

Analyzing Plan Diagrams of Data Query Optimizers

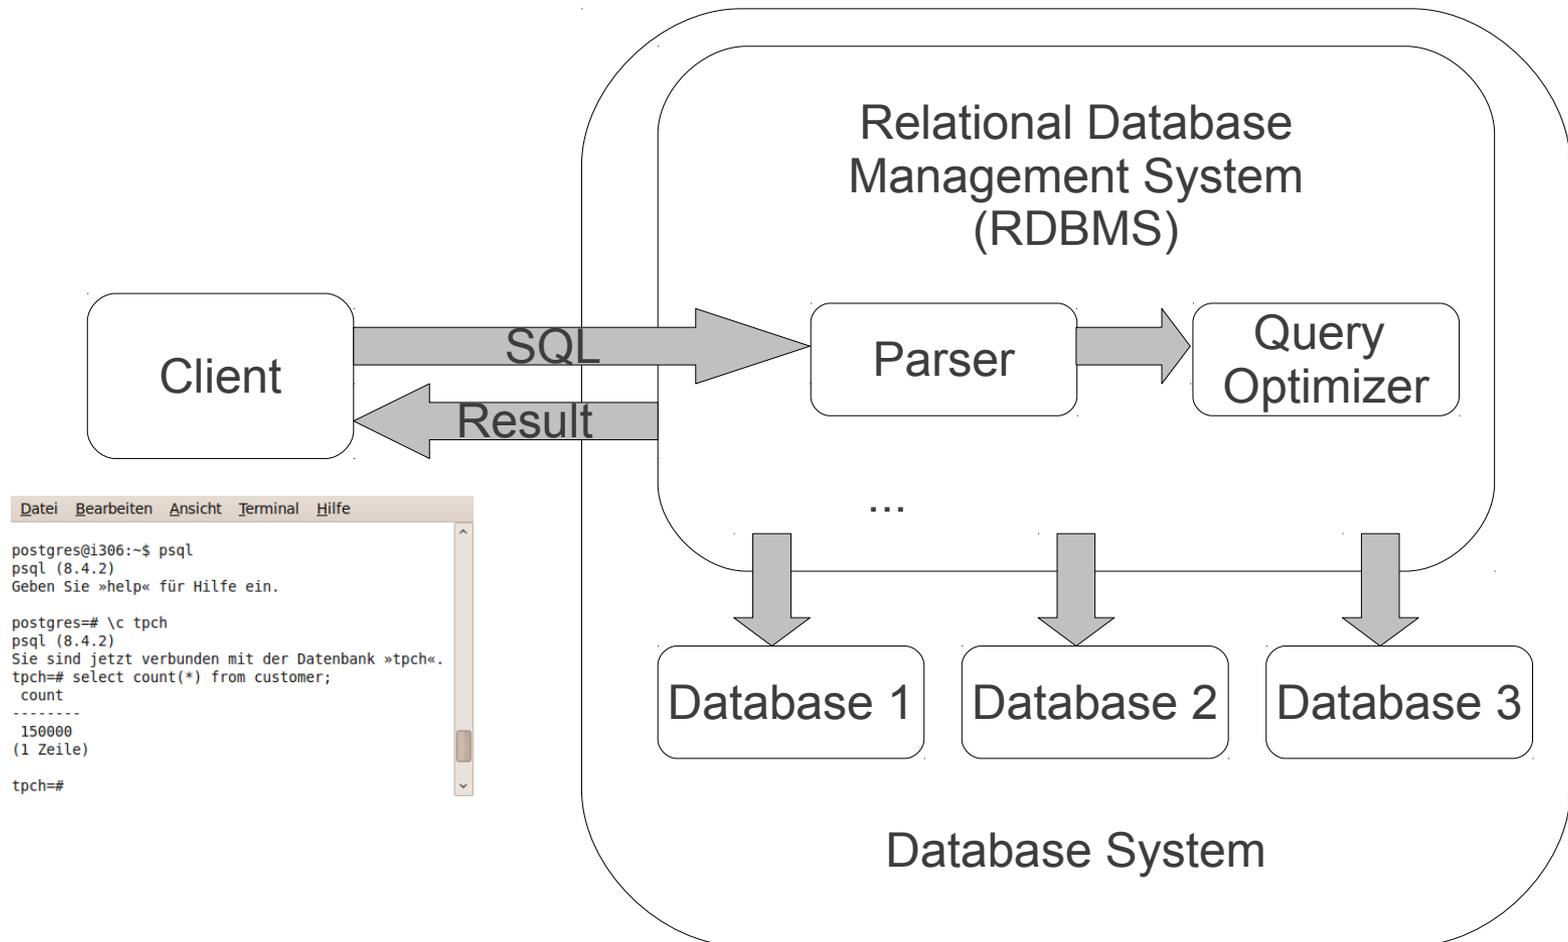
INSTITUT FÜR PROGRAMMSTRUKTUREN UND DATENORGANISATION (IPD), FAKULTÄT FÜR INFORMATIK

Paper

- Analyzing Plan Diagrams of Database Query Optimizers
- written in 2005
- Naveen Reddy, Jayant R. Haritsa
- Database Systems Lab, SERC/CSA
- Indian Institute of Science, Bangalore 560012, INDIA

