

Factorized databases II

Jakub Zavodny (University of Oxford, UK)

DAMOL

DATA ANALYSIS AND MODELING LAB

Palacky University, Olomouc, Czech Republic



europa
european
social fund in the
czech republic



EUROPEAN UNION



MINISTRY OF EDUCATION,
YOUTH AND SPORTS



OP Education
for Competitiveness

INVESTMENTS IN EDUCATION DEVELOPMENT

Factorised Databases 2.

Factorised Database Engine

Jakub Závodný, University of Oxford

based on joint work with PhD supervisor Dan Olteanu

Factorised Relations

$$\begin{aligned} & (\langle \text{Tesco} \rangle \times (\langle \text{food} \rangle \cup \langle \text{clothes} \rangle) \cup \langle \text{Lidl} \rangle \times \langle \text{food} \rangle) \times (\langle \text{CzechRep} \rangle \cup \langle \text{Slovakia} \rangle) \cup \\ & (\langle \text{Tesco} \rangle \cup \langle \text{Lidl} \rangle \cup \langle \text{Aldi} \rangle) \times \langle \text{food} \rangle \times (\langle \text{Poland} \rangle \cup \langle \text{UK} \rangle) \cup \\ & (\langle \text{Tesco} \rangle \cup \langle \text{Aldi} \rangle) \times \langle \text{books} \rangle \times (\langle \text{UK} \rangle \cup \langle \text{Spain} \rangle) \end{aligned}$$

Representations of relations that

- compactly encode combinations of groups of values.
- allow for enumerating tuples with constant delay.
- allow query evaluation *without expanding* to flat relations
 - ▶ Using factorisations can boost query evaluation (processing a smaller representation).

Key Results

1. How compact can factorised query results be?

- Asymptotic size bounds for factorisations of query results.
- Parameter $s(Q)$ characterising the factorisability of results of Q .

2. Can such factorisations speed up query evaluation?

- Compute factorisations directly from query and input data.
- Optimisation and evaluation techniques for queries with selections, projections, joins, aggregates and ordering on factorisations.
- FDB: a main-memory engine using factorised relations.


Key Challenges

2. Can such factorisations speed up query evaluation?

Factorisation Trees

Uniform nesting structures \longleftrightarrow defined by *factorisation trees* (f-trees).

$$\bigcup_{\text{Chains}} \left(\langle \text{chain} \rangle \times \left(\bigcup_{\text{Products}} \langle \text{product} \rangle \right) \times \left(\bigcup_{\text{Countries}} \langle \text{country} \rangle \right) \right) \longleftrightarrow$$



$$\bigcup_{\text{Chains}} \left(\langle \text{chain} \rangle \times \left(\bigcup_{\text{Products}} \langle \text{product} \rangle \times \left(\bigcup_{\text{Countries}} \langle \text{country} \rangle \right) \right) \right) \longleftrightarrow$$



Factorisation Trees

F-tree **defines** the factorisation of a relation.



+

A	B	C
x	1	2
x	1	3
x	2	2
x	2	3
y	4	4

⇓

$$\langle x \rangle \times (\langle 1 \rangle \cup \langle 2 \rangle) \times (\langle 2 \rangle \cup \langle 3 \rangle) \\ \cup \langle y \rangle \times \langle 4 \rangle \quad \times \langle 4 \rangle$$

Example: Evaluating a Selection on a Factorised Relation

Evaluate $\sigma_{B=C}R(A, B, C)$:

A	B	C
x	1	2
x	1	3
x	2	2
x	2	3
y	4	4

\mapsto

A	B = C
x	2
y	4

$$\langle x \rangle \times (\langle 1 \rangle \cup \langle 2 \rangle) \times (\langle 2 \rangle \cup \langle 3 \rangle) \cup \\ \langle y \rangle \times \langle 4 \rangle \quad \times \langle 4 \rangle$$

\mapsto

$$\langle x \rangle \times \langle 2 \rangle \cup \\ \langle y \rangle \times \langle 4 \rangle$$



\mapsto



Example: Evaluating a Selection on a Factorised Relation

To evaluate $\sigma_{B=C}$:

