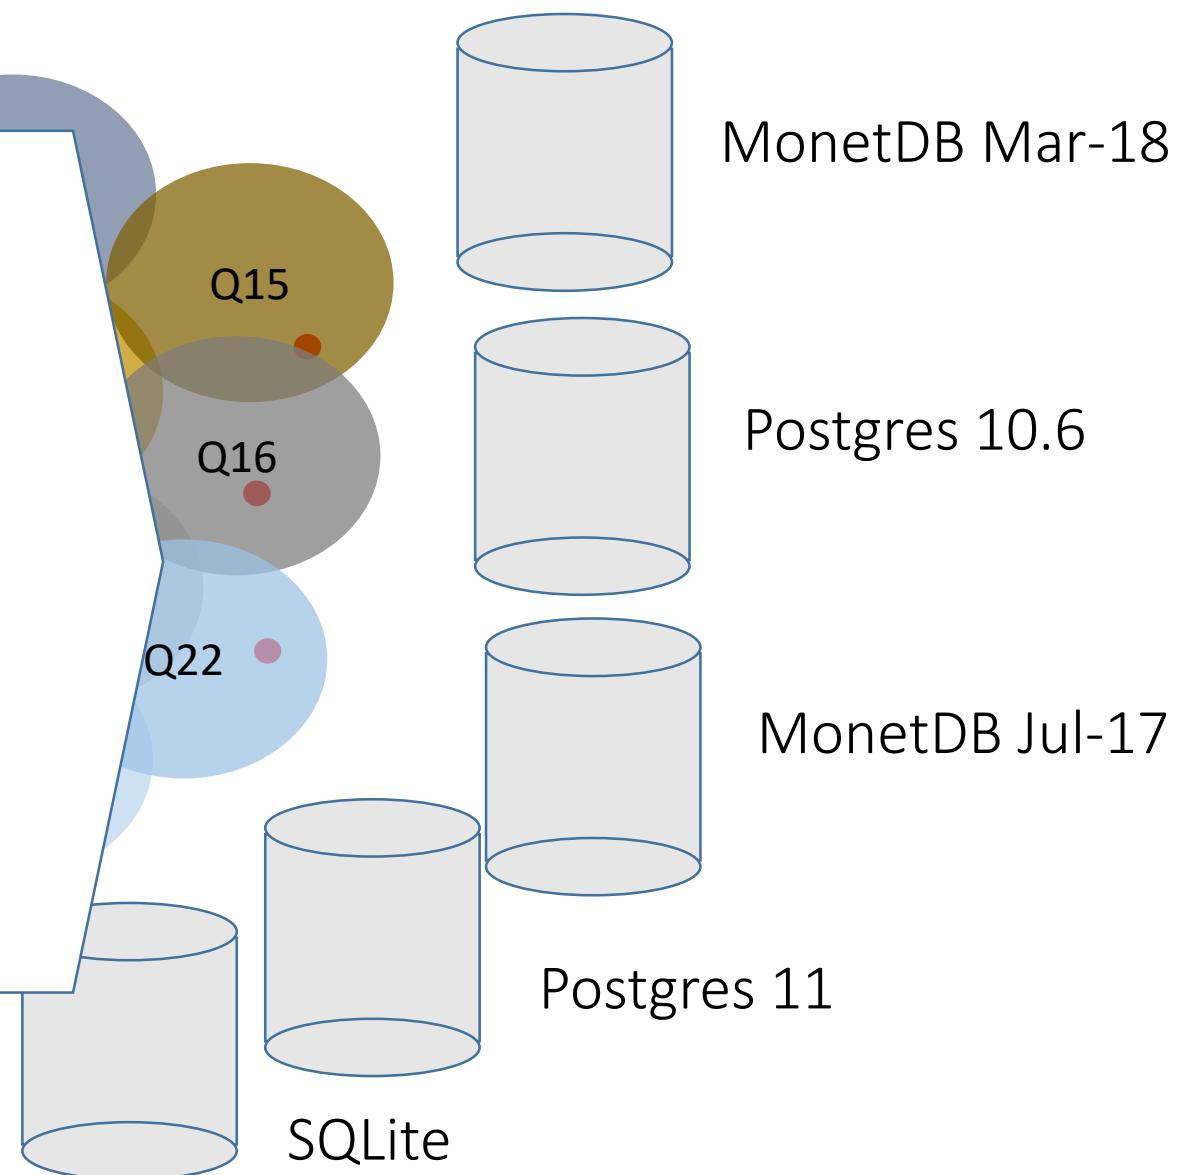
```
-- Query 6  
  
select  
    sum(l_extendedprice * l_discount) as revenue  
from  
    lineitem  
where  
    l_shipdate < date '1994-01-01' + interval '1' year
```



