DPconv: Super-Polynomially Faster Join Ordering

Mihail Stoian, Andreas Kipf

Data Systems Lab @UTN

@SIGMOD'25, June 25, 2025



Join Ordering Recap

SQL Query

```
SELECT *

FROM R_1, R_2, R_3, R_4

WHERE R_1.a = R_2.b

AND R_2.c = R_3.d

AND R_3.e = R_4.f
```

Join Ordering Recap

SQL Query \implies Query Graph

SELECT * FROM R_1 , R_2 , R_3 , R_4 WHERE R_1 .a = R_2 .b AND R_2 .c = R_3 .d AND R_3 .e = R_4 .f



Join Ordering Recap



SELECT * FROM R_1 , R_2 , R_3 , R_4 WHERE R_1 .a = R_2 .b AND R_2 .c = R_3 .d AND R_3 .e = R_4 .f



SIGMOD'79 ⇒ DPsize

Access Path Selection in a Relational Database Management System

P. Griffiths Selinger M. M. Astrahan D. D. Chamberlin R. A. Lorie T. G. Price

IBM Research Division, San Jose, California 95193

ABSTRACT: In a high level query and data manipulation language such as SQL, requests are stated non-procedurally, without reference to access paths. This paper describes how System R chooses access paths for both simple (single relation) and complex queries (such as joins), given a user specification of desired data as a boolean expression of predicates. System R is an experimental database management system developed to carry out research on the relational model of data. System R was designed and built by members of the IBM San Jose Research Laboratory. access path for each table in the SQL statement. Of the many possible choices, the optimizer chooses the one which minimizes "total access cost" for performing the entire statement.

This paper will address the issues of access path selection for queries. Retrieval for data manipulation (UPDATE, DELETE) is treated similarly. Section 2 will describe the place of the optimizer in the processing of a SQL statement, and section 3 will describe the storage component access paths that are available on a single physically stored table. In section 4 the















Cost Functions



Cost Functions



Cost Functions





$\textbf{Exact}~\mathbf{C}_{\max}$

 $O(3^{n})$

















DPconv



Approximate C_{out} $O(3^{-1})$ super-polynomial $O^*(2^n \log W / \epsilon)$

Beyond Theory



- Bellman's optimality principle.
- Given set $S \subseteq \{1, ..., n\}$:
 - \circ c(S) = Join cardinality of S.
 - DP[S] = Optimal cost to join relations in S.
- \Rightarrow Solution: DP[{1, ..., n}].

- Bellman's optimality principle.
- Given set $S \subseteq \{1, ..., n\}$:
 - \circ c(S) = Join cardinality of S.
 - DP[S] = Optimal cost to join relations in S.
- \Rightarrow Solution: DP[{1, ..., n}].



- Bellman's optimality principle.
- Given set $S \subseteq \{1, ..., n\}$:
 - \circ c(S) = Join cardinality of S.
 - DP[S] = Optimal cost to join relations in S.
- \Rightarrow Solution: DP[{1, ..., n}].
- Recursion for C_{out}:

DP[*S*] =



$(\mathrm{DP}[T] + \mathrm{DP}[S \setminus T])$

- Bellman's optimality principle.
- Given set $S \subseteq \{1, ..., n\}$:
 - \circ c(S) = Join cardinality of S.
 - DP[S] = Optimal cost to join relations in S.
- \Rightarrow Solution: DP[{1, ..., n}].
- Recursion for C_{out}:



$DP[S] = \min_{T \subseteq S} (DP[T] + DP[S \setminus T])$

- Bellman's optimality principle.
- Given set $S \subseteq \{1, ..., n\}$:
 - \circ c(S) = Join cardinality of S.
 - DP[S] = Optimal cost to join relations in S.
- \Rightarrow Solution: DP[{1, ..., n}].
- Recursion for C_{out}:



