AnEfficientSQL-basedRDFQueryingScheme

EugeneInseokChongSouripriyaDasGeorgeEadonJagannathanSrinivasan

Oracle

OneOracleDrive,Nashua,NH03062,USA

Abstract

Devising a scheme for efficient and scalable querying of Resource Description Framework (RDF) data has been an active area of current research. However, most approaches define new languages for querying RDF data, which has the following shortcomings: 1) They are difficult to integrate with SQL queries used in database applications, and 2) They incur inefficiency as data has to be transformed from SQL to the corresponding language data format. This paper proposes a SQL based scheme that avoids these problems.Specifically, it introduces a SQL table function RDF_MATCH to query RDF data. The results of RDF MATCH table function can be further processed by SQL's rich querying capabilities and seamlessly combined with queries on traditional relational data. Furthermore, the RDF_MATCH table function invocation is rewritten as a SQL query, thereby avoiding run-time table function procedural overheads. It also enables optimization of rewrittenqueryinconjunctionwiththerestofthe query. The resulting query is executed efficiently by making use of B-tree indexes as well as specialized subject-property materialized views. This paper describes the functionality of the RDF_MATCH table function for querying RDF data, which can optionally include user-defined rulebases, and discusses its implementation in OracleRDBMS.Italsopresentsanexperimental study characterizing the overhead eliminated by avoiding procedural code at runtime. characterizing performance under various input conditions, and demonstrating scalability using

Proceedingsofthe31 stVLDBConference, Trondheim,Norway,2005 80 million RDF triples from UniProt protein and annotation data.

1. Introduction

ResourceDescriptionFramework(RDF)[1]isalangu age forrepresentinginformation(metadata)aboutresou rcesin the World Wide Web. The resources are not limited t o web pages but can also include things that can be identified on web. The specification of metadata in the generic RDF format makes it suitable for automatic consumptionbyadiversesetofapplications.

The RDF data represented as a collection of < subjec t, property, object> triples, can easily be stored in а relational database. The paper addresses the issue of efficiently querying such RDF data. For querying RD F data, most approaches define yet another query lang uage, which in turn issues SQL to process user requests. In contrast, this paper proposes a SQL-based scheme fo querying RDF data. Specifically, it proposes an RDF MATCH table function with the following functionality:

- The ability to search for an arbitrary pattern agai nst the RDF data including inferencing based on RDFS [3]rules, and
- The ability to include a collection of user-defined rulesasanoptionaldatasource.



Figure1:RDFdataforreviewersmodel

Toillustratethebasicfunctionality,considerRDF data about research paper reviewers. The RDF classes and the triple instances are shown in Figure 1. Assuming the eRDF data is stored in the data base as the model 'review ers', user

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distribute d for direct commercial advantage, the VLDB copyright notice and the title of the publication and its date appear, and notice is give nthat copying is by permission of the Very Large Data Base Endowment. To copy otherwise, or to republish, requires a fee and/or s pecial permission from the Endowment

can issue the following query to find reviewers who students with a geless than 25:

are

```
SELECT t.r reviewer
FROM TABLE(RDF_MATCH(
    '(?r ReviewerOf ?c)
    (?r rdf:type Student)
    (?r Age ?a)',
    RDFModels('reviewers'),
    NULL, NULL)) t
WHERE t.a < 25;</pre>
```

ThevariousargumentstoRDF_MATCHareasfollows:

- The first argument captures the graph pattern to search for. It uses SPARQL-like syntax [13] and variablesareprefixed with a '?' character.
- The second argument specifies the model(s) to be queried.
- The third argument specifies the rulebases (if any) Here the NULL argument indicates absence of rulebases.
- The fourth argument specifies user-defined namespacealiases(ifany).HeretheNULLargument indicates that no user-defined aliases are used, however default aliases such as rdf: are always available.

By processing RDF data using SQL the regular databa se tables can be queried in a single query along with data. For example, a user can join results of RDF q uery with a traditional employees tables aytofind the emailid of the faculty reviewers:

```
SELECT t.r reviewer, e.emailid emailid
FROM TABLE(RDF_MATCH(
    '(?r ReviewerOf ?c)
    (?r rdf:type Faculty)',
    RDFModels('reviewers'),
    NULL, NULL)) t, employees e
WHERE t.r = e.name;
```

Providing RDF querying capability in SQL would enab le ntics applications to easily process domain-specific sema storedasRDFdatainarelationaldatabase.Thisb ecomes evenmoreimportantespeciallyinthecontextofse mantic web applications, since RDF is an important buildin g block of the semantic web [4]. Also, in future, if the vast amount of data stored in relational databases is ma de available as RDF triples [1], then the RDF MATCH function can be used to query such data. Furthermor e, applications that need to handle large volumes of metadata such as portals and e-marketplaces can als 0 benefitfromthisfunctionality.

The proposed SQL-based RDF querying scheme involving the RDF_MATCH table function has been designed to meet most of the requirements identifie d in [5], including RDF Graph pattern matching, limiting resultingsubgraphs, and returning results that may contain subgraphs from the graphentailed by input RDF grap hs. It also handles RDFS and user-defined rules in a seaml ess manneralong with the models.

Withregardstoimplementation, the key aspects of our approach areas follows:

- The RDF_MATCH functionality is introduced as a SQL table function. Although this avoids kernel changes, the approach is not suitable for class of queries that return large result sets, where the overhead of returning result via the RDBMS table function infrastructure tends to dominate the query costs(seeSection4.1fordetails).Tocircumvent this problem an extension to RDBMS table function infrastructureisimplemented, whichallowsarewr ite of table function with a SQL query. With this extension, processing of RDF_MATCH table functional languagerun-timesystemotherthantheSQLengine.
- TheRDFdatatriplesarestoredafternormalization in two tables, namely IdTriples (triples in the identi fier format) and UriMap (uritoidentifier mapping). Oth er storage organizations are possible but are not considered in this paper. This would be addressed in future work. The storage representation supports handling literal of multiple datatypes as well as supports multiple representations of same literal value.
- The core implementation of RDF_MATCH query translates to a self-join query on IdTriples table. To efficiently execute this query, a set of B-tree ind exes and materialized views are defined on the IdTriples andUriMapTables.
- Aclassofmaterialized viewscalled *subject-property matrix* materialized join views (SPMJVs) is introduced to avoid the inefficiency resulting from storing heterogenous data in the canonical triple format table. Also, the statistics collected on the SPMJVs serve as statistics for the corresponding portionsofthetripletable.
- Our approach relies on the RDBMS cost-based optimizer for optimizing the resulting query (that is, after rewrite of table function invocation). This approachhastheadvantagethatRDBMSoptimizeris leveraged. However, a shortcoming is that optimizer can generate sub-optimal plans. We plan to address this problem in future by enhancing the optimizert o betterhandletheclassofself-joinqueries.
- Rulebases are in general handled by generating SQL queries, which may optionally involve table functions. Also we support the notion of indexing rulebases, which allows pre-computing and storing the data derived by applying rulebases to specified RDFmodels.

