

DPconv:

# Super-Polynomially Faster Join Ordering

---

Mihail Stoian, Andreas Kipf

UTN

@Gray Systems Lab

January 28, 2025

# It's not just about join ordering..



Nature

<https://www.nature.com/articles> ⋮

## Quantum supremacy using a programmable ...

by F Arute · 2019 · Cited by 9186 — Therefore, in order to claim quantum supremacy we need a quantum processor that executes the program with sufficiently low error rates. Building ...

# It's not just about join ordering..



Nature  
https://ww  
Quantum s  
by F Arute · 20  
quantum proce

acy we need a  
ding ...

# It's not just about join ordering..



First super-polynomial speedups for einsum optimization.

## Problem Statement

- Fundamental problem in query optimization.
- Statement: Given a SQL query, find the order in which to join the relations.
- Goal: Minimize query time or memory usage.

# Join Ordering

## Problem Statement

- Fundamental problem in query optimization.
- Statement: Given a SQL query, find the order in which to join the relations.
- Goal: Minimize query time or memory usage.
- Cost functions – **minimize**:
  - $C_{out}$ : Sum of the intermediate sizes  $\approx$  “I want fast queries”.

## Problem Statement

- Fundamental problem in query optimization.
- Statement: Given a SQL query, find the order in which to join the relations.
- **Goal:** Minimize query time or memory usage.
- Cost functions – **minimize:**
  - $C_{out}$ : Sum of the intermediate sizes  $\approx$  “I want fast queries”.
  - $C_{max}$ : Maximum intermediate size  $\approx$  “Please avoid disk spilling”.

# Join Ordering

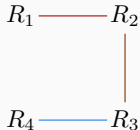
## Problem Statement

- Fundamental problem in query optimization.
- Statement: Given a SQL query, find the order in which to join the relations.
- **Goal:** Minimize query time or memory usage.
- Cost functions – **minimize:**
  - $C_{\text{out}}$ : Sum of the intermediate sizes  $\approx$  “I want fast queries”.
  - $C_{\text{max}}$ : Maximum intermediate size  $\approx$  “Please avoid disk spilling”.
- Research question since  $\sim 50$  years: *How fast can we get?*

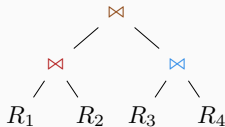


# Join Ordering: Notation

```
SELECT *  
FROM R1, R2, R3, R4  
WHERE R1.a = R2.b  
AND R2.c = R3.d  
AND R3.e = R4.f
```



Query graph



Join tree

# Join Ordering: Dynamic Programming

## Dynamic Programming

- Use Bellman's optimality principle.
- Dynamic programming table  $DP[S]$ : The optimal cost to join the relations in the set  $S$ .

---

<sup>1</sup>For  $C_{\max}$ , replace both  $+$  by  $\max$ .

# Join Ordering: Dynamic Programming

## Dynamic Programming

- Use Bellman's optimality principle.
- Dynamic programming table  $DP[S]$ : The optimal cost to join the relations in the set  $S$ .  
⇒ Solution value:  $DP[\{1, \dots, n\}]$ .
- Optimize  $DP[S]$  as

$$DP[S] = c(S) + \min_{T \subseteq S} (DP[T] + DP[S \setminus T]).^1$$

---

<sup>1</sup>For  $C_{\max}$ , replace both  $+$  by  $\max$ .

# Join Ordering: Dynamic Programming

## Dynamic Programming

- Use Bellman's optimality principle.
- Dynamic programming table  $DP[S]$ : The optimal cost to join the relations in the set  $S$ .  
⇒ Solution value:  $DP[\{1, \dots, n\}]$ .
- Optimize  $DP[S]$  as

$$DP[S] = c(S) + \min_{T \subseteq S} (DP[T] + DP[S \setminus T]).^1$$

## Running Time Analysis

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

---

<sup>1</sup>For  $C_{\max}$ , replace both  $+$  by  $\max$ .

## Exact Algorithms

- Selinger et al. (1979):  $\text{DPsizeLinear}^2 \rightarrow \Theta(2^n n)$ .

---

<sup>2</sup>Left-deep only.

## Exact Algorithms

- Selinger et al. (1979):  $\text{DPsizeLinear}^2 \rightarrow \Theta(2^n n)$ .
- Vance and Maier (1996):  $\text{DPsub} \rightarrow \Theta(3^n)$ .

---

<sup>2</sup>Left-deep only.

## Exact Algorithms

- Selinger et al. (1979):  $\text{DPsizeLinear}^2 \rightarrow \Theta(2^n n)$ .
- Vance and Maier (1996):  $\text{DPsub} \rightarrow \Theta(3^n)$ .
- Moerkotte and Neumann (2006):  $\text{DPccp} \rightarrow \Theta(\#\text{ccp})$ .

---

<sup>2</sup>Left-deep only.

## Exact Algorithms

- Selinger et al. (1979):  $\text{DPsizeLinear}^2 \rightarrow \Theta(2^n n)$ .
- Vance and Maier (1996):  $\text{DPsub} \rightarrow \Theta(3^n)$ .
- Moerkotte and Neumann (2006):  $\text{DPccp} \rightarrow \Theta(\#\text{ccp})$ .
- Haffner and Dietrich (2023):  $A^* \rightarrow O(\#\text{ccp})$ .

