



Guilherme Fetter Damasio

University of Ontario Institute of Technology and IBM Centre for Advanced Studies







Outline

- Introduction
- Relational Database
- Graph Database
- Our Research





Introduction

Data is everywhere





© Dopyright 2013 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice.





Introduction







Introduction

- How can we store all this data?
 - Database



Edgar Codd invents the Relational Data Model, and its first order theory. IBM team implements System R. 1970





- Relational Database
 - Highly structured data
 - Represents information in tables with rows (tuples) and columns (attributes)
- References to tuples in other table by referring to their primary-key (PK) attributes via foreign-key (FK) columns

PERSON						
ID	NAME	CAR_ID				
1	John	1				
5	Brad	5				
1		<u> </u>				
PK		FK				

CAR							
ID	TYPE	MAX_NUM_OF_PAS					
1	Mercedes	2					
3	BMW	5					
<u> </u>	·	·					
PK							





- Structured Query Language (SQL)
 - Standard language for storing, manipulating and retrieving data in databases
 - Very-high-level language
 - Query optimization
- SQL example: find everyone who has a BMW car with max number of passengers = 5

```
SELECT NAME FROM "PERSON"

WHERE (

PERSON."CAR_ID" IN (

SELECT ID FROM CAR

WHERE (

CAR."TYPE" = 'BMW' AND

CAR."MAX_NUM_OF_PAS" = 5)

)

)

)
```





- Joins are computed at query time by matching primary- and foreign-keys of the many rows of the to-be-joined tables
- These operations are compute- and memory-intensive and have an exponential cost
- A lot of the data available is naturally represented as graphs
 - Data is related to resources
 - Example: Social network data







How can we improve the performance for this type of data?











- Relationships are first-class citizens of the graph data model
- Graph databases enable us to build sophisticated models that map closely to our problem domain
- Each node (entity or attribute) in the graph database model directly and physically contains a list of relationship-records that represent its relationships to other nodes
- Whenever you run the equivalent of a JOIN operation, the database just uses this list and has direct access to the connected nodes, eliminating the need for a expensive search / match computation





- Support a very flexible and fine-grained data model
 - Manage rich domains in an easy and intuitive way







- Relational model to Graph model
 - Each entity table is represented by a label on nodes
 - Each row in a entity table is a node
 - Columns on those tables become node properties
 - Foreign keys become relationships
- Connections between nodes is directly link in such a way that relating data becomes a simple matter of following connections
- Helps companies to make sense of the masses of connect data that exist today





- Resource Description Framework (RDF)
 - W3C specification
 - Standard model for data interchange on the web
- The language description model allows the creation of statements about a resource either by defining its relationships with other resources or by defining its attributes
- Model is composed by triples
 - Subject
 - Predicate
 - Object
- Subjects, predicates and objects are normally composed by a URI and a node identification





Example:

A person called "John" that has a car "BMW", where this car has maximum number of passengers equals to "5".







- RDF can be represented in different formats:
 - N-triples, turtle, RDF/XML, etc.

N-triples representation of the previous example:

<http://myLocal/person/John> <http://myLocal/hasName> "John" . <http://myLocal/person/John> <http://myLocal/hasCarType> <http://myLocal/car/BMW> . <http://myLocal/car/BMW> <http://myLocal/hasType> "BMW" . <http://myLocal/car/BMW> <http://myLocal/hasMaxNumberOfPassengers> "5"^^<http://www.w3.org/2001/XMLSchema#integer> .

- RDF can also be represented by quadruples
 - Graphs





- There is no direct way to represent properties in relationships, but it can be accomplished by *Blank Nodes*
 - Resource for which a URI or literal is not given
- Can make use of ontologies
 - Vocabularies in a standard format
 - Establishes the relationship between variables
- Ex: Integration of data coming from different publishers
 - The data can be imported into a common RDF model, eg, by using converters to the publishers' databases. However, one database may use the term "author", whereas the other may use the term "creator"
- To make the integration complete, and extra definition should be added to the RDF data
 - Describing the fact that the relationship described as "author" is the same as "creator"





