

# OPTIMIZATION OF SPARQL BY USING CORESPARQL

Jinghua Groppe, Sven Groppe and Jan Kolbaum

*IFIS, University of Luebeck, Ratzeburger Allee 160, D-23538 Luebeck, Germany*

**Keywords:** Semantic Web, RDF, SPARQL, coreSPARQL, Query Optimization, Query Rewriting.

**Abstract:** SPARQL is becoming an important query language for RDF data. Query optimization to speed up query processing has been an important research topic for all query languages. In order to optimize SPARQL queries, we suggest a core fragment of the SPARQL language, which we call the coreSPARQL language. coreSPARQL has the same expressive power as SPARQL, but eliminates redundant language constructs of SPARQL. SPARQL engines and optimization approaches will benefit from using coreSPARQL, because fewer cases need to be considered when processing coreSPARQL queries and the coreSPARQL syntax is machine-friendly. In this paper, we present an approach to automatically transforming SPARQL to coreSPARQL, and develop a set of rewriting rules to optimize coreSPARQL queries. Our experimental results show that our optimization of SPARQL speeds up RDF querying.

## 1 INTRODUCTION

The Semantic Web uses the Resource Description Framework (RDF) (Beckett, 2004) as its data format to describe information in the web. RDF provides a model and syntax for describing data, but it does not provide querying functionalities. A number of RDF querying languages have been developed, e.g. RQL, N3, Versa, SeRQL, Triple and RDQL. When (Haase et al., 2004.) compares these six languages, SPARQL (Prud'hommeau and Seaborne, 2007) has not emerged. SPARQL was first proposed on 12th October 2004 and became an official W3C Recommendation on 15th January 2008. Many RDF stores support or plan to support SPARQL, e.g. Jena (Wilkinson et al., 2003) and Sesame (Broekstra et al., 2002). SPARQL becomes increasingly important as an RDF query language.

The optimization of queries has been an active research topic for improving the performance of query processing. An important optimization technique is rewriting of queries. While query rewriting has been extensively studied in the relational databases and XML areas, there is no complete and thorough work on rewriting of SPARQL queries. Therefore, we focus on the rewriting and simplification of SPARQL queries. In this paper we develop a core fragment of the SPARQL language to simplify SPARQL, which we name coreSPARQL, and a set of rules to optimize coreSPARQL queries.

SPARQL supports a large number of different language constructs, which brings flexibility of expressiveness, but also redundancy of expressions. For example, the three expressions of SPARQL in Figure 1 have the same semantics. Redundant expressive power increases the difficulties of query processing. It is also obvious that the syntax for Expression 1 is user-friendly, but Expression 3 is more easily to be interpreted by a machine.

Expression 1	Expression 2	Expression 3
(1 {?x 3}).	[] rdf:first 1; rdf:rest _:b. _:b rdf:first {\$x 3}; rdf:rest rdf:nil.	_:b1 rdf:first 1. _:b1 rdf:rest _:b2. _:b2 rdf:first _:b3. _:b3 ?x 3. _:b2 rdf:rest rdf:nil.

Figure 1: Three SPARQL expressions with same semantics.

In order to reduce the number of cases, which must be considered when processing SPARQL queries, and in order to make SPARQL queries more machine-processable, we suggest the coreSPARQL language, which is a core fragment of the SPARQL language. coreSPARQL possesses the same expressive power as SPARQL, but does not contain redundant

---

This work is funded by the German Research Foundation (DFG) project GR 3435/1-1 LUPOSDATE.

language constructs of SPARQL and only allows machine friendly syntax. We develop an approach, which automatically transforms SPARQL queries to coreSPARQL queries.

SPARQL queries written by users or generated by applications are often un-optimized, and thus sub-optimal. Sub-optimal queries impact query processing performance. Based on coreSPARQL, we develop a set of simplification rules to rewrite coreSPARQL queries, and transform a sub-optimal query into an optimal query by eliminating redundant parts and optimizing sub-expressions. Our performance study shows that after our optimization, SPARQL can be processed more efficiently, and the transformation of SPARQL to coreSPARQL has a low overhead. Due to the limitation of space, we do not present our experiment results in this paper.