This scheme has been implemented in Oracle RDBMS using Oracle's table function infrastructure. In ad dition, the implementation uses Oracle's B-tree indexes, function-basedindexes, and materialized views.

Performance experiments conducted using RDF data for WordNet, the lexical database for English langu [8], and UniProtprotein and annotation data[14]v alidate the feasibility of this scheme and demonstrate that it scaleswellforlargedatasets.

Thekeycontributionsofthepaperare:

- A SQL-based scheme for querying RDF data. No changes are made to SQL. Instead, RDF_MATCH table function is introduced. User can leverage all of the SQL capability to process the result of the RDF_MATCH tablefunction.
- An efficient and scalable SQL based implementation of RDF_MATCH table function, including querying ondataderived by applying rule bases..
- AnextensiontoRDBMStablefunctioninfrastructure that eliminates bulk of the runtime overheads for class of table functions that can be expressed as a SQLquery.
- A study characterizing RDF query performance as wellasidentifyingoverheadsinvariouscomponents

1.1 RelatedWork

For querying RDF data, a number of query languages have been developed. This includes RDQL [9], RDFQL [10], RQL[11], SPARQL[13], SquishQL[1], and RSQL [15]. These are declarative query languages with qu ite a fewsimilaritiestoSQL.However,theschemepropos edin this paper differs from all of the above in that it allows SQLitselftobeusedtoqueryRDFdatabyintroduc inga table function. The main advantage of this SQL-base d schemeisthatitallowsleveragingtherichfuncti onalityof SQLandefficientlycombininggraphquerieswithqu eries againsttraditionaldatabasetables.

Withrespecttohandlingrulebases,ourschemeisq uite similar to RDFQL where one or more data sources or rulebasescanbespecified.

With regards to implementation, query languages suc h as RQL and SquishQL try to push as much of the functionality as possible to underlying database by formulating SQL queries against tables storing RDF data. Our approach for implementing the RDF_MATCH table function is somewhat similar. However, it is tightl у integratedwiththeSQLengineandwiththetablef unction SQL rewrite functionality further optimization is p ossible such as filter condition pushdown. Our approach us es materializedviewstospeedupqueriesonRDFdata. This is in addition to the typical database indexes one can defineonthetablesstoringRDFdata.

1.2 OrganizationofthePaper

Section 2 describes key concepts of supporting RDF_MATCH based queries. Section 3 discusses the designandimplementationoftheRDF_MATCHfunction on top of Oracle RDBMS. Section 4 discusses an RDBMS table function infrastructure enhancement that t can eliminate bulk of the mapping overhead for RDF_MATCH queries. Section 5 describes the performance experiments conducted using RDF data for r

WordNet and UniProt. Section 6 concludes with a summaryandfuturework.

2. KeyConcepts

This section gives the terminology used in the rest of the paper, outlines the requirements for querying RDFd ata, and describes how the SQL based RDF querying scheme meets these requirements.

2.1Terminology

RDF can be used to capture domain semantics. The basic unit of information is a *fact* expressed as a <subject, property (predicate), object> *triple*. For example, the fact, 'John's age is 24', can be represented by <'John', 'age', '24'>triple. A collection of triples, typically pertaining to adomain or sub-domain, constitutes an *RDF model*.

Triples in a model can be classified as *schema* triples and *data* triples. Schema triples, specified using RDFS, describe the schema-related information (for exampl e, <'age', 'rdfs:domain', 'Person'>), whereas data tr iples describe the instance data. Note that a triple's su property arealways URIs while its object can be a URI or literal.

An RDF model is also referred to as *RDF graph*, where each triple forms a <property>edge that conn ects the<subject>nodetothe<object>node.

An RDF data set can optionally include one or more *rulebases*, each containing a collection of *rules*. A *rule* when applied to a model yields additional triples. An RDFmodelaugmented with a rulebase is equivalent to the original set of triples plus the triples inferred b y applying the rulebase to the model.

2.2ASQLbasedRDFQueryingScheme

Atable function is introduced to SQL as described below to satisfy most of the requirements described in [5] by a SQL-based RDF querying scheme.

RDF_MATCH Table Function: For data stored in a database, an RDF_MATCH table function is introduced with the ability to search for an arbitrary graph p attern against the RDF data, including inferencing based o n RDFS and user-defined rules. The signature of the t able functionisas follows:

RDF_MATCH (
Pattern	VARCHAR,
Models	RDFModels,
RuleBases	RDFRules,
Aliases	RDFAliases
)	

RETURNS AnyDataSet;

The first parameter captures the graph pattern to b e matched. It is specified as a collection of one or some <Subject, Property, Object> triple patterns. Typica lly, each triple pattern contains some variables. Variab les alwaysstartwitha '?'character.

Among the remaining parameters, the first two species fy a list of RDF models and (optional) rule bases, which h

together constitute the RDF data to be queried. The last parameterspecifiesaliasesfornamespaces.

The result returned by RDF_MATCH is a table of rows. Each resulting row contains values (bindings) for the variables used in the graph pattern. Substituti ng the values in the graph pattern would identify the correspondingmatchingsubgraph.

The exact definition of the result table, that is, the set of columns and their data types, varies depending u pon the graph pattern used in an RDF_MATCH invocation. (Use of the AnyDataSet data type allows us to define RDF MATCHwiththisflexibility.)Specifically,for each variable in a given graph pattern, the result table has a column with the same name as the variable (without the starting (?'). These columns are used for returning the lexicalvaluesforthecorrespondingvariables.In addition, the result table has additional columns (of form <variable>\$type)forreturningdatatypeinformationfor each variable that may be bound to literals as well .Note that based upon current RDF restrictions (that subj ects and predicates in triples must be URIs and not lite rals), only those variables that do not appear as subject and predicate components of triples in a graph pattern canbe boundtoliteralvalues.