 $DP[S] = c(S) + \min_{T \subseteq S} (DP[T] + DP[S \setminus T])$

- Bellman's optimality principle.
- Given set $S \subseteq \{1, ..., n\}$:
 - \circ c(S) = Join cardinality of S.
 - DP[S] = Optimal cost to join relations in S.
- \Rightarrow Solution: DP[{1, ..., n}].
- Recursion for C_{out}:



$\mathrm{DP}[S] = c(S) + \min_{T \subseteq S} \left(\mathrm{DP}[T] + \mathrm{DP}[S \setminus T] \right)$

• Running time:

$$\sum_{S\subseteq [n]} 2^{|S|}$$

- Bellman's optimality principle.
- Given set $S \subseteq \{1, ..., n\}$:
 - \circ c(S) = Join cardinality of S.
 - DP[S] = Optimal cost to join relations in S.
- \Rightarrow Solution: DP[{1, ..., n}].
- Recursion for C_{out}:



$\mathrm{DP}[S] = c(S) + \min_{T \subseteq S} (\mathrm{DP}[T] + \mathrm{DP}[S \setminus T])$

• Running time:

$$\sum_{S\subseteq [n]} 2^{|S|} = \sum_{k=0}^{n} \binom{n}{k} 2^{k}$$

- Bellman's optimality principle.
- Given set $S \subseteq \{1, ..., n\}$:
 - \circ c(S) = Join cardinality of S.
 - DP[S] = Optimal cost to join relations in S.
- \Rightarrow Solution: DP[{1, ..., n}].
- Recursion for C_{out}:



 $\mathrm{DP}[S] = c(S) + \min_{T \subseteq S} \left(\mathrm{DP}[T] + \mathrm{DP}[S \setminus T] \right)$

• Running time:

$$\sum_{S\subseteq [n]} 2^{|S|} = \sum_{k=0}^{n} \binom{n}{k} 2^{k} = (1+2)^{n} = 3^{n}$$

FOURIER MEETS MÖBIUS: FAST SUBSET CONVOLUTION

ANDREAS BJÖRKLUND, THORE HUSFELDT, PETTERI KASKI, AND MIKKO KOIVISTO

ABSTRACT. We present a fast algorithm for the subset convolution problem: given functions f and g defined on the lattice of subsets of an *n*-element set N, compute their subset convolution f * g, defined for all $S \subseteq N$ by

$$(f * g)(S) = \sum_{T \subseteq S} f(T)g(S \setminus T),$$

FOURIER MEETS MÖBIUS: FAST SUBSET CONVOLUTION

ANDREAS BJÖRKLUND, THORE HUSFELDT, PETTERI KASKI, AND MIKKO KOIVISTO

ABSTRACT. We present a fast algorithm for the subset convolution problem: given functions f and g defined on the lattice of subsets of an *n*-element set N, compute their subset convolution f * g, defined for all $S \subseteq N$ by

$$(f * g)(S) = \sum_{T \subseteq S} f(T)g(S \setminus T),$$

FOURIER MEETS MÖBIUS: FAST SUBSET CONVOLUTION

ANDREAS BJÖRKLUND, THORE HUSFELDT, PETTERI KASKI, AND MIKKO KOIVISTO

ABSTRACT. We present a fast algorithm for the subset convolution problem: given functions f and g defined on the lattice of subsets of an *n*-element set N, compute their subset convolution f * g, defined for all $S \subseteq N$ by

$$(f * g)(S) = \sum_{T \subseteq S} f(T)g(S \setminus T),$$

(+, ×)

FOURIER MEETS MÖBIUS: FAST SUBSET CONVOLUTION

ANDREAS BJÖRKLUND, THORE HUSFELDT, PETTERI KASKI, AND MIKKO KOIVISTO

ABSTRACT. We present a fast algorithm for the subset convolution problem: given functions f and g defined on the lattice of subsets of an *n*-element set N, compute their subset convolution f * g, defined for all $S \subseteq N$ by

$$(f * g)(S) = \sum_{T \subseteq S} f(T)g(S \setminus T),$$
(+, ×)