**Related Work.** (Pérez et al., 2006) suggests several rules for rewriting AND, UNION and OPTIONAL expressions in SPARQL queries. The purpose of the rewriting is constructing a critical fragment of UNION-free graph pattern expressions for the study of evaluation complexity.

(Bernstein et al., 2007), (Groppe et al., 2007a), (Broekstra et al., 2002) and (Groppe et al., 2009) reorder triple patterns in order to reduce the size of intermediate results. (Groppe et al., 2007a) pushes a filter expression upward if all the variables in the filter expression has already been bound. (Bernstein et al., 2007) reorders triple patterns according to their selectivity, which is estimated based on schemas. (Broekstra et al., 2002) and (Groppe et al., 2009) both observe that the number of variables might impact the sizes of the intermediate resultant data. (Broekstra et al., 2002) reorders the triple patterns according to the number of variables, while (Groppe et al., 2009) considers the number of the new variables, which have not been bound so far, because the occurred variables are bound with the result of previous triple patterns.

An amount of work contributes to the rewriting of relational algebra, and develops a number of equivalency rules (Arasu et al., 2006) (Chaudhuri, 1998) (Ioannidis, 1996) (Jarke and Koch, 1984). Some of our and other equivalency rules for rewriting SPARQL queries are adapted from the equivalency rules for relational algebra, e.g. the rules for comparison operators.

Several contributions are dedicated to the transformation of SPARQL queries to SQL queries, and the storage of RDF data in relational databases, and thus use proven database technologies, e.g. (Chong et al., 2005), (Chebotko et al., 2007) and (Cyganiak, 2005).

(Groppe et al., 2007b), (Weiss et al., 2008) and (Groppe et al., 2009) suggest different indices for fast data access. (Groppe et al., 2009) develops a new approach to compute join of triple patterns by dynamically restricting triple patterns.

## 2 RDF AND SPARQL

Figure 2 presents an example of RDF data and of a SPARQL query.

RDF data is a set of triples of the form Subject Predicate Object, which are RDF terms, e.g. IRIs, literals or blank nodes. Figure 2 provides an example of RDF data with 3 triples. SPARQL selects RDF data based on graph pattern matching, where the core component of SPARQL graph patterns is a set of triple patterns  $s \ p \ o$ .  $s \ p \ o$  corresponds to the subject (s), predicate (p) and object (o) of a RDF triple, but they can be variables as well as RDF terms. A triple pattern matches a subset of the RDF data, where the RDF terms in the triple pattern correspond to the ones in the RDF data. The query result of a triple pattern consists of pairs of variables with their bound values, i.e. corresponding RDF terms in the matched subset of the RDF data. The result of a set of triple patterns is the join of the result of each triple pattern.

Book.rdf	Book.sparql
@prefix ex: <http://book/>	prefix ex: <http://book/>
ex:book1 ex:title "XML".	SELECT ?y, ?z
ex:book2 ex:title "Index".	WHERE { ?x ex:title ?y.
ex:book2 ex:pages 90.	?x ex:pages ?z.}

Figure 2: RDF data and SPARQL query.

The SPARQL query Book.sparql in Figure 2 consists of the SELECT clause and the WHERE clause. The SELECT clause identifies the variables to appear in the query results, and the WHERE clause contains two triple patterns, which identify the constraints on RDF data. The triple pattern  $?x \ ex:title \ ?y$  matches the first two triples of Book.rdf, such that its result is  $\{<?x=ex:book1, ?y="XML">, <?x=ex:book2, ?y="Index">\}$ . The triple pattern  $?x \ ex:pages \ ?z$  matches the last triple of Book.rdf, such that the result is  $\{<?x=ex:book2, ?z=90>\}$ . The two triple patterns impose a join over the common variable  $?x$ , such that the result of the two triple patterns is  $\{<?x=ex:book2, ?y="Index", ?z=90>\}$ . The final query result is  $\{<?y="Index", ?z=90>\}$ .