SPARQL

- SPARQL is a recursive acronym for SPARQL Protocol and RDF Query Language
 - Semantic query language for databases able to retrieve data stored in RDF format
 - Support by W3C
- SPARQL query that returns every person who has a BMW with maximum number of passengers as 5

```
PREFIX predicate: <<u>http://myLocal/</u>>
SELECT ?name
WHERE
{
    ?x predicate:hasName ?name ?name ?name ?name ?roar predicate:hasCarType ?car .
    ?car predicate:hasType "BMW" .
    ?car predicate:hasMaxNumberOfPassengers 5
}
```





SPARQL

Property paths

- elt* \rightarrow zero or more occurrences of *elt*
- elt+ \rightarrow one or more occurrences of *elt*
- elt? \rightarrow A path of zero or one *elt*
- elt1 / elt2 \rightarrow A sequence path of *elt1*, followed by *elt2*
- $^{\text{elt}} \rightarrow$ Inverse path (object to subject)

```
PREFIX predicate: <<u>http://myLocal/</u>>
SELECT ?a ?b ?c
WHERE
```

```
{
```

- ?a predicate:hasType "NLJOIN" .
- ?a predicate:hasLeftChild+ ?b .
- ?b predicate:hasType "TBSCAN" .
- ?b predicate:hasGenericChild ?c .
- ?c predicate:hasType "BASEOB" .







Graph Database Pros

- Flexibility
 - The data captured can be easily changed and extended for additional attributes and objects
- Data relationship exploration
 - No join index lookup (just follow connections)
 - Traverse millions of nodes per second
 - "Which supplier provided the products owned by this group of customers?"
- Indexing
 - Graph databases are naturally indexed by relationships (the strength of the underlying model), providing faster access compared to relational data for data





Graph Database Cons

- Not efficient at
 - Processing high volumes of transactions
 - Handling queries that span the entire database
- For most common graph databases, you have to store all the data on one server
- Still fairly new compared to relational software, which has now existed for a full generation
 - It takes time to build a solid database market after all, regardless of data model
- Some graph databases only offer the graph model, but the underlying implementation is backed by a traditional





Relational Database x Graph Database

- Which one is better?
 - Depends

Graph Database

If your domain entities have relationships to other entities, and your queries rely on exploring those relationships

Relational Database

If you have large-volume analytics queries typical of data warehousing





"It is important to note the consequence of using graph databases. The query latency in a graph is proportional to how much of the graph you choose to explore in a query, and is not proportional to the amount of data stored."

Jim Webber, author of *Graph Databases*





Query Performance Problem Determination with Knowledge Base in Semantic Web System OptImatch

Guilherme Fetter Damasio

Jaroslaw Szlichta Piotr Mierzejewski Calisto Zuzarte

University of Ontario Institute of Technology and IBM Centre for Advanced Studies







Outline

- I. Background
- II. Motivation
- III. Architecture
- IV. Transforming Diagnostic Data
- V. Searching for Problem Patterns
- VI. Providing Recommendations with Knowledge Base
- VII. Experimental Study
- VIII. Current Work





Outline

I. Background

- II. Motivation
- III. Architecture
- IV. Transforming Diagnostic Data
- V. Searching for Problem Patterns
- VI. Providing Recommendations with Knowledge Base
- VII. Experimental Study
- VIII. Current Work





Background

DB2

- Relational Database
 Management System
- Query tuning
- Access Plan or Query Execution Plan (QEP)

QEP

- Contains a graphical format portion represented in a tree structure
- Contains a text format portion







