

DrillBeyond: Processing Multi-Result Open World SQL Queries

Julian Eberius*, [Maik Thiele](#), Katrin Braunschweig and Wolfgang Lehner
Technische Universität Dresden, Germany

Motivation

DOMAINS OF...

- ..information retrieval and
- database systems...

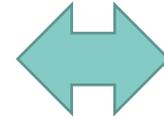
...HAVE BEEN TRADITIONALLY KEPT SEPARATE

DIFFERENCES IN...

- | | | | |
|--------------------------------|------------------------|--------|----------------------------------|
| ▪ type of data that is managed | (structured | versus | unstructured) |
| ▪ query language used | (fully specified query | versus | keyword query) |
| ▪ nature of the query result | (exact single answer | versus | ranked list of possible answers) |
| ▪ usage scenarios | (analytic scenarios | versus | information gathering) |

WE EXPLORE A NEW WAY...

- ..of merging the two paradigms for use in ad-hoc and self-service analytics



DRILLBEYOND = NOVEL HYBRID DBMS/IR SYSTEM

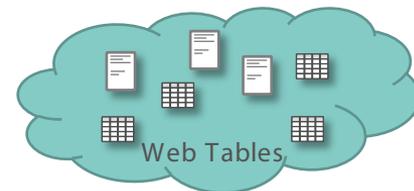
- Blurring the lines between type of data managed, query language used, and nature of the query result

```
select n_name, gdp, avg(o_totalprice)
from nation, customer, orders
where
  n_nationkey=c_nationkey
  andc_custkey=o_custkey
  and gdp > 10.0
group by n_name, gdp
order by gdp desc
```

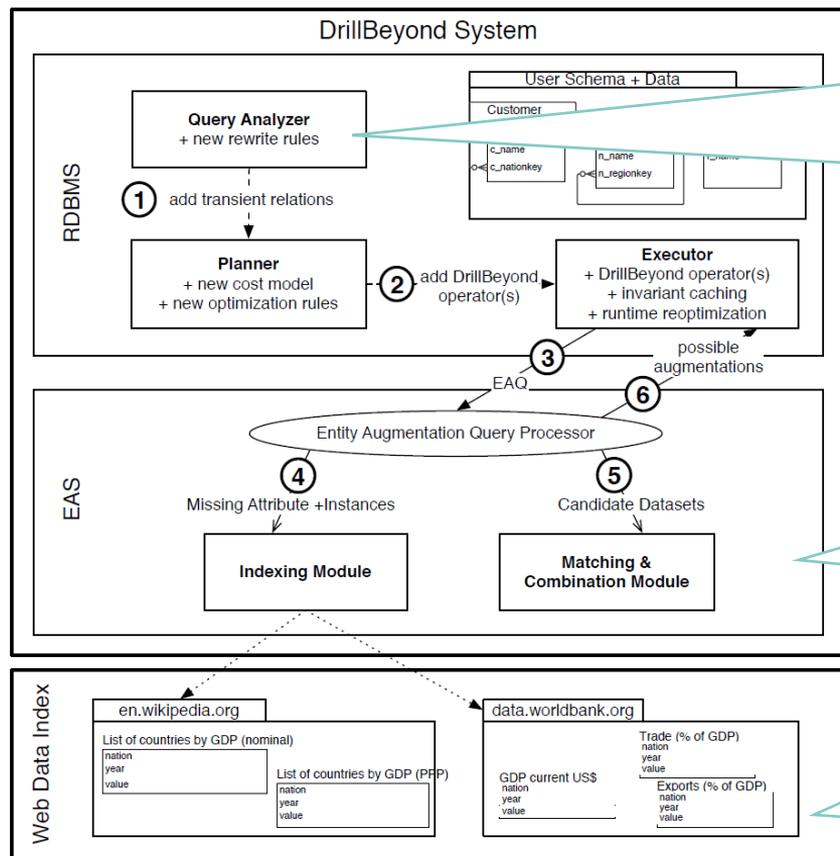


MIX OF...

