IN2090 – Databaser og datamodellering

13 – Spørreprosessering

Leif Harald Karlsen
leifhka@ifi.uio.no

Universitetet i Oslo

Spørreprosessering: Oversikt

```
SELECT p.name, c.name
  FROM products p INNER JOIN
       categories c USING (cid);
```

$$\pi_{p.name,c.name}\left(\rho_{p}\left(products\right) \bowtie_{p.cid=c.cid} \rho_{c}\left(categories\right)\right)$$

```
                                 QUERY PLAN
-------------------------------------------------------------------------------
Hash Join  (cost=1.18..3.26 rows=77 width=65)
  Hash Cond: (p.category_id = c.category_id)
  ->  Seq Scan on products p  (cost=0.00..1.77 rows=77 width=19)
  ->  Hash  (cost=1.08..1.08 rows=8 width=50)
        ->  Seq Scan on categories c  (cost=0.00..1.08 rows=8 width=50)
(5 rows)
```

```
      name      |      name
----------------+----------------
 Chai           | Beverages
 Chang          | Beverages
 Aniseed Syrup  | Condiments
 Chef Anton's Cajun Seasoning | Condiments
 Chef Anton's Gumbo Mix | Condiments
 Grandma's Boysenberry Spread | Condiments
 Uncle Bob's Organic Dried Pears | Produce
```

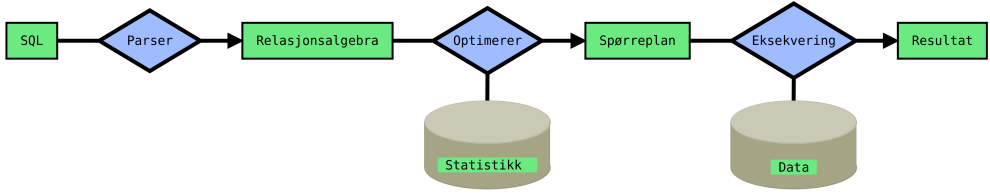

- En spørring går gjennom flere steg før den til slutt blir evaluert over dataene
- Disse stegene sørger for at spørringen blir besvart så effektivt som mulig

Fra SQL til relasjonell algebra

```
SELECT p.name, c.name
  FROM products p INNER JOIN
       categories c USING (cid);
```

$$\pi_{p.name,c.name}\left(\rho_{p}\left(products\right) \bowtie_{p.cid=c.cid} \rho_{c}\left(categories\right)\right)$$

```
                                QUERY PLAN
-------------------------------------------------------------------------
 Hash Join  (cost=1.18..3.26 rows=77 width=65)
   Hash Cond: (p.category_id = c.category_id)
   ->  Seq Scan on products p  (cost=0.00..1.77 rows=77 width=19)
   ->  Hash  (cost=1.08..1.08 rows=8 width=50)
         ->  Seq Scan on categories c  (cost=0.00..1.08 rows=8 width=50)
(5 rows)
```

```
          name          |    name
------------------------+------------
 Chai                   | Beverages
 Chang                  | Beverages
 Aniseed Syrup          | Condiments
 Chef Anton's Cajun Seasoning | Condiments
 Chef Anton's Gumbo Mix | Condiments
 Grandma's Boysenberry Spread | Condiments
 Uncle Bob's Organic Dried Pears | Produce
```

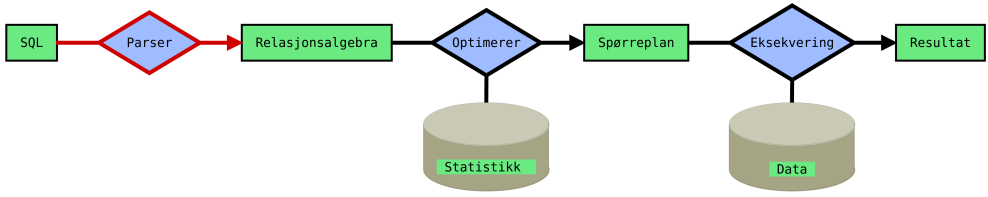

- Det første som skjer er at spørringen sjekkes syntaktisk (f.eks. tabellene og kolonnene finnes i databasen, typene er riktige, osv.)
- Deretter oversettes spørringen til et spørre-tre over relasjons algebraen