Example: Consider the following query (from Section 1) to find student reviewers who are less than 25 year old:

TheRDF_MATCHinvocationreturnsthefollowingtabl e:

r	С	c\$type	a	a\$type
John	IDBC2005	URI	24	xsd:int

Notethat, sincevariable ?rappears assubject component, the value of column risal ways a URI and hence there is noneed for an additional column to return its data type.

SQL constructs may be used to extend the above query, to do aggregation, grouping, and ordering, f or example:

```
SELECT t.c conf,
        COUNT(*) row_count, AVG(t.a) avg_age
FROM TABLE(RDF_MATCH(.....)) t
GROUP BY t.c
ORDER BY avg_age;
```

Though the above examples show querying RDF data only, userscanalsoquerytheassociated RDF schem a, for example, to obtain domains and ranges for a propert y. Thus, the key benefits of using a table function fo r querying RDF data is that the standard SQL construc ts can be used for further processing of the results. This includes iterating over the results, constrainingt he results using WHERE clause predicates, grouping the results using GROUP BY clause, sorting the results using ORDER BY clause, and limiting the results by using the ROWNUM clause. Also, the SQL set operations can be used to combine result sets of two or more invocati ons of RDF_MATCH. With the table function SQL rewrite functionality discussed in Section 4.2 the optimize r will be able to optimize the whole SQL query including f ilter condition pushdown.

Rule and Rulebases: A rule is identified by a name and the rulebase to which it belongs. A rule consists o faleft hand side (LHS) pattern for the antecedents, an opt ional filter condition that further restricts the subgrap hs matched by the LHS, an optional list of namespace aliases, and a right hand side (RHS) pattern for th e consequents. For example, the rule that "chairpers on of a conference is also a reviewer of the conference" is represented as follows:

('rb',				rulebas	e name
'Chai	rpersonRule'	,		rule na	ame
`(?r	ChairPerson	Of	?c)′,	LHS pat	tern
NULL,				filter	condition
NULL,				aliases	5
`(?r	ReviewerOf	?c)	()	RHS pat	tern

The following query will return both John and Mary as reviewers. The latter is implicitly inferred by app lying rulebase rbtothe reviewersmodel.

```
SELECT t.r reviewer
FROM TABLE(RDF_MATCH(
    '(?r ReviewerOf ?c)',
    RDFModels ('reviewers'),
    RDFRules('rb'), NULL)) t;
```

Auser can create rulebases and add rules by using APIs. Once the rulebases are created and populated, they can be specified in a RDF_MATCH query. Note that the RDFS rulebase (named rdfs) is created by the system and is implicitly available for users.

3.DesignandImplementation

This section describes the design and implementatio n of the SQL-based RDF Querying Scheme. This scheme is implemented on top of Oracle RDBMS. Although the descriptionhereassumesOracleRDBMS,theschemec an besupportedinanyRDBMSthatsupportstablefunct ions, materializedjoinviews,andB-treeindexes.

3.1RDFDataStorageandMultipleDataType Hanlding

The RDF data must be stored compactly and the storage format should be suitable for efficient que ry processing. In our scheme, RDF data is stored (afte r normalization) in two tables: IdTriples (ModelID SubjectID, PropertyID, ObjectID, ...) and UriMap (UriID, UriValue, ...). This normalization is critical because URIs (or literals) are typically repeated. Also, it enables efficient query processing due to the compa ct size. Given an RDF triple, its three URIs (or liter als) are first mapped to corresponding identifiers using tab 1e UriMap. If no mapping is found for a URI (or litera 1), a new unique UriID is generated and the new mapping i S

inserted into the UriMap table. A tuple comprising the ModelID (for the RDF model) and the three UriIDs is thenstored into the IdTriple stable.

Auserviewiscreatedontheunderlyingtableshol ding RDF data, which presents only selective portions (a t model granularity) of the RDF data to the users bas ed on their privileges. Also, the RDF_MATCH function is executed with invoker's privileges. Thus, this sche me limits each invoker's access via RDF_MATCH query to only the appropriate portion of the RDF data.

TypedliteralsarestoredintheUriMaptablewith their type. To support matching between multiple representations for the same value, such as the int eger123 and the float 12.3E+1, each literal is mapped to a canonical literal. Literals that represent the same value willmaptothesamecanonicalliteral, and aliter almaybe its own canonical literal. The canonical literal ID (which isused when joining on the object column) and the exact literalID(which is used when returning the object tothe user) are both stored in IdTriples. For simplicity, queries in this paper are written as if there was a single objectID columninIdTriples.

The first literal entered for a value becomes the canonical literal. To support mapping other equi-va lued literals to this canonical literal, there is a flag in UriMap to indicate that the literal is a canonical literal . Further. the pre-defined datatypes are partitioned into fami lies, where all types in a family are associated with a s ingle value space. For example, float and integer types b oth belong to the numeric family. For each type family, there isafunctiontoconverttheUriMaplexicalreprese ntations into a canonical form, such as a native database ty pe.A function-based index for this purpose is defined on UriMap, soacanonical form can efficiently be mapp edto thecorrespondingcanonicalliteralduringquerying

3.2RDF_MATCHTableFunction

The RDF_MATCH functionality is implemented as a SQL table function using Oracle's table function interfaces[6].

RDFQueryProcessing

- **Compile Time Processing** : At compile time, the form of the table result, namely the set of columns is determined. The kernel passes information regarding columns referenced in the outer SQL query to the ta function. This allows for optimization of table fun queries based on columns referenced in the SQL quer y containingtablefunctioninvocation.

- Execution Time Processing : Based on the input arguments, namely, pattern, models, rules, and alia ses, a SQL query is generated against the IdTriples and Ur iMap tables. Figure 2 shows the various layers of implementation. There are two types of implementati on: conventional procedural processing (discussed in Se ction 4.1) and a new declarative rewrite-based processing (discussed in Section 4.2). Due to overheads with t he procedural implementation, RDF_MATCH function uses therewrite-based implementation.