Background

				1					1	
				0.00655366	←	Estimate nun	nber of rows		0.00655366	
				~NLJOIN	←	Туре			^NLJOIN	
				(133)	←				(162)	
				142.555	←	Cumulative t	otal cost		142.555	
				25.1664	←				25.1664	
			/	+		-\	U COSI	/	+	\
			0.00664213		0	.986681		0.00664213		0.986681
			NLJOIN			IXSCAN		NLJOIN		IXSCAN
			(134)			(153)		(163)		(182)
			136.875			16.9821		136.875		16.9821
			24.1664			3		24.1664		3
		/	+	\		1	/		\	1
		0.00021628		30.7108	9.	73094e+06	0.00021628		30.7108	9.73094e+06
		NLJOIN		IXSCAN	INDE	X: CITICOM	NLJOIN		IXSCAN	INDEX: CITICOM
		(135)		(152)	HIER	TO CARD 00	(164)		(181)	HIER TO CARD 00
		131.176		17.0017		Q128	131.176		17.0017	Q45
		23.1664		3		See Providence	23.1664		3	1000
	/	+	\	1		/	+	\	1	
	0.000283536		0.762794	9.8623e+06		0.000283536		0.762794	9.8623e+06	
	NLJOIN		IXSCAN	INDEX: CITIC	MO	NLJOIN		IXSCAN	INDEX: CITIC	MO
	(136)		(151)	CARDS 2HIER MA	P 5P	(165)		(180)	CARDS 2HIER MA	P 5P
	125.497		16.9809	Q130	1007 Refe	125.497		16.9809	Q46	8. 7 T . 861
	22.1664		3			22.1664		3	62	
/	+	\	1	1		+	\	1		
0.00035736		0.793418	321134	0.00035	736		0.793418	321134		
NLJOIN		IXSCAN	INDEX: CITICO	M NLJOI	N		IXSCAN	INDEX: CITICO	M	
(137)		(150)	HIERLOOK7 04	(166)		(179)	HIERLOOK7 04		
119.819		16.9812	Q131	119.8	19		16.9812	Q47		
21.1664		3	0.000	21.16	64		3	10-10-10-10-10-1		
1	2			<i>,</i> , , , , , , , , , , , , , , , , , ,		0	1			8





Background

8.	333	33
TΒ	SCA	N
(2)
54	77.	33
51	16.	67

2)	TBSCAN: (Ta	ble Scan)	
	Cumulative	Total Cost:	5477.33
	Cumulative	CPU Cost:	1.52058e+009
	Cumulative	I/O Cost:	5116.67
	Cumulative	Re-Total Cost:	2690.61
	Cumulative	Re-CPU Cost:	7.83255e+008
	Cumulative	Re-I/O Cost:	0
	Cumulative	First Row Cost:	5477.33
	Estimated	Bufferpool Buffer	s: 0
	Arguments:		
	MAXPAGES:	(Maximum pages for	r prefetch)
	ALL		
	PREFETCH:	(Type of Prefetch)
	NONE		
	SCANDIR :	(Scan Direction)	
	FORWAR	D	
	SPEED : SLOW	(Assumed speed of	scan, in sharing structures)
	THROTTLE: FALSE	(Scan may be thro	ttled, for scan sharing)
	VISIBLE : FALSE	(May be included	in scan sharing structures)
	WRAPPING: FALSE	(Scan may start a	nywhere and wrap)





Outline

I. Background

II. Motivation

- III. Architecture
- IV. Transforming Diagnostic Data
- V. Searching for Problem Patterns
- VI. Providing Recommendations with Knowledge Base
- VII. Experimental Study
- VIII. Current Work





Motivation

- World's most valuable data remains in relational databases
- Complex analytic queries are a part of business operations
- Database systems are becoming more sophisticated
- Currently, tools for query performance problem determination have limitations, so performance analysis is often best done by manually analysis





Motivation

- Existing tools provide running recommendations for specific known problems
 - IBM[®] Optim Query Tuner[®], IBM Optim Workload Tuner[®], etc.
- Some characteristics of query execution plan (QEPs) are not easily found by the use of search tools like "grep"
- Query performance problem determination is valuable inside IBM support of business clients and database optimizer development organization
- Optimizations and performance tuning strategies becomes necessary to maintain the usability of the database





Motivation

- Incorporating OptImatch many optimization problems could be automatically identified and resolved
- Text graph version of a snippet of a QEP from IBM DB2
- Problem: NLJOIN has inner stream of type table scan (TBSCAN)
 - Costly
 - NLJOIN scans entire inner TABLE each rows from the outer
- Solution: create an index of the target table of the TBSCAN

5751.8	(19860.9				
SJOTN		NLJOIN				
<u>141</u>)		<u>(142</u>) 16246.59				
	1000000	4909.624				
	/	++	\			
	19.13	2	4043			
	FETC	FETCH (143)				
	(14					
	26.0884		15771			
	2.6	2.624				
	/+	\	1			
	19.12	1228	812130			
	IXSCAN	SALES FACT	CUST DIM			
	(<u>144</u>) 11708.7	Q2	Q1			
	5250 I					
	9.18948	e+07				
	IDX1 Q	2				
		A.				