# The Solution, TPC-H as a start

```
-- Query 6

select
    sum(l_extendedprice * l_discount) as revenue
from
    lineitem
where
    l_shipdate >= date '1994-01-01' and
    l_shipdate < date '1994-01-01' + interval '1' year and
    l_discount between .06 - 0.01 and .06 + 0.01 and
    l_quantity < 24
```



# SQLscalpel compiles the query into a grammar

```
-- Query 6

select
    sum(l_extendedprice * l_discount) as revenue
from
    lineitem
where
    l_shipdate >= date '1994-01-01' and
    l_shipdate < date '1994-01-01' + interval '1' year and
    l_discount between .06 - 0.01 and .06 + 0.01 and
    l_quantity < 24
```

tpch\_q6:  
SELECT \${projection} FROM \${table}  
WHERE \${pred} \${predlist}\*  
  
projection:  
sum(l\_extendedprice \* l\_discount) as revenue  
  
table:  
lineitem  
  
pred:  
l\_shipdate >= date '1994-01-01'  
l\_shipdate < date '1994-01-01' + interval '1' year  
l\_discount between 0.06 - 0.01 and 0.06 + 0.01  
l\_quantity < 24  
  
predlist:  
AND \${pred}

# SQLscalpel enumerates templates

```
SELECT ${projection} FROM ${table}  
WHERE ${pred}
```

```
SELECT ${projection} FROM ${table}  
WHERE ${pred} AND ${pred}
```

```
SELECT ${projection} FROM ${table}  
WHERE ${pred} AND ${pred} AND ${pred}
```

```
SELECT ${projection} FROM ${table}  
WHERE ${pred} AND ${pred} AND ${pred} AND ${pred}
```

tpch\_q6:

```
SELECT ${projection} FROM ${table}  
WHERE ${pred} ${predlist}*  
      
```

projection:

```
sum(l_extendedprice * l_discount) as revenue  
      
```

table:

```
lineitem  
      
```

pred:

```
l_shipdate >= date '1994-01-01'  
l_shipdate < date '1994-01-01' + interval '1' year  
l_discount between 0.06 - 0.01 and 0.06 + 0.01  
l_quantity < 24  
      
```

predlist:

```
AND ${pred}  
      
```

tag	templates	space	tag	templates	space
Q1	40	9207	Q12	8484	162918
Q2	58160	6354837405	Q13	16	81
Q3	240	29295	Q14	6	21
Q4	28	81	Q15	40	372
Q5	108	96579	Q16	608	25515
Q6	4	15	Q17	26	81
Q7	>100K	–	Q18	576	43659
Q8	480	5478165	Q19	>100K	–
Q9	1512	3528441	Q20	320	3339.0
Q10	384	722925	Q21	18464	4255065
Q11	162	7203	Q22	156	777

**Figure 6: TPC-H query space**

Maintain a query pool

Pick promising candidates

Keep a workflow database

Support analysis

The logo for SQL scalpel, featuring the word "SQL" in large gray letters, followed by "scalpel" in blue italicized letters, with a horizontal gray line extending from the end of "scalpel".