Figure2:RDF_MATCHImplementationOverview

Consider as an example use of the RDF_MATCH table functioninthefollowingSQLquery:

```
SELECT t.r reviewer, t.c conf, t.a age
FROM TABLE(RDF_MATCH(
         '(?r ReviewerOf ?c)
         (?r rdf:type Student)
         (?r Age ?a)',
```

First, aliases specified (if any) are substituted w ith the namespaces to expand all alias-encoded URIs. Next, URIs and literals, such as 'ReviewerOf', 'rdf:type' are convertedintoUriIDsusinglookupsontheUriMapt able:

```
FROM IdTriples t1, IdTriples t2. IdTriples t3
WHERE t1.PropertyID = 14 AND t2.PropertyID = 11
AND t2.ObjectID = 4 AND t3.PropertyID = 29
```

Then a self-join query is generated based on matchi ng variablesacrosstriples(e.g. '?r')inthepattern :

```
WHERE ... t1.SubjectID = t2.SubjectID AND
t2.SubjectID = t3.SubjectID
```

Next, the internal IDs are joined with the Uri Mapt able to generate the join result in the URI (and literal) f or mat:

Note that 'r' is a URI, so there is no type associated, whereas 'c' and 'a' have a date type (c\$type, a\$type) associated.

The models argument is used to restrict the IdTriples table based on the corresponding model identifiers in the above self-join SQL query. After the transformation phase, the generated single SQL query is optimized and executed to obtain results.

In addition to the table function arguments, kernel implicitly provides information regarding the colum ns (which correspond to the variables in the graph pat tern) referenced elsewhere in the original SQL query. The RDF_MATCH implementation has been optimized to compute values only for these columns. This avoids additional joins with UriMap table to get the corresponding UriValue. In general, a query withn triple patternandmvariableswillresultinaquerywith (n+m-1) joins, assuming mvariables are projected (hence m joins with UriMap table). Experiment IV (Section 5.6) demonstrates the performance benefits of this optimization.

The rules argument contains a list of rulebases to be applied to infernew triples. This is discussed below.

Rule Processing: To handle rules, the RDF MATCH function replaces references to the IdTriples table in the generated SQL with subqueries or table functions th at vieldtherelevantexplicitandinferredtriples.S ubqueries are used whenever the required inferencing can be d one conveniently within a SQL query (i.e., without expl icitly materializing intermediate results). These subqueri es generally take the form of a SQL UNION with one UNION component for each rule that yields a relevan t triple, plus one component to select the explicit t riples. Table functions will be used when the subquery appr oach isnotfeasible.

To support the RDFS inference rules, we must RDFS compute a transitive closure for the two transitive properties: rdfs:subClassOf (rule rdfs11) and hese rdfs:subPropertyOf(rulerdfs5).InOracleRDBMS,t transitive closures can be computed using hierarchi cal queries with the START WITH and CONNECT BY NOCYCLE clauses. Note that CONNECT BY NOCYCLE queries can handle graphs that contain cycl es bygeneratingtherowinspiteoftheloopinuser data.The remaining RDFS rules can be implemented with simple SQLqueries.

To ensure that RDFS inferencing can be done within
single SQL query, the user is prohibited from exten
the built-in RDFS vocabulary. This means, for examp
that there cannot be a property that is a sub-prope
the rdfs:subPropertyOf property, nor can there be a
user-
defined rule that yields rdfs: domain triples.a
ding
the
rty of
user-

User-defined rules can be classified as follows bas ed upontheextentofrecursion, if any, in the rule:

- Non-recursive rules: The antecedents cannot be inferredbythegivenrule,oranyrulethatdepend son thegivenrule'sconsequents.
- Simple recursive rules: These rules are used to associate transitivity and symmetry characteristics withuser-definedproperties.
- Rulesthatusearbitraryrecursionunliketheother two categories.

Non-recursive user-defined rules can be evaluated u sing SQL (join) queries by formulating the FROM and WHERE clauses based upon the antecedents and the SELECT clause based on the consequents of the rule soas to return the inferred triples. Note that the triples that match the antecedents of a user-defined rule could themselves be inferred, so the FROM clause may reference subqueries to find inferred triples. The ChairpersonRulegiveninSection2.2wouldtranslat einto SQLasfollows:

```
SELECT ...
FROM (
-- (?x ChairpersonOf ?c) => (?x ReviewerOf ?c)
SELECT tl.SubjectID, 14 PropertyID, tl.ObjectID
FROM IdTriples t1
WHERE tl.PropertyID = 56
UNION
-- explicit ReviewerOf triples
SELECT tl.SubjectID, tl.PropertyID, tl.ObjectID
FROM IdTriples t1
WHERE tl.PropertyID = 14
) tl;
```

Simplerecursiverulesinvolvingtransitivityands ymmetry can be evaluated as follows. Symmetry can be easily handled with a simple SQL query. However, handling transitivitywithasingleSQLqueryrequiressome typeof hierarchical query (e.g., using the START WITH and CONNECT BY NOCYCLE clauses in Oracle RDBMS), asinthecaseoftransitiveRDFSrules.

Suppose the user's query is:

```
RDF_MATCH(
`(?a rdf:type Male)
(?a AncestorOf ?b)',
```

There is a user-defined rule to make AncestorOf transitive, and for simplicity we assume that the R DFS rulebase is not used. So after translation we have a join between IdTriples(fortherdf:typetriple) and as ubquery which computes the transitive closure using CONNECT BY(forAncestorOf):

The third class of rules involving arbitrary recurs ionis the most complicated, and it has not been addressed in the current implementation. These rules will be evaluat ed using table functions, because an unknown number of passes over the intermediate results are required to find all inferred triples.

3.3SpeedingupRDF_MATCHQueries

The speed up is achieved by creating materialized j oin views(MJVs)andcreatingappropriateB ⁺-treeindexeson them,andindexingRDFdataandrulebases.Eachof these isdescribedindetailbelow.

Generic Materialized Join Views: The query generated by RDF_MATCH table function involves a self-join of IdTriples table if the same variable is used in more than one triple of the search pattern. Depending on how many triples are specified, a multi-way join needs to be executed. Since the join cost is a major portion of the total

processingtime, materialized joinviews can be def inedto speed up RDF_MATCH processing. The row size of IdTriples table is small and hence the materialized join view can be a good candidate for reducing the join cost. In general, six materialized two-way join views, na mely ioins between SubjectID-SubjectID, SubjectID-PropertyID, SubjectID-ObjectID, PropertyID-PropertyID, PropertyID-ObjectID, and ObjectID-ObjectID can be defined as long as the storage requirement is met. Most useful materialized join v iews for typical queries, however, are joins between SubjectID-ObjectID, SubjectID-SubjectID, and ObjectID-ObjectID.Notethattheindividualmaterialized join views could be created for a subset of databa sedon theworkloadcharacteristics.