- **Merge nodes** B, C in the tree




- **Merge-join corresponding unions** over B and C in linear time.

$$\begin{array}{l} \langle x \rangle \times (\langle 1 \rangle \cup \langle 2 \rangle) \times (\langle 2 \rangle \cup \langle 3 \rangle) \cup \\ \langle y \rangle \times \langle 4 \rangle \quad \quad \quad \times \langle 4 \rangle \end{array} \quad \mapsto \quad \begin{array}{l} \langle x \rangle \times \langle 2 \rangle \cup \\ \langle y \rangle \times \langle 4 \rangle \end{array}$$

Query Evaluation: Merge Selection

In general, we can evaluate $\sigma_{B=C}$ if B and C are **siblings**:

- **Merge nodes** B, C in the tree



- **Merge-join corresponding unions** over B and C in linear time.

$$\dots(\langle b_1 \rangle \times B_1 \cup \langle b_2 \rangle \times B_2 \cup \dots \cup \langle b_n \rangle \times B_n) \times \\ (\langle c_1 \rangle \times C_1 \cup \langle c_2 \rangle \times C_2 \cup \dots \cup \langle c_m \rangle \times C_m) \dots$$

\mapsto


$$\dots(\langle b_{i_1} = c_{j_1} \rangle \times B_{i_1} \times C_{j_1} \cup \dots \cup \langle b_{i_k} = c_{j_k} \rangle \times B_{i_k} \times C_{j_k}) \dots$$

Call this the **merge operator**.

Query Evaluation: Absorb Selection

Also, we can evaluate $\sigma_{A=B}$ if A is **ancestor** of B :

- **Absorb node B** into A in the tree




- **Filter unions** over B in linear time.

$$\begin{aligned} & \dots \langle a \rangle \times \left(\bigcup \dots (\langle b_1 \rangle \times B_1 \cup \langle b_2 \rangle \times B_2 \cup \dots \cup \langle b_n \rangle \times B_n) \dots \right) \\ \mapsto & \dots \langle a \rangle \times \left(\bigcup \dots (\langle b_i = a \rangle \times B_i) \dots \right) \quad \mapsto \quad \langle a \rangle \times \left(\bigcup \dots B_i \dots \right) \end{aligned}$$

Call this the **absorb operator**.

Query Evaluation: Selection


How to join $A = B$ when A and B are not siblings or descendants?



Query Evaluation: Restructuring Operators

How to join $A = B$ when A and B are not siblings / descendants?

Restructure the factorisation.



Restructuring operators

- **Swap with parent** (re-group terms)
- **Push up** (factor out common expressions)

in quasilinear time in $(|\text{input}| + |\text{output}|)$.

Query Evaluation

Operators on factorisations **defined by** operators on f-trees:

Equality selection operators

- **Merge** siblings
- **Absorb** descendant

Restructuring operators

- **Swap** with parent
- **Push up**

+

- **Project** away a leaf attribute
- **Cartesian product** = simple tree concatenation, constant time!

Any SPJ query can be evaluated by a sequence of these operators
(called an **f-plan** for the query).

Query Evaluation: Example

PlaysFor		CompetesIn		LeagueStadium	
player	team	team	league	league	stadium
Messi	Barcelona	Barcelona	Primera	Primera	CampNou
Villa	Barcelona	Barcelona	Champions	Champions	CampNou
Cech	Chelsea	Chelsea	Premier	Champions	Wembley
Torres	Chelsea	Chelsea	Champions	Premier	Stamford
van Persie	Arsenal	Arsenal	Premier	Premier	Wembley

BasedIn		IsIn	
team	city	stadium	city
Barcelona	Barcelona	CampNou	Barcelona
Chelsea	London	Wembley	London
Arsenal	London	Stamford	London

Which players can I see on which stadiums?

Which stadiums and which teams are in the same city?


Query Evaluation: Example

Which players can I see on which stadiums?

Which stadiums and which teams are in the same city?