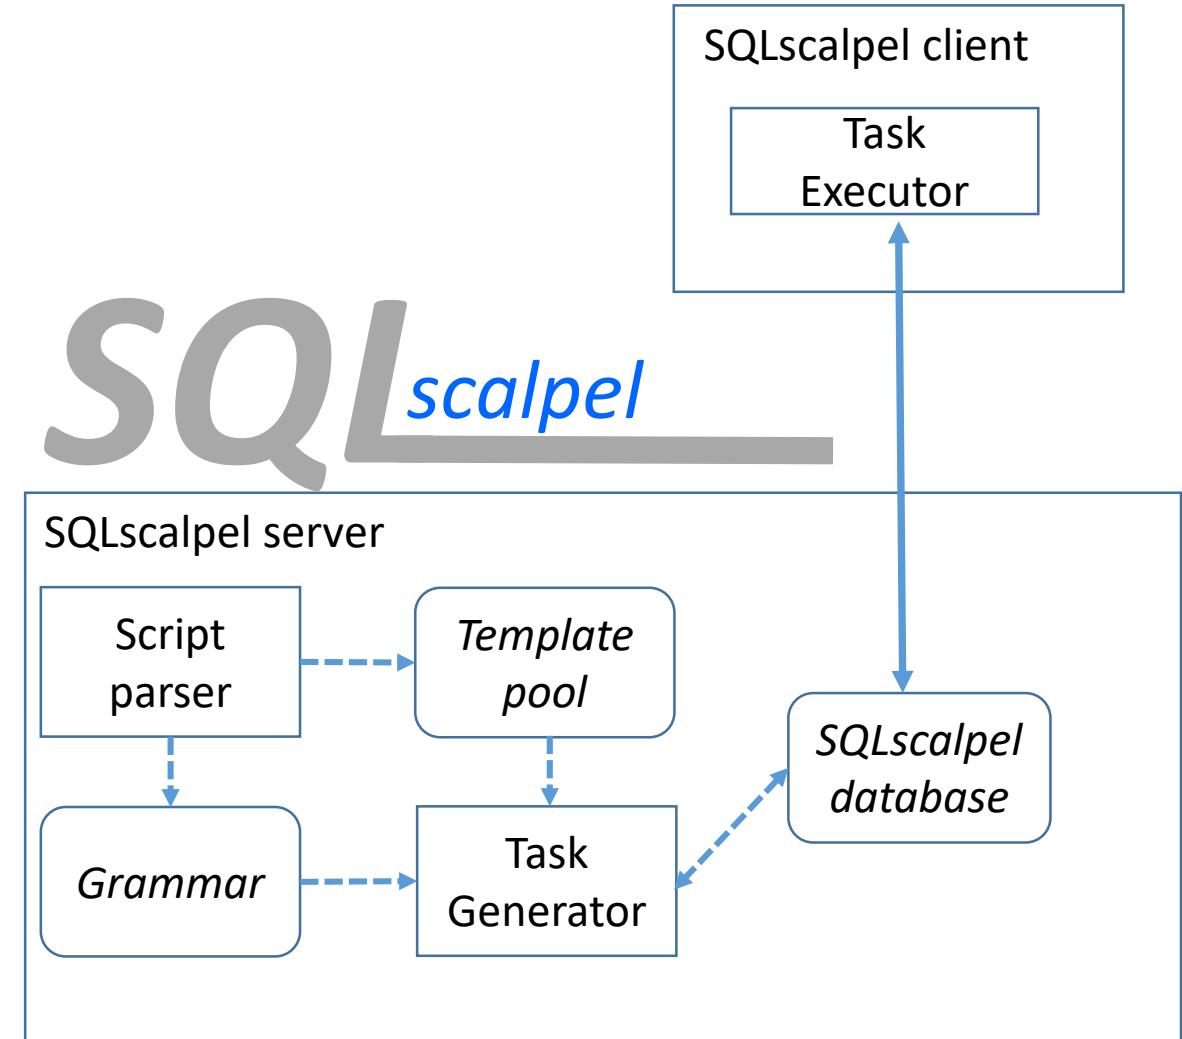
# SQLscalpel architecture

Maintain a query pool

Pick promising candidates

Keep a workflow database

Support analysis



Maintain a query pool



Pick promising candidates

Keep a workflow database

Support analysis

SQLscalpel architecture

### Strategies:

- Baseline
- Random
- Alter a lexical term
- Expand a query
- Prune a query
- Ditch manually

## SQLscalpel architecture

Maintain a query pool

Pick promising candidates →

Keep a workflow database

Support analysis

Strategies:

- FCFS
- Manual steering
- Simulated annealing
- Biased terms

SCALscalpel X +

localhost:5000/scalpel/TPC-H/q01

Projects Products Platforms

SQL scalpel TPC-H q01 Story Edit Scalpel Config Queries Queue Results History Terms Scatter Comment

## The Scalpel grammar

A Scalpel grammar is a concise description of a collection of test cases. It is described using a grammar composed of rules identified by an identifier following by a colon. Each rule is followed by a series of alternative text snippets to construct the test case. Each snippet in the grammar is only used once in a single test case. The grammar rules can be embedded as references \${name} or \${[name]} in the snippets.