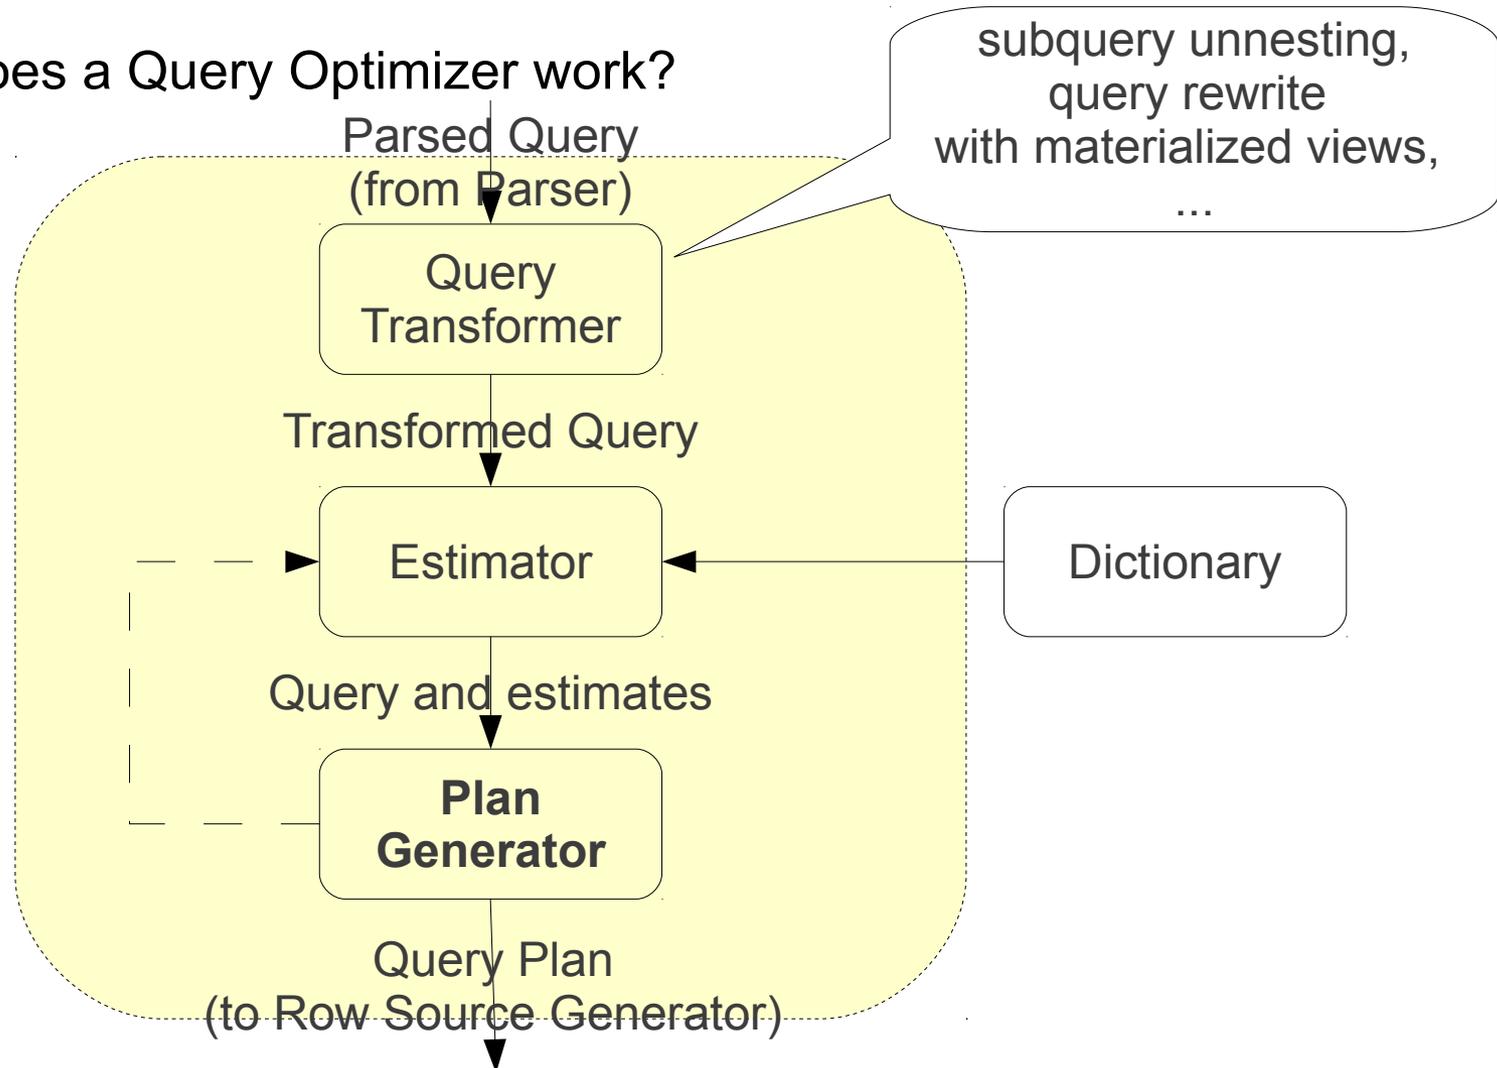
Relational Database Management System

- Where is a Query Optimizer located?



Query Optimizer

- How does a Query Optimizer work?



Query Optimizer

- Identifies the most efficient strategy to execute SQL queries
- Returns query plan
- Part of the RDBMS
- Based on estimated number of rows of relevant relations
- Relevant join conditions are given in the query

Query Plan

- A strategy to execute SQL queries
- Query Plan does not affect the result
- Example: (id is FK of Aid)
- select o,p from A inner join B on (A.id=B.Aid) where B.p=0;

- Plan P1: AB

- A=3, B=6/3=2 => cost(AB)=6

A	id	o
	1	X
	2	Y
	3	Z

B	id	Aid	p
	1	1	0
	2	1	1
	3	2	2
	4	2	0
	5	3	0
	6	3	1

- Plan P2: BA

- B=3 (using where), A=1 (primary key) => cost(BA)=3

How to analyze a query optimizer?

1. Send SQL statement EXPLAIN SELECT ... FROM ... WHERE ...
2. Optimizer returns query plan
3. This can be repeated by varying the where condition(s) systematically until the selectivity space is covered
4. Corresponding plan and cost for each set of conditions can be visualized



Query plan costs

A	id	o
	1	X
	2	Y
	3	Z

B	id	Aid	p
	1	1	0
	2	1	1
	3	2	2
	4	2	0
	5	3	0
	6	3	1

```
test=# explain select o,p
      from A inner join B on (A.id=B.Aid)
      where B.p=0;
```

QUERY PLAN

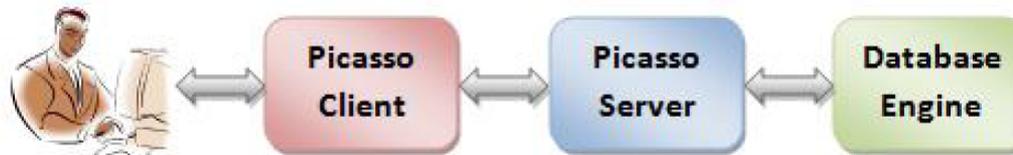
```
Hash Join  (cost=1.07..2.18 rows=3 width=6)
  Hash Cond: (b.aid = a.id)
  -> Seq Scan on b  (cost=0.00..1.07 rows=3 width=8)
      Filter: (p = 0)
  -> Hash  (cost=1.03..1.03 rows=3 width=6)
      -> Seq Scan on a  (cost=0.00..1.03 rows=3 width=6)
```

- Decision support benchmark
 - database schema, data
 - queries for testing (Q1 through Q22)
- We use same data set as the authors of the paper:

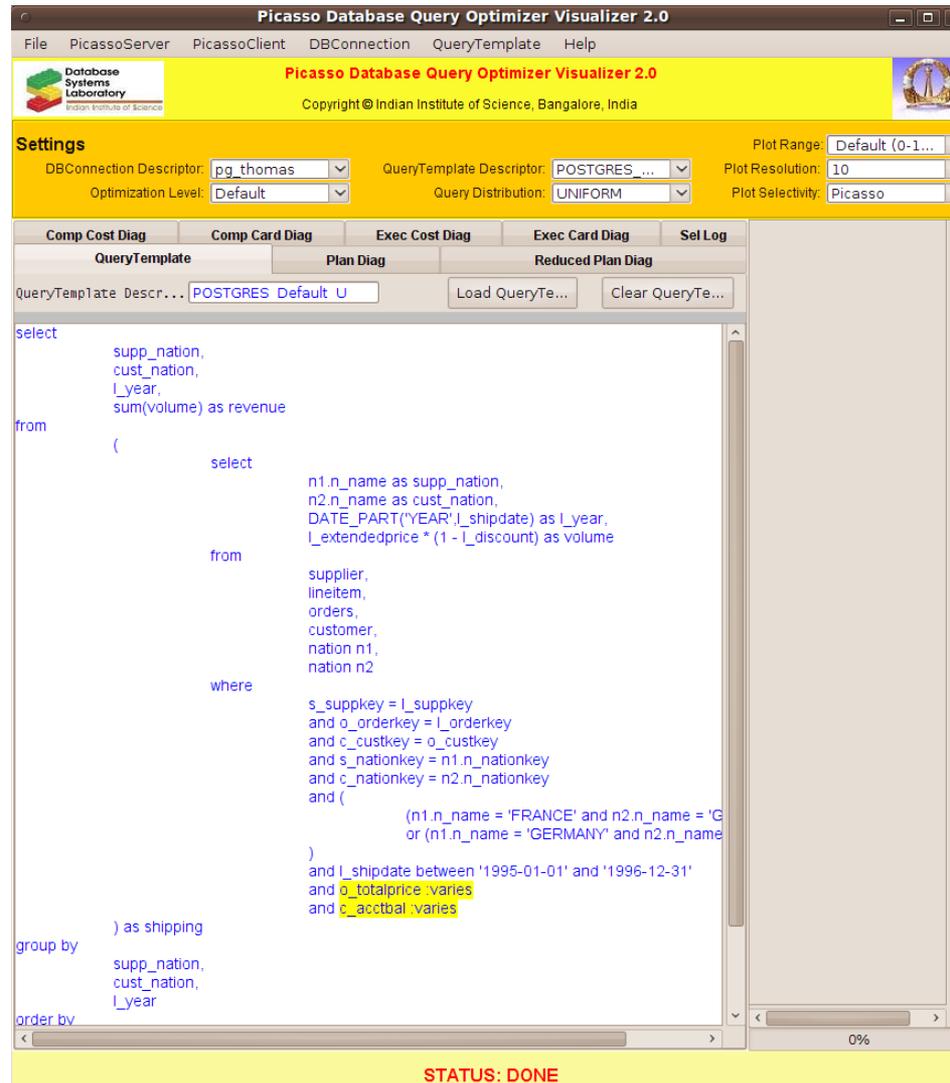
relation	#rows
part.tbl	200000
nation.tbl	25
customer.tbl	150000
region.tbl	5
partsupp.tbl	800000
supplier.tbl	10000
orders.tbl	1500000
lineitem.tbl	6001215