FOURIER MEETS MÖBIUS: FAST SUBSET CONVOLUTION

ANDREAS BJÖRKLUND, THORE HUSFELDT, PETTERI KASKI, AND MIKKO KOIVISTO

ABSTRACT. We present a fast algorithm for the subset convolution problem: given functions f and g defined on the lattice of subsets of an *n*-element set N, compute their subset convolution f * g, defined for all $S \subseteq N$ by

$$(f * g)(S) = \sum_{T \subseteq S} f(T)g(S \setminus T),$$
(+, ×)

FOURIER MEETS MÖBIUS: FAST SUBSET CONVOLUTION

ANDREAS BJÖRKLUND, THORE HUSFELDT, PETTERI KASKI, AND MIKKO KOIVISTO

ABSTRACT. We present a fast algorithm for the subset convolution problem: given functions f and g defined on the lattice of subsets of an *n*-element set N, compute their subset convolution f * g, defined for all $S \subseteq N$ by

$$(f * g)(S) = \sum_{T \subseteq S} f(T)g(S \setminus T),$$

where addition and multiplication is carried out in an arbitrary ring. Via Möbius transform and inversion, our algorithm evaluates the subset convolution in $O(n^22^n)$ additions and multiplications, substantially improving upon the straightforward $O(3^n)$ algorithm. Specifically, if the input func-

O Ring vs. **C** Semi-Ring: Intuition

- Ring example: $(+, \times)$.
 - \circ 2 + 3 = 5.
 - 5 **+** (-2) = 3.

O Ring vs. **C** Semi-Ring: Intuition

- Ring example: (+, ×).
 - \circ 2 + 3 = 5.
 - 5 **+** (-2) = 3.
- Semi-ring example: (**min**, +), (min, max).
 - min(2, 3) = 3.
 - How to reverse?

• Find a mapping:

$$(+, \times) \Rightarrow (\min, +)$$

• Find a mapping:

$$(+, \times) \Rightarrow (\min, +)$$

Represent a value v as x^v

• Find a mapping:

$$(+, \times) \Rightarrow (\min, +)$$

• Substitution Represent a value **v** as $x^{v} \rightarrow e.g.$, **5** is represented as x^{5} .

• Motivation: Find a mapping:

$$(+, \times) \Rightarrow (\min, +)$$

• **Provide a set of a set of**

"+" becomes product:
$$x^{a+b} = x^a x^b$$

"min" = extracting the lowest monomial

• Motivation: Find a mapping:

$$(+, \times) \Rightarrow (\min, +)$$

• **Q** Represent a value **v** as $x^{v} \rightarrow e.g.$, **5** is represented as x^{5} .

"+" becomes product:
$$x^{a+b} = x^a x^b$$

"min" = extracting the lowest monomial

- **Drawback**: Running time becomes *pseudo-polynomial*.
 - O(W) overhead, where W is the largest value.

Takeaway

• It is possible to map a semi-ring to a ring. **Drawback**:

$$O(2^n n^2) \implies O^*(2^n W).$$

• There is a case where we can avoid the O(W) factor.

Avoiding pseudo-polynomiality

• Ad-hoc dynamic programming:

 $DP[S] = \max \{ DP[T], DP[S \setminus T] \}$



• Ad-hoc dynamic programming:

DP[S] =

 $\min_{\mathcal{T}\subset \mathcal{S}} \max \{ \mathrm{DP}[\mathcal{T}], \mathrm{DP}[\mathcal{S} \setminus \mathcal{T}] \}$



• Ad-hoc dynamic programming:

$$DP[S] = \max \left\{ c(S), \min_{T \subset S} \max \{ DP[T], DP[S \setminus T] \} \right\}$$



• Ad-hoc dynamic programming:

$$\mathrm{DP}[S] = \max\left\{c(S), \min_{T \subset S} \max\left\{\mathrm{DP}[T], \mathrm{DP}[S \setminus T]\right\}\right\}$$

• *Provide the Provide the Pro*

$$\left.\begin{array}{c}\mathsf{C}(\mathsf{S})\\\mathsf{M}\\\mathsf{T}\\\mathsf{T}\\\mathsf{S}\setminus\mathsf{T}\end{array}\right.$$

• Ad-hoc dynamic programming:

$$DP[S] = \max \left\{ c(S), \min_{T \subset S} \max \{ DP[T], DP[S \setminus T] \} \right\}$$

The DP does not create new values!