SPARQL provides rich capabilities to select and filter data, and we refer the interested reader to

(Prud'hommeaus and Seaborne, 2007) for a complete description of SPARQL.

### 3 CORESPARQL

SPARQL allows redundant language constructs and supports abbreviated syntax. The redundancy brings the flexibility of expressiveness and abbreviations bring the simplification of expressions, but they do not increase the expressive power of the language. That a SPARQL query can be expressed in different forms increases the number of cases to be processed; the abbreviated syntaxes are not machine-friendly. In order to make SPARQL queries more machine-processable, and to reduce the number of cases, which must be considered when processing SPARQL queries, we abstract a subset from the SPARQL language, and name the subset *coreSPARQL*.

#### 3.1 Defining coreSPARQL

In Definition 1, we describe coreSPARQL in terms of the common and different properties with SPARQL. Figure 3 demonstrates several SPARQL and corresponding coreSPARQL components.

component	SPARQL	coreSPARQL
triple pattern	s1 p1 o1; p2 \$x.	s1 p1 o1. s1 p2 ?x.
blank node []	[ p o].	_:b p o.
group graph pattern	{ {s1 p1 o1} s2 p2 o2. }	{ s1 p1 o1. s2 p2 o2. }
&& operator	Filter(A && B).	Filter(A). Filter(B).

Figure 3: SPARQL and corresponding coreSPARQL components.

**Definition 1 (coreSPARQL).** coreSPARQL is a core fragment of SPARQL. A coreSPARQL query is also a SPARQL query. coreSPARQL has the same expressive power as SPARQL, but allows only machine-friendly syntax, and eliminates many redundant language constructs. Especially, in coreSPARQL,

- all triple patterns are only in the form: s p o;
- a group graph pattern cannot directly nest another group graph pattern;
- variable names start only with ?;
- blank nodes [] are not allowed;
- RDF collections of the form (...) are not allowed;

- neither prefixed IRIs nor IRIs, which are relative to a BASE-declaration, are allowed.
- the keyword a is not allowed;
- the && operator is not allowed. □

#### 3.2 Transforming SPARQL to CORESPARQL

SPARQL provides user-friendly syntax to write RDF queries, and coreSPARQL queries are easy to program. Therefore, the next task for us is to find a way to automatically transform SPARQL queries to coreSPARQL queries. We develop a set of transformation rules, such that a SPARQL query can be transformed into a coreSPARQL query by recursive application of these rules, i.e. if the expression of a left-hand side of a rule occurs in a SPARQL query, it is replaced with the right-hand side of the rule.

We use the following notation to describe these rules: we write s (s1, s2,...), p (p1, p2, ...), o (o1, o2,...) for the subject, predicate, and object of a triple pattern, os (os1, os2, ...) for a list of objects, e.g. os = o1, o2, o3, ..., om, where m ≥ 1, and pos (pos1, pos2, ...) for predicate-object-lists, e.g., pos = p1 os1; p2 os2; ...; pm osm, where m ≥ 1. A blank node [] is replaced by a blank node label, e.g. \_:b, where b must be not used elsewhere in the query. Note that some patterns in the following rules may be not supported by SPARQL. Such patterns are intermediate results of the transformation, and will be translated to standard language constructs after the transformation.

- Rule 1: eliminating Object-Lists:
  - 1.1 s1 p1 o1, os. => s1 p1 o1. s1 p1 os.
- Rule 2: eliminating Predicate-Object-Lists:
  - 2.1 s1 p1 os1; pos. => s1 p1 os1. s1 pos.
- Rule 3: eliminating blank nodes [].
  - 3.1 [] => \_:b
  - 3.2 [ pos ]. => \_:b pos.
  - 3.3 [ pos ] p1 os1. => \_:b pos. \_:b p1 os1.
  - 3.4 s1 p1 [ pos ]. => s1 p1 \_:b. \_:b pos.
- Rule 4: eliminating RDF collections ( ), where e (e1, e2,...) is an element of the collection, i.e. a variable, a literal, a blank node, or a collection. Here, we introduce a variant of the collection, e.g. (e)<sub>s=:b</sub>. to restrict that the blank node, which is allocated for the collection (e), must be \_:b.
  - 4.1 (e) pos. => \_:b rdf:first e.  
\_:b rdf:rest rdf:nil.  
\_:b pos.
  - 4.2 (e). => \_:b rdf:first e. \_:b rdf:rest rdf:nil.
  - 4.3 (e1 e2 e3...). => \_:b rdf:first e1.  
\_:b rdf:rest (e2 e3...).





