

The (Potential) Perks of Integrating Provenance Support into Database Engines

Boris Glavic, Gopal Shankar¹

Illinois Institute of Technology

DBGroup



GCASR 2013, May 3, 2013

¹Master student

- 1 Motivation
- 2 Research Directions
- 3 Improving Performance with Provenance-aware Operators
- 4 Conclusions



Given a piece of data

- How do we know ...
 - which data it is derived from?
 - which transformations (SQL) where used to create it?
 - who created it?
 - ...

Example

		result	
		shop	rev
t ₁		Migros	125
t ₂		Coop	25



Given a piece of data

- How do we know ...
 - which data it is derived from?
 - which transformations (SQL) where used to create it?
 - who created it?
 - ...

Example

Compute the **revenue** for each **shop** as **sum** of **prices** of **items** sold

Example

result		
	shop	rev
t ₁	Migros	125
t ₂	Coop	25

```
SELECT shop,  
       sum(price) AS rev  
FROM sales, items  
WHERE itemId = id  
GROUP BY shop
```

sales		items		
	shop	itemId	id	price
s ₁	Migros	1	i ₁	100
s ₂	Migros	3	i ₂	10
s ₃	Coop	3	i ₃	25



Use Cases

- Debugging (tracking the sources of errors)
- Propagating annotations
- Gain deeper understanding of data and transformations
 - Estimate quality, trust
- Access Control (PBAC)
- Supporting technology (e.g., Probabilistic databases, Creating test databases)

Application Domains

- Complex database queries, e.g., datawarehousing
- E-science and curated databases
- Data integration

Annotation Propagation

- Annotate inputs
- Propagate these annotations during query processing
 - Combine/filter annotations based on data-dependencies of operators

Example

- Annotate inputs with tuple IDs (singleton set)
- Join: Set union
- Duplicate Removal: Set union

Example

	Result		
t	<table><thead><tr><th>a</th></tr></thead><tbody><tr><td>1</td></tr></tbody></table>	a	1
a			
1			



```
SELECT DISTINCT R.a  
FROM R, S  
WHERE R.a = S.c;
```

	R			S	
	a	b		c	d
r ₁	1	2	s ₁	1	2
r ₂	3	6	s ₂	1	3



Example

- Annotate inputs with tuple IDs (singleton set)
- Join: Set union
- Duplicate Removal: Set union

$$\begin{aligned} & (\{r_1\} \cup \{s_1\}) \cup (\{r_1\} \cup \{s_2\}) \\ & = \{r_1, s_1, s_2\} \end{aligned}$$

Example

	Result
	a
$\{r_1, s_1, s_2\}$	1



```
SELECT DISTINCT R.a
FROM R, S
WHERE R.a = S.c;
```

	R		S	
	a	b	c	d
r_1	1	2	1	2
r_2	3	6	1	3



Relational encoding of provenance

- 1 Standard relations
 - Extend query result tuples with complete input tuples in provenance (or only ID columns)
 - Duplicate tuples to fit in provenance
- 2 New data types (DBMS extensibility mechanism)
 - Use new/existing complex data type to represent provenance

Example (Provenance)

$\{ \langle r_1, s_1 \rangle, \langle r_1, s_2 \rangle \}$

Example

	Result		
	<table><thead><tr><th>a</th></tr></thead><tbody><tr><td>1</td></tr></tbody></table>	a	1
a			
1			
t			



```
SELECT DISTINCT R.a  
FROM R, S  
WHERE R.a = S.c;
```


	R			S	
	a	b		c	d
r ₁	1	2	s ₁	1	2
r ₂	3	6	s ₂	1	3



Relation with Provenance + Query Result Data

- **Data:** result tuple + all tuples from input tables in provenance
 - Result tuple might have to be duplicated!
- **Schema:** + input attributes (renamed)

Example (Query Result)



Relation with Provenance + Query Result Data

- **Data:** result tuple + all tuples from input tables in provenance
 - Result tuple might have to be duplicated!
- **Schema:** + input attributes (renamed)

Example (Provenance Representation)

R		S			Result + Provenance				
a	b	c	d	q	a	P(a)	P(b)	P(c)	P(d)
1	2	1	2	\rightarrow	1	1	2	1	2
3	4	1	3		1	1	2	1	3



Relation with Provenance + Query Result Data

- **Data:** result tuple + all tuples from input tables in provenance
 - Result tuple might have to be duplicated!
- **Schema:** + input attributes (renamed)