---

<sup>2</sup>Left-deep only.



## Exact Algorithms

- Selinger et al. (1979):  $\text{DPsizeLinear}^2 \rightarrow \Theta(2^n n)$ .
- Vance and Maier (1996):  $\text{DPsub} \rightarrow \Theta(3^n)$ .
- Moerkotte and Neumann (2006):  $\text{DPccp} \rightarrow \Theta(\#\text{ccp})$ .
- Haffner and Dietrich (2023):  $A^* \rightarrow O(\#\text{ccp})$ .

## Approximation Algorithms

- Many good heuristics, yet no theoretical guarantees.

---

<sup>2</sup>Left-deep only.

# Join Ordering: Status Quo

## Exact Algorithms

- Selinger et al. (1979):  $\text{DPsizeLinear}^2 \rightarrow \Theta(2^n n)$ .
- Vance and Maier (1996):  $\text{DPsub} \rightarrow \Theta(3^n)$ .
- Moerkotte and Neumann (2006):  $\text{DPccp} \rightarrow \Theta(\#\text{ccp})$ .
- Haffner and Dietrich (2023):  $A^* \rightarrow O(\#\text{ccp})$ .

## Approximation Algorithms

- Many good heuristics, yet no theoretical guarantees.
- For a good reason: it's *hard* [1]:
  - NP-hard to approximate the optimal cost  $K$  within  $2^{\Theta(\log^{\delta-1} K)}$  for any  $\delta > 0$ .

---

<sup>2</sup>Left-deep only.

TL;DR

Join Ordering DP = Subset Convolution

## Our Results: Formally

---

Let  $W$  be the largest join cardinality. Then,

---

## Our Results: Formally

---

Let  $W$  be the largest join cardinality. Then,

- $C_{\text{out}}$ : Can be solved in  $O^*(2^n W)$  time.<sup>3</sup>

---

<sup>3</sup> $O^*$  hides polynomial factors in  $n$ .

## Our Results: Formally

Let  $W$  be the largest join cardinality. Then,

- $C_{\text{out}}$ : Can be solved in  $O^*(2^n W)$  time.<sup>3</sup>
- $C_{\text{out}}$ : Can be  $(1 + \varepsilon)$ -approximated in  $O^*(2^n \log W/\varepsilon)$  time.

---

<sup>3</sup> $O^*$  hides polynomial factors in  $n$ .

## Our Results: Formally

Let  $W$  be the largest join cardinality. Then,

- $C_{\text{out}}$ : Can be solved in  $O^*(2^n W)$  time.<sup>3</sup>
- $C_{\text{out}}$ : Can be  $(1 + \varepsilon)$ -approximated in  $O^*(2^n \log W / \varepsilon)$  time.
- $C_{\text{max}}$ : Can be solved in  $O(2^n n^3)$  time (does not depend on  $W$ ).

---

<sup>3</sup> $O^*$  hides polynomial factors in  $n$ .

# Our Results: Formally

Let  $W$  be the largest join cardinality. Then,

- $C_{\text{out}}$ : Can be solved in  $O^*(2^n W)$  time.<sup>3</sup>
- $C_{\text{out}}$ : Can be  $(1 + \varepsilon)$ -approximated in  $O^*(2^n \log W / \varepsilon)$  time.
- $C_{\text{max}}$ : Can be solved in  $O(2^n n^3)$  time (does not depend on  $W$ ).

## Implications

- $C_{\text{max}}$ : [Super-polynomial speedup](#).

---

<sup>3</sup> $O^*$  hides polynomial factors in  $n$ .



# Our Results: Formally

Let  $W$  be the largest join cardinality. Then,

- $C_{\text{out}}$ : Can be solved in  $O^*(2^n W)$  time.<sup>3</sup>
- $C_{\text{out}}$ : Can be  $(1 + \varepsilon)$ -approximated in  $O^*(2^n \log W / \varepsilon)$  time.
- $C_{\text{max}}$ : Can be solved in  $O(2^n n^3)$  time (does not depend on  $W$ ).

## Implications

- $C_{\text{max}}$ : Super-polynomial speedup.
- $C_{\text{out}}$ : First  $(1 + \varepsilon)$ -approximation algorithm.

---

<sup>3</sup> $O^*$  hides polynomial factors in  $n$ .

# Our Results: Formally

Let  $W$  be the largest join cardinality. Then,

- $C_{\text{out}}$ : Can be solved in  $O^*(2^n W)$  time.<sup>3</sup>
- $C_{\text{out}}$ : Can be  $(1 + \varepsilon)$ -approximated in  $O^*(2^n \log W / \varepsilon)$  time.
- $C_{\text{max}}$ : Can be solved in  $O(2^n n^3)$  time (does not depend on  $W$ ).

## Implications

- $C_{\text{max}}$ : Super-polynomial speedup.
- $C_{\text{out}}$ : First  $(1 + \varepsilon)$ -approximation algorithm.
- Beyond databases: einsum optimization.

---

<sup>3</sup> $O^*$  hides polynomial factors in  $n$ .

# Our Results: Formally

Let  $W$  be the largest join cardinality. Then,

- $C_{\text{out}}$ : Can be solved in  $O^*(2^n W)$  time.<sup>3</sup>
- $C_{\text{out}}$ : Can be  $(1 + \varepsilon)$ -approximated in  $O^*(2^n \log W / \varepsilon)$  time.
- $C_{\text{max}}$ : Can be solved in  $O(2^n n^3)$  time (does not depend on  $W$ ).