processing time and query cost. A query is unsatisfiable, if it contains conflicting constraints. For example, two constraints  $\text{IsIRI}(?x)$  and  $\text{Filter}(?x = \text{"http://example.com"})$  contradict each other:  $\text{IsIRI}(?x)$  requires that  $?x$  is an IRI, but  $\text{Filter}(?x = \text{"http://example.com"})$  requires that  $?x$  is a string.

In order to optimize queries and improve the evaluation performance, we develop a set of equivalency rules to detect conflicting and redundant constraints. By recursive application of these rules, a coreSPARQL query can be optimized to a more simple expression, or even to an empty expression, i.e. the query is unsatisfiable.

We use the rewriting rules in

Table 1 to simplify coreSPARQL queries, where  $C$  ( $C1, C2, \dots$ ) represents a literal,  $G$  a graph pattern or the query pattern, i.e. the outer-most graph pattern, and  $E$  ( $E1, E2, \dots$ ) an expression. Additionally, we introduce a new graph pattern: *void graph pattern*, denoted by  $\perp$ . Contrary to the empty group pattern  $\{\}$  in SPARQL, which matches any RDF graph, a void graph pattern  $\perp$  does not match any RDF graph. If a SPARQL query is simplified to the void graph pattern, the query is unsatisfiable. Note that  $\perp$  is an intermediate result during simplification, and any satisfiable SPARQL expressions will not contain  $\perp$  after optimization.

## 5 CONCLUSIONS

We suggest the coreSPARQL language, which is a core fragment of SPARQL, but has the same expressiveness as SPARQL. Optimization approaches, SPARQL engines and all applications, which process SPARQL queries, benefit from coreSPARQL, because coreSPARQL possesses machine-friendly syntax and thus is easy to program, contains less language constructs and thus reduces the number of cases to be considered.

We develop a set of transformation rules to translate SPARQL queries to coreSPARQL queries, and a set of rewriting rules to further optimize coreSPARQL queries. We develop a prototype of our approach, which shows that our optimization speeds up SPARQL query processing.

Table 1: Rewriting rules for optimizing coreSPARQL queries.