Example (Provenance Representation)


R		S			Result + Provenance				
a	b	c	d	q	a	P(a)	P(b)	P(c)	P(d)
1	2	1	2	\rightarrow	1	1	2	1	2
3	4	1	3		1	1	2	1	3



Relation with Query Result Data + Extra Attribute

- **Data:** result tuple + new attribute to store provenance
 - Value of new attribute: complete provenance for this output
- **Schema:** + new provenance attribute


Example (Query Result)



Relation with Query Result Data + Extra Attribute


- **Data:** result tuple + new attribute to store provenance
 - Value of new attribute: complete provenance for this output
- **Schema:** + new provenance attribute

Example (Provenance Representation)



Query rewrite - Standard relations

- Take original query q and rewrite into q^+ :
 - Attach provenance to inputs (on-the-fly or join existing)
 - Propagate through operations
 - \Rightarrow Compute original results + provenance





Advantages

- Reuse of existing DB technology
- DBMS independence (to some extent)


Disadvantages

- Unnecessarily complicated and inefficient computation
 - Impossible to propagate provenance through some operators
 - \Rightarrow need to duplicate parts of the query
- Hard to benefit from overlap and structure in provenance
- Complex queries that are hard to optimize



Query rewrite - Single Value

- Take original query q and rewrite into q^+ :
 - Attach provenance to inputs (on-the-fly or join existing)
 - User-defined functions (UDF) to combine propagated provenance





Advantages

- Reuse of existing DB technology
- DBMS independence (to some extent)

Disadvantages

- Unnecessarily complicated and inefficient computation
 - Impossible to propagate provenance through some operators
 - DBMS not optimized for very large attribute values
 - Optimizer has to handle provenance computations as black-boxes
- Hard to benefit from overlap and structure in provenance



Implications of loose integration of provenance

- Severe performance limitations
 - SQL and DBMS not designed for provenance \Rightarrow complex queries
 - Different data-flow (mostly append only, less reduction) \Rightarrow relational execution engine not a good fit
 - Provenance structure predictable and overlapping \Rightarrow no exploited
- \Rightarrow hinder broad adaptation
 - Access Control (**performance critical**)
 - Supporting technology, e.g., probabilistic data (**performance critical**)

What can be gained by full integration of provenance in DBMS engine?

- Provenance-aware operators
- Provenance optimized data-flow during query processing
- Performance optimized compression techniques

- 1 Motivation
- 2 Research Directions**
- 3 Improving Performance with Provenance-aware Operators
- 4 Conclusions



Tight Integration of Provenance Computation in the DBMS core

- Provenance-aware operators
 - Simplify queries \Rightarrow easier to optimize
 - Direct propagation without “rerouting” \Rightarrow less redundancy
- Compression schemes exploiting provenance structure
 - Faster compression/decompression \Rightarrow applicable during provenance computation
 - Execute operations on compressed provenance directly
 - Approximate provenance
- Investigate non-traditional data-flows for provenance
 - Cache provenance between queries
 - Share provenance between operators \Rightarrow divert from traditional operator base processing
 - ...



Tight Integration of Provenance Computation in the DBMS core

- **Provenance-aware operators**
 - Simplify queries \Rightarrow easier to optimize
 - Direct propagation without “rerouting” \Rightarrow less redundancy
- Compression schemes exploiting provenance structure
 - Faster compression/decompression \Rightarrow applicable during provenance computation
 - Execute operations on compressed provenance directly
 - Approximate provenance
- Investigate non-traditional data-flows for provenance
 - Cache provenance between queries
 - Share provenance between operators \Rightarrow divert from traditional operator base processing
 - ...



Rational

- Overlap in provenance
- Avoid generating large provenance for intermediates if not used in final provenance
 - Performance gain!

Approach

- Use declarative descriptions as placeholders for actual provenance
- Only expand when necessary
- Suited to present provenance to user?



Example

```
SELECT count(*), R.b
  (SELECT R.a
   FROM R
   WHERE EXISTS (SELECT * FROM S WHERE S.c > 5) AND R.a < 10)
GROUP BY R.b
HAVING count(*) > 1000
```

- Each output will have all tuples from S with $S.c > 5$ in its provenance
- Propagation will attach this potentially large set of tuples
 - to every tuple from R
 - many of these tuples may belong to groups that will be filtered eventually
 - \Rightarrow unnecessary work



Example

```
SELECT count(*), R.b
  (SELECT R.a
   FROM R
   WHERE EXISTS (SELECT * FROM S WHERE S.c > 5) AND R.a < 10)
GROUP BY R.b
HAVING count(*) > 1000
```

Provenance

a	Prov
1	$\{r_1, s_1, s_3, s_5, \dots\}$
13	$\{r_5, s_1, s_3, s_5, \dots\}$
...	...