## Implications

- $C_{\text{max}}$ : Super-polynomial speedup.
- $C_{\text{out}}$ : First  $(1 + \varepsilon)$ -approximation algorithm.
- Beyond databases: einsum optimization.
  - Used in quantum circuit simulation.
  - Speedup over the state-of-the-art algorithm in `opt_einsum`.

---

<sup>3</sup> $O^*$  hides polynomial factors in  $n$ .

# Overview: Part I

## General Framework

- Join ordering DP:

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

# Overview: Part I

## General Framework

- Join ordering DP:

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

- Subset convolution in the  $(\min, \otimes)$  semi-ring:

$$(f \star g)(S) = \min_{T \subseteq S} (f(T) \otimes g(S \setminus T)).$$

# Overview: Part I

## General Framework

- Join ordering DP:

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

- Subset convolution in the  $(\min, \otimes)$  semi-ring:

$$(f \star g)(S) = \min_{T \subseteq S} (f(T) \otimes g(S \setminus T)).$$

- In other words:

$$\text{DP}[S] = c(S) \otimes (\text{DP} \star \text{DP})(S).$$

# Overview: Part I

## General Framework

- Join ordering DP:

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

- Subset convolution in the  $(\min, \otimes)$  semi-ring:

$$(f \star g)(S) = \min_{T \subseteq S} (f(T) \otimes g(S \setminus T)).$$

- In other words:

$$\text{DP}[S] = c(S) \otimes (\text{DP} \star \text{DP})(S).$$

- Instantiations:

# Overview: Part I

## General Framework

- Join ordering DP:

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

- Subset convolution in the  $(\min, \otimes)$  semi-ring:

$$(f \star g)(S) = \min_{T \subseteq S} (f(T) \otimes g(S \setminus T)).$$

- In other words:

$$DP[S] = c(S) \otimes (DP \star DP)(S).$$

- Instantiations:
  - $C_{\text{out}}$ :  $(\min, +)$ .



# Overview: Part I

## General Framework

- Join ordering DP:

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

- Subset convolution in the  $(\min, \otimes)$  semi-ring:

$$(f \star g)(S) = \min_{T \subseteq S} (f(T) \otimes g(S \setminus T)).$$

- In other words:

$$DP[S] = c(S) \otimes (DP \star DP)(S).$$

- Instantiations:
  - $C_{\text{out}}$ :  $(\min, +)$ .
  - $C_{\text{max}}$ :  $(\min, \max)$ .

## 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

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

## Overview: Part II

---

### Fast Subset Convolution

- Evaluating convolutions in semi-rings is hard –  $O(3^n)$ .

## Overview: Part II

---

### Fast Subset Convolution

- Evaluating convolutions in semi-rings is hard –  $O(3^n)$ .
- Convolution in the  $(+, \times)$  ring is much faster –  $O(2^n n^2)$ :

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

# Overview: Part II

## Fast Subset Convolution

- Evaluating convolutions in semi-rings is hard –  $O(3^n)$ .
- Convolution in the  $(+, \times)$  ring is much faster –  $O(2^n n^2)$ :

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

- **Good news:** We can map a semi-ring into a ring via *polynomials*.
  - Represent a value  $v$  as  $x^v$ .
  - **Intuition:**
    - “+” becomes product:  $x^{a+b} = x^a x^b$ .
    - “min” is just taking the lowest monomial.

# Overview: Part II

## Fast Subset Convolution

- Evaluating convolutions in semi-rings is hard –  $O(3^n)$ .
- Convolution in the  $(+, \times)$  ring is much faster –  $O(2^n n^2)$ :

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

- **Good news:** We can map a semi-ring into a ring via *polynomials*.
  - Represent a value  $v$  as  $x^v$ .
  - **Intuition:**
    - “+” becomes product:  $x^{a+b} = x^a x^b$ .
    - “min” is just taking the lowest monomial.
- **Bad news:** Running time becomes pseudopolynomial:  $O^*(2^n W)$ , where  $W$  is the largest value in  $f$  and  $g$ .

# Overview: Part II

## Fast Subset Convolution

- Evaluating convolutions in semi-rings is hard –  $O(3^n)$ .
- Convolution in the  $(+, \times)$  ring is much faster –  $O(2^n n^2)$ :

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

- **Good news:** We can map a semi-ring into a ring via *polynomials*.
  - Represent a value  $v$  as  $x^v$ .
  - **Intuition:**
    - “+” becomes product:  $x^{a+b} = x^a x^b$ .
    - “min” is just taking the lowest monomial.
- **Bad news:** Running time becomes pseudopolynomial:  $O^*(2^n W)$ , where  $W$  is the largest value in  $f$  and  $g$ .

Where to go?

## DPconv: Takeaway

- *No* pseudopolynomial factor when we optimize for  $C_{\max}$ .

---

<sup>4</sup>Or in  $O^*(2^{\frac{3n}{2}}/\sqrt{\epsilon})$ -time, however this is not yet practical.