- ..relational queries on a database and...
- ..top- k keyword- based searches (= Entity Augmentation Queries)



DrillBeyond – System Architecture



QUERY ANALYZER

- Maps SQL query tokens to the database catalog
- For unrecognized tokens (e.g. "gdp") we introduce transient metadata
- Query is rewritten to include an additional join with a transient relation

ENTITY AUGMENTATION SYSTEM

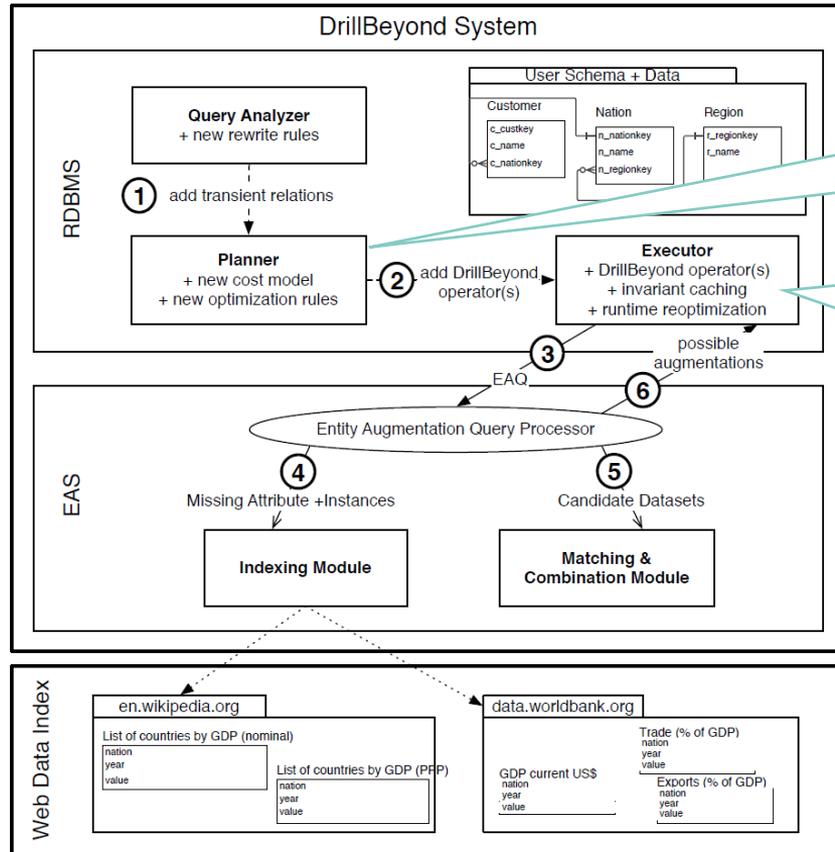
- Implements the entity augmentation processing

WEB DATA INDEX

- Index of Web tables
- Generic system exposing an interface for keyword-based document search

talk yesterday

DrillBeyond – System Architecture (2)



QUERY PLANNER

- Minimize the overhead of creating multiple result variants

EXECUTOR

- Implements the repeated execution of operator trees using the DrillBeyond operator ω

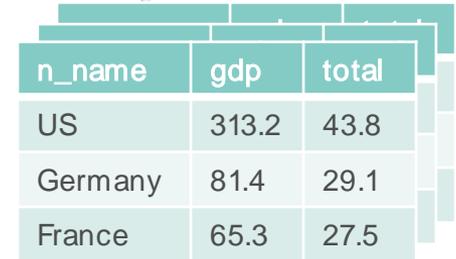
System Integration Challenges

MULTI-SOLUTION QUERY PROCESSING

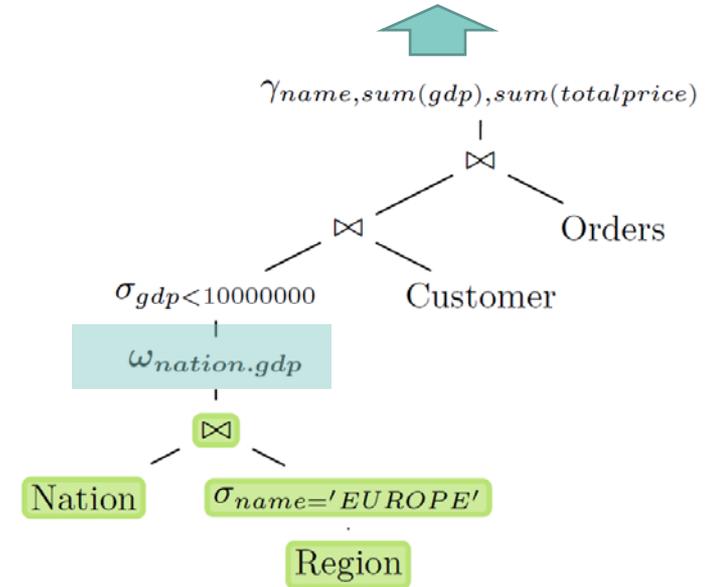
- Top-k query result versions for a single user query
- Naïve way: process the query k times → query runtime increased by factor k
- Goals:
 - Minimizes duplicate work between query executions
 - **Maximize invariant parts** of multiple executions