Picasso

- Picasso
 - Sends explain queries against the database
 - Stores and visualizes query plan and costs
- Picasso Server
 - Connects to Database (SQL-Server, Oracle, Sybase, postgresql, ...)
 - Multiuser
- Picasso Client
 - Connects to PICASSO Server
 - GUI for PICASSO Server



Run test with Picasso



Picasso Database Query Optimizer Visualizer 2.0

File PicassoServer PicassoClient DBConnection QueryTemplate Help

Database Systems Laboratory
Indian Institute of Science
Copyright © Indian Institute of Science, Bangalore, India

Settings

DBConnection Descriptor: pg_thomas QueryTemplate Descriptor: POSTGRES_... Plot Range: Default (0-1...
Optimization Level: Default Query Distribution: UNIFORM Plot Resolution: 10 Plot Selectivity: Picasso

Comp Cost Diag Comp Card Diag Exec Cost Diag Exec Card Diag Sel Log

QueryTemplate Plan Diag Reduced Plan Diag

QueryTemplate Descr... POSTGRES Default U Load QueryTe... Clear QueryTe...

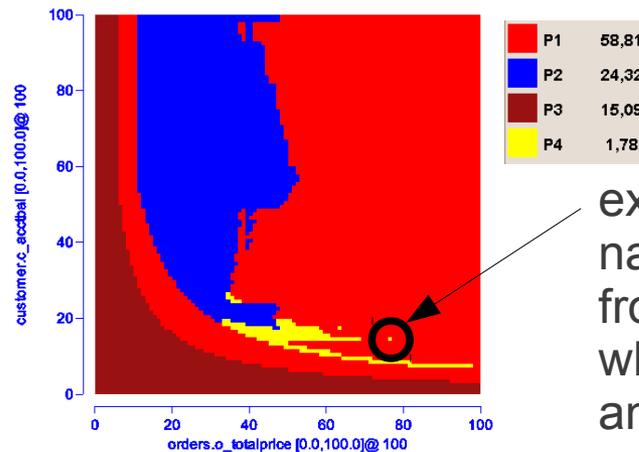
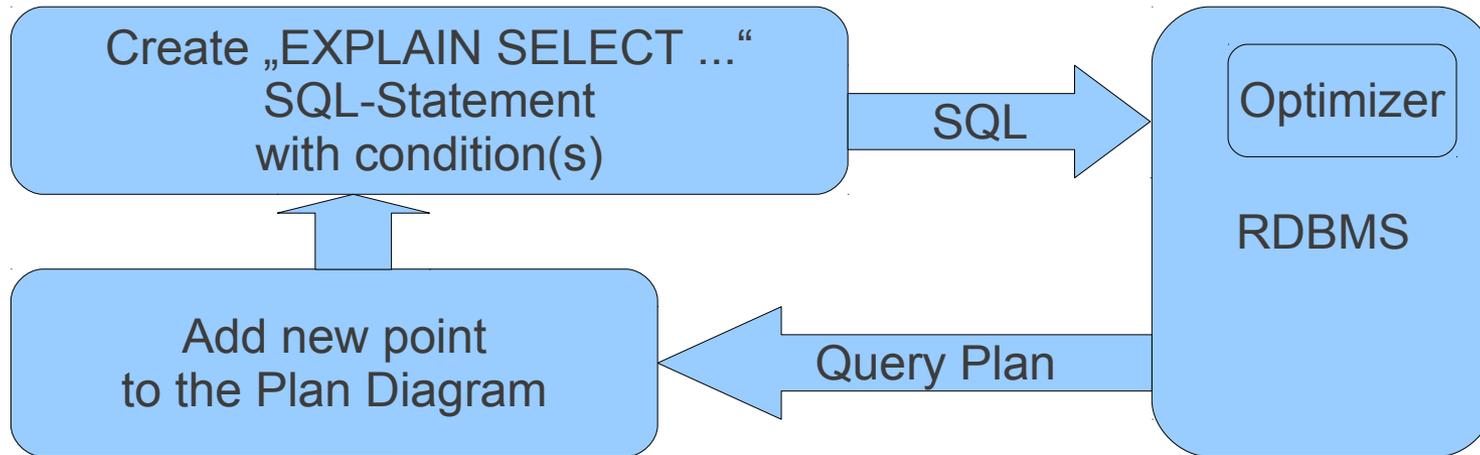
```
select
  supp_nation,
  cust_nation,
  l_year,
  sum(volume) as revenue
from
  (
    select
      n1.n_name as supp_nation,
      n2.n_name as cust_nation,
      DATE_PART('YEAR',l_shipdate) as l_year,
      l_extendedprice * (1 - l_discount) as volume
    from
      supplier,
      lineitem,
      orders,
      customer,
      nation n1,
      nation n2
    where
      s_suppkey = l_suppkey
      and o_orderkey = l_orderkey
      and c_custkey = o_custkey
      and s_nationkey = n1.n_nationkey
      and c_nationkey = n2.n_nationkey
      and (
        (n1.n_name = 'FRANCE' and n2.n_name = 'GERMANY')
        or (n1.n_name = 'GERMANY' and n2.n_name = 'FRANCE')
      )
      and l_shipdate between '1995-01-01' and '1996-12-31'
      and o_totalprice > 100000
      and o_acctbal > 100000
    ) as shipping
group by
  supp_nation,
  cust_nation,
  l_year
order by
  revenue desc
```