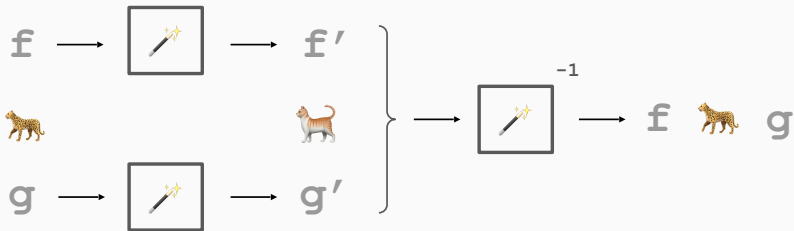
### DPconv: Takeaway

- No pseudopolynomial factor when we optimize for  $C_{\max}$ .
- No pseudopolynomial factor when we move to  $(1 + \varepsilon)$ -approximation  
→  $C_{\text{out}}$  can be  $(1 + \varepsilon)$ -approximated in  $O^*(2^n \log W/\varepsilon)$ .<sup>4</sup>

---

<sup>4</sup>Or in  $O^*(2^{\frac{3n}{2}}/\sqrt{\varepsilon})$ -time, however this is not yet practical.

# Fast Convolutions: Intuition



# Fast Subset Convolution

- Same principle: Transform the functions by a magic box.
- Our magic box – **zeta transform**:

$$\zeta f(S) = \sum_{T \subseteq S} f(T),$$

for a subset  $S \subseteq [n]$ .

# Fast Subset Convolution

- Same principle: Transform the functions by a magic box.
- Our magic box – **zeta transform**:

$$\zeta f(S) = \sum_{T \subseteq S} f(T),$$

for a subset  $S \subseteq [n]$ .

- Naive evaluation:  $O(3^n)$ .

# Fast Subset Convolution

- Same principle: Transform the functions by a magic box.
- Our magic box – **zeta transform**:

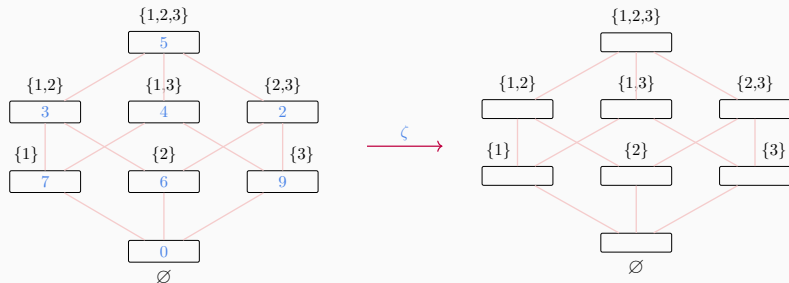
$$\zeta f(S) = \sum_{T \subseteq S} f(T),$$

for a subset  $S \subseteq [n]$ .

- Naive evaluation:  $O(3^n)$ .
- Clever evaluation:  $O(2^n n)$ .

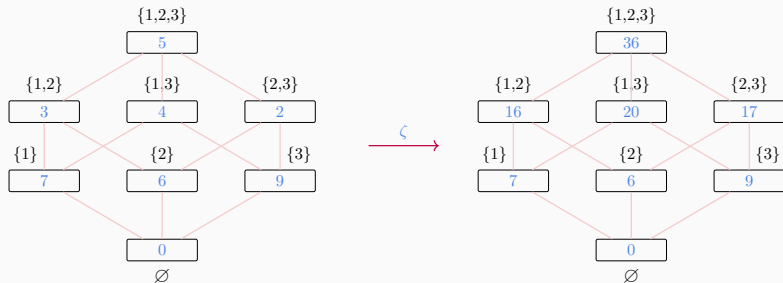
# Zeta Transform: Hands-On

- You have 1 minute to fill up:



**Figure 1:** How to compute the zeta transform of a set function.

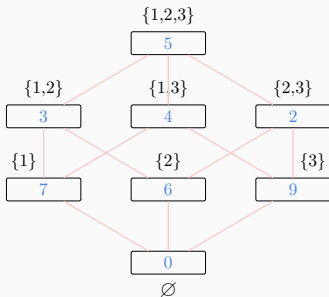
# Zeta Transform: Hands-On



**Figure 2:** How to compute the zeta transform of a set function.

# Fast Zeta Transform: Pseudocode

```
zeta(f):  
  for d in range(n):  
    for S in range(2**n):  
      if S & 2**d:  
        f[S] += f[S ^ 2**d]
```





# Inverse Zeta Transform

---

- "Wait, you also need the inverted magic box, right?"

# Inverse Zeta Transform

---

- "Wait, you also need the inverted magic box, right?"
- The inverse of the zeta transform:

$$\mu f(S) = \sum_{T \subseteq S} (-1)^{|S \setminus T|} f(T).$$

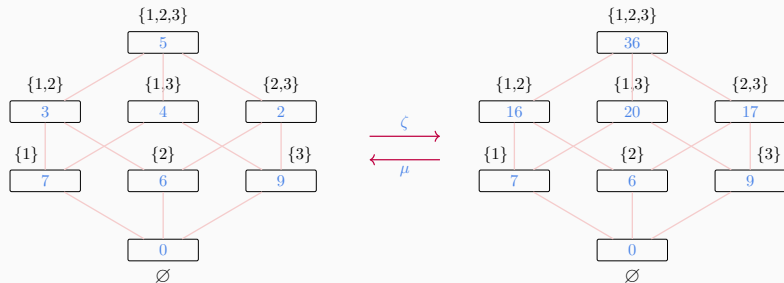
# Inverse Zeta Transform

- "Wait, you also need the inverted magic box, right?"
- The inverse of the zeta transform:

$$\mu f(S) = \sum_{T \subseteq S} (-1)^{|S \setminus T|} f(T).$$

- In other words:  $f = \mu\zeta f = \zeta\mu f$ .