Outline

- I. Background
- II. Motivation

III. Architecture

- IV. Transforming Diagnostic Data
- V. Searching for Problem Patterns
- VI. Providing Recommendations with Knowledge Base
- VII. Experimental Study
- VIII. Current Work





Architecture







Outline

- I. Background
- II. Motivation
- III. Architecture

IV. Transforming Diagnostic Data

- V. Searching for Problem Patterns
- VI. Providing Recommendations with Knowledge Base
- VII. Experimental Study
- VIII. Current Work





Transforming Diagnostic Data

Parser

- Define the property structure and property name for the RDF
- RDF
 - Represents the QEP
 - RDF is supported by DB2
 - RDF does not enforce specific schema, but it can be enforced by specifying predicates and establishing the relationship between LOLEPOPs

Subject	Predicate	Object
		\neg

<http://explainPlan/PlanPop/2> <http://explainPlan/PlanPred/hasPopType> "NLJOIN" .
<http://explainPlan/PlanPop/2> <http://explainPlan/PlanPred/hasOuterInputStream> <http://explainPlan/PlanPop/3> .
<http://explainPlan/PlanPop/2> <http://explainPlan/PlanPred/hasInnerInputStream> <http://explainPlan/PlanPop/5> .
<http://explainPlan/PlanPop/3> <http://explainPlan/PlanPred/hasPopType> "FETCH" .
<http://explainPlan/PlanPop/3> <http://explainPlan/PlanPred/hasEstimatedCardinality> "19.12" .
<http://explainPlan/PlanPop/5> <http://explainPlan/PlanPred/hasEstimatedCardinality> "19.12" .
<http://explainPlan/PlanPop/5> <http://explainPlan/PlanPred/hasEstimatedCardinality> "4043.0".
<http://explainPlan/PlanPop/5> <http://explainPlan/PlanPred/hasIonInputLeg> "INNER" .
<http://explainPlan/PlanPop/5> <http://explainPlan/PlanPred/hasIonInputStreamPop> <http://explainPlan/PlanBaseObject/CUST_DIM> .
<http://explainPlan/PlanBaseObject/CUST_DIM> <http://explainPlan/PlanPred/hasEstimateCardinality> "812130.0"</http://explainPlan/PlanPred/hasEstimateCardinality> "812130.0"</http://explainPlan/PlanPred/hasEstimateCardinality> "812130.0"</http://explainPlan/PlanPred/hasEstimateCardinality> "812130.0"</http://explainPlan/PlanPred/hasEstimateCardinality> "812130.0"</http://explainPlan/PlanPred/hasEstimateCardinality> "812130.0"</http://explainPlan/PlanPred/ha





Transforming Diagnostic Data

SPARQL

- RDF language
- Property paths
- Supported by DB2

Handlers

- Result (Retrieval)
- Internal (Filtering)
- Relationship
- Blank Node

PREFIX predURI: http://explainPlan/PlanPred/ SELECT (?pop1 AS ?TOP) (?pop2 AS ?ANY2) (?pop4 AS ?BASE4) WHERE [?pop1 predURI:hasPopType "NLJOIN" . ?pop1 predURI:hasOuterInputStream ?BNodeOfpop2_to_pop1 . ?BNodeOfpop2_to_pop1 predURI:hasOuterInputStream ?pop2. ?pop2 predURI:hasOutputStream ?BNodeOfpop2_to_pop1. ?BNodeOfpop2_to_pop1 predURI:hasOutputStream ?pop1 . ?pop1 predURI:hasInnerInputStream ?BNodeOfpop3_to_pop1. ?BNodeOfpop3_to_pop1 predURI:hasInnerInputStream ?pop3. ?pop3 predURI:hasOutputStream ?BNodeOfpop3_to_pop1. ?BNodeOfpop3_to_pop1 predURI:hasOutputStream ?pop1 . ?pop3 predURI:hasPopType "TBSCAN". ?pop3 predURI:hasEstimateCardinality ?internalHandler1. FILTER (?internalHandler1 > 100). ?pop3 predURI:hasInputStream ?BNodeOfpop4_to_pop3. ?BNodeOfpop4_to_pop3 predURI:hasInputStream ?pop4 . ?pop4 predURI:hasOutputStream ?BNodeOfpop4_to_pop3. ?BNodeOfpop4_to_pop3 predURI:hasOutputStream?pop3. ?pop4 predURI:isABaseObj ?internalHandler2. } ORDER BY ?pop1