STATUS: DONE

Run test with Picasso

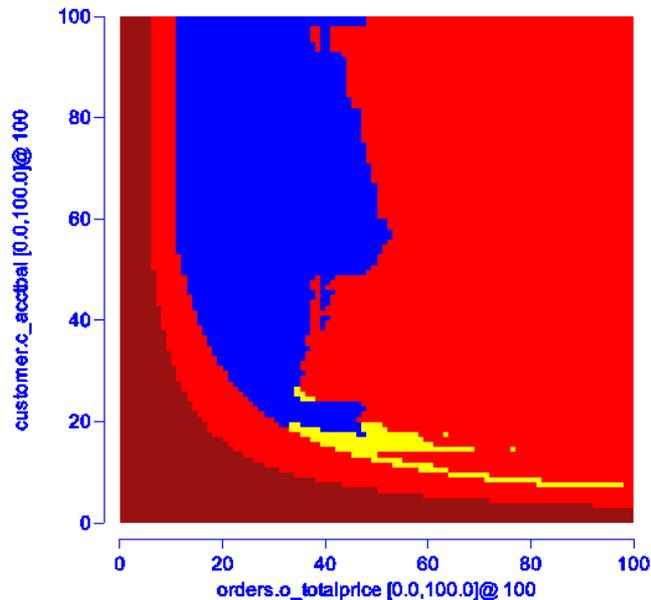
- Picasso varies WHERE conditions according to plot resolution:
 - **and o_totalprice :varies**
 - **and c_acctbal :varies**
- **2** varying conditions \Rightarrow **2**-dimensional selectivity space (2D)
- Picasso runs EXPLAIN-Statements to determine the query plan chosen by the optimizer.
- **2D** and Resolution=100 $\Rightarrow 100^2 = 10,000$ EXPLAIN-Statements

Diagram creation process

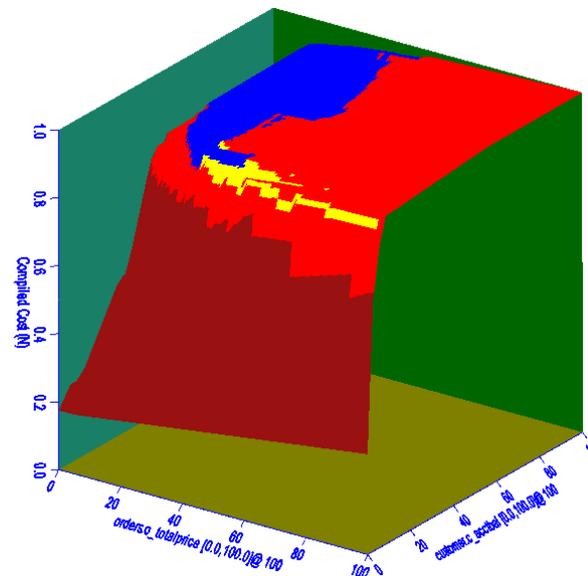


explain select
 name, price
 from customer, order
 where [...]
 and o_totalprice <= 98679
 and c_acctbal <= 5809;
 result: P4

Plan Diagram and Cost Diagram



Plan Diagram:
2D-visualization of execution plans



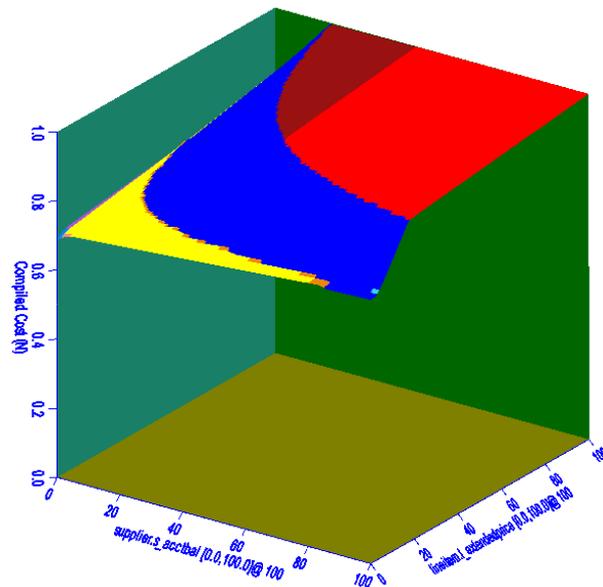
Cost Diagram:
3D-visualization of **estimated** costs.
Estimation is done by optimizer.

	P1	58,81
	P2	24,32
	P3	15,09
	P4	1,78

↑
Plan

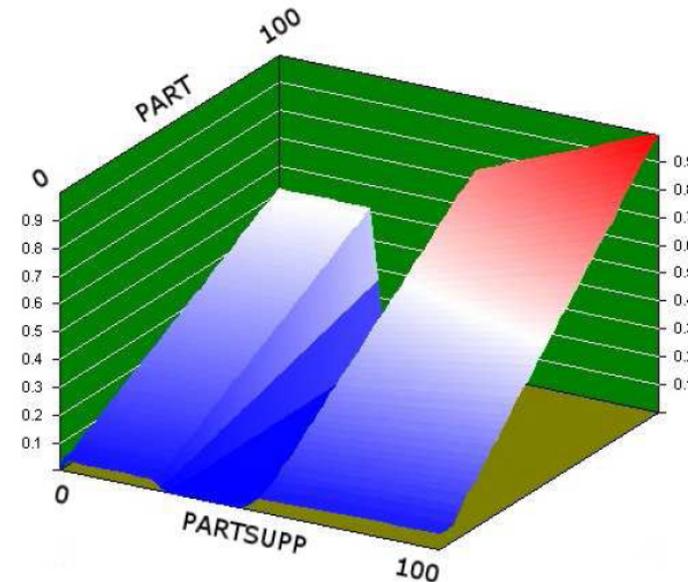
↑
Coverage of selectivity space in %

Cost domination Principle



Monotonic cost behaviour:
 Increasing selectivity space
 increases estimated costs

Cost domination principle:
 When increasing selectivity space, estimated costs must increase monotonically.

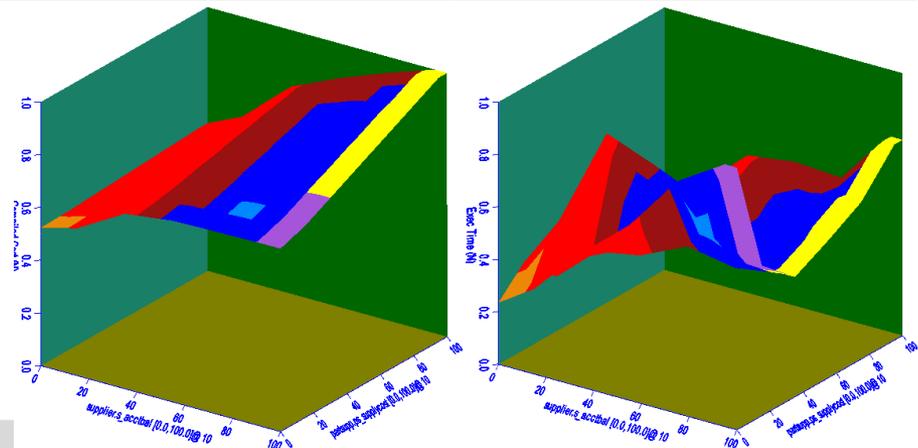


Non-monotonic cost behaviour:
 Increasing selectivity space does not
 always increase estimated costs