[Modify the scalpel grammar.](#) [Manage the dialect translation table.](#) [Show dialect translation table.](#)

```
query:
    select ${l_select_term} ${select_expr}* ${l_from} ${l_where} group by ${l_group_by_term} ${group_by_expr}* order by ${l_order_by_term}
group_by_expr:
    , ${l_group_by_term}
select_expr:
    , ${l_select_term}
order_by_expr:
    , ${l_order_by_term}
l_from:
    from lineitem
l_where:
    where l_shipdate <= date '1998-12-01' - interval '90' day ( 3 )
l_group_by_term:
    l_returnflag
    l_linenumber
l_select_term:
    l_returnflag
    l_linenumber
    sum ( l_quantity ) as sum_qty
    sum ( l_extendedprice ) as sum_base_price
    sum ( l_extendedprice * ( 1 - l_discount ) ) as sum_disc_price
    sum ( l_extendedprice * ( 1 - l_discount ) * ( 1 + l_tax ) ) as sum_charge
    avg ( l_quantity ) as avg_qty
    avg ( l_extendedprice ) as avg_price
    avg ( l_discount ) as avg_disc
    count ( * ) as count_order
l_order_by_term:
    l_returnflag
    l_linenumber
```

Firefox prevented this page from automatically reloading.



Projects Products Platforms

TPC-H ▾  
q01 ▾

Story

Config

Scalpel

Queries

Queue

Results

History

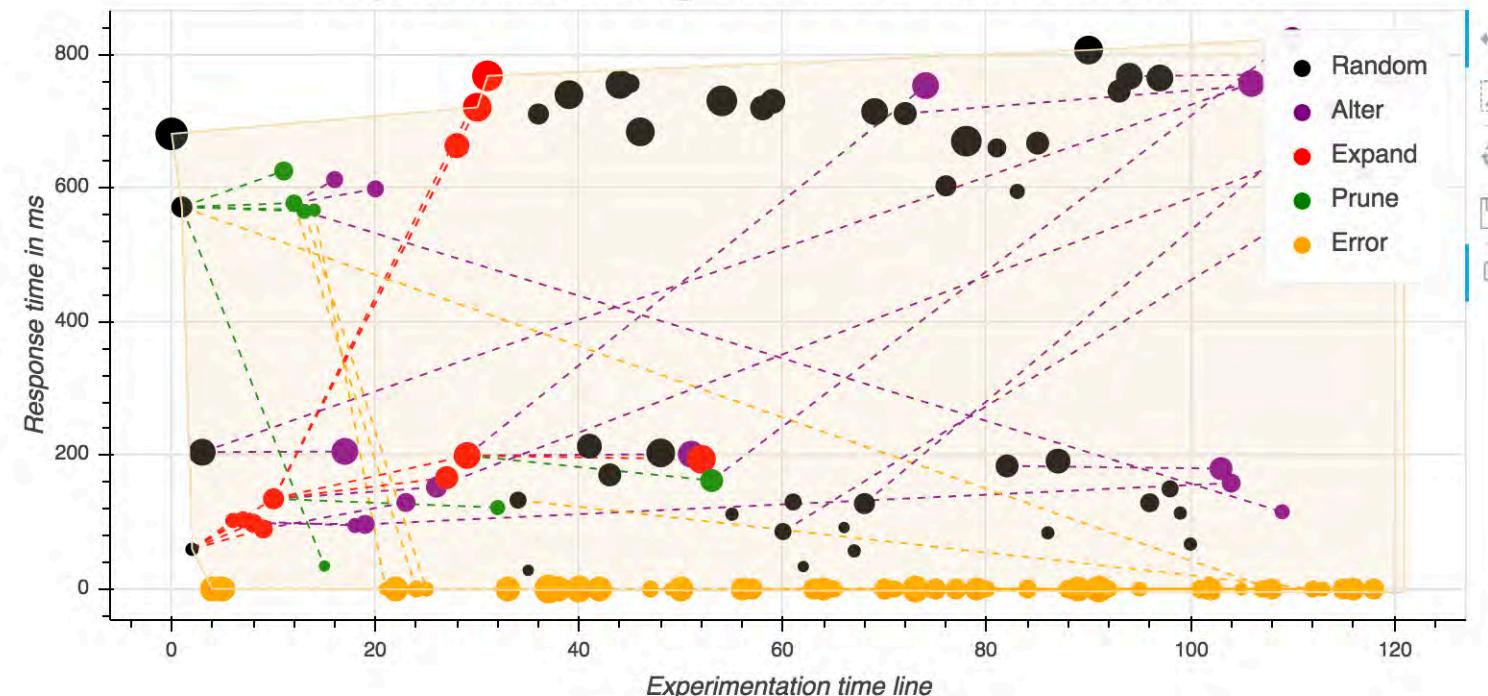
Terms

## Scalpel experiment lineage

The exploration steps through the query space are visualized as a map. It simply shows the performance of the queries in milliseconds and the lineage relationships between them.