# Inverse Zeta Transform: Example



**Figure 3:** Inverting the zeta transform.

## First attempt

1. Magic box:  $\zeta f, \zeta g$ .

## First attempt

1. Magic box:  $\zeta f, \zeta g$ .
2. Compute  $\zeta f \cdot \zeta g$ .

## First attempt

1. Magic box:  $\zeta f, \zeta g$ .
2. Compute  $\zeta f \cdot \zeta g$ .
3. Inverted magic box:  $\mu(\zeta f \cdot \zeta g)$ .

# Fast Subset Convolution

## First attempt

1. Magic box:  $\zeta f, \zeta g$ .
2. Compute  $\zeta f \cdot \zeta g$ .
3. Inverted magic box:  $\mu(\zeta f \cdot \zeta g)$ .

Unfortunately, this only computes

$$\sum_{\substack{U, V \subseteq S \\ U \cup V = S}} f(U)g(V).$$



# Fast Subset Convolution

## First attempt

1. Magic box:  $\zeta f, \zeta g$ .
2. Compute  $\zeta f \cdot \zeta g$ .
3. Inverted magic box:  $\mu(\zeta f \cdot \zeta g)$ .

Unfortunately, this only computes

$$\sum_{\substack{U, V \subseteq S \\ U \cup V = S}} f(U)g(V).$$

*How can we fix it?*

## Optimizing $C_{\max}$

- Ad-hoc dynamic programming:

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

## Optimizing $C_{\max}$

- Ad-hoc dynamic programming:

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

- **Note:** The DP does not create any new values!  
→ We can binary search the optimal value  $DP[\{1, \dots, n\}]$ .

## Optimizing $C_{\max}$

- Ad-hoc dynamic programming:

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

- **Note:** The DP does not create any new values!  
→ We can binary search the optimal value  $DP[\{1, \dots, n\}]$ .
- Fix a join cardinality  $\gamma$  and define a new cardinality function:

$$c'(S) = \text{if } c(S) \leq \gamma \text{ then } 1 \text{ else } 0.$$

## Optimizing $C_{\max}$

- Ad-hoc dynamic programming:

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

- **Note:** The DP does not create any new values!  
→ We can binary search the optimal value  $DP[\{1, \dots, n\}]$ .
- Fix a join cardinality  $\gamma$  and define a new cardinality function:

$$c'(S) = \text{if } c(S) \leq \gamma \text{ then } 1 \text{ else } 0.$$

- Now just run the  $(+, \times)$  dynamic programming with  $c'$ :

$$DP[S] = c'(S) + \sum_{T \subset S} DP[T] \cdot DP[S \setminus T].$$

## Optimizing $C_{\max}$

- Ad-hoc dynamic programming:

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

- **Note:** The DP does not create any new values!  
→ We can binary search the optimal value  $DP[\{1, \dots, n\}]$ .
- Fix a join cardinality  $\gamma$  and define a new cardinality function:

$$c'(S) = \text{if } c(S) \leq \gamma \text{ then } 1 \text{ else } 0.$$

- Now just run the  $(+, \times)$  dynamic programming with  $c'$ :

$$DP[S] = c'(S) + \sum_{T \subset S} DP[T] \cdot DP[S \setminus T].$$

- If  $DP[\{1, \dots, n\}] > 0$ , then  $\gamma$  is feasible.

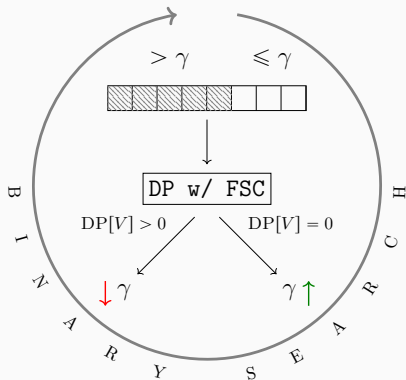
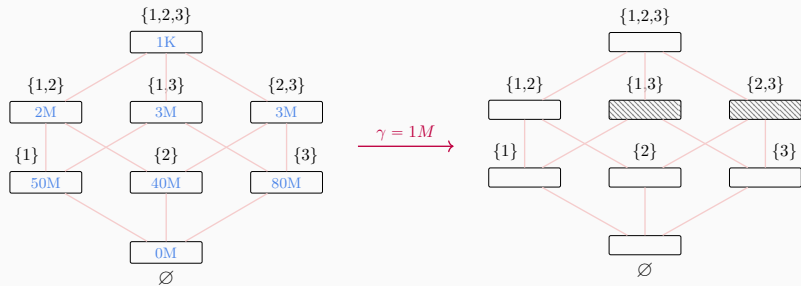


Figure 4: How DPconv optimizes  $C_{\max}$ .