PlaysFor \bowtie_{team} CompetesIn \bowtie_{league} LeagueStadium

BasedIn \bowtie_{city} IsIn



$$\begin{aligned} & (\langle \text{FCBarcelona} \rangle \times (\langle \text{Messi} \rangle \cup \langle \text{Villa} \rangle)) \times \\ & \quad \times (\langle \text{Primera} \rangle \times \langle \text{CampNou} \rangle \cup \\ & \quad \quad \langle \text{Champions} \rangle \times (\langle \text{CampNou} \rangle \cup \langle \text{Wembley} \rangle)) \cup \\ & \langle \text{Chelsea} \rangle \times (\langle \text{Cech} \rangle \cup \langle \text{Torres} \rangle) \times \\ & \quad \times (\langle \text{Premier} \rangle \times (\langle \text{Stamford} \rangle \cup \langle \text{Wembley} \rangle)) \cup \\ & \quad \quad \langle \text{Champions} \rangle \times (\langle \text{CampNou} \rangle \cup \langle \text{Wembley} \rangle) \cup \\ & \langle \text{Arsenal} \rangle \times \langle \text{van Persie} \rangle \times \\ & \quad \times (\langle \text{Premier} \rangle \times (\langle \text{Stamford} \rangle \cup \langle \text{Wembley} \rangle)) \end{aligned}$$
$$\begin{aligned} & (\langle \text{Barcelona} \rangle \times \langle \text{FCBarcelona} \rangle) \\ & \quad \times \langle \text{CampNou} \rangle \cup \\ & \langle \text{London} \rangle \times (\langle \text{Chelsea} \rangle \cup \langle \text{Arsenal} \rangle) \\ & \quad \times (\langle \text{Wembley} \rangle \cup \langle \text{Stamford} \rangle) \end{aligned}$$

Which players *from the same city* can I see on which stadiums?

Query Evaluation: Example



Which players *from the same city* can I see on which stadiums?

$(\text{PlaysFor} \bowtie_{\text{team}} \text{CompetesIn} \bowtie_{\text{league}} \text{LeagueStadium}) \bowtie_{\text{team, stadium}} (\text{BasedIn} \bowtie_{\text{city}} \text{IsIn})$


$$\begin{aligned} & (\langle \text{FCBarcelona} \rangle \times (\langle \text{Messi} \rangle \cup \langle \text{Villa} \rangle)) \times \\ & \quad \times (\langle \text{Primera} \rangle \times \langle \text{CampNou} \rangle \cup \\ & \quad \quad \langle \text{Champions} \rangle \times (\langle \text{CampNou} \rangle \cup \langle \text{Wembley} \rangle)) \cup \\ \langle \text{Chelsea} \rangle \times (\langle \text{Cech} \rangle \cup \langle \text{Torres} \rangle) \times \\ & \quad \times (\langle \text{Premier} \rangle \times (\langle \text{Stamford} \rangle \cup \langle \text{Wembley} \rangle)) \cup \\ & \quad \quad \langle \text{Champions} \rangle \times (\langle \text{CampNou} \rangle \cup \langle \text{Wembley} \rangle)) \cup \\ \langle \text{Arsenal} \rangle \times \langle \text{van Persie} \rangle \times \\ & \quad \times (\langle \text{Premier} \rangle \times (\langle \text{Stamford} \rangle \cup \langle \text{Wembley} \rangle)) \end{aligned} \times \left((\langle \text{Barcelona} \rangle \times \langle \text{FCBarcelona} \rangle \right. \\ \left. \times \langle \text{CampNou} \rangle \cup \right. \\ \left. \langle \text{London} \rangle \times (\langle \text{Chelsea} \rangle \cup \langle \text{Arsenal} \rangle) \right. \\ \left. \times (\langle \text{Wembley} \rangle \cup \langle \text{Stamford} \rangle) \right)$$