Eliminating the same components:	
<ul style="list-style-type: none"> <li><math>G \ G \Rightarrow G</math></li> </ul>	
Constant propagation:	
<ul style="list-style-type: none"> <li><math>\text{Filter}(?x=C). \text{Filter}(\dots?x\dots) \Rightarrow \text{Filter}(?x=C). \text{Filter}(\dots C\dots)</math>, if <math>?x</math> is not the parameter of a bound function.</li> </ul>	
E.g. $\text{Filter}(?x=10). \text{Filter}(?x>5) \Rightarrow \text{Filter}(?x=10). \text{Filter}(10>5)$ .	
E.g. $\text{Filter}(?x=\text{"work"}). \text{Filter}(\text{Lang}(?x)=\text{"EN"}) \Rightarrow \text{Filter}(\text{Lang}(\text{"work"}) = \text{"EN"})$ .	
Variable binding:	
<ul style="list-style-type: none"> <li><math>\text{Filter}(\text{bound}(?x)). \text{Filter}(\dots?x\dots) \Rightarrow \text{Filter}(\dots?x\dots)</math>, if <math>?x</math> of <math>\text{Filter}(\dots?x\dots)</math> is neither a parameter of a bound function nor inside an operand of <math>\parallel</math>.</li> <li><math>\text{Filter}(!\text{bound}(?x)). \text{Filter}(\dots?x\dots) \Rightarrow \perp</math>, if <math>?x</math> of <math>\text{Filter}(\dots?x\dots)</math> is neither a parameter of a bound function nor inside an operand of <math>\parallel</math>.</li> </ul>	
E.g. $\text{Filter}(\text{bound}(?x)). \text{Filter}(?x>10) \Rightarrow \text{Filter}(?x>10)$ .	
E.g. $\text{Filter}(!\text{bound}(?x)). \text{Filter}(?x=\text{"red"}) \Rightarrow \perp$ .	
Functions IsIRI, IsBlank, IsLiteral:	
<ul style="list-style-type: none"> <li><math>\text{IsIRI}(C) \Rightarrow \begin{cases} \text{true, if } C \text{ is an IRI;} \\ \text{false, if } C \text{ is not a IRI.} \end{cases}</math></li> </ul>	
E.g. $\text{IsIRI}(\text{mailto:alice@work.example}) \Rightarrow \text{true}$	
E.g. $\text{IsIRI}(\text{"mailto:alice@work.example"}) \Rightarrow \text{false}$	
The rules for the functions IsURI, IsBlank, IsLiteral are analogous to this one.	
Functions LangMatches, Regex:	
<ul style="list-style-type: none"> <li><math>\text{LangMatches}(C, L) \Rightarrow \begin{cases} \text{true, if } C \text{ matches } L; \\ \text{false, if } C \text{ does not match } L \end{cases}</math></li> </ul>	
E.g. $\text{LangMatches}(\text{"work@EN"}, \text{"EN"}) \Rightarrow \text{true}$	
E.g. $\text{LangMatches}(\text{"work"}, \text{"EN"}) \Rightarrow \text{false}$	
The rules for the function Regex are analogous to this one	
Function Lang:	
<ul style="list-style-type: none"> <li><math>\text{Lang}(C1@C2) \Rightarrow C2</math></li> </ul>	
E.g. $\text{Lang}(\text{"work"@EN}) \Rightarrow \text{"EN"}$	
Function Filter:	
<ul style="list-style-type: none"> <li><math>\text{Filter}(\text{false}) \Rightarrow \perp</math></li> <li><math>\text{Filter}(\text{true}) \Rightarrow \{\}</math></li> <li><math>C1 \text{ op } C2 \Rightarrow \begin{cases} \text{true, if } C1 \text{ op } C2 = \text{true;} \\ \text{false, if } C1 \text{ op } C2 = \text{false;} \end{cases}</math></li> </ul>	
E.g. $\text{Filter}(1>10) \Rightarrow \text{Filter}(\text{false}) \Rightarrow \perp$	
Elimination of $\perp$ and $\{\}$	
<ul style="list-style-type: none"> <li><math>G \text{ Optional } \perp \Rightarrow G</math></li> <li><math>G \text{ UNION } \{\} \Rightarrow G</math></li> <li><math>G \ \{\} \Rightarrow G</math></li> <li><math>G \text{ Graph } n \ \{\} \Rightarrow G</math></li> </ul>	<ul style="list-style-type: none"> <li><math>G \text{ Optional } \{\} \Rightarrow G</math></li> <li><math>G \text{ UNION } \perp \Rightarrow G</math></li> <li><math>G \ \perp \Rightarrow \perp</math></li> <li><math>\text{Graph } n \ \perp \Rightarrow \perp</math></li> </ul>
where $n$ is a variable or an IRI	
E.g. $\text{Graph } ?g \ \perp \Rightarrow \perp$	
E.g. $\{s \ p \ o\} \ \perp \Rightarrow \perp$	

Table 1: Rewriting rules for optimizing coreSPARQL queries (cont.).