The materialized join views are incrementally maintained on demand by the user using the DBMS_MVIEW.REFRESH API.AproceduralAPIisprovided to analyze IdTriples table to estimate the size of various materialized views, based on which a user can defin e a subsetofmaterialized views.

Subject-Property Matrix Materialized Join Views: To minimize the query processing overheads that are in herent in the canonical triples-based representation of RD F, subject-property matrix based materialized join views can be used. These materialized views can be designed u sing the following basic ideas:

- Foragroupofsubjects, choose a set of single-val ued properties that occur together. These can be direct properties of these subjects or nested properties. A propertyp 1isadirectpropertyofsubjectx 1 if there is $atriple(x_1, p_1, x_2)$. Approperty misanested property ofsubjectx 1 if there is a set of triples such as, (x 1,p 1, x_2), ..., (x_m, p_m, x_{m+1}) , where m >1. For example, if we have a set of triples, (John, address, addr1), (addr1, zip, 03062), then zip is a nested property of John.
- Create a (subject-property matrix) materialized joi n view each of whose rows contains values of these propertiesforasubjectinthegroup.

Queryperformancecanbeimprovedsignificantlythr ough the use of such materialized join views because an umber of joins can be eliminated. For example, Table 1s hows a sample RDF data and Table 2 shows a matrix material ized join view created for subjects who are students with their direct property age and nested property city (named in the view as studies At to denote the city where his/her university is located).

This subject-property matrix can be exploited by an RDBMS optimizer to process an RDF query using the followingquerypatterntoretrievethe ageand studiesAt infoforeachs tudent:

`(?r	rdf:type	Student)
(?r	enrolledAt	?u)
(?r	age	?a)
(?u	city	?city)'

andretrievingvaluesofvariables ?r, ?city,and ?a.

Table1:StudentInfoRDFData

Subject	Property	Object
John	rdf:type	Student
John	EnrolledAt	Univ1
John	Age	24
Pam	rdf:type	Student
Pam	EnrolledAt	Univ2
Pam	Age	22
Univ1	UnivName	NYU
Univl	City	New York
Univ2	City	Los Angeles

Table2:StudentMatrix

Subject	StudiesAt	Age
John	New York	24
Pam	Los Angeles	22

This query will normally require a 4-way self-joinonthe IdTriples table (leaving out the conversion betweenIDs and URIs, for simplicity). However, by using thematrix in Table 2, the query can be processed by simplyselecting all the rows from the materialized join view.Thus, self-joins can be completely eliminated inthis case.This can lead to significant speed-up inquery processing.

Ingeneral, for the type of queries shown above a query requiring an n-way join could potentially be processed using a matrix with m-properties with (n - m) joins.

In typical usage of such matrices, each subject in the group will have one value for each of the chosen properties. Usage may involve sparseness to somee xtent to allow expanding the group of subjects to include those subjects that may have no values for a few of the properties in the selected subset.

It may be noted that use of these matrices as materialized join views for performance gain needs to be evaluated against the workload for potential benefits versus the space overhead incurred for additionals to rage. The issue of which views to materialize is dependen tupon the search pattern and it is up to the user to decide the which is frequents earch pattern.

The problems of obtaining property-specific statist ics foratriplestore with heterogeneous data can be matrix with the use of statistics computed on the matrix materialized views because those can serve as stati stics for the corresponding portions of the vertical table e.

Finally, Jena2's [12] property tables (clustering multiple properties) are in many ways similar to su bject-property matrices. The main differences include the following:

- Subject-property matrix is an auxiliary structure, not aprimary storage structure. So, these matrices may dropped or redefined as necessary without requiring a datare-loading.
- The definition of subject-property matrix allows us e of nested properties and hence allows more ways of

creating useful materialized views for optimizing performanceofavarietyofqueriesinaworkload.

Indexing **Rulebases:** Rulebases specified in RDF_MATCHqueryareapplied, by default, during que ry , if a processing to the specified list of models. However rulebase is used frequently then that rulebase can be indexed using a set of APIs provided for this purpo se. Indexing a rulebase for an RDF model refers to precomputing the triples that can be inferred with res pectto the specified model. These pre-computed triples are stored in a separate table and are used subsequentl y duringRDF_MATCHqueryprocessingtospeedupquery execution. In general, a pre-computing may need to be done for a combination of models and rulebases, tha tis, applying a set of rules from the union of rulebases to a triplesfromtheunionofasetofRDFmodels.

However, these pre-computed results cannot be used directly to process RDF_MATCH queries that referenc e additional rulebases or models. Currently for such cases, all inferencing must be done at query execution tim e. Notice that inferencing can only add triples to the graph, sothepre-computed triples are always valid for th elarger set of rulebases and models, though the pre-compute d results are not necessarily complete. We plan to ex plore handling these cases by analyzing the rulebases and models so we can avoid re-computing portions of the precomputedresultsthatarecomplete.

Indexing RDF Data: As mentioned earlier, the core processing involves performing self-joins on IdTrip les table. Thus, creating the right set of indexes on IdTriples is critical for performance improvement. There are typically two types of query patterns: 1) given ap joining subject with subject, or object with object as shown below:

`(?r	ReviewerOf	?c)
(?r	Age	?a)'
`(?r	ReviewerOf	?c)
(?c	rdf:type	Conference)'

or

SincepropertyistypicallyspecifiedasaURIvalu e,index key with property as the first column may allow pru ning the search space to a single range in the B-tree in dex. Further, having all the three columns (namely PropertyID, SubjectID, and ObjectID) as part of the key may allow index-only access provided the additi onal storage space required for three column indexes can be accommodated. Based upon these observations, we hav e usedtwothreecolumnindexeswiththefollowingke ysin allofourperformanceexperimentsdescribedinSec tion5 ObjectID> and : < PropertyID, SubjectID, <PropertyID, ObjectID, SubjectID>. Use of keyprefix compression in indexes allowed reducing the storagespacerequiredfortheindexes.

The choice for indexes may depend on the actual RDF data and workload characteristics. We need to explo further to see how any algorithm for choosing index may need to be customized to exploit constraints su row for matsused for RDF triples storage and typica queries that involvemulti-ways elfjoins.

4.MinimizingOverheadsbyan EnhancementtoRDBMS

Thissection discusses an enhancement to Oracle RDB MS table function infrastructure that can minimize tab le function processing overheads.