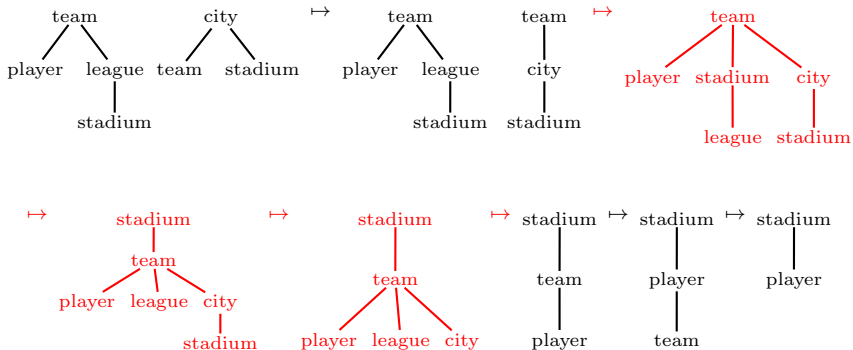
Now perform selections $\text{team} = \text{team}$ and $\text{stadium} = \text{stadium}$.

Query Evaluation: Example F-plan



$\langle \text{CampNou} \rangle \times (\langle \text{Messi} \rangle \cup \langle \text{Villa} \rangle) \cup$
 $\langle \text{Wembley} \rangle \times (\langle \text{Cech} \rangle \cup \langle \text{Torres} \rangle \cup \langle \text{van Persie} \rangle) \cup$
 $\langle \text{Stamford} \rangle \times (\langle \text{Cech} \rangle \cup \langle \text{Torres} \rangle \cup \langle \text{van Persie} \rangle)$

Query Evaluation: Different Example F-plan



$\langle CampNou \rangle \times (\langle Messi \rangle \cup \langle Villa \rangle) \cup$
 $\langle Wembley \rangle \times (\langle Cech \rangle \cup \langle Torres \rangle \cup \langle van Persie \rangle) \cup$
 $\langle Stamford \rangle \times (\langle Cech \rangle \cup \langle Torres \rangle \cup \langle van Persie \rangle)$

Query Optimisation: Challenges

How to choose a good f-plan for a given query?

New challenges:

- New degrees of freedom:
Restructuring between joins.
- New (conflicting) objectives:
Fast execution vs. compact representation of the result?
- How to estimate factorisation size in advance?

Query Optimisation: Approaches

Exhaustive search:

- Explore the entire space of allowed operators, find cheapest path.

Greedy search:

- Always choose the cheapest query operator.
- Always choose the cheapest restructuring to enable query operators.

Metrics for factorisation size and evaluation time:

- Asymptotic size measure $s(\mathcal{T})$.
- Estimates from statistical catalogue information about the data.

FDB

We implement all these ideas in

FDB = Prototype DB engine using factorisations at storage level.
(In-memory, written in C++)

Experimentally compared with

RDB = Lightweight relational engine.
(In-memory, written in C++)

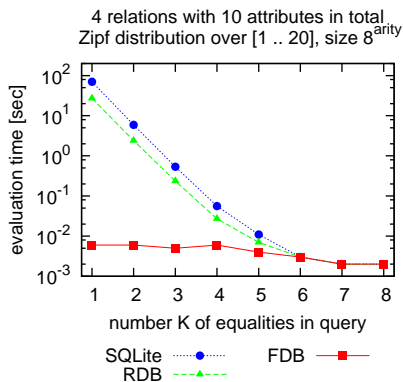
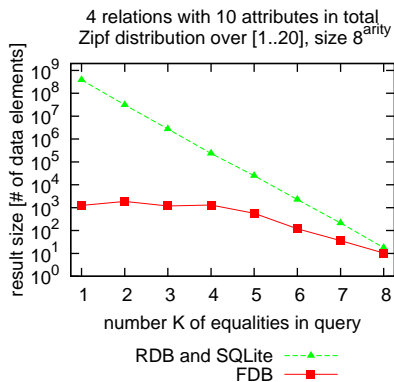
SQLite

PostgreSQL

Experimental Evaluation

FDB vs. RDB vs. SQLite


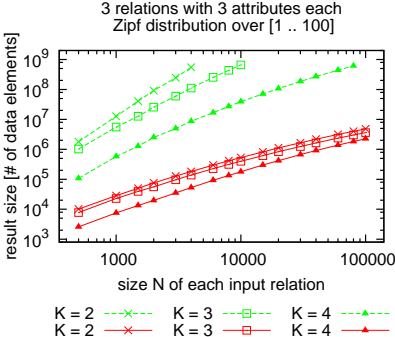
Queries with K equalities. Input is flat relations.