Estimated costs vs. execution costs

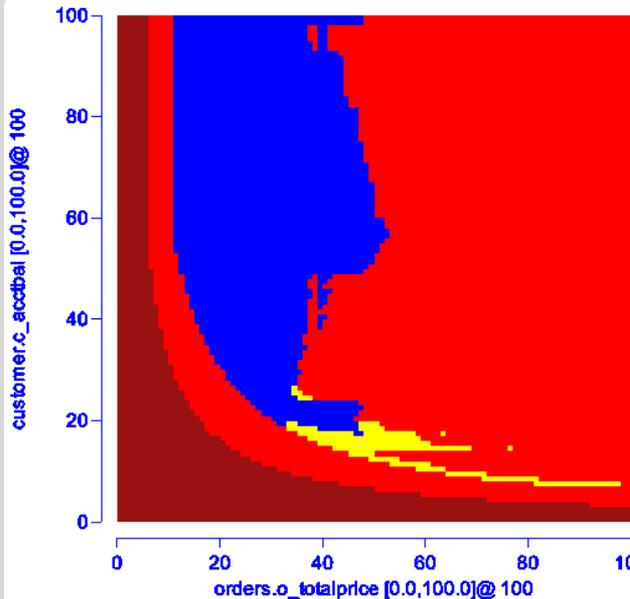
	Estimated costs	Execution costs
SQL Statement for each point in diagram	explain select ...	select ...
Measurement	#rows, ...	execution time
Creation time for a 10x10 diagram	10s	4h 10min
Monotonicity	Estimated cost diagram should always be monotonic	Execution cost diagram might not be monotonic, e.g. if estimations were not precise

=> Even if we have Monotonic estimated cost behaviour, we want to reduce #Plans and #Segments.

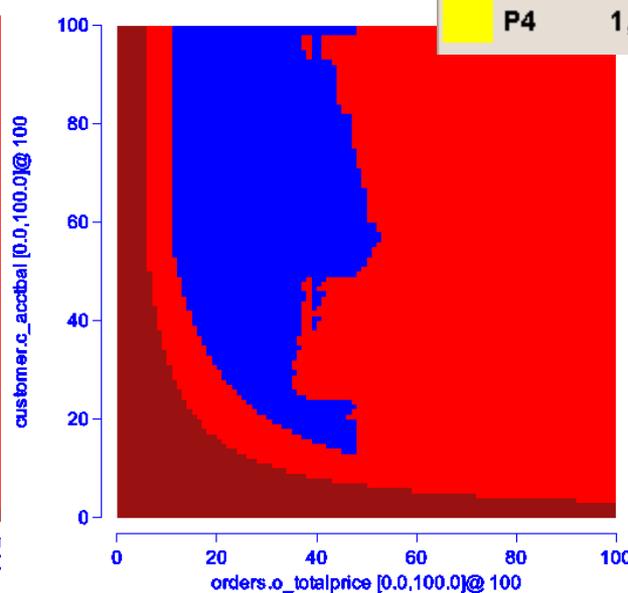


Reduced Plan Diagram

■	P1	58,81
■	P2	24,32
■	P3	15,09
■	P4	1,78

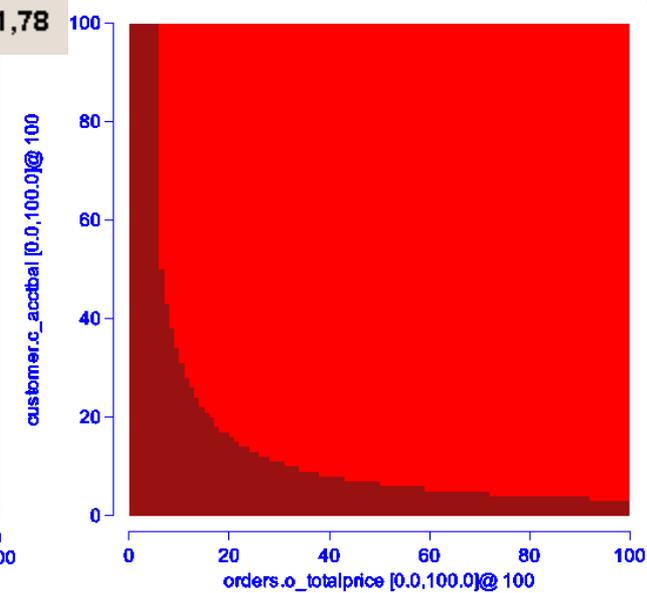


Plan Diagram



Reduced Plan Diagram

Maximum allowed Cost Increase:
5%



10%

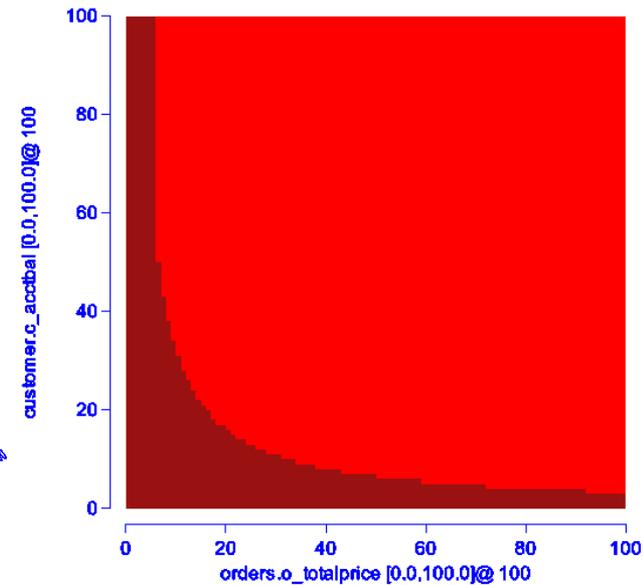
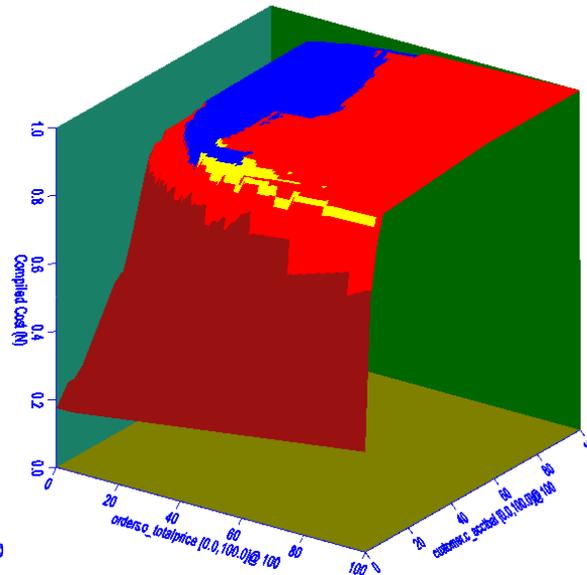
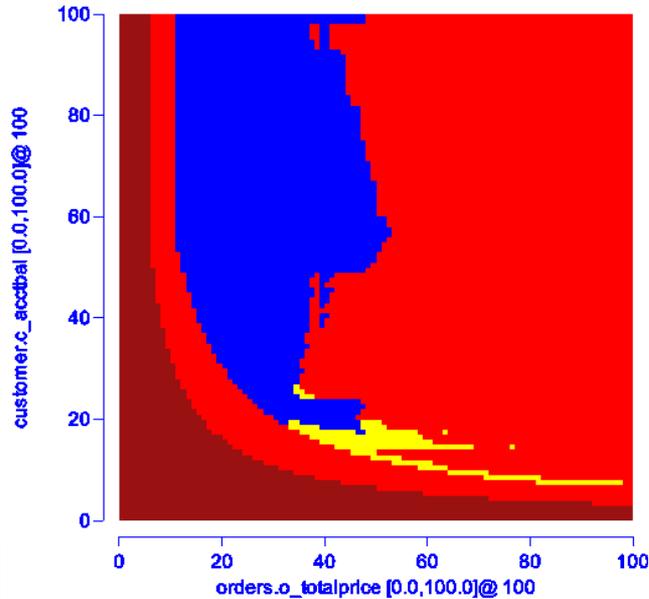
Reduced Plan Diagram:

Try to reduce #Plans without increasing **estimated** costs more than X% **in each** point

(X% is called Costgreedy Reduction or plan optimality tolerance threshold.)

