

Noname manuscript No.
(will be inserted by the editor)

A Survey of RDF Stores & SPARQL Engines for Querying Knowledge Graphs

Waqas Ali · Muhammad Saleem · Bin Yao ·
Aidan Hogan · Axel-Cyrille Ngonga Ngomo

Received: date / Accepted: date

Abstract Recent years have seen the growing adoption of non-relational data models for representing diverse, incomplete data. Among these, the RDF graph-based data model has seen ever-broadening adoption, particularly on the Web. This adoption has prompted the standardization of the SPARQL query language for RDF, as well as the development of a variety of local and distributed engines for processing queries over RDF graphs. These engines implement a diverse range of specialized techniques for storage, indexing, and query processing. A number of benchmarks, based on both synthetic and real-world data, have also emerged to allow for contrasting the performance of different query engines, often at large scale. This survey paper draws together these developments, providing a comprehensive review of the techniques, engines and benchmarks for querying RDF knowledge graphs.

Keywords Knowledge Graph · Storage · Indexing · Query Processing · SPARQL · Benchmarks

W. Ali
SEIEE, Shanghai Jiao Tong University, Shanghai, China
E-mail: waqasali@sjtu.edu.cn

M. Saleem
AKSW, University of Leipzig, Leipzig, Germany
E-mail: saleem@informatik.uni-leipzig.de

B. Yao
SEIEE, Shanghai Jiao Tong University, Shanghai, China
E-mail: yaobin@cs.sjtu.edu.cn

A. Hogan
DCC, University of Chile & IMFD, Santiago, Chile
E-mail: ahogan@dcc.uchile.cl

A.-C. Ngonga Ngomo
University of Paderborn, Paderborn, Germany
E-mail: axel.ngonga@upb.de

1 Introduction

The Resource Description Framework (RDF) is a graph-based data model where triples of the form (s, p, o) denote directed labeled edges $s \xrightarrow{p} o$ in a graph. RDF has gained significant adoption in the past years, particularly on the Web. As of 2019, over 5 million websites publish RDF data embedded in their webpages [27]. RDF has also become a popular format for publishing knowledge graphs on the Web, the largest of which – including Bio2RDF, DBpedia, PubChemRDF, UniProt and Wikidata – contain billions of triples. These developments have brought about the need for optimized techniques and engines for querying large RDF graphs. We refer to engines that allow for storing, indexing and processing joins over RDF as *RDF stores*.

A variety of query languages have been proposed for RDF, but the