Oversettelse: Eksempel

```
SELECT p.product_name, d.unit_price
FROM categories AS c JOIN
  products AS p USING (category_id) JOIN
  order_details AS d USING (product_id)
WHERE c.category_name = 'Beverages'
  AND d.discount >= 0.25;
```

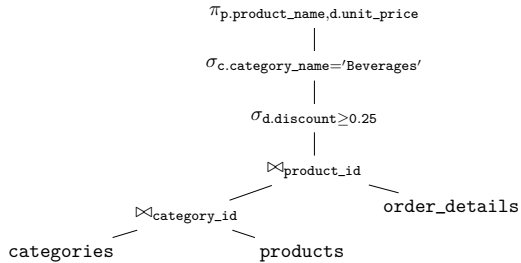

$$\pi_{\texttt{p.product\_name,d.unit\_price}}\Big(\sigma_{\texttt{c.category\_name}='\texttt{Beverages}'}\Big(\sigma_{\texttt{d.discount}\geq 0.25}\big(\texttt{categories} \bowtie_{\texttt{category\_id}} \texttt{products} \bowtie_{\texttt{product\_id}} \texttt{order\_details}\big)\Big)\Big)$$

Ulike spørringer – Likt resultat

```
SELECT p.name, c.name
  FROM products p INNER JOIN
       categories c USING (cid);
```

$$\pi_{p.name,c.name}\left(\rho_{p}\left(products\right) \bowtie_{p.cid=c.cid} \rho_{c}\left(categories\right)\right)$$

```
                                  QUERY PLAN
------------------------------------------------------------------------------
Hash Join  (cost=1.18..3.26 rows=77 width=65)
  Hash Cond: (p.category_id = c.category_id)
  ->  Seq Scan on products p  (cost=0.00..1.77 rows=77 width=19)
  ->  Hash  (cost=1.08..1.08 rows=8 width=50)
        ->  Seq Scan on categories c  (cost=0.00..1.08 rows=8 width=50)
(5 rows)
```

```
          name          |     name
------------------------+--------------
 Chai                   | Beverages
 Chang                  | Beverages
 Aniseed Syrup          | Condiments
 Chef Anton's Cajun Seasoning | Condiments
 Chef Anton's Gumbo Mix | Condiments
 Grandma's Boysenberry Spread | Condiments
 Uncle Bob's Organic Dried Pears | Produce
```

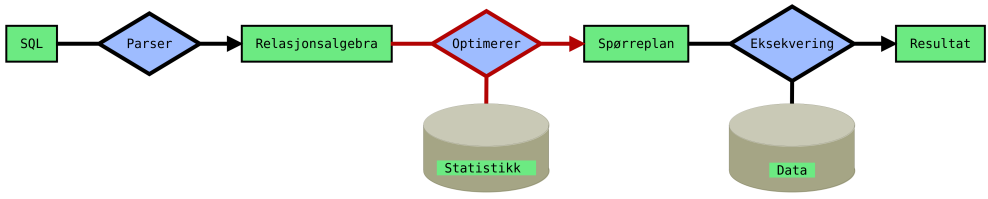

- Spørringen uttrykt i relasjonell algebra kan manipuleres algebraisk
- Dette brukes for å genere forskjellige men ekvivalente spørringer
- Altså, spørringer som gir samme svar, men ser forskjellige ut
- Forskjellige spørringer kan ha ulik kompleksitet
- De ulike spørringene skal så (i neste steg) bli tilordnet en ca. kostnad
- Vi vil så velge den spørringen som er billigst å eksekvere