## Optimizing $C_{\max}$ : Example



**Figure 5:** The cardinalities  $> \gamma$  are ignored.



# Optimizing $C_{\text{out}}$

- Ad-hoc DP:  $O(3^n)$ .
- $C_{\text{out}}$  is working in the  $(\min, +)$  semi-ring.
  - Using the FSC framework, we obtain a  $O^*(2^n W)$ -time algorithm.

# Optimizing $C_{\text{out}}$

- Ad-hoc DP:  $O(3^n)$ .
- $C_{\text{out}}$  is working in the  $(\min, +)$  semi-ring.
  - Using the FSC framework, we obtain a  $O^*(2^n W)$ -time algorithm.
- **However**, this is not so practical in the context of join ordering.
  - Join cardinalities can be *very* large.

# Optimizing $C_{\text{out}}$

- Ad-hoc DP:  $O(3^n)$ .
- $C_{\text{out}}$  is working in the  $(\min, +)$  semi-ring.
  - Using the FSC framework, we obtain a  $O^*(2^n W)$ -time algorithm.
- **However**, this is not so practical in the context of join ordering.
  - Join cardinalities can be *very* large.
- Two open questions in algorithm design research:

# Optimizing $C_{\text{out}}$

- Ad-hoc DP:  $O(3^n)$ .
- $C_{\text{out}}$  is working in the  $(\min, +)$  semi-ring.
  - Using the FSC framework, we obtain a  $O^*(2^n W)$ -time algorithm.
- **However**, this is not so practical in the context of join ordering.
  - Join cardinalities can be *very* large.
- Two open questions in algorithm design research:
  - Is  $O(3^n)$  the best we can do for  $(\min, +)$  subset convolution?

# Optimizing $C_{\text{out}}$

- Ad-hoc DP:  $O(3^n)$ .
- $C_{\text{out}}$  is working in the  $(\min, +)$  semi-ring.
  - Using the FSC framework, we obtain a  $O^*(2^n W)$ -time algorithm.
- **However**, this is not so practical in the context of join ordering.
  - Join cardinalities can be *very* large.
- Two open questions in algorithm design research:
  - Is  $O(3^n)$  the best we can do for  $(\min, +)$  subset convolution?
  - Is the pseudopolynomial  $O(W)$  factor actually needed?

## Joint Optimization: $C_{\text{cap}} = C_{\text{max}} + C_{\text{out}}$

### Motivation $C_{\text{cap}}$

- Optimizing for  $C_{\text{max}}$  alone may lead to slow plans.  
→ Also enforce that we have an optimal  $C_{\text{out}}$  plan.

## Joint Optimization: $C_{\text{cap}} = C_{\text{max}} + C_{\text{out}}$

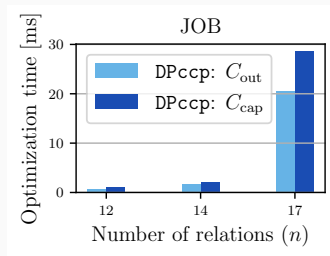
### Motivation $C_{\text{cap}}$

- Optimizing for  $C_{\text{max}}$  alone may lead to slow plans.  
→ Also enforce that we have an optimal  $C_{\text{out}}$  plan.
- First run DPccp with  $C_{\text{max}}$ , then  $C_{\text{out}}$  with a threshold.

# Joint Optimization: $C_{\text{cap}} = C_{\text{max}} + C_{\text{out}}$

## Motivation $C_{\text{cap}}$

- Optimizing for  $C_{\text{max}}$  alone may lead to slow plans.  
→ Also enforce that we have an optimal  $C_{\text{out}}$  plan.
- First run DPccp with  $C_{\text{max}}$ , then  $C_{\text{out}}$  with a threshold.
- **However**, this is slower:



**Figure 6:** Initial overhead in optimizing  $C_{\text{cap}}$  on JOB



## Joint Optimization: $C_{\text{cap}} = C_{\text{max}} + C_{\text{out}}$

---

### (Very) Simple Solution

- Replace DP<sub>ccp</sub> in  $C_{\text{max}}$  optimization by DP<sub>conv</sub>.

# Approximating $C_{\text{out}}$ ?

---

Recall:  $C_{\text{out}}$  takes  $O^*(2^n W)$ -time with our framework.

**How to dissolve the pseudopolynomial factor?**

# Approximating $C_{\text{out}}$ ?

---

Recall:  $C_{\text{out}}$  takes  $O^*(2^n W)$ -time with our framework.

## How to dissolve the pseudopolynomial factor?

- We will approximate the optimal value of  $C_{\text{out}}$  within  $1 + \varepsilon$ .

# Approximating $C_{\text{out}}$ ?

Recall:  $C_{\text{out}}$  takes  $O^*(2^n W)$ -time with our framework.

## How to dissolve the pseudopolynomial factor?

- We will approximate the optimal value of  $C_{\text{out}}$  within  $1 + \varepsilon$ .
- Example: The optimal plan has  $C_{\text{out}} = 400K$  and  $\varepsilon = 1\%$ .  
→ Then we obtain a plan of cost  $\leq 404K$ .

# Approximating $C_{\text{out}}$ ?

Recall:  $C_{\text{out}}$  takes  $O^*(2^n W)$ -time with our framework.