Comparison operators:	
<ul style="list-style-type: none"> <li>• <math>\text{FILTER}(?x \text{ op1 } C1). \text{FILTER}(?x \text{ op2 } C2). \Rightarrow</math> <math display="block">\left\{ \begin{array}{l} \text{FILTER}(?x \text{ op1 } C1), \text{ if } ((\text{op1}=\text{op2} \wedge (C1 \text{ op1 } C2) \wedge \text{op1} \in \{&lt;, &lt;=, &gt;=, &gt;\}) \vee ((C1 \text{ op1 } C2) \wedge C1 \neq C2 \wedge \\ \text{op1}, \text{op2} \in \{&lt;, &lt;=, &gt;=, &gt;\}) \vee \text{op1}, \text{op2} \in \{&gt;=, &gt;\}) \vee (C1=C2 \wedge ((\text{op1}='&lt;' \wedge \text{op2}='&lt;=') \vee \\ \text{op1}='&gt;' \wedge \text{op2}='&gt;='))) \\ \text{FILTER}(?x \text{ op2 } C2), \text{ if } ((\text{op1}=\text{op2} \wedge (C2 \text{ op1 } C1) \wedge \text{op1} \in \{&lt;, &lt;=, &gt;=, &gt;\}) \vee ((C2 \text{ op2 } C1) \wedge C1 \neq C2 \wedge \\ \text{op1}, \text{op2} \in \{&lt;, &lt;=, &gt;=, &gt;\}) \vee (C1=C2 \wedge ((\text{op2}='&lt;' \wedge \text{op1}='&lt;=') \vee \\ \text{op2}='&gt;' \wedge \text{op1}='&gt;='))) \\ \text{FILTER}(\text{false}), \text{ if } ((C1 &gt; C2 \wedge \text{op1} \in \{&gt;, &gt;=, &gt;\}) \wedge \text{op2} \in \{&lt;, &lt;=, &lt;\}) \vee (C1 &lt; C2 \wedge \text{op1} \in \{&lt;, &lt;=, &lt;\}) \wedge \text{op2} \in \{&gt;, &gt;=, &gt;\}) \vee \\ (C1=C2 \wedge \text{op1} \neq \text{op2} \wedge (\text{op1}, \text{op2} \in \{=, !=\} \vee \text{op1}, \text{op2} \in \{&lt;, &gt;\})) \\ \text{FILTER}(?x \text{ op1 } C1). \text{FILTER}(?x \text{ op2 } C2), \text{ otherwise.} \end{array} \right.</math> </li> </ul>	
E.g. $\text{Filter}(?x > 10). \text{Filter}(?x > 30). \Rightarrow \text{Filter}(?x > 30); \text{Filter}(?x > 30). \text{Filter}(?x < 10). \Rightarrow \text{Filter}(\text{false}).$	
Operators $\parallel, !$ and $\neg$	
<ul style="list-style-type: none"> <li>• <math>E \parallel \text{true} \Rightarrow \text{true}</math>      • <math>\text{false} \parallel \text{false} \Rightarrow \text{false}</math>      • <math>E \parallel E \Rightarrow E</math>      • <math>!(A1 \text{ op } A2) \Rightarrow A1 \neg(\text{op}) A2</math></li> <li>• <math>?x \text{ op1 } C1 \parallel ?x \text{ op2 } C2 \Rightarrow</math> <math display="block">\left\{ \begin{array}{l} ?x \text{ op1 } C1, \text{ if } ((\text{op1}=\text{op2} \wedge (C2 \text{ op1 } C1) \wedge \text{op1} \in \{&lt;, &lt;=, &gt;=, &gt;\}) \vee ((C2 \text{ op1 } C1) \wedge C1 \neq C2 \wedge (\text{op1}, \text{op2} \in \{&lt;, &lt;=, &gt;=, &gt;\}) \vee \\ \text{op1}, \text{op2} \in \{&gt;=, &gt;\}) \vee (C1=C2 \wedge ((\text{op1}='&lt;' \wedge \text{op2}='&lt;') \vee (\text{op1}='&gt;' \wedge \text{op2}='&gt;'))) \\ ?x \text{ op2 } C2, \text{ if } ((\text{op1}=\text{op2} \wedge (C1 \text{ op1 } C2) \wedge \text{op1} \in \{&lt;, &lt;=, &gt;=, &gt;\}) \vee ((C1 \text{ op2 } C2) \wedge C1 \neq C2 \wedge (\text{op1}, \text{op2} \in \{&lt;, &lt;=, &gt;=, &gt;\}) \vee \\ \text{op1}, \text{op2} \in \{&gt;=, &gt;\}) \vee (C1=C2 \wedge ((\text{op2}='&lt;' \wedge \text{op1}='&lt;') \vee (\text{op2}='&gt;' \wedge \text{op1}='&gt;'))) \\ \text{Bound}(?x), \text{ if } (\text{op1}=\neg(\text{op2}) \wedge C1=C2), \\ ?x \text{ op1 } C1 \parallel ?x \text{ op2 } C2, \text{ otherwise} \end{array} \right.</math> </li> <li>• <math>\neg(=) \Rightarrow !=</math>      • <math>\neg(!=) \Rightarrow =</math>      • <math>\neg(&lt;) \Rightarrow &gt;=</math>      • <math>\neg(&lt;=) \Rightarrow &gt;</math>      • <math>\neg(&gt;) \Rightarrow &lt;=</math>      • <math>\neg(&gt;=) \Rightarrow &lt;</math></li> </ul>	