4.1 RDFQueryProcessingComponents

The RDF query processing time using RDF_MATCH table function (t_{total}), without the kernel enhancement discussedinSection4.2, can be represented as fol lows:

 $t_{total} = t_{core} + t_{sql2proc} + t_{proc2canonical} + t_{canonical2sql}$ Here t_{core} represents the core processing time, that is, the cost of SQL query that performs the self-joins on IdTriples table and any additional joins with UriMa p table. Once the results are computed, they are copi edinto variables of the table function procedure (t_{sql2proc}), and subsequently it is converted to canonical format (t_{proc2canonical}) so it can be returned to via RDBMS table function infrastructure, and finally transformed ba ck (t_{canonical2sql}) so it can be consumed by the outer SQL query.

The component, t_{total} - t_{core}, is dependent on the result computed by table function (note: not on the overal 1 result) and hence it will dominate the query costs when the table function results etsize is large. The Experiment I (described in Section 5.3) demonstrates the overhea ds incurred for varying number of result rows. To avo id this overhead an enhancement to RDBMS is implemented as discussed below.

4.2 ANewTableFunctionInterface

The following extension of RDBMS table function infrastructure is implemented, that would allow a s imple rewriteoftablefunction with a SQL query.

As an alternative to the current TableStart(), TableFetch(), and TableClose() interfaces, RDBMS shouldsupportanewtablefunctioninterface:

TableRewriteSQL(arg1, ..., argn) RETURNS VARCHAR;This function takes the arguments specified in thetablefunction and generates a SQL string. For table functiontionsdefined using this interface, RDBMS table functioninfrastructuredoesthefollowingprocessing:

- InvokethecorrespondingroutinetogeneratetheSQ L string,
- Substitute the generated SQL string into the origin al SQL query, and
- Reparseandexecutetheresultingquery.

Theneteffectissame as if the user typed in the generated SQL query in place of the table function. However, general function mechanism cannot be used here beca use of the FROM clause. It has to be the table function

Suppose the RDF_MATCH table function be defined using the TableRewriteSQL() interface. Consider t he followingquery:

```
SELECT t.a age
FROM TABLE(RDF_MATCH(
    '(?r Age ?a)',
    RDFModels('reviewers'),
    NULL, NULL)) t
WHERE t.a < 25;</pre>
```

Theresultingqueryafterrewritingthetablefunct ionisas follows:

SELECT t.a age

```
FROM (SELECT ul.UriValue a, ul.Type a$type
FROM IdTriples tl, UriMap ul
WHERE tl.PropertyID = 29 AND tl.ModelID = 1
AND ul.UriID = tl.SubjectID) t
WHERE t.a < 25;</pre>
```

Note that the subquery in bold font is the SQL frag ment that is returned from TableRewriteSQL() for the above RDF_MATCH invocation. Now, the whole SQL query is optimized and executed. For example, the filter condition is pushed inside the subquery for further optimizat ion.

The advantage of such a scheme is that it avoids the overhead of copying the results into table function variables, as well as eliminates the table function infrastructure overhead of transforming the result to canonical form and re-transforming it back to present in the appropriate datatype format. However, such as cheme is applicable only when the table function can be defined declaratively using SQL (as is the case for RDF_MATCH).

5.PerformanceStudy

This section describes the performance experiments conducted using RDF_MATCH table function.

5.1ExperimentalSetup

The experiments are conducted using Oracle10 g Release 1(10.1.0.2.0)onaRedHatEnterpriseLinuxAS3s ystem withone3.06GHzPentium4CPUand2048MBofmain memory.Adatabasebuffercacheof256MB,sharedp ool of 256 MB, and database block size of 8 KB is used.

The timings reported below are the mean result from tenormore trials with warmcaches.

5.2Dataset

The experiments I through IV are conducted using an RDF representation of WordNet [11], a lexical datab ase for the English language, which organizes English w ords into synonym sets, categorizes these synonym sets according to part of speech (noun, verb, etc.), and enumerates linguistic relationships (antonymOf, similarTo, etc.) between these synonym sets. In the RDF representation, each part of speech is modeled as an rdfs:Class, and each linguistic relationship is mod eled as

anrdf:Property.ThisRDFSchemaforWordNetissho wn inFigure3.



Figure3:WordNetRDFSchema

The hyponymOf property is used to denote that the subject represents a specialization of the object. For example,skyscraperisahyponymofbuilding.

Table 3. Property and Resource Statistics of WordNet

Property	Count		Resources (explicitrdf:type)	Count
WordForm	174,002		Verb	12,127
Rdf:type	99,653		Noun (6,025
glossaryEntry	99,642	1	djectiveSatellite 1	0,912
hyponymOf	78,445	A	djective 7,	003
SimilarTo	21,858	A	dverb 3,	575
Others	26		Others	11
Total	473,626		Total	99,653

The relevant logical statistics for the experimenta 1 configuration is shown in Table 3. The logical stat istics can be computed simply with the RDF MATCH table function. For example, to find number of resources typed as'verb', auseruseRDF MATCHtable function with the pattern `(?w rdf:type wn:verb)'. This type of quervisexpectedtorunefficientlyasitresults inasingle table query. For example, the above query took less than 0.01 seconds.

Thedataisstoredinthenormalizedformintwota bles, namely, IdTriples table of size 14 MB and UriMap ta ble of size 34 MB. The indexes on IdTriples table and UriMaptableareofsize22MBand26MBrespective ly.

Experiments V and VI use large-scale UniProt data with80milliontriples(seeSection5.7formore details).

5.3ExperimentI:OverheadEstimation

This experiment characterizes the benefit of the TableRewriteSQL()enhancementdescribedinSection 4. Fourconfigurationsaretested:

- 1) RDF_MATCH with the current table function interface (TableStart(), TableFetch(), and TableClose()). Execution time of this table function correspondstothet total terminSection4.1.
- 2) SQL query equivalent to RDF_MATCH with the enhanced interface (TableRewriteSQL()). Execution

core term in time of this query corresponds to the t Section4.1.

t

e

n

t

:

- Table function (using the current interface) that 3) fetches from a SOL query, but does not return any rows.TheSQLqueryissimpleanditsexecutiontim is negligible. Execution time of this table functio correspondstothet sql2procterminSection4.1.
- Table function (using the current interface) that 4) returns rows, but does not execute any SQL. Execution time of this table function corresponds t 0 $t_{\text{proc2canonical}}+t_{\text{canonical2sql}}$ inSection4.1.