⇒ We can binary search DP[{1, ..., n}].



- Fix a cardinality **U**.
- Rescale the join cardinalities:
 - $\circ \quad \leq \boldsymbol{U} \rightarrow \boldsymbol{1}.$
 - $\circ > \mathbf{U} \rightarrow 0.$

- Fix a cardinality **U**.
- Rescale the join cardinalities:
 - $\circ \quad \leq \boldsymbol{U} \rightarrow \boldsymbol{1}.$
 - $\circ \quad > {\bm U} \to 0.$



- Fix a cardinality **U**.
- Rescale the join cardinalities:
 - $\circ \quad \leq \boldsymbol{U} \rightarrow \boldsymbol{1}.$
 - $\circ \quad > {\bm U} \to 0.$



- Fix a cardinality **U**.
- Rescale the join cardinalities:
 - $\circ \quad \leq \boldsymbol{U} \to \boldsymbol{1}.$
 - $\circ \quad > \boldsymbol{U} \to \boldsymbol{0}.$
- Run the DP in the $(+, \times)$ ring:

$\mathrm{DP}[S] = \qquad \mathrm{DP}[T] \cdot \mathrm{DP}[S \setminus T]$

- Fix a cardinality **U**.
- Rescale the join cardinalities:
 - $\circ \quad \leq \boldsymbol{U} \rightarrow \boldsymbol{1}.$
 - $\circ \quad > \mathbf{U} \to \mathbf{0}.$
- Run the DP in the $(+, \times)$ ring:

$$DP[S] = \sum_{T \subset S} DP[T] \cdot DP[S \setminus T]$$

- Fix a cardinality **U**.
- Rescale the join cardinalities:
 - $\circ \quad \leq \boldsymbol{U} \rightarrow \boldsymbol{1}.$
 - $\circ > \mathbf{U} \rightarrow 0.$
- Run the DP in the $(+, \times)$ ring:

$$DP[S] = c(S) + \sum_{T \subset S} DP[T] \cdot DP[S \setminus T]$$

- Fix a cardinality **U**.
- Rescale the join cardinalities:
 - $\circ \quad \leq \boldsymbol{U} \to \boldsymbol{1}.$
 - $\circ \quad > \boldsymbol{U} \to \boldsymbol{0}.$
- Run the DP in the $(+, \times)$ ring:

$\mathrm{DP}[S] = c(S) + \sum_{T \subset S} \mathrm{DP}[T] \cdot \mathrm{DP}[S \setminus T]$

• If $DP[\{1, ..., n\}] > 0 \Rightarrow U$ is feasible.

- Fix a cardinality **U**.
- Rescale the join cardinalities:
 - $\circ \leq \mathbf{U} \rightarrow 1.$
 - $\circ \quad > \boldsymbol{U} \to \boldsymbol{0}.$
- Run the DP in the (+, ×) ring:

$\mathrm{DP}[S] = c(S) + \sum_{T \subset S} \mathrm{DP}[T] \cdot \mathrm{DP}[S \setminus T]$

- If $DP[\{1, ..., n\}] > 0 \Rightarrow U$ is feasible.
- Running time: $O(log(2^n) \cdot 2^n n^2) = O(2^n n^3)$.

The End?

• DPconv does not take into account the sparsity of the graph.

 \Rightarrow We need a sparse subset convolution.

• *Polynomial-space* join ordering?

 \Rightarrow Preliminary results for acyclic query graphs.