Ulike spørringer: Eksempel

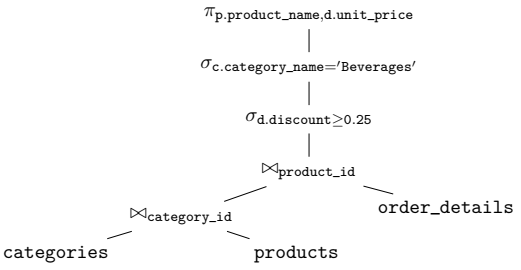

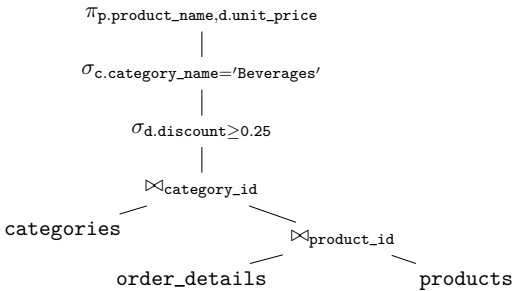

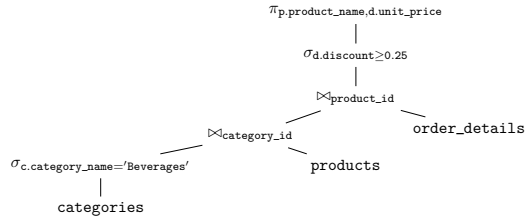

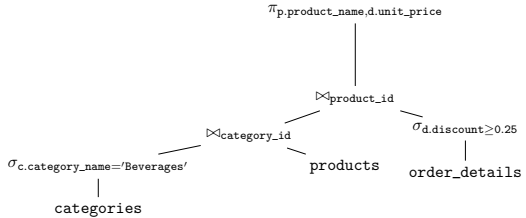

Fra spørring til kostnad

```
SELECT p.name, c.name
  FROM products p INNER JOIN
       categories c USING (cid);
```

$$\pi_{p.name,c.name}\left(\rho_{p}\left(products\right) \bowtie_{p.cid=c.cid} \rho_{c}\left(categories\right)\right)$$

```
                                  QUERY PLAN
-------------------------------------------------------------------------------
Hash Join  (cost=1.18..3.26 rows=77 width=65)
  Hash Cond: (p.category_id = c.category_id)
  ->  Seq Scan on products p  (cost=0.00..1.77 rows=77 width=19)
  ->  Hash  (cost=1.08..1.08 rows=8 width=50)
        ->  Seq Scan on categories c  (cost=0.00..1.08 rows=8 width=50)
(5 rows)
```

```
          name          |     name
------------------------+--------------
 Chai                   | Beverages
 Chang                  | Beverages
 Aniseed Syrup          | Condiments
 Chef Anton's Cajun Seasoning | Condiments
 Chef Anton's Gumbo Mix | Condiments
 Grandma's Boysenberry Spread | Condiments
 Uncle Bob's Organic Dried Pears | Produce
```

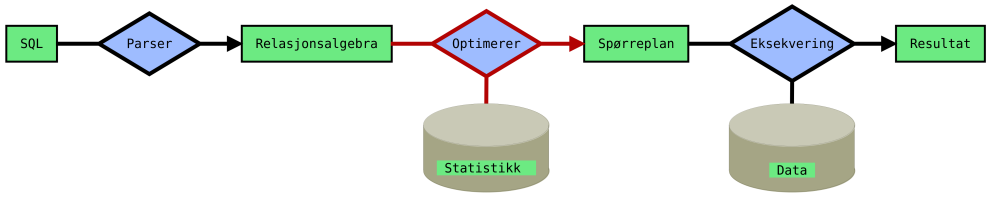

- De ulike spørringene blir så tilordnet en kostnad
- Kostnadsevalueringen bruker statistikk over databasen
- F.eks. antall rader i hver tabell, antall ulike verdier i hver kolonne, osv.
- Bruker her også skranker (f.eks. `UNIQUE`, `CHECK`) og indeksstrukturer
- Høyere kostnad betyr lengre eksekveringstid
- Databasen velger så den spørringen med lavest kostnad