TPC-H Query 7

P1	58,81
P2	24,32
P3	15,09
P4	1,78



Plan Diagram

- 4 Plans
- P4 covers 1,78% only
- Many small segments
- Not smooth

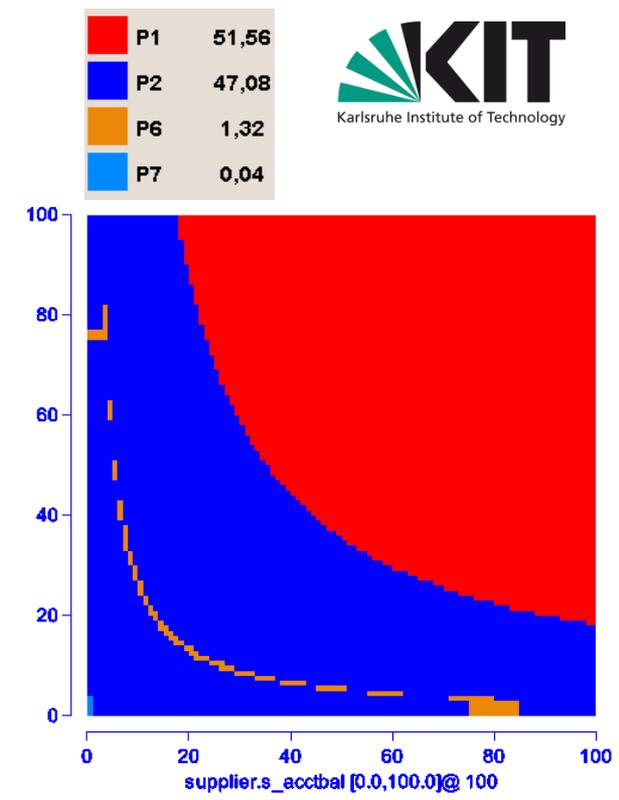
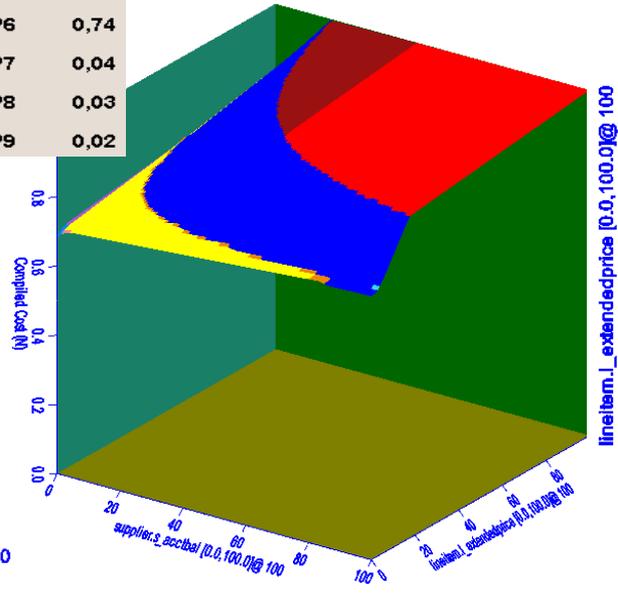
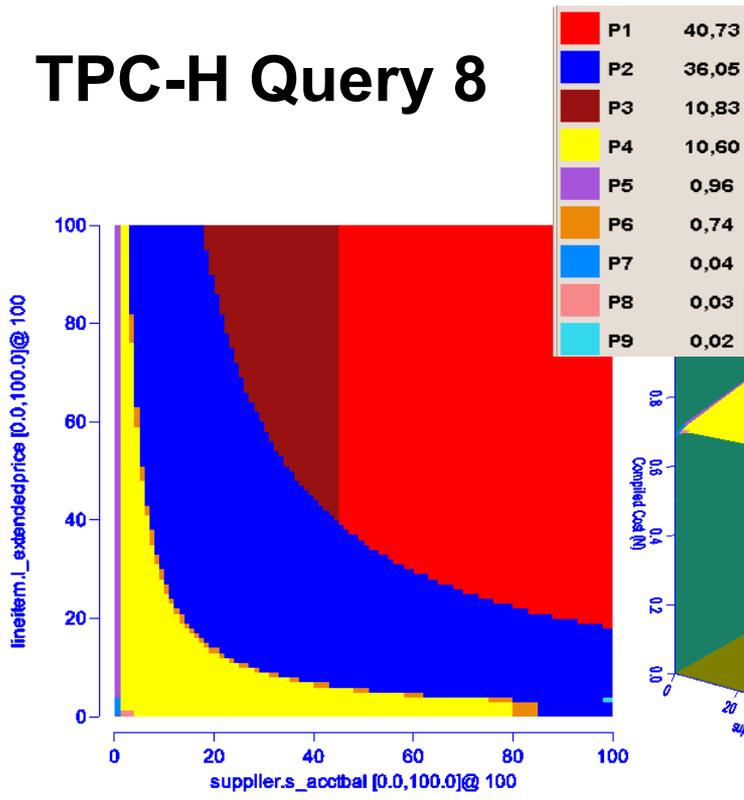
Cost Diagram