Experimental Evaluation

FDB vs. RDB vs. SQLite

Scaling behaviour with increasing data size. Input is flat relations.



Recent Development: Aggregation, Ordering

“Find the total price of items sold.”

```
SELECT SUM(price) FROM Sales;
```

“Find the average price of items sold to each customer.”

```
SELECT customer,AVG(price) FROM Sales  
GROUP BY customer;
```

“Find the customers with top 10 sales.”

```
SELECT customer,SUM(price) AS spending FROM Sales  
GROUP BY customer  
ORDER BY spending  
LIMIT 10;
```

Aggregation on Factorisations

Aggregation can also be done in linear time, using recursion:

The number of tuples in a factorisation F , $\text{COUNT}(F)$, is defined as:

- $\text{COUNT}(\langle a \rangle) = 1$.
- $\text{COUNT}(F_1 \cup \dots \cup F_k) = \text{COUNT}(F_1) + \dots + \text{COUNT}(F_k)$.
- $\text{COUNT}(F_1 \times \dots \times F_k) = \text{COUNT}(F_1) \cdot \dots \cdot \text{COUNT}(F_k)$.

Sum of all values of attribute A in a factorisation F , $\text{SUM}_A(F)$, is defined as:

- $\text{SUM}_A(\langle a \rangle) = a$,
if the value $\langle a \rangle$ is of attribute A .
- $\text{SUM}_A(F_1 \cup \dots \cup F_k) = \text{SUM}_A(F_1) + \dots + \text{SUM}_A(F_k)$.
- $\text{SUM}_A(F_1 \times \dots \times F_k) = \text{SUM}_A(F_1) \cdot \text{COUNT}(F_2) \cdot \dots \cdot \text{COUNT}(F_k)$,
where wlog values of attribute A are in expression F_1 .

Aggregation: Example

Orders			Pizzas		Items	
customer	date	pizza	pizza	item	item	price
Mario	Monday	Capricciosa	Margherita	base	base	6
Mario	Tuesday	Margherita	Capricciosa	base	ham	1
Pietro	Friday	Hawaii	Capricciosa	ham	mushrooms	1
Lucia	Friday	Hawaii	Capricciosa	mushrooms	pineapple	2
Mario	Friday	Capricciosa	Hawaii	base		
			Hawaii	ham		
			Hawaii	pineapple		


Itemized sales: $\text{Sales} = \text{Orders} \bowtie \text{Pizzas} \bowtie \text{Items}$.

Revenue per customer: $\varpi_{\text{customer}, \text{sum}_{\text{price}}}(\text{Sales})$.

```
SELECT customer, SUM(price) FROM Sales GROUP BY customer;
```

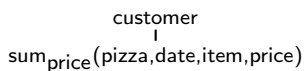

Aggregation: Example

$Sales = Orders \bowtie Pizzas \bowtie Items$


$$\begin{aligned} & \langle Capricciosa \rangle \times (\langle Monday \rangle \times \langle Mario \rangle \cup \langle Friday \rangle \times \langle Mario \rangle) \\ & \quad \times (\langle base \rangle \times \langle 6 \rangle \cup \langle ham \rangle \times \langle 1 \rangle \cup \langle mushrooms \rangle \times \langle 1 \rangle) \\ \cup & \langle Hawaii \rangle \times \langle Friday \rangle \times (\langle Lucia \rangle \cup \langle Pietro \rangle) \\ & \quad \times (\langle base \rangle \times \langle 6 \rangle \cup \langle ham \rangle \times \langle 1 \rangle \cup \langle pineapple \rangle \times \langle 2 \rangle) \\ \cup & \langle Margherita \rangle \times \langle Tuesday \rangle \times \langle Mario \rangle \times \langle base \rangle \times \langle 6 \rangle \end{aligned}$$