[Highlight results for just a few terms.](#)

Response time lineage sf1/MonetDB\_Mar18/localhost



TPC-H ▾  
q01 ▾

Story

Config

Scalpel

Queries

Queue

Results

History

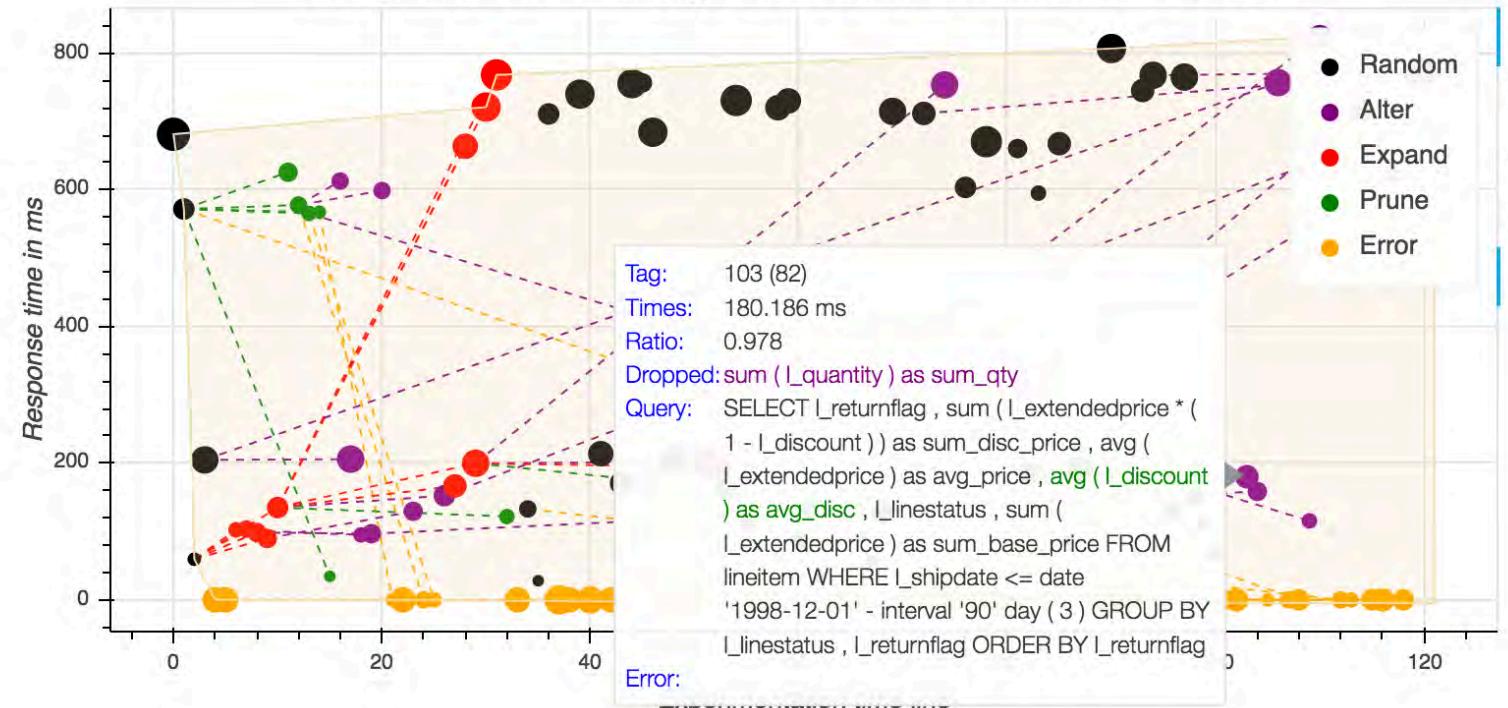
Terms

## Scalpel experiment lineage

The exploration steps through the query space are visualized as a map. It simply shows the performance of the queries in milliseconds and the

Highlight results for just a few terms.