QUERY PLANNING

- Web tables are not fully known at plan-time
→ no selectivity information or data types at plan-time
- Goal: plan queries with relations only known at run-time



n_name	gdp	total
US	313.2	43.8
Germany	81.4	29.1
France	65.3	27.5



DrillBeyond Operator ω

INIT()

- Initializes state

```
function INIT  
  state  $\leftarrow$  'collecting'  
  tuplestore  $\leftarrow$   $\emptyset$   
  augMap  $\leftarrow$  HashMap()  
  n  $\leftarrow$  0
```

NEXT()

- Produces augmented tuples in three phase
 - Collect()
 - Augment()
 - Project()

```
function NEXT  
  if state = 'collecting' then  
    COLLECT()  
    AUGMENT()  
    state  $\leftarrow$  'projecting'  
  return PROJECT()
```

DrillBeyond Operator ω (2)

COLLECT()

- Pulls and stores all tuples
- Blocking

```
function COLLECT
  while true do
    t  $\leftarrow$  NEXT(childPlan)
    if t = NULL then
      break
    tuplestore  $\leftarrow$  t
    augKey  $\leftarrow$  TEXTATTRS(t)
    if augKey  $\notin$  augMap then
      augMap[augKey]  $\leftarrow$   $\emptyset$ 
```

AUGMENT()

- Pass all entries from the EAS into a hashtable

```
function AUGMENT
  augReq  $\leftarrow$  ( $\forall k \in$  augMap | augMap[k] =  $\emptyset$ )
  for all augKey, [augValues...]  $\in$  SEND(augReq) do
    augMap[augKey]  $\leftarrow$  [augValues...]
```

DrillBeyond Operator ω (3)

PROJECT()

- Produces output tuples by replaying the stored tuples and filling the augmentation attribute

RESCAN()

- Called when subtrees have to be re-executed

NEXTVARIANT()

- Produces the multi-variant query results

function PROJECT

$t \leftarrow \text{NEXT}(\text{tuplestore})$

if $t = \text{NULL}$ **then return** NULL

$\text{augKey} \leftarrow \text{TEXTATTRS}(t)$

$t[\text{augAttr}] = \text{augMap}[\text{augKey}][n]$ **return** t

function RESCAN

$\text{state} \leftarrow \text{'collecting'}$

$\text{tuplestore} \leftarrow \emptyset$

function NEXTVARIANT

RESCAN(tuplestore)

$n \leftarrow n + 1$

Why Blocking?

EXAMPLE

- Send each tuple to the augmentation system at its own (tuple-at-a-time)
- Augmentation system may choose ds_1 , ds_2 and ds_4

country	gdp
Russia	?
UK	?
USA	?

MORE CONSISTENT: DS_1 AND DS_3

FOR QUALITY REASONS THE DRILLBEYOND OPERATOR...

- ..needs to be a blocking operator
- Realized in the “collecting” state in function Next()
- Consumes tuples from underlying operators until these are exhausted
- Hands them over to the augmentation system