Spørreplaner

```
SELECT p.name, c.name
  FROM products p INNER JOIN
       categories c USING (cid);
```

$$\pi_{p.name,c.name}\left(\rho_{p}\left(products\right) \bowtie_{p.cid=c.cid} \rho_{c}\left(categories\right)\right)$$

```
                                QUERY PLAN
-------------------------------------------------------------------------
 Hash Join  (cost=1.18..3.26 rows=77 width=65)
   Hash Cond: (p.category_id = c.category_id)
   ->  Seq Scan on products p  (cost=0.00..1.77 rows=77 width=19)
   ->  Hash  (cost=1.08..1.08 rows=8 width=50)
         ->  Seq Scan on categories c  (cost=0.00..1.08 rows=8 width=50)
(5 rows)
```

```
          name          |    name
------------------------+------------
 Chai                   | Beverages
 Chang                  | Beverages
 Aniseed Syrup          | Condiments
 Chef Anton's Cajun Seasoning | Condiments
 Chef Anton's Gumbo Mix | Condiments
 Grandma's Boysenberry Spread | Condiments
 Uncle Bob's Organic Dried Pears | Produce
```

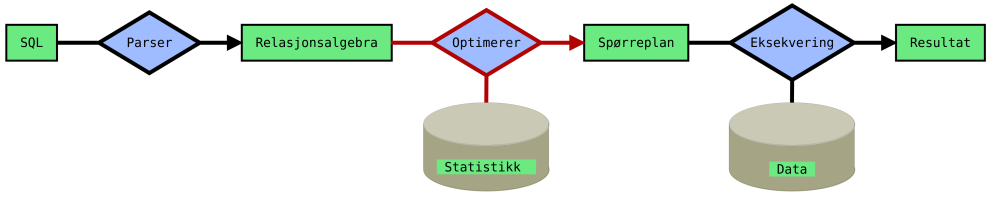

- Det siste som skjer i dette trinnet er at det blir laget en spørreplan for den valgte spørringen
- Dette er en mer detaljert plan for hvordan spørringen skal eksekveres

EXPLAIN

- Av og til kan det være nyttig å få se denne spørreplanen
- F.eks. dersom man lurer på hvordan spørringen vil bli eksekvert
- Eller dersom man ønsker et ca. estimat på hvor komplisert spørringen blir å eksekvere
- Dette kan gjøres ved å skrive `EXPLAIN` foran spørringen
- Spørringen blir da ikke eksekvert
- (Merk: Ikke pensum å kunne forstå spørreplaner!)

EXPLAIN: Eksempel

```
psql=> EXPLAIN SELECT p.product_name, d.unit_price
FROM categories AS c JOIN
  products AS p USING (category_id) JOIN
  order_details AS d USING (product_id)
WHERE c.category_name = 'Beverages'
  AND d.discount >= 0.25;
                                    QUERY PLAN
-------------------------------------------------------------------------------------
 Hash Join  (cost=3.32..43.04 rows=20 width=21)
   Hash Cond: (d.product_id = p.product_id)
   ->  Seq Scan on order_details d  (cost=0.00..38.94 rows=154 width=6)
         Filter: (discount >= '0.25'::double precision)
   ->  Hash  (cost=3.20..3.20 rows=10 width=19)
         ->  Hash Join  (cost=1.11..3.20 rows=10 width=19)
               Hash Cond: (p.category_id = c.category_id)
               ->  Seq Scan on products p  (cost=0.00..1.77 rows=77 width=21)
               ->  Hash  (cost=1.10..1.10 rows=1 width=2)
                     ->  Seq Scan on categories c  (cost=0.00..1.10 rows=1 width=2)
                           Filter: ((category_name)::text = 'Beverages'::text)
(11 rows)
```