## REFERENCES

- Arasu, A., Babu, S., Widom, J., 2006. The CQL continuous query language: semantic foundations and query execution. *VLDB Journal*, 15(2): 121-142.
- Beckett, D. (editor), 2004. RDF/XML Syntax Specification (Revised), *W3C Recommendation*, 10th February 2004.
- Bernstein, A., Stocker, M., Kiefer, C., 2007. SPARQL Query Optimization Using Selectivity Estimation. *ISWC'07*.
- Broekstra, J., Kampman, A., van Harmelen., 2002. Sesame: A Generic Architecture for Storing and Querying RDF and RDF Schema. *ISWC'02*.
- Chaudhuri, S., 1998. An Overview of Query Optimization in Relational Systems, In *ACM PODS'98*.
- Chebotko, A., Lu, S., Fotouhi, F., 2007. Semantics Preserving SPARQL-to-SQL Translation. Technical report TR-DB-112007-CLF.
- Chong, E. I., Das S., Eadon G., Srinivasan J., 2005. An Efficient SQL-based RDF Querying Scheme, *VLDB'05*.
- Cygniak, R., 2005. A relational algebra for SPARQL. Technical report HPL-2005-170.
- Groppe, S., Groppe, J., Kukulenz, D., Linnemann, V., 2007a. A SPARQL Engine for Streaming RDF Data,

*3rd International Conference on Signal-Image Technology & Internet-Based Systems (SITIS'07)*.

- Groppe, J., Groppe, S., Ebers, S., Linnemann, V., 2009. Efficient Processing of SPARQL Joins in Memory by Dynamically Restricting Triple Patterns. *ACM SAC'09*.
- Groppe, S., Groppe, J., Linnemann, V., 2007b. Using an Index of Precomputed Joins in order to Speed Up SPARQL Processing, *ICEIS'07*.
- Haase, P., Broekstra, J., Eberhart, A., Volz, R., 2004. A Comparison of RDF Query Languages. in *ISWC'04*.
- Ioannidis, Y. E., 1996. Query optimization, In *ACM Computing Surveys*, Vol. 28, No. 1.
- Jarke, M., Koch, J., 1984. Query Optimization in Database Systems, In *ACM Computing Surveys*, Vol. 16, No. 2.
- Pérez, J. Arenas, M., Gutierrez C., 2006. Semantics and Complexity of SPARQL. *ISWC'06*.
- Prud'hommeaux E., Seaborne A., 2007. SPARQL Query Language for RDF, *W3C Recommendation*, 15 Jan. 2007.
- Weiss, C., Karras, P., Bernstein, A., 2008. Hexastore: Sextuple Indexing for Semantic Web Data Management, *VLDB'08*.
- Wilkinson, K., Sayers, C., Kuno, H. A., Reynolds, D. 2003. Efficient RDF Storage and Retrieval in Jena2. In *SWDB'03 co-located with VLDB'03*.