{"id":"hasEstimateCardinality", _____ "value":"1","sign":"<"}</pre> Pop1 predURI:hasEstimateCardinality ?internalHandler1 FILTER (?internalHandler1 < 1)</p>

{"id":"hasEarlyOutFlag","value":"LEFT"} --> ?pop1 predURI:hasEarlyOutFlag "LEFT"





Outline

- I. Background
- II. Motivation
- III. Architecture
- IV. Transforming Diagnostic Data

V. Searching for Problem Patterns

- VI. Providing Recommendations with Knowledge Base
- VII. Experimental Study
- VIII. Current Work





Searching for Problem Patterns

- Web-based graphical interface
- Upload Files
- Properties and relationship between nodes
- Pattern #1 example (High read cost)
 - (i) is of type "NLJOIN"
 - (ii) outer input stream of type "ANY" and cardinality > 1
 - (iii) has inner input stream of type "TBSCAN"
 - (iv) inner input stream has cardinality > 100
 - (v) "TBSCAN" has generic input stream of type "BASE OB"
- Solution
 - Create an index in the table read by the TBSCAN

Select Pop Type : BASE OB ✓ Estimate Cardinality: 100 > Actual Cardinality: Cardinality Discrepancy Factor: Name: Schema: Distinct Row Values: Tablespace Overhead:	NLJOIN.1 ANY.2 TESCAN:3
✓ Estimate Cardinality: 100 > ▼ Actual Cardinality: ✓ ✓ Cardinality Discrepancy Factor: ✓ ✓ Name: ✓ ✓ Schema: ✓ ✓ Buffer Poll Pages: ✓ ✓ Distinct Row Values: ✓ ✓ Tablespace Overhead: ✓ ✓	ANY2 TESCAN:3
✓ Estimate Cardinality: 100 > • Actual Cardinality: • • • Cardinality Discrepancy Factor: • • • Name: • • • • Schema: • • • • Buffer Poll Pages: • • • • Distinct Row Values: • • • • Tablespace Overhead: • • • •	ANY:2 TESCAN:3
Actual Cardinality: Cardinality Discrepancy Factor: Name: Schema: Buffer Poll Pages: Distinct Row Values: Tablespace Overhead:	ANY.2 TESCAN:3
Cardinality Discrepancy Factor: Vame: Schema: Shffer Poll Pages: Distinct Row Values: Tablespace Overhead: V 	ANY:2 TBSCAN:3
Name: Schema: Buffer Poll Pages: Distinct Row Values: Tablespace Overhead: V	
Schema: Buffer Poll Pages: Distinct Row Values: Tablespace Overhead:	
Buffer Poll Pages: * Distinct Row Values: * Tablespace Overhead: *	
Distinct Row Values:	
Tablespace Overhead:	BASE OB:4
	BROE OD.4
Tablespace Transfer Rate:	
Prefetch Page Count:	
Index Leaf Pages:	
Index Tree Levels:	
Index Full Key Card:	
Index Page Density:	
Column In Index:	





Searching for Problem Patterns

Query R	esult			0		
******* Time to Number Files M Files M ******	Analysi of Files atching atching *******	************* s: 0 hr, 0 m Analized: 2 Pattern: 2 Recommendati ***********	**************************************	**************************************	f 198 NLJ (162 490	60.9 OIN 142) 46.59 9.624
Result	for the :	file: 1000sq	1_3_ADVITO_HOTEL	_REPORT.sql.explain	19.12 FETCH	4043 TBSCAN
TOP	ANY2	TBSCAN3	BASE4		(143)	(145)
142	143	37	COUNTRY_CCRS		26.0884 2.624	4907
====== Result	for the	===== file: 1000sq	1_7CARD_HOLDER	_DEMOGRAPHICS_REPORT.sql.explain	19.12 1228 IXSCAN SALES_ (144) Q2 11708.7 5250	FACT CUST_DI Q1
TOP	ANY2	TBSCAN3	BASE4]	9,18948e+07	
4	5	47	COUNTRY_CCRS		IDX1 Q2	
Result TOP 4	for the ANY2	file: 1000sq TBSCAN3 47	1_7CARD_HOLDER BASE4 COUNTRY_CCRS	_DEMOGRAPHICS_REPORT.sql.explain	11708.7 5250 9.18948e+1 IDX1 Q2	07