Evaluering

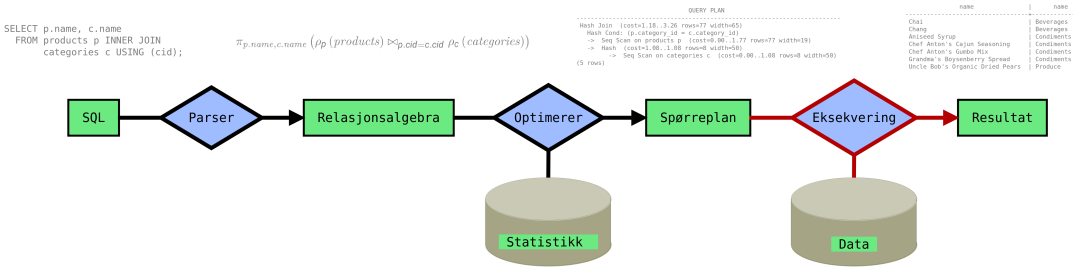

- Til slutt evalueres spørringen over databasen
- Databasen har så svært effektive algoritmer for joins, oppslag, sorteing, osv.
- Merk: Databasen trenger kun én algoritme per operator i den (utvidede) relasjonelle algebraen

ANALYZE

- Dersom vi ønsker å vite hvor lang tid en spørring faktisk tar å eksekvere, samt detaljert analyse av hver del av spørreplanen kan vi bruke `EXPLAIN ANALYZE`
- Får da også informasjon om minnebruk
- Da vil spørringen bli eksekvert, og databasen samler så nøyaktig informasjon om eksekveringen
- Dersom en spørring tar lang tid kan dette bruker for å finne ut hvilken del av spørringen som er komplisert
- Kan også brukes for å finne manglende indeksstrukturer

ANALYZE: Eksempel

```
psql=> EXPLAIN ANALYZE SELECT p.product_name, d.unit_price
FROM categories AS c JOIN
  products AS p USING (category_id) JOIN
  order_details AS d USING (product_id)
WHERE c.category_name = 'Beverages'
  AND d.discount >= 0.25;
                                                              QUERY PLAN
---------------------------------------------------------------------------------------------------------------------------------------
 Hash Join  (cost=3.32..43.04 rows=20 width=21) (actual time=0.130..1.066 rows=32 loops=1)
   Hash Cond: (d.product_id = p.product_id)
   ->  Seq Scan on order_details d  (cost=0.00..38.94 rows=154 width=6) (actual time=0.031..0.887 rows=154 loops=1)
         Filter: (discount >= '0.25'::double precision)
         Rows Removed by Filter: 2001
   ->  Hash  (cost=3.20..3.20 rows=10 width=19) (actual time=0.085..0.085 rows=12 loops=1)
         Buckets: 1024  Batches: 1  Memory Usage: 9kB
         ->  Hash Join  (cost=1.11..3.20 rows=10 width=19) (actual time=0.034..0.077 rows=12 loops=1)
               Hash Cond: (p.category_id = c.category_id)
               ->  Seq Scan on products p  (cost=0.00..1.77 rows=77 width=21) (actual time=0.008..0.022 rows=77 loops=1)
               ->  Hash  (cost=1.10..1.10 rows=1 width=2) (actual time=0.016..0.016 rows=1 loops=1)
                     Buckets: 1024  Batches: 1  Memory Usage: 9kB
                     ->  Seq Scan on categories c  (cost=0.00..1.10 rows=1 width=2) (actual time=0.008..0.012 rows=1 loops=1)
                           Filter: ((category_name)::text = 'Beverages'::text)
                           Rows Removed by Filter: 7
 Planning Time: 0.567 ms
 Execution Time: 1.146 ms
(17 rows)
```

Takk for nå!

Lykke til med forberedelsene til eksamen!