ds_1 American Countries

country	gdp m USD
USA	X
Canada	Y

ds_3 World GDPs

country	gdp Mil. \$
UK	X
Russia	Z

ds_2 European Economy

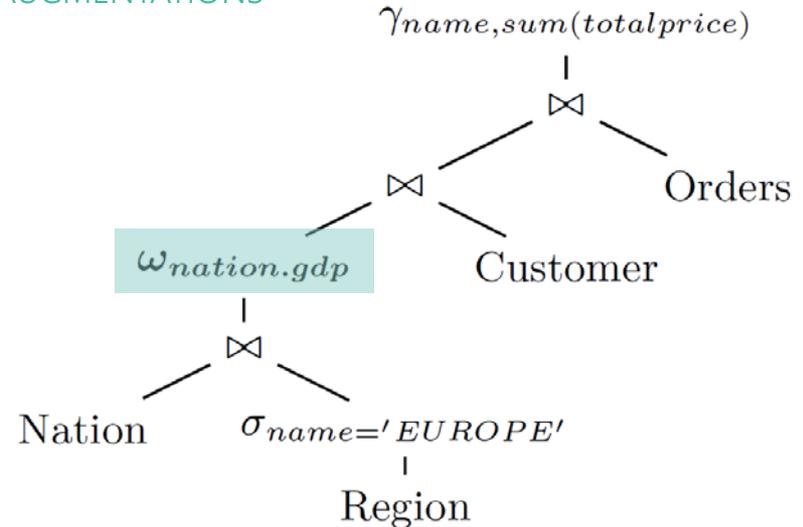
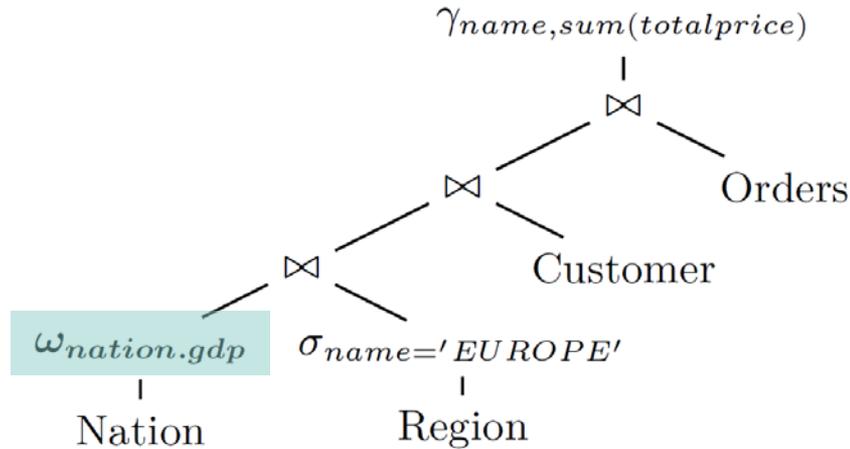
country	gdp EUR
Russia	X
France	Y

ds_4 European Economy

country	gdp EUR
UK	X

Lower Placement Bound

SAME QUERY, TWO DIFFERENT QUERY PLANS \rightarrow TWO DIFFERENT AUGMENTATIONS



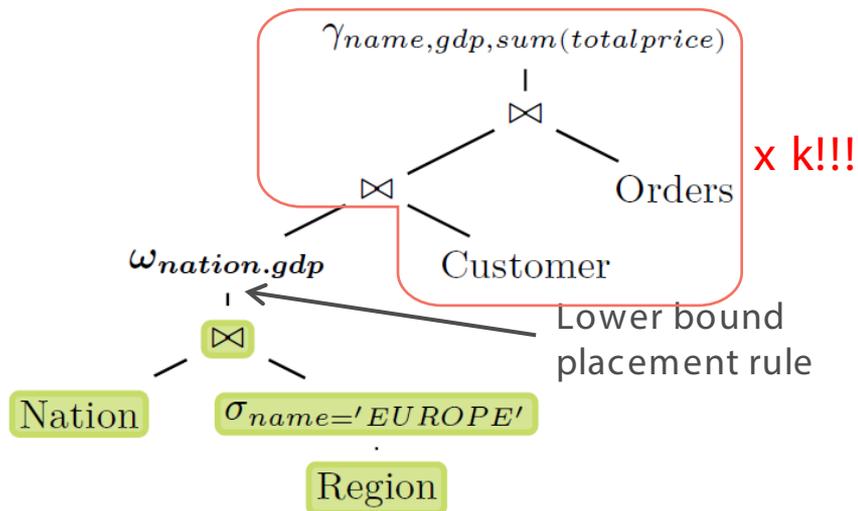
LOWER PLACEMENT BOUND

- ω can only be placed when all filters on R, e.g., joins and predicates, have been applied
- In other words: apply ω to the minimum number of entities in R