Response time lineage sf1/MonetDB\_Mar18/localhost



[TPC-H ▾](#)  
[q01 ▾](#) Story Config Scalpel Queries Queue Results Dashboard

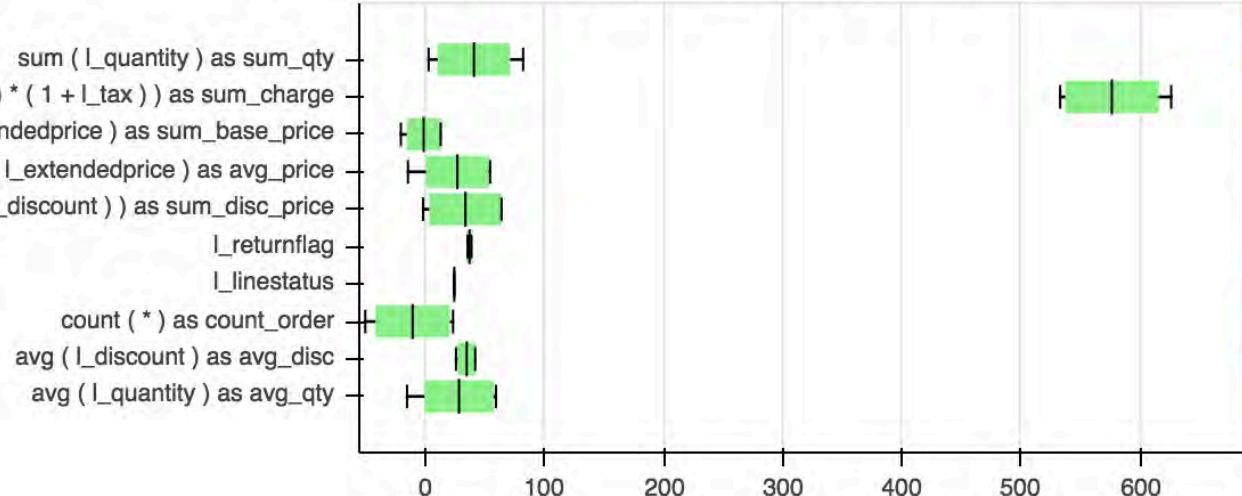
## Scalpel terms analysis

The lexical terms impact on a query is visualized below. It is calculated by taking two queries which differ only in a single term, then sum up the results and calculate the mean and stddev, and also show the outliers.

The measured performance can drop below 0 due to caching effects of tables over a sequence of similar queries.

[Pre-filter the result table](#)