Figure 4 shows the query processing time for these components as the number of rows returned is varied from the bottom, Core SQL, SQL to Proc, Proc to SQL andOtherinthatorder.



Figure4:RDF MATCHQueryProcessing Components(standarddeviation $\sigma \leq 0.0838$)

The results demonstrate that t $_{sql2proc}$ and t $_{proc2canonical}$ + $t_{canonical 2 sql}$ are linear in the number of rows returned, and that these overheads dominate the core SQL processi ng time when a large number of rows are returned. The enhanced table function interface avoids this per-r ow overhead, and therefore it is preferred over the cu rrent table function interface. In all of the remaining performance experiments, we run the queries with th e enhancedRDF MATCHtablefunctioninterface.

5.4 Experiment II: Varying Number of Triples in the SearchPattern

As the number of triples in the RDF MATCH search pattern increases, RDF MATCH performs an increasing cterize number of self-joins on the Triples table. To chara how the varying number of self-joins impacts performance, queries are run to find 'hyponymOf'pa thsof varying length. For example, the query to find twotriple 'hyponymOf'pathsis:

```
SELECT AVG(LENGTH(a))
FROM TABLE (RDF_MATCH (
   `(?a
          wn:hyponymOf
                           ?b)
           wn:hyponymOf
                           ?c)',
    (?b
  RDFModels('WordNet'),
  NULL, NULL));
```

Thequeries are run without materialized views, and witha generic SubjectID-ObjectID materialized view, as

described in Section 3.3. Figure 5 shows the query processing time as the number of triples in the sea rch pattern varies. Note that the number of matches dec lines asthenumberoftriplesincrease, from 78, 445 matc hesfor the one-triple query to 45,619 matches for the sixtriple query.



Figure5:RDF MATCHPerformanceForVarious Searches (σ ≤0.0881)

As expected, processing time increases with the number of triples due to corresponding increase in the number of self-joins. The materialized view general ly improves performance, except for 1-triple and 5-tri ple case. For 1-triple case, no benefit is expected, as the resultingquerydoesnotinvolveanyself-joins.Fo rthe5triple case, the benefit derived due to usage of materialized view is offset because the optimizer c hooses asub-optimalplan.

5.5ExperimentIII: VaryingFilterConditions

This experiment characterizes the impact of SQL predicates that filter the results found by RDF MAT CH. Thefollowingsearchpatternisusedforthisexper iment:

	-		-
`(?c0		wn:wordForm	?word)
(?c0		wn:wordForm	?syn1)
(?c0		wn:wordForm	?syn2)
(?c1		wn:wordForm	?syn1)
(?c2		wn:wordForm	?syn2)
(?c1		rdf:type	wn:Adverb)
(?c2		rdf:type	wn:Verb)'

This query is executed with four different equality filters (e.g., word='clear') and four different range fil ters(e.g., (word >= 'bat' AND word < 'bounce')) to yield approximately 350, 1050, 2000, and 3125 matches wit h each type of filter. Figure 6 shows the query proce ssing time for these filters. Note that this query finds 79,885 matchesin8secondswhenthereisnofilterpredic ate.As expected, less selective filters require greater pr ocessing time. Notice that equality filters are more efficie nt than range filters. This is because the equality filter is implemented with a single lookup in the UriMap tabl eto findtheUriIDfortheliteralgiveninthefilter. Incontrast, range predicates require a join between the IdTripl esand UriMap table to get the values needed for filter evaluation.



Figure6:RDF_MATCHPerformancewithFilter Conditions (equality: σ ≤0.0029;range: σ ≤0.0619)

5.6ExperimentIV:VaryingProjectionList

This experiment characterizes the benefit of the projection list optimization done by RDF_MATCH. The followingsearch pattern is used for this experimen t:



Figure7:RDF_MATCHPerformanceForVarying ProjectionLists (σ ≤0.0724)

Thissearchpattern, which involves 4 variables and yields 1,470 matches, is used in queries with varying sets variables referenced in the SELECT list. Figure 7 s the query processing time as the projection lists a re changed.

The projection list optimization eliminates joins w ith the UriMap table for variables that are not referen ced outside of RDF_MATCH. It is clear that large performancegainsarepossible from this optimizati on.

5.7ExperimentV:Large-ScaleRDFData

RDF_MATCH This experiment characterizes performance for querying large-scale data. UniProt proteinandannotationdatainRDFformat[14]isu sedfor this experiment. To study scalability we created se veral datasets using varying subsets (from 10 million to 80 million triples) of the UniProt data. The largest d ataset, corresponding to 5.2 GB of RDF/XML data, occupies 2 .5 GB for IdTriplestable, 1.7GB for UriMaptable, 3. 6GB for IdTriples indexes, and 1.2 GB for UriMap indexe S. SixqueriesadaptedfromexamplesgivenwiththeUn iProt data (shown in Table 4) are then run against these datasets.

Table4.QueriesadaptedfromUniProtsamplequerie s

Description	Pattern	Projection	Result limit
Q1:Displaytherangesof	6triples 5vars	3vars	15000rows
Q2:Listproteinswith publicationsbyauthors	5triples 5vars	3vars	10rows
withmatchingnames Q3:Countthenumberof timesapublicationbya specificauthoriscited	3triples 2vars	0vars	32rows
Q4:Listresourcesthat arerelatedtoproteins annotatedwithaspecific keyword	3triples 2vars	1var	3000rows
Q5:Listgenesassociated withhumandiseases	7triples 5vars	3vars	750rows
Q6:Listrecently modifiedentries	2triples 2vars 1rangepred.	2vars	8000rows

Each query includes a ROWNUM predicate to limit the number of result rows so that the number of mat remains constant even as the dataset size changes. aggregate functions are used in the SELECT list to the overhead of returning multipler owstotheclie nt.

The RDF_MATCH search pattern for Query 1, for example, is as follows:

```
SELECT AVG(LENGTH(protein)), AVG(LENGTH(begin)),
       AVG(LENGTH(end))
FROM TABLE(RDF_MATCH(
   `(?p
                rdf:type
                                up:Protein)
    (?p
                up:annotation
                                ?a)
                rdf:type
    (?a
                up:Transmembrane_Annotation)
    (?a
                up:range
                                ?range)
    (?range
                up:begin
                                ?begin)
    (?range
                up:end
                                ?end)'
   RDFModels('UniProt')
                          NULL, NULL))
```

WHERE rownum <= 15000;

Execution times (in seconds) for these queries (see Table 5) remain almost the same even as datasets ize chan ges.