Outline

- I. Background
- II. Motivation
- III. Architecture
- IV. Transforming Diagnostic Data
- V. Searching for Problem Patterns
- VI. Providing Recommendations with Knowledge Base
- VII. Experimental Study
- VIII. Current Work





- Populated with predetermined problem patterns and associated query plan recommendations
- Problem pattern \rightarrow SPARQL query with result handlers
- Use of alias to specify a result handler
- Use of handlers tagging

PREFIX predURI: <http://explainPlan/PlanPred/> SELECT (?pop1 AS ?TOP) (?pop2 AS ?ANY2) (?pop4 AS ?BASE4)

- Static and dynamic components
- Use of "@" sign to represent a result handler





- Use of 1 handler: @TOP
- Use of multiple handlers: [@TOP, @BASE4]
- Limit number of occurrences: [@TOP, @BASE4]:1
 List only 1 result
- Use of helper functions
- List columns: @TOP.listColumns("PREDICATE")
 List columns from an alias handler in the predicate
- Example: "Create index on table @BASE4 on columns @TOP.listColumns("INPUT")"





Query R	esult			
****** Time to Number of Files Mo Files Mo ******	Analysis Analysis of Files atching F atching F	s: 0 hr, 0 m Analized: 2 Pattern: 2 Recommendati	********************** in, 2 sec, 284 m 0 on: 2 **********	**************************************
Result	for the f	file: 1000sq	1_3_ADVITO_HOTEL	_REPORT.sql.explain
TOP	ANY2	TBSCAN3	BASE4	
9	10	37	COUNTRY_CCRS	
<pre>{Q2.DB_({Q2.MER({Q2.TRAI </pre>	CR_FLAG}, CH_ACCEPT NS_DATE}, for the f	<pre>{Q2.MCC(A) FOR_ID}, {Q2 {Q2.TXNSKE {Q2.TXNSKE } ;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;</pre>	<pre>), {Q2.MERCHANT_ .MERCH_COUNTRY}, Y}, {Q3.SHORT_NA 1_7CARD_HOLDER</pre>	_DEMOGRAPHICS_REPORT.sql.explain
TOP	ANY2	TBSCAN3	BASE4	
4	5	47	COUNTRY_CCRS	
RECOMME TAG: (RECOM (221.AD (221.AD (221.CA (221.CA (221.CA (221.DE (221.HA)	NDATIONS: MENDATION DENDATION DRESS_LIN DUNT_OVEF RDH_LAST_ ST_CENTEF PARTMENT] ERARCHY_N WEL. (021	I: ia tabela {C IE_1}, {Q21 LIMIT}, {Q NAME}, {Q21 DESCRIPTIO , {Q21.DISC IUM}, {Q21.I .OPEN DATE}	OUNTRY_CCRS} nas ADDRESS_LINE_2}, 21.AMOUNT_PAST_D .CARD_ACTIVATION N}, {Q21.COUNTRY RETIONARY_CODE}, D_VERIFICATION_C . {Q21.OTHER ACC	<pre>colunas {Q21.ACCOUNT_NUMBER}, {Q21.ACCOUNT_TYPE}, {Q21.ADDRESS_LINE_3}, {Q21.AMOUNT_OF_LAST_PAYMENT}, JE_CYCLE1}, {Q21.BILL_TYPE}, {Q21.CARDH_FIRST_NAME}, _DATE}, {Q21.CITY}, {Q21.COST_CENTER}, _NUM}, {Q21.CREDIT_LIMIT}, {Q21.CREDIT_RATING}, {Q21.EMAIL}, {Q21.EMPLOYEE_ID}, {Q21.FIRST_NAME}, DDE}, {Q21.LAST_NAME}, {Q21.LAST_PAY_DATE}, </pre>





- OptImatch can provide advanced guidance with variety of recommendations
- Pattern #2 (estimation of the execution cost by optimizer)
 - (i) LOLEPOP of type index Scan (IXSCAN) or table scan (TBSCAN)
 - (ii) has cardinality < 0.001
 - (iii) has a generic input stream of type Base Object (BASE OB)
 - (iv) the generic input stream has cardinality > 100000
- Recommendation:
 - Create column group statistics (CGS) on equality local predicate columns and CGS on equality join predicate columns of the Base Object.