Term impact for sf1/MonetDB\_Mar18/localhost



SCALscalpel +

localhost:5000/scatter

SQL *scalpel* Projects Products Platforms

TPC-H ▾  
q01 ▾

Story

Edit

Scalpel

Config

Queries

Queue

Results

History

Terms

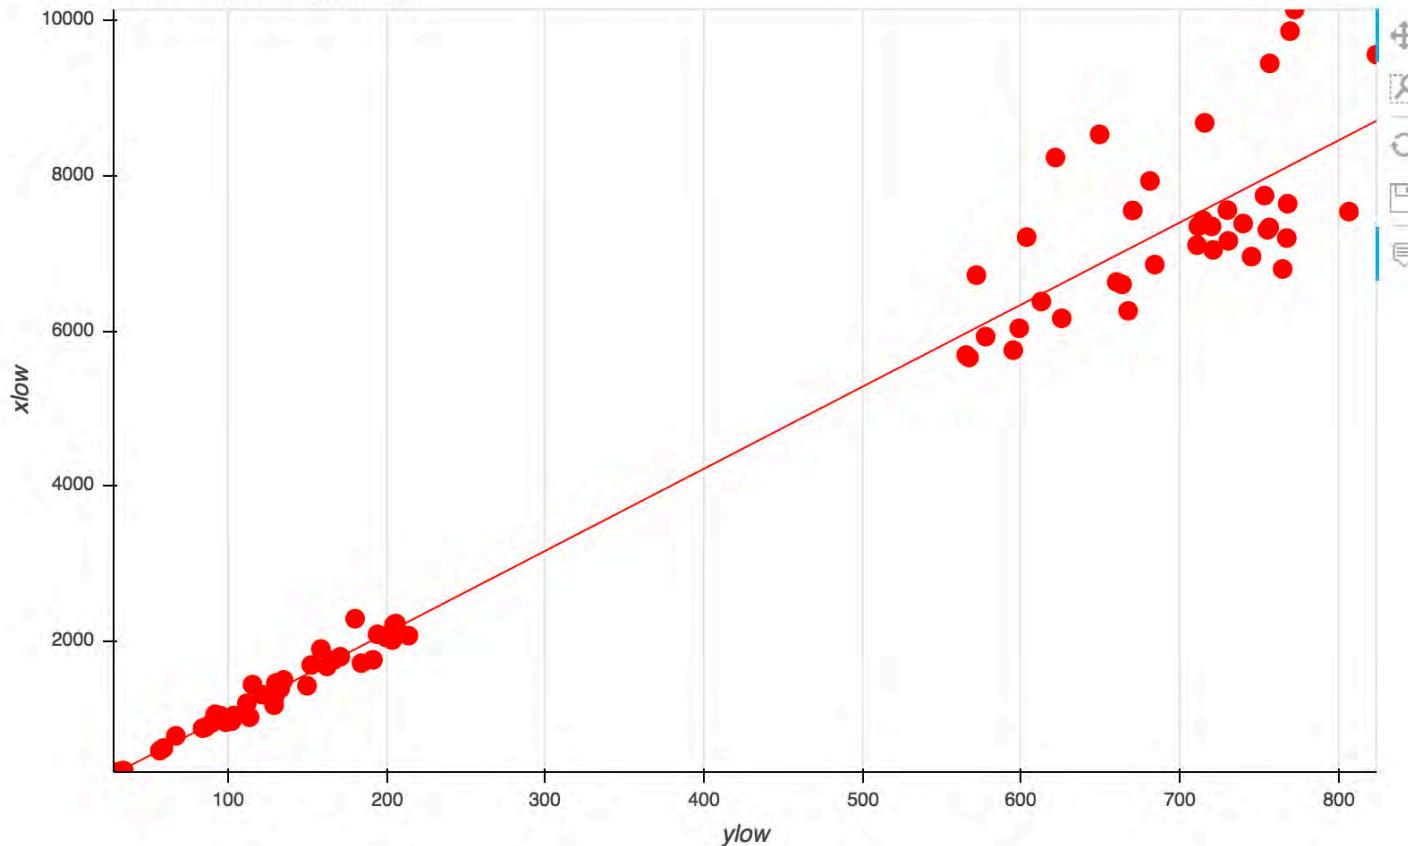
Scatter

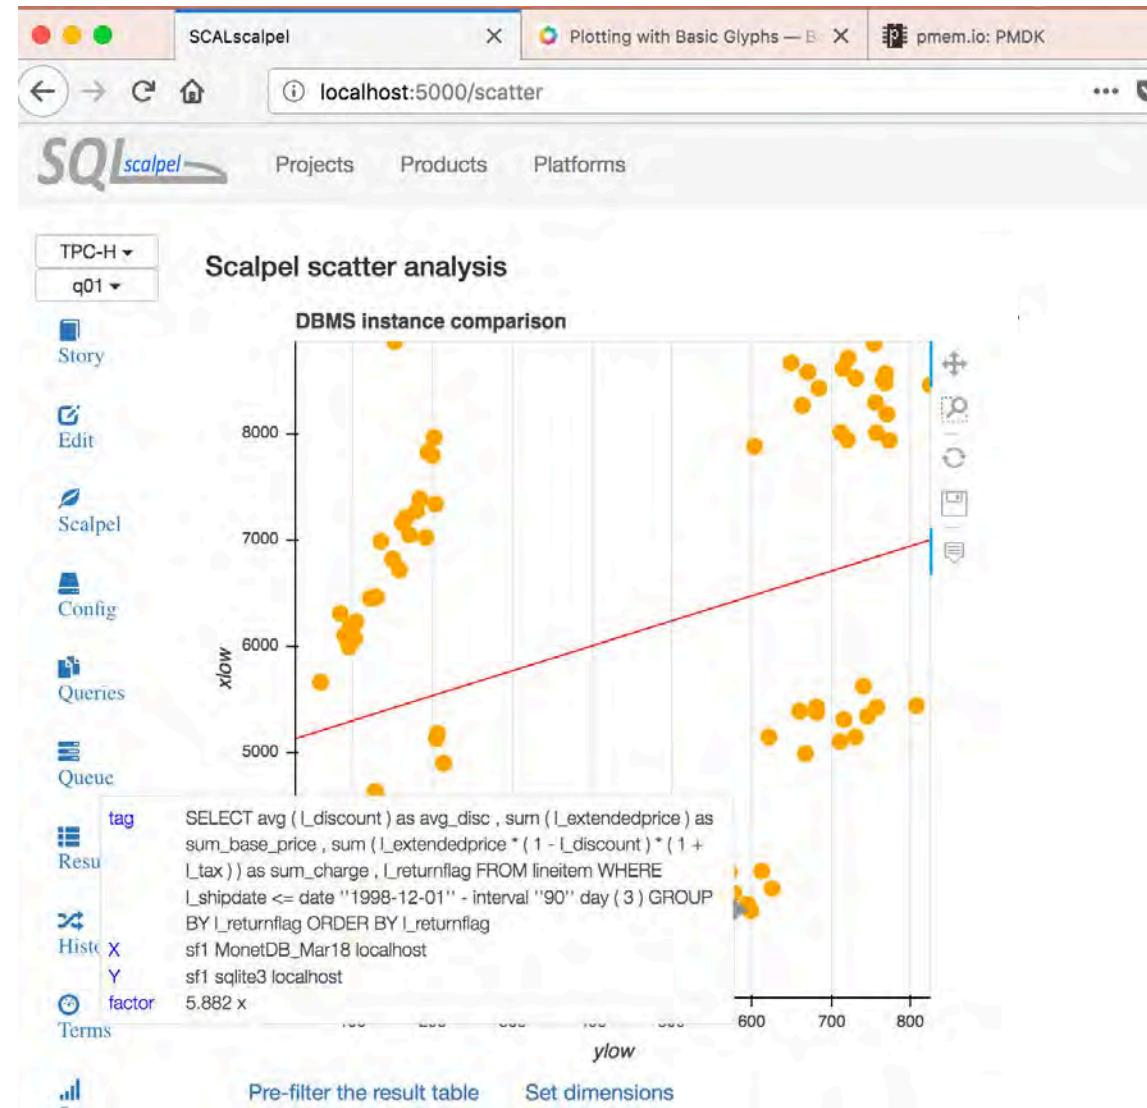
## Scalpel scatter analysis

Compare all queries executed against two database/dbms/platform combinations.

Pre-filter the result table Set dimensions

DBMS instance comparison





Handle SQL dialects

Easily generate  
erroneous queries

Brings the target down  
using Cartesian results

tpch\_q6:

```
SELECT ${projection} FROM ${table} WHERE ${pred} ${predlist}*  
    
```

projection:

```
sum(l_extendedprice * l_discount) as revenue  
    
```

table:

lineitem

lineitem, regions

pred:

```
l_shipdate >= date '1994-01-01'  
    
```

```
l_shipdate < date '1994-01-01' + interval '1' year  
    