Table5.RDF_MATCHPerformanceScalability

	Q1	Q2	Q3	Q4	Q5 (Q6
10MTriples	0.86	< 0.01 <	< 0.01 0.0	03 0.1	8 0.46	-
20MTriples	0.95	< 0.01 <	< 0.01 0.0	0.1	9 0.47	
40MTriples	0.96	< 0.01 <	< 0.01 0.0	03 0.1	8 0.47	
80MTriples	1.03	< 0.01 <	< 0.01 0.0	03 0.2	0 0.49	
Maximum σ	.054	0.002	0.002	.011	.065	0.07

This shows that RDF_MATCH based query performance is scalable, that is, retrieval cost performance is scalable, that is, retrieval cost performance is scalable.

5.8ExperimentVI:Subject-PropertyMJVs

ToseepotentialbenefitsfromuseofSubject-Prope rty MJVs(SPMJVs),weusedthefollowingquerypattern againstthe80MtripleUniProtdataset:

Samota	leooninpieonin	olulubet.
`(?s	up:name	?n)
(?s	rdf:type	up:Protein)
(?s	up:curated	true)
(?s	up:created	?cre)
(2ª	un:modified	2mod) /

An SPMJV was created for rdf:type, up:curated, up:created, and up:modified properties. This SPMJV contained489,695rowsandoccupied39MB;therewa sa single B ⁺tree index on the subject, which occupied 19 MB.

Twoqueries were tested: (#1) aCOUNT (*) query, and (#2) a query that selects ?n, ?cre, and ?mod. Each query was posed with and without use of the SPMJV. The results in Table 6 shows that this can lead to sign ificant performance benefits.

Table6:RDF_MATCHPerformancewithand

withoutSPMJVS		
Query		Time (sec)
#1 w/o	SPMJV	4.87
#1 w/	SPMJV	1.79
#2 w/o	SPMJV	13.68
#2 w/	SPMJV	9.05

6.ConclusionsandFutureWork

The paper proposed a SQL based scheme for querying RDF data. Specifically, the RDF_MATCH table function n is introduced with the ability to perform pattern-b ased match against RDF data (graph) that can optionally include triples inferred by applying RDFS or user-d effined rules. Users can do further processing (iterate ove r, constrain using filter conditions, limit the result s, etc.) using standard SQL constructs.

The RDF_MATCH table function itself is implemented by generating a SQL query against table s holdingRDFdata.Forefficientqueryprocessing,g eneric and subject-property matrix materialized join views , and indexes (on RDF data and rulebases) are used. Furthermore, a kernel enhancement is implemented th at eliminates RDF_MATCH table function run-time processingoverheads.

The experimental study conducted using RDF data for WordNet and UniProt demonstrates that the SQL based schemeisefficient and scalable.

We expect that providing RDF querying capability as part of SQL will enable a database system to suppor t wider range of applications as well as facilitate b uilding semantically rich applications. The RDF querying capability can also be used in conjunction with dat a mining techniques on RDF data collected from divers e applicationstodiscoverinterestingsemanticrelat ionships.

In future, we plan to consider alternate storage representations for RDF triples. A promising storag e representation is partial normalization, where only the namespaces are normalized. That is, URIs are repres ented by the (namespace identifier, URI suffix). Also, we plan to enhance RDBMS optimizer to improve its capabilit ies in optimizing the class of self-join queries that t ypically occur while querying RDF data. The selection of sui table join method, join order, and subject-property matri х materializedjoinviewsiscriticalingeneratinga noptimal plan. Allowing users to specify hints to influence the optimizationprocesswillalsobeexplored.

Acknowledgments

WethankJayBanerjeeforhisusefulcommentsonea rlier draftsofthispaper.

References

- [1] *RDFPrimer*.W3CRecommendation,10February 2004, <u>http://www.w3.org/TR/rdf-primer</u>.
- [2]L.Miller,A.Seaborne,A.Reggiori.Three ImplementationsofSquishQL,aSimpleRDFQuery Language. *FirstInternationalSemanticWeb Conference(ISWC2002)*, June2002.
- [3] *RDFVocabularyDescriptionLanguage1.0:RDF Schema*.W3CRecommendation10February2004, <u>http://www.w3.org/TR/rdf-schema</u>.
- [4]T.Berners-Lee,J,Handler,OLassila.TheSema ntic Web. *ScientificAmerican* ,May2001.
- [5] RDFDataAccessUseCasesandRequirements .W3C WorkingDraft,2June2004, http://www.w3.org/TR/rdf-dawg-uc/.
- [6] PipelinedandParallelTableFunction. DataCartridge Developer'sGuideRelease2 (9.2) PartNo.A96595-01, OracleCorporation, March2002.
- [7]R.G.Bello,etal.MaterializedViewsinOracl e.In Proceedingsofthe24 thInt.Conf.onVeryLargeData Bases,(1998),659-664.
- [8] WordNet,TheLexicalDatabaseforEnglish Language, http://www.cogsci.princeton.edu/~wn.
- [9] RDQL-AQueryLanguageforRDF ,W3CMember Submission9January2004, <u>http://www.w3.org/Submission/2004/SUBM-RDQL-</u>20040109.
- [10] *RDFQLDatabaseCommandReference* , <u>http://www.intellidimension.com/default.rsp?topic=/pa</u> ges/rdfgateway/reference/db/default.rsp.
- [11]G.Karvounarakis,S.Alexaki,V.Christophides , D. Plexousakis,M.Scholl.RQL:ADeclarativeQuery LanguageforRDF. WWW2002,May7-11,2002, Honolulu,Hawaii,USA.
- [12]KevinWilkinson,CraigSayers,andHarumiKuno EfficientRDFStorageandRetrievalinJena2. First InternationalWorkshoponSemanticWeband Databases,pp.131-151,2003.
- [13] SPARQLQueryLanguageforRDF ,W3CWorking Draft,12October2004, <u>http://www.w3.org/TR/2004/WD-rdf-sparql-query-20041012/</u>.
- [14]UniProtDataSet, <u>http://www.isb-sib.ch/~ejain/rdf/</u>.[15]LiMa,ZhongSu,YuePan,LiZhang,TaoLiu.
- RStar:AnRDFStorageandQueryingSystemfor EnterpriseResourceManagement. *CIKM*,pp.484-491, 2004.