- Pattern #3 (poor join order)
 - (i) LOLEPOP of type JOIN
 - (ii) has a descendant outer input stream of type JOIN
 - (ii) has a descendant inner input stream of type JOIN
 - (iii) the descendant outer input stream join is a Left Outer Join
 - (iv) descendant inner input stream join is a Left Outer Join
- Recommendation:
 - Rewrite the query from the following structure (T1 LOJ T2) ... JOIN ... (T3 LOJ T4) to ((T1 LOJ T2).... JOINT3) LOJ T4 as the rewritten query is more efficient
- With all provided patterns and recommendations, a user, with no particular knowledge or training, can run a general test of all predetermined problem patterns against a given query workload





Outline

- I. Background
- II. Motivation
- III. Architecture
- IV. Transforming Diagnostic Data
- V. Searching for Problem Patterns
- VI. Providing Recommendations with Knowledge Base
- VII. Experimental Study
- VIII. Current Work





Experimental Study

- Focuses on 3 objectives:
 - Effectiveness using real IBM customer query workload.
 - Scalability and performance:
 - Sizes of the query workload,
 - Number of LOLEPOPs
 - Number of recommendations in the knowledge base.
 - Comparative study with manual search for patterns by experts
 - Time and precision
- Patterns used through experimental study:
 - Pattern #1
 - Pattern #2
 - Pattern #3





Performance and Scalability Study Size of Query Workload



- Search increases linearly with the number of QEP files
- For 1000 QEPs, time to perform the search is less than 30 seconds
- Pattern #3 takes longer because it is more complex, containing descendant nodes (using recursion to analyze)





Performance and Scalability Study Number of LOLEPOPs



- For patterns #1 and #2 the search increases linearly with the number LOLEPOPs
- Even large and complex queries (with around 500 LOLEPOPs) can be processed efficiently by our tool (less than 1400 milliseconds)





Performance and Scalability Study Number of Recommendations in Knowledge Base



- OptImatch has linear dependency over the number of recommendations in the knowledge base
- Process a 1000 query workload against 1000 problem patterns and recommendations in around 3hr30min





Comparative User Study



- OptImatch drastically reduces the time to search for a pattern
- OptImatch does not only perform significantly faster than a manual search but it also guarantees correctness.





Outline

- I. Background
- II. Motivation
- III. Architecture
- IV. Transforming Diagnostic Data
- V. Searching for Problem Patterns
- VI. Providing Recommendations with Knowledge Base
- VII. Experimental Study
- VIII. Current Work





Current Work

- Extension of OptImatch
 - Automatically discover, based on previously knowledge, of the best QEP to be applied to a given query
 - Automatically recommendations such as changing database configuration in order to improve performance
- Knowledge Base
 - Information about QEP and queries
- Matching similarities





Current Work

- Normalization
 - Reduction of data redundancy
- 3 main pieces:
 - BaseData (tables and columns)
 - LinkData (sql) \rightarrow Ex: tables being joined, predicates used
 - Knowledge Base (our current RDF)

BaseData		LinkData
FableColumnIndex	SQL	SQL Table
rd Type Column[] lumn[] IsNull Length ages Length ages	NumOfTables Table[] Predicate[] BestPlanID	NumOfTables Table[] Predicate[] BestPlanID

	Knowledge	Base Record		Knowled	lge Base
Query	SPARQL	GUIDELINE	RDF_ID	ID	RDF
			ID	ID	BestPlanRDF





Current Work

- Old OptImatch: 894 ms
- New OptImatch: 1.52 ms



Average Time (Query 1)





Publications

EDBT

 G. Damasio, P. Mierzejewski, J. Szlichta and C. Zuzarte: Query Performance Problem Determination with Knowledge Base in Semantic Web System OptImatch. EDBT 2016, 515-526

ICDE

 G. Damasio, P. Mierzejewski, J. Szlichta and C. Zuzarte: OptImatch: Semantic Web System for Query Problem Determination. IEEE ICDE 2016, 1334-1337