## How to dissolve the pseudopolynomial factor?

- We will approximate the optimal value of  $C_{\text{out}}$  within  $1 + \varepsilon$ .
- Example: The optimal plan has  $C_{\text{out}} = 400K$  and  $\varepsilon = 1\%$ .  
→ Then we obtain a plan of cost  $\leq 404K$ .
- Running time:  $O^*(2^n \log W/\varepsilon)$ .

# Approximate $(\min, +)$ Subset Convolution

## Definition

Given two set functions  $f, g$ , approximate their  $(\min, +)$  subset convolution:

$$(f \star g)(S) \leq \tilde{h}(S) \leq (1 + \varepsilon)(f \star g)(S)$$

for all  $S \subseteq [n]$ , with  $\varepsilon > 0$ .

# Approximate $(\min, +)$ Subset Convolution

## Definition

Given two set functions  $f, g$ , approximate their  $(\min, +)$  subset convolution:

$$(f \star g)(S) \leq \tilde{h}(S) \leq (1 + \varepsilon)(f \star g)(S)$$

for all  $S \subseteq [n]$ , with  $\varepsilon > 0$ .

## How to employ it

- Fix a relative error  $\delta > 0$ .

# Approximate $(\min, +)$ Subset Convolution

## Definition

Given two set functions  $f, g$ , approximate their  $(\min, +)$  subset convolution:

$$(f \star g)(S) \leq \tilde{h}(S) \leq (1 + \varepsilon)(f \star g)(S)$$

for all  $S \subseteq [n]$ , with  $\varepsilon > 0$ .

## How to employ it

- Fix a relative error  $\delta > 0$ .
- Hence, the total error is  $(1 + \delta)^{n-1}$ .



# Approximate $(\min, +)$ Subset Convolution

## Definition

Given two set functions  $f, g$ , approximate their  $(\min, +)$  subset convolution:

$$(f \star g)(S) \leq \tilde{h}(S) \leq (1 + \varepsilon)(f \star g)(S)$$

for all  $S \subseteq [n]$ , with  $\varepsilon > 0$ .

## How to employ it

- Fix a relative error  $\delta > 0$ .
- Hence, the total error is  $(1 + \delta)^{n-1}$ .
- To obtain  $(1 + \varepsilon)$ -approximation, simply set  $\delta := \Theta(\varepsilon/(n - 1))$ .

## Benchmarks

- Clique queries with cardinalities  $c(S)$  s.t.  $c(S) \leq c(S_1) \cdot c(S_1)$ .

## Benchmarks

- Clique queries with cardinalities  $c(S)$  s.t.  $c(S) \leq c(S_1) \cdot c(S_1)$ .
- When comparing to the  $A^*$ -based optimizer (SIGMOD'23), we take their original clique queries since their algorithm performance is sensitive to the cardinalities.

## Competitors

- $\text{DP}_{\text{sub}}[\text{max}, \text{out}]$ :  $\Theta(3^n)$  – independent of the query shape.

## Benchmarks

- Clique queries with cardinalities  $c(S)$  s.t.  $c(S) \leq c(S_1) \cdot c(S_1)$ .
- When comparing to the  $A^*$ -based optimizer (SIGMOD'23), we take their original clique queries since their algorithm performance is sensitive to the cardinalities.

## Competitors

- $\text{DP}_{\text{sub}}[\text{max}, \text{out}]$ :  $\Theta(3^n)$  – independent of the query shape.
- $\text{DP}_{\text{ccp}} = \text{DP}_{\text{sub}}$  on clique queries.

## Benchmarks

- Clique queries with cardinalities  $c(S)$  s.t.  $c(S) \leq c(S_1) \cdot c(S_1)$ .
- When comparing to the  $A^*$ -based optimizer (SIGMOD'23), we take their original clique queries since their algorithm performance is sensitive to the cardinalities.

## Competitors

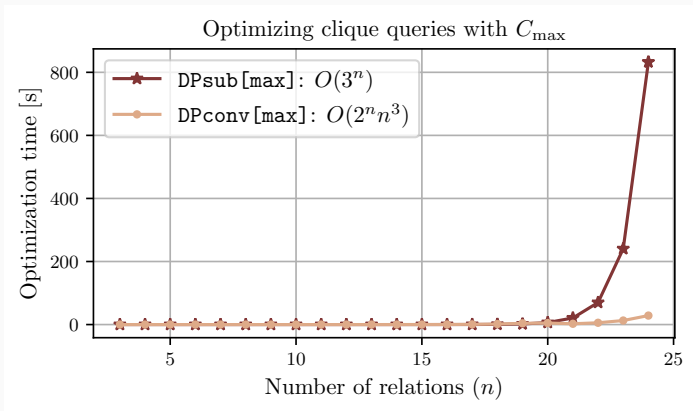
- $DP_{sub}[\max, out]$ :  $\Theta(3^n)$  – independent of the query shape.
- $DP_{ccp} = DP_{sub}$  on clique queries.
- $A^*[out]$ :  $O(\#ccp)$  – adapts to the cardinalities.

## Benchmarks

- Clique queries with cardinalities  $c(S)$  s.t.  $c(S) \leq c(S_1) \cdot c(S_1)$ .
- When comparing to the  $A^*$ -based optimizer (SIGMOD'23), we take their original clique queries since their algorithm performance is sensitive to the cardinalities.