Aggregation: Example


$\varpi_{customer, \text{sum}_{price}}(\text{Sales})$



$\langle \text{Capricciosa} \rangle \times (\langle \text{Monday} \rangle \times \langle \text{Mario} \rangle \cup \langle \text{Friday} \rangle \times \langle \text{Mario} \rangle)$
 $\times (\langle \text{base} \rangle \times \langle 6 \rangle \cup \langle \text{ham} \rangle \times \langle 1 \rangle \cup \langle \text{mushrooms} \rangle \times \langle 1 \rangle)$
 $\cup \langle \text{Hawaii} \rangle \times \langle \text{Friday} \rangle \times (\langle \text{Lucia} \rangle \cup \langle \text{Pietro} \rangle)$
 $\times (\langle \text{base} \rangle \times \langle 6 \rangle \cup \langle \text{ham} \rangle \times \langle 1 \rangle \cup \langle \text{pineapple} \rangle \times \langle 2 \rangle)$
 $\cup \langle \text{Margherita} \rangle \times \langle \text{Tuesday} \rangle \times \langle \text{Mario} \rangle \times \langle \text{base} \rangle \times \langle 6 \rangle$

Aggregation: Example


$\varpi_{customer, \text{sum}_{price}}(\text{Sales})$



$\langle \text{Capricciosa} \rangle \times (\langle \text{Monday} \rangle \times \langle \text{Mario} \rangle \cup \langle \text{Friday} \rangle \times \langle \text{Mario} \rangle)$
 $\times (\langle \text{base} \rangle \times \langle 6 \rangle \cup \langle \text{ham} \rangle \times \langle 1 \rangle \cup \langle \text{mushrooms} \rangle \times \langle 1 \rangle)$
 $\cup \langle \text{Hawaii} \rangle \times \langle \text{Friday} \rangle \times (\langle \text{Lucia} \rangle \cup \langle \text{Pietro} \rangle)$
 $\times (\langle \text{base} \rangle \times \langle 6 \rangle \cup \langle \text{ham} \rangle \times \langle 1 \rangle \cup \langle \text{pineapple} \rangle \times \langle 2 \rangle)$
 $\cup \langle \text{Margherita} \rangle \times \langle \text{Tuesday} \rangle \times \langle \text{Mario} \rangle \times \langle \text{base} \rangle \times \langle 6 \rangle$

Aggregation: Example


$\varpi_{customer, \text{sum}_{price}}(\text{Sales})$



$\langle \text{Capricciosa} \rangle \times (\langle \text{Monday} \rangle \times \langle \text{Mario} \rangle \cup \langle \text{Friday} \rangle \times \langle \text{Mario} \rangle)$
 $\times (\langle \text{base} \rangle \times \langle 6 \rangle \cup \langle \text{ham} \rangle \times \langle 1 \rangle \cup \langle \text{mushrooms} \rangle \times \langle 1 \rangle)$
 $\cup \langle \text{Hawaii} \rangle \times \langle \text{Friday} \rangle \times (\langle \text{Lucia} \rangle \cup \langle \text{Pietro} \rangle)$
 $\times (\langle \text{base} \rangle \times \langle 6 \rangle \cup \langle \text{ham} \rangle \times \langle 1 \rangle \cup \langle \text{pineapple} \rangle \times \langle 2 \rangle)$
 $\cup \langle \text{Margherita} \rangle \times \langle \text{Tuesday} \rangle \times \langle \text{Mario} \rangle \times \langle \text{base} \rangle \times \langle 6 \rangle$

Aggregation: Example


$\varpi_{customer, \text{sum}_{price}}(\text{Sales})$



$\langle \text{Capricciosa} \rangle \times (\langle \text{Monday} \rangle \times \langle \text{Mario} \rangle \cup \langle \text{Friday} \rangle \times \langle \text{Mario} \rangle) \times \langle 8 \rangle$
 $\cup \langle \text{Hawaii} \rangle \times \langle \text{Friday} \rangle \times (\langle \text{Lucia} \rangle \cup \langle \text{Pietro} \rangle) \times \langle 9 \rangle$
 $\cup \langle \text{Margherita} \rangle \times \langle \text{Tuesday} \rangle \times \langle \text{Mario} \rangle \times \langle 6 \rangle$