Maximize Invariant Sub-Trees

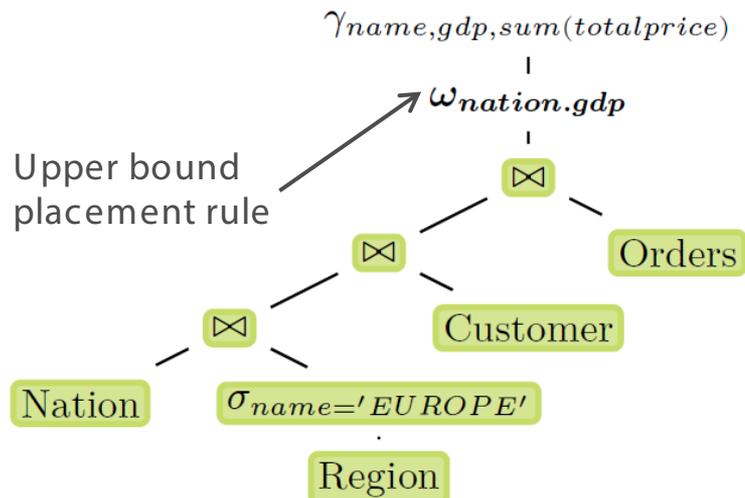
OPTIMAL PLAN...

- ..with respect to a single query execution



OPTIMAL PLAN...

- ..with respect to multiple query executions
- Upper bound placement rule: Place ω not earlier than the augmented values need to be accessed



Maximize Invariant Sub-Trees (2)

OBSERVATION

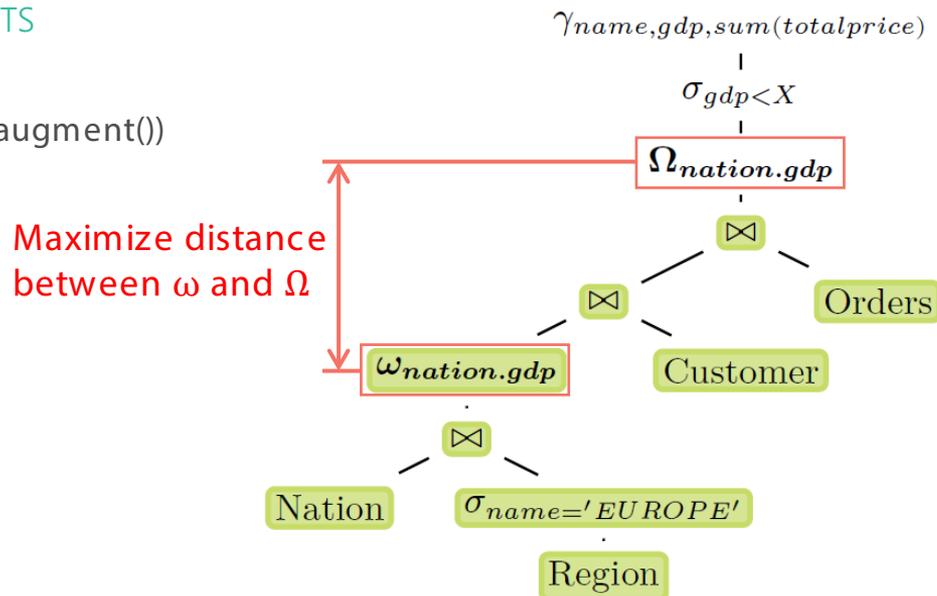
- We can separate the input part of ω from the actual projection of augmented values

SPLIT THE DRILLBEYOND OPERATOR INTO TWO PARTS

- ω
 - Performing the augmentation (collect() and augment())
 - Projects placeholders
 - Placed at the lower placement bound
- Ω
 - Performing the projection of values
 - Dereferences to the correct value array
 - Placed at the upper placement bound