## Competitors

- $\text{DP}_{\text{sub}}[\text{max}, \text{out}]$ :  $\Theta(3^n)$  – independent of the query shape.
- $\text{DP}_{\text{ccp}} = \text{DP}_{\text{sub}}$  on clique queries.
- $A^*[\text{out}]$ :  $O(\#\text{ccp})$  – adapts to the cardinalities.
- $\text{DP}_{\text{conv}}[\text{max}]$ :  $\Theta(2^n n^3)$  – independent of the query shape.



**Figure 7:** Optimizing for  $C_{\max}$  on clique queries

Evaluation:  $C_{cap} = C_{max} + C_{out}$

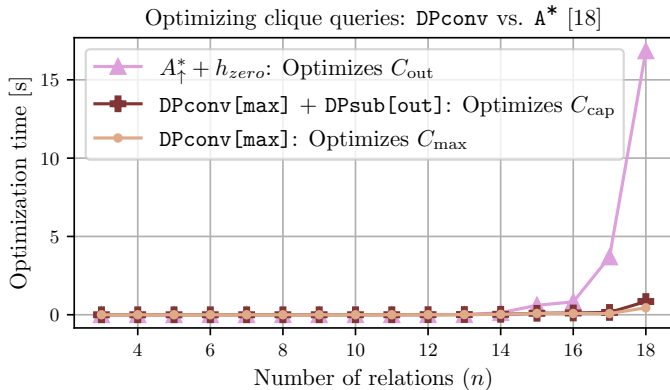
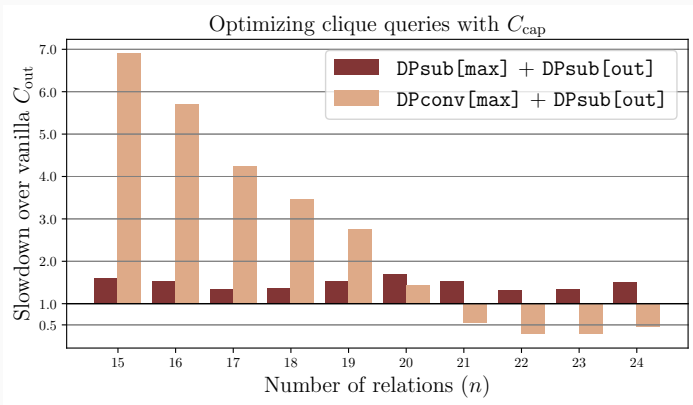


Figure 8: Optimizing for  $C_{out}$  and  $C_{cap}$  on clique queries



**Evaluation:**  $C_{\text{cap}} = C_{\text{max}} + C_{\text{out}}$



**Figure 9:** Optimizing for  $C_{\text{cap}}$  on clique queries

# Open Problems

---

- Currently, `DPconv` is agnostic to the query shape.
  - For  $C_{\max}$ , it *always* runs in  $O(2^n n^3)$ -time.
  - We would need a sparse subset convolution.
  - Sparse = Few connected subgraphs.

# Open Problems

---

- Currently, DPconv is agnostic to the query shape.
  - For  $C_{\max}$ , it *always* runs in  $O(2^n n^3)$ -time.
  - We would need a sparse subset convolution.
  - Sparse = Few connected subgraphs.
- Polynomial-*space* join ordering.
  - All exact join order optimizers take exponential *space*.
  - Example: Steiner tree can be solved in polynomial space.

# Open Problems

---

- Currently, DPconv is agnostic to the query shape.
  - For  $C_{\max}$ , it *always* runs in  $O(2^n n^3)$ -time.
  - We would need a sparse subset convolution.
  - Sparse = Few connected subgraphs.
- Polynomial-*space* join ordering.
  - All exact join order optimizers take exponential *space*.
  - Example: Steiner tree can be solved in polynomial space.
- Jointly optimize the memory of concurrent queries using  $C_{\max}$ .
  - AutoWLM [2] can predict query's memory requirements.

# Outlook

- This summer: Internship @GSL in Barcelona with Tiemo Bang.





S. Chatterji, S. S. K. Evani, S. Ganguly, and M. D. Yemmanuru.

**On the Complexity of Approximate Query Optimization.**

In L. Popa, S. Abiteboul, and P. G. Kolaitis, editors, *Proceedings of the Twenty-first ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems, June 3-5, Madison, Wisconsin, USA*, pages 282–292. ACM, 2002.



G. Saxena, M. Rahman, N. Chainani, C. Lin, G. Caragea, F. Chowdhury, R. Marcus, T. Kraska, I. Pandis, and B. M. Narayanaswamy.

**Auto-WLM: Machine Learning Enhanced Workload Management in Amazon Redshift.**

In S. Das, I. Pandis, K. S. Candan, and S. Amer-Yahia, editors, *Companion of the 2023 International Conference on Management of Data, SIGMOD/PODS 2023, Seattle, WA, USA, June 18-23, 2023*, pages 225–237. ACM, 2023.

## Evaluation: Costs

- CEB benchmark (13,644 queries).
- 2,873 queries:
  - $C_{\text{out}}$  has 6.8% larger  $C_{\text{max}}$ .
  - $C_{\text{max}}$  loses 22.8% in  $C_{\text{out}}$ .
  - $C_{\text{cap}}$  loses only 9.5% in  $C_{\text{out}}$  while maintaining optimal  $C_{\text{max}}$ .