Using Declarative Summarization

a	Prov
1	$\{r_1, \sigma_{c>5}(S)\}$
13	$\{r_5, \sigma_{c>5}(S)\}$
...	...





- 1 Motivation
- 2 Research Directions
- 3 Improving Performance with Provenance-aware Operators**
- 4 Conclusions



Data-flow and access patterns

- Provenance computation \neq Regular queries
 - Tuple breadth constantly increasing
 - Coupling of input data with intermediate data
 - Overlap

Direct propagation through operators not always possible

- Adding additional rows to input of an aggregation
- \Rightarrow change result!





Approach

- Operators optimized for provenance specific tasks
- Support direct propagation of provenance
- Implement and evaluate new operators in **Perm**
 - Perm is a relational provenance system based on query rewrite
 - Implemented as an extension to Postgres

Rationale

- Less complex queries
- Specialized code
- Less redundancy in computation



Problem

- Cannot propagate provenance through aggregation directly
 - Would change aggregation results
- Compute aggregation provenance
 - Compute the original aggregation
 - Compute provenance for query below aggregation
 - Join these two
- Complexity
 - Each level of aggregation
 - Doubles query size
 - Adds additional join
 - ⇒ Exponential in levels of aggregations!




Case study: AggProject - A Provenance-aware Aggregation Operator

Solution

- Develop a new type of aggregation that caches input provenance and attaches it to aggregation results





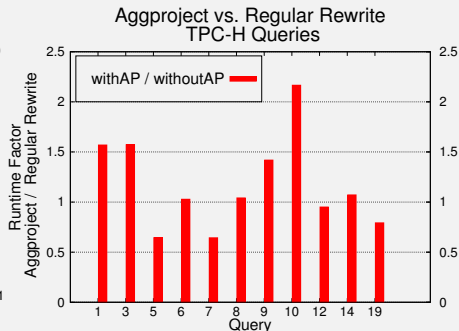
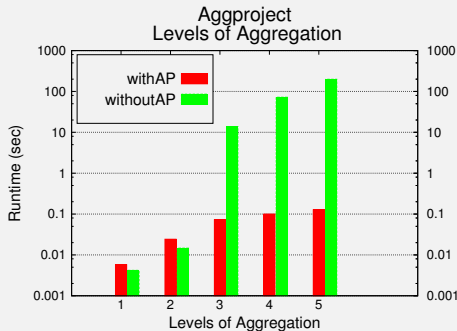
Implementation in Perm

- Add `AGGPROJECT` as new clause to SQL
- Physical `aggproject` operator using aggregation by sorting
- “Hack” integration with query optimizer
- Integration with provenance query rewrites
- Support multiple levels of aggregation



Setup

- TPC-H Benchmark Dataset (1GB)
- Perm vs. Perm with Aggproject



AggProject

- Develop cost-model
- Implemented Hash-based variant
 - In-memory
 - With disk caching

Provenance-aware Operators

- Set operations
- Nested subqueries



- 1 Motivation
- 2 Research Directions
- 3 Improving Performance with Provenance-aware Operators
- 4 Conclusions**



- Provenance become more and more important
- State-of-the-art of relational provenance systems unsatisfactory
 - Query rewrite, custom data types, stored procedures
 - Seriously performance limitations
 - Reduces adaptation for use cases that benefit from provenance
- Integration of provenance in DBMS engine core
 - Potential for addressing the performance bottleneck



- Continue provenance-aware operators approach
- Efficient provenance representations and declarative summarizations
 - Develop model and equivalences
 - Cost model and performance analysis
- Non-traditional data flows
 - Caching provenance
 - Operators shared provenance data structures
 - Indexing



- **Homepage:** `http://www.cs.iit.edu/~glavic/`
- **DBGroup:** `http://www.cs.iit.edu/~dbgroup/`
- **Perm:**
 - Project page:
`http://www.cs.iit.edu/~dbgroup/research/perm.php`
 - Svn repository: `https://permdbms.svn.sourceforge.net/svnroot/permdbms/trunk`