→ MINIMIZING AUGMENTATION OPERATOR COST

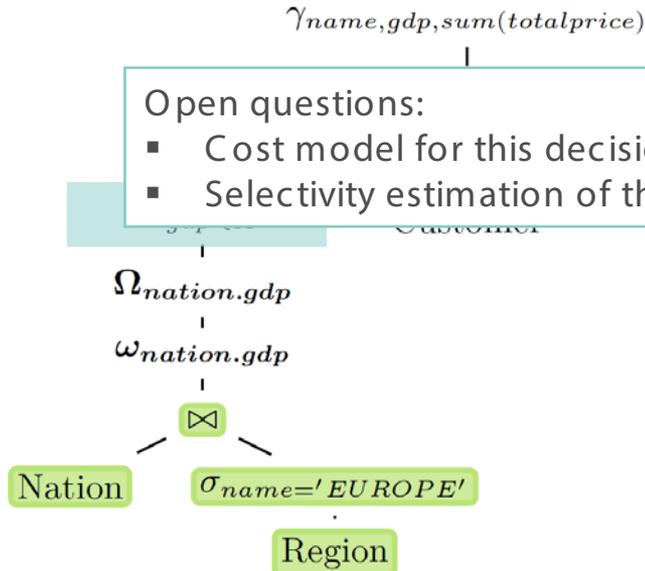
→ MAXIMIZING SIZE OF INVARIANT QUERY PARTS



Dynamic Selection Pull-Up

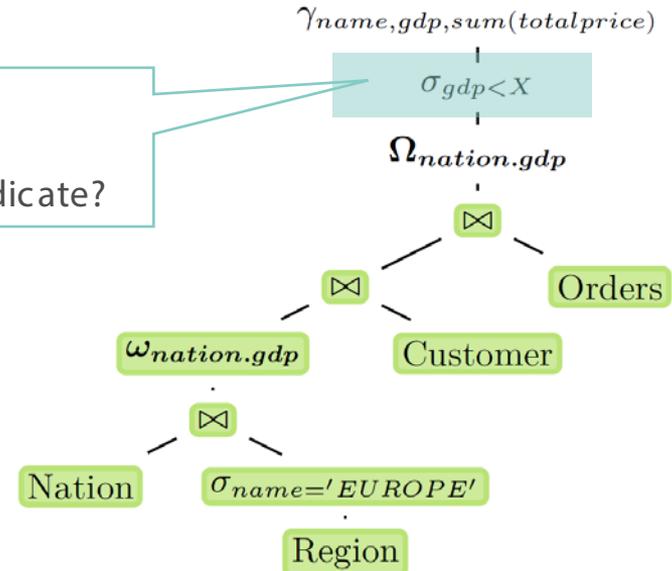
SELECTION PUSH-DOWN

- Smaller intermediate join result...
- ..but more varying nodes in the plan



SELECTION PULL-UP

- Place Ω as describe before
- Larger invariant subplan...
- ..but increased join results



Dynamic Selection Pull-Up (2)

COST MODEL

- Binary decision → decide for the minimum cost of the subtree above ω

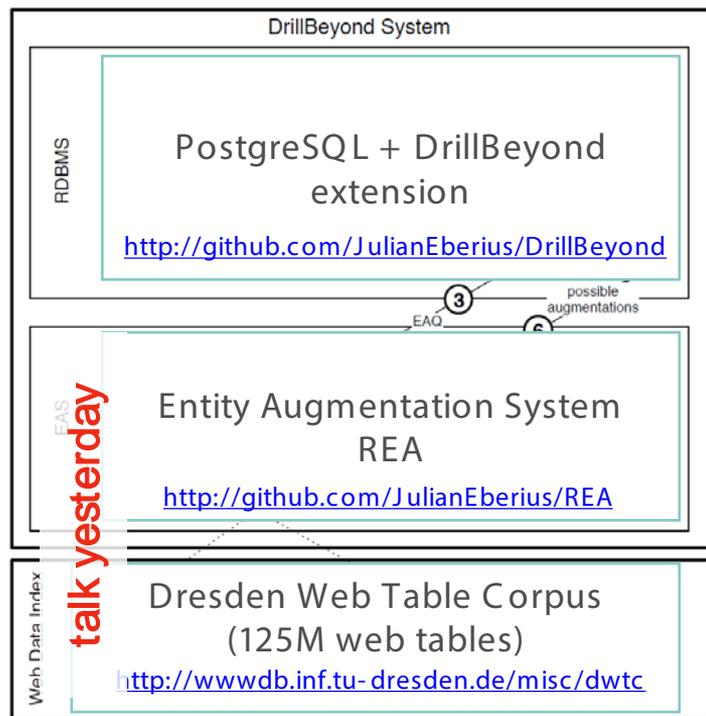
$$C_k(\omega, \top) = \begin{cases} C_1(\omega, \Omega) + \sum_{i=1}^k s_i \cdot C_1(\omega, \top) & (\text{Pull - up}) \\ \sum_{i=1}^k s_i \cdot C_1(\omega, \top) & (\text{no Pull - up}) \end{cases}$$

- $C(x, y)$: cost of the subplan from operator x to y
- s_i
 - Selectivity of the i^{th} data source for the attribute in question
 - Not known at plan-time → determined in the `augment()` phase



Evaluation

IMPLEMENTATION



TEST DATABASE

- Variation of TPC-H replacing generic identifiers with real-world entities, e.g. real company names for the Supplier relation

TEST QUERIES

- Subset of TPC-H queries in which dimension tables are used
- Added arbitrary where-clauses to one of the dimensions: "relation.X > Y"

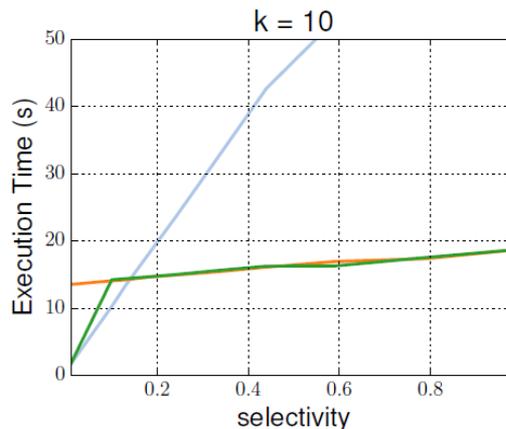
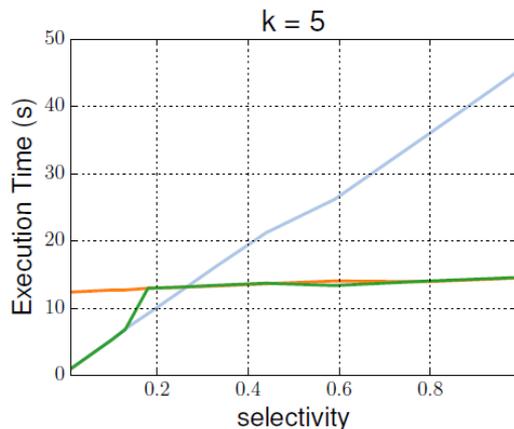
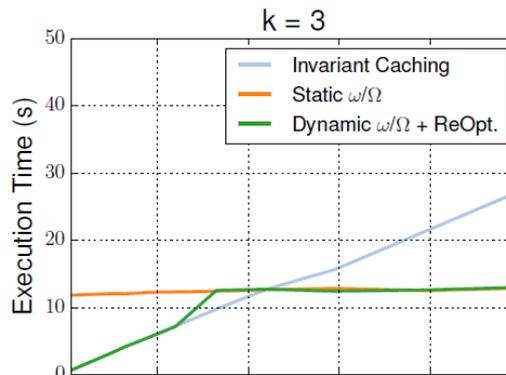
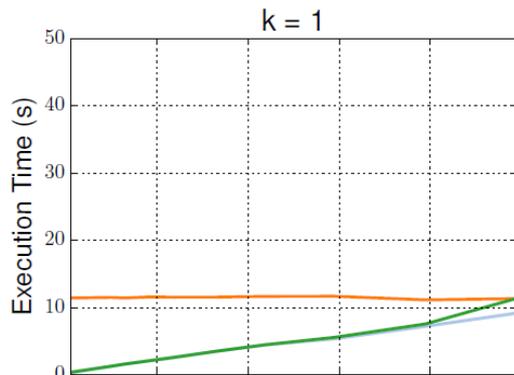
PARAMETERS

- Number of augmentations $k \in (1, 3, 5, 10)$
- Selectivity s ranging between 0.01 and 0.99 in ten steps

Overall Performance



Performance of query 9, by selectivity



Summary

DRILLBEYOND A RDBMS / IR HYBRID SYSTEM

- Integration of top-k entity augmentation system into a RDBMS
- Multiple alternative query result solving the uncertainty and ambiguity of the automated integration



NEW OPERATORS: DRILLBEYOND ω/Ω

- Implemented in traditional interface functions: Init(), Next() and ReScan()
- Own cost model

NEW OPTIMIZATION STRATEGIES

- Invariant Caching
- Maximization of invariant sub-trees
- Selection push-down versus selection pull-up (runtime reoptimization)