Aggregation: Example


$\varpi_{customer, sum_{price}}(\text{Sales})$



$\langle Lucia \rangle \times \langle Hawaii \rangle \times \langle Friday \rangle \times \langle 9 \rangle$
 $\cup \langle Mario \rangle \times (\langle Capricciosa \rangle \times (\langle Monday \rangle \cup \langle Friday \rangle)) \times \langle 8 \rangle \cup$
 $\quad \langle Margherita \rangle \times \langle Tuesday \rangle \times \langle 6 \rangle$
 $\cup \langle Pietro \rangle \times \langle Hawaii \rangle \times \langle Friday \rangle \times \langle 9 \rangle$

Aggregation: Example

$\varpi_{customer, \text{sum}_{price}}(\text{Sales})$




$\langle Lucia \rangle \times \langle 9 \rangle \cup \langle Mario \rangle \times \langle 22 \rangle \cup \langle Pietro \rangle \times \langle 9 \rangle$

Experimental Evaluation

FDB vs. SQLite vs. PostgreSQL

Aggregation queries on $Sales = Orders \bowtie Items \bowtie Packages$.



$Q_1 = \varpi_{\text{package, date, customer; sum(price)}(Sales)$

$Q_2 = \varpi_{\text{customer; revenue} \leftarrow \text{sum(price)}(Sales)$

$Q_3 = \varpi_{\text{date, package; sum(price)}(Sales)$


$Q_4 = \varpi_{\text{package; sum(price)}(Sales)$

$Q_5 = \varpi_{\text{sum(price)}(Sales)$

Experimental Evaluation

FDB vs. SQLite vs. PostgreSQL

Aggregation queries on $Sales = Orders \bowtie Items \bowtie Packages$.
Performance with increasing data size.




$Q_2 = \varpi_{\text{customer}; \text{revenue} \leftarrow \text{sum}(\text{price})}(\text{Sales})$
 $Q_3 = \varpi_{\text{date, package}; \text{sum}(\text{price})}(\text{Sales})$

Experimental Evaluation

FDB vs. SQLite vs. PostgreSQL

Aggregation queries on Orders, Items and Packages.



$Q_1 = \varpi_{\text{package, date, customer; sum(price)}(Sales)$

$Q_2 = \varpi_{\text{customer; revenue} \leftarrow \text{sum(price)}(Sales)$

$Q_3 = \varpi_{\text{date, package; sum(price)}(Sales)$

$Q_4 = \varpi_{\text{package; sum(price)}(Sales)$

$Q_5 = \varpi_{\text{sum(price)}(Sales)$

Experimental Evaluation

Factorisations **can boost query performance** for many-to-many data.

FDB outperforms relational engines **SQLite** and **PostgreSQL**:

- if input is factorised.
- if output can be factorised.
- for flat to flat aggregation.
(Unless they are hand-optimised to push aggregates past joins.)
- for flat to flat queries with `ORDER BY` and `LIMIT`.

Factorised Relations

$$\begin{aligned} & (\langle \text{Tesco} \rangle \times (\langle \text{food} \rangle \cup \langle \text{clothes} \rangle) \cup \langle \text{Lidl} \rangle \times \langle \text{food} \rangle) \times ((\langle \text{CzechRep} \rangle \cup \langle \text{Slovakia} \rangle) \cup \\ & (\langle \text{Tesco} \rangle \cup \langle \text{Lidl} \rangle \cup \langle \text{Aldi} \rangle) \times \langle \text{food} \rangle \times ((\langle \text{Poland} \rangle \cup \langle \text{UK} \rangle) \cup \\ & (\langle \text{Tesco} \rangle \cup \langle \text{Aldi} \rangle) \times \langle \text{books} \rangle \times ((\langle \text{UK} \rangle \cup \langle \text{Spain} \rangle)) \end{aligned}$$

Representations of relations that

- compactly encode combinations of groups of values.
- allow for enumerating tuples with constant delay.
- allow query evaluation *without expanding* to flat relations
 - ▶ Using factorisations can boost query evaluation (processing a smaller representation).

Thank you!