```

```
l_discount between 0.06 - 0.01 and 0.06 + 0.01  
    
```

```
l_quantity < 24  
    
```

predlist:

```
AND ${pred}  
    
```

- SQLscalpel prototype is up and running
  - Full-stack infrastructure
  - Drivers for MonetDB, MySQL, SQLite, Postgres
- Functional enhancements
  - Multi-partitioned tables
  - Built-in forums
- Future projects
  - The story

Questions



Q1: SELECT count(\*) FROM nation WHERE nation.n\_name='BRAZIL'

Q2: SELECT count(\*) FROM nation WHERE nation.n\_name='BRAZIL' AND nation.n\_regionkey=1

$$\frac{\frac{T_A(Q_2)}{T_B(Q_2)}}{\frac{T_A(Q_1)}{T_B(Q_1)}} \rightarrow \frac{T_A(Q_2)}{T_B(Q_2)} \frac{T_B(Q_1)}{T_A(Q_1)}$$

Q1: SELECT count(\*) FROM nation WHERE nation.n\_name='BRAZIL'

Q2: SELECT count(\*) FROM nation WHERE nation.n\_name='BRAZIL' AND nation.n\_regionkey=1

$$\frac{T_A(Q_2)}{T_B(Q_2)} \quad \frac{T_B(Q_1)}{T_A(Q_1)}$$