- monotonic

Reduced Plan Diagram

- 2 Plans only
- Max Increase $\leq 6.73\%$
- Avg Increase $\leq 0.72\%$

TPC-H Query 8



Plan Diagram

- 9 Plans
- 80% covered by 3 Plans
- Many small segments
- Not smooth

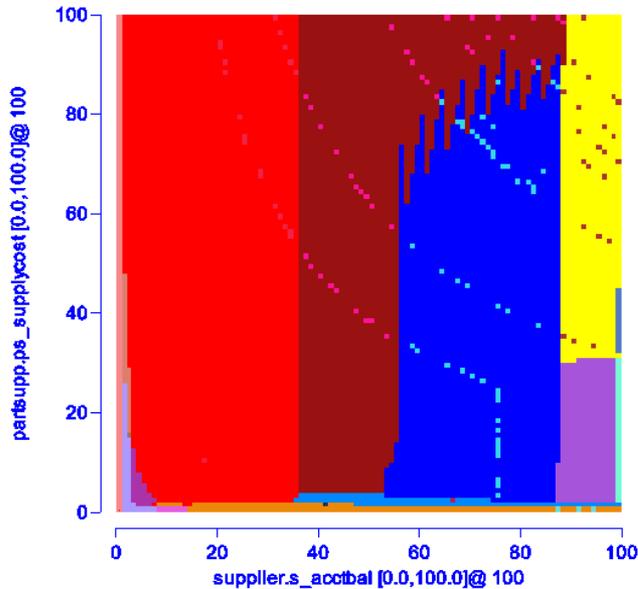
Cost Diagram

- monotonic

Reduced Plan Diagram

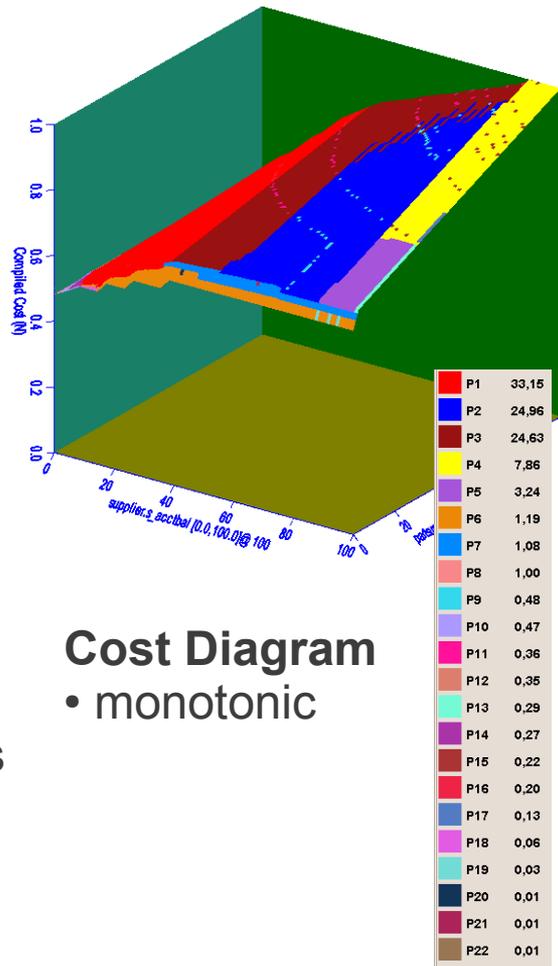
- 4 Plans only
- Max Increase $\leq 9.8\%$
- Avg Increase $\leq 0.41\%$

TPC-H Query 9



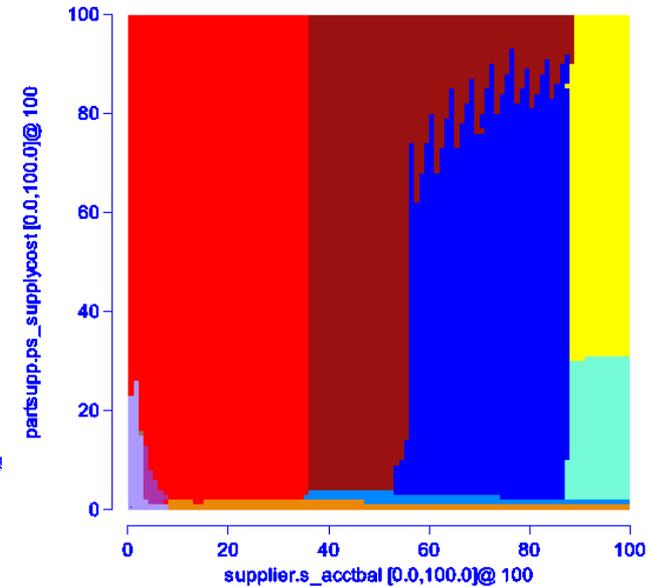
Plan Diagram

- 22 Plans
- 80% covered by 3 Plans
- Many small segments
- Not smooth
- Highly instable



Cost Diagram

- monotonic



Reduced Plan Diagram

- 10 Plans
- Max Increase $\leq 9.15\%$
- Avg Increase $\leq 0.15\%$

Clarifications

- #Plans does not necessarily correspond with execution time of one query. We do **not** rate performance!
- But: We consider the **stability** of an optimizer.

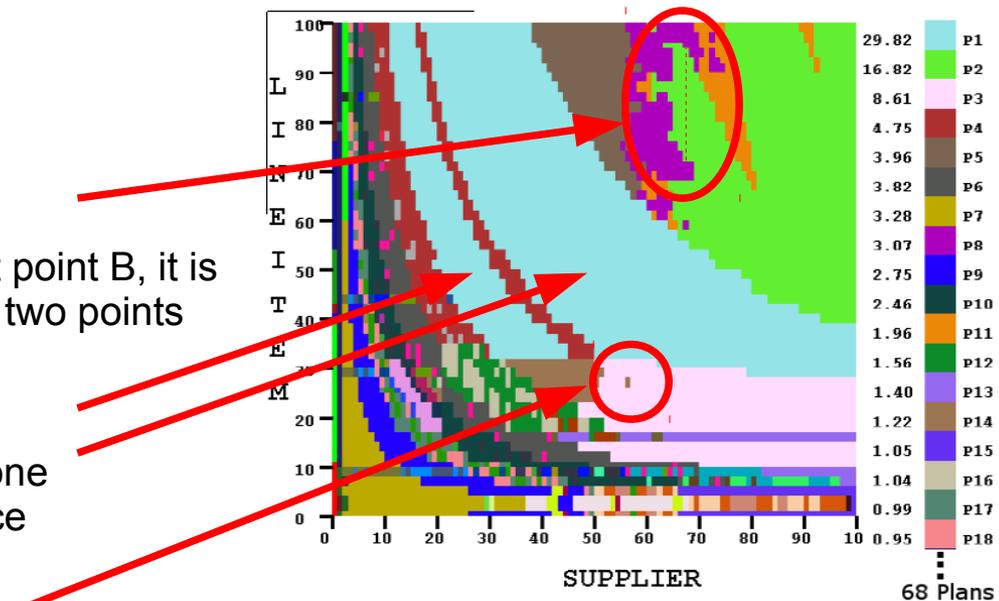
Hypothesis

- Modern query optimizers make extremely fine-grained plan choices
- smaller plans occupy less than 1% of selectivity space
- Optimizers are over-sophisticated
- Processing overheads with query optimization could be lowered

Parametric query optimization (PQO)

- PQO
 - apriori identify the optimal set of plans at compile time
 - at run time, use plan corresponding to selectivity parameter(s).
- PQO Assumptions:
- Plan Convexity
 - If Plan P is optimal at point A and at point B, it is also optimal at all points joining the two points
- Plan Uniqueness
 - An optimal plan P appears at only one contiguous region in the entire space
- Plan Homogeneity
 - An optimal plan P is optimal within the entire region enclosed by its plan boundaries.

But: today's optimizers can look like that:

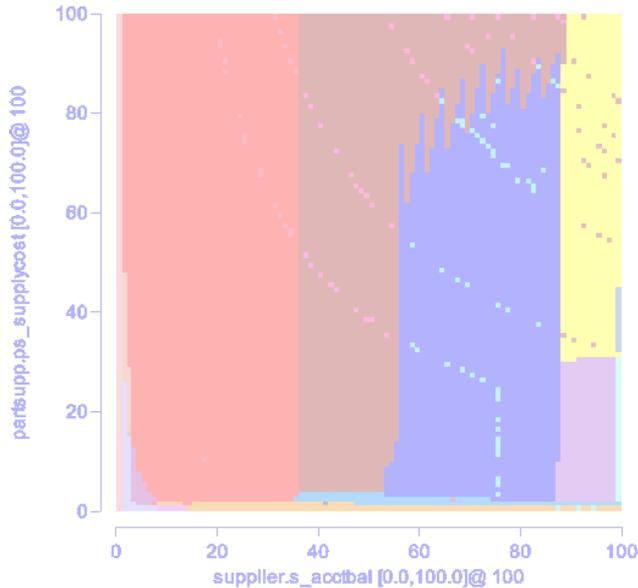


Better: reduce #plans for PQO

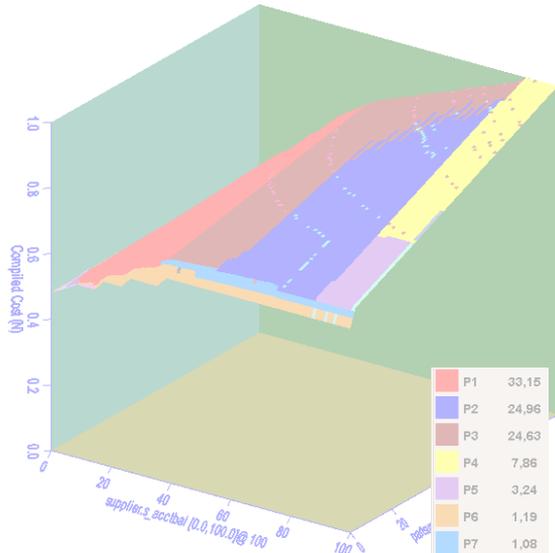
Conclusion

- Often highly intricate diagrams
- Typically 80% of space covered by 20% of plans (postgresql: 37%)
- Cardinality of plan diagram can be reduced, without materially affecting the query cost
- Assumptions of parametric query optimization literature do not hold in practice. (Reduction needed)
- Needed:
 - Mechanism for pruning the plan search space
 - Directly reduce #plans by optimizer

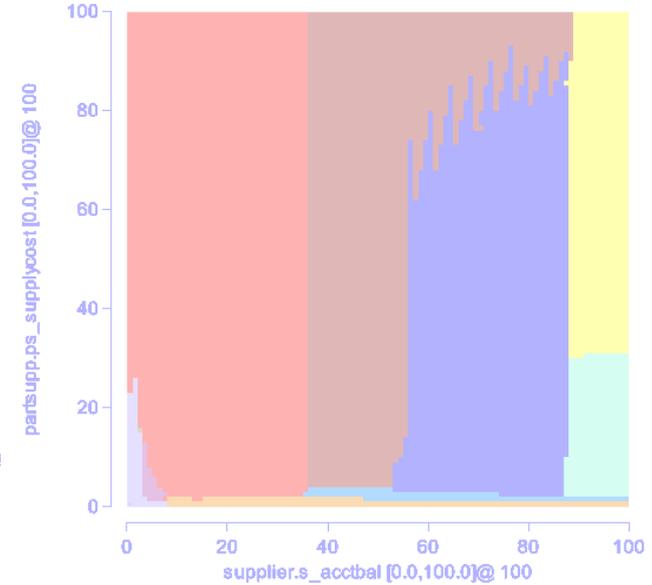
Discussion



Plan Diagram



Cost Diagram



Reduced Plan Diagram

Do you have any questions?



P1	33,15
P2	24,96
P3	24,63
P4	7,86
P5	3,24
P6	1,19
P7	1,08
P8	1,00
P9	0,48
P10	0,47
P11	0,36
P12	0,35
P13	0,29
P14	0,27
P15	0,22
P16	0,20
P17	0,13
P18	0,06
P19	0,03
P20	0,01
P21	0,01
P22	0,01

Comparison of Database Analyzers

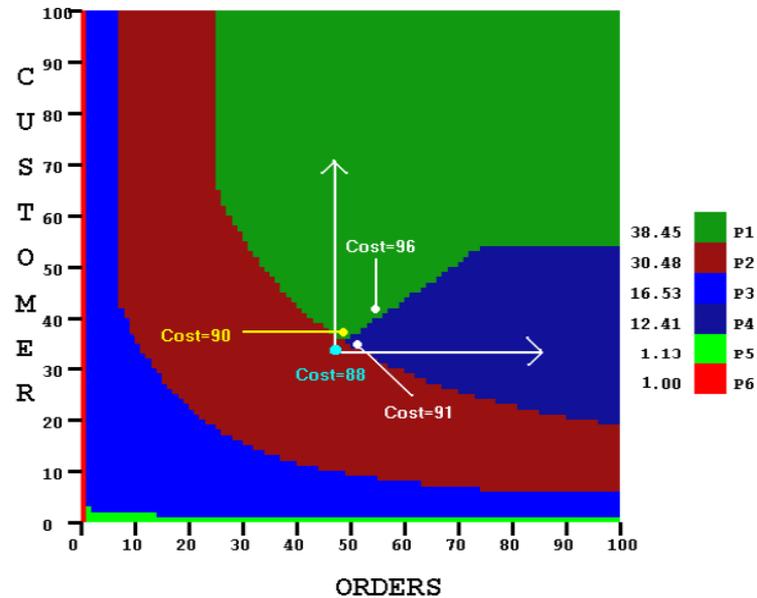
TPC-H Query	#Plans	%plans for 80%	Gini Index
2	2	100,00%	0,47
5	8	37,50%	0,65
7	4	50,00%	0,57
8	9	33,33%	0,68
9	22	13,64%	0,76
10	9	22,22%	0,48
18	7	14,29%	0,24
21	4	25,00%	0,28
postgresql	8,13	37,00%	0,52

Measurement of statistical dispersion

OptA	28,7	17,00%	0,79
OptB	24,5	23,00%	0,72
OptC	28,8	16,00%	0,8

- OptX: Commercial DBs tested by IIS
 - postgresql: Same dataset, similar reference system
- postgresql uses smaller #Plans to cover space.
- postgresql needs higher percentage of plans to cover 80% of selectivity space.
- postgresql's Gini Index is close to 0.5 which is better than 0.7 or 0.8.

Generation of Reduced Plan Diagram

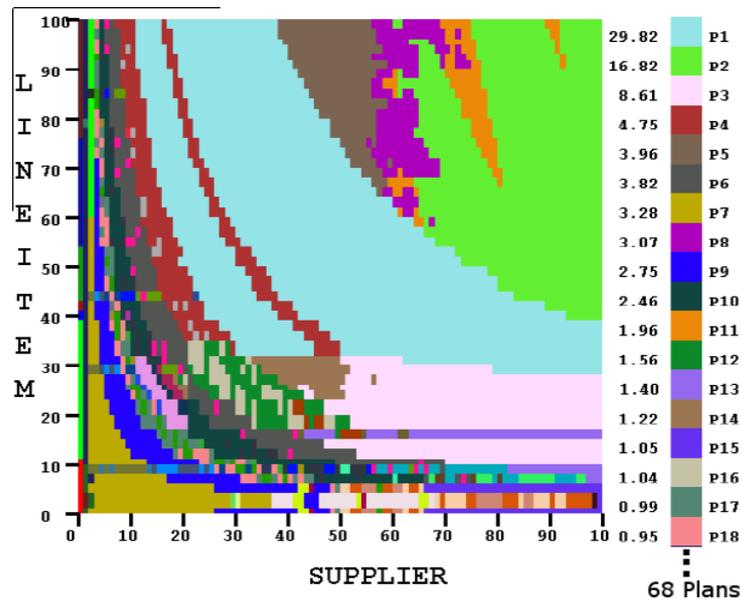


- P1, P4 are upper bounds of P2
- We move point to P1 which has lowest cost (90)
- If all Points of a plan are swit

Query plan costs

- The order of the operations in a query plan is important:
- $\text{cost}(AB) \neq \text{cost}(BA)$
- $\text{cost}(ABC) \neq \text{cost}(A) * \text{cost}(BC)$
- $\text{cost}(ABC) \neq \text{cost}(AB) * \text{cost}(C)$

Complex patterns



- Rapidly alternating choices
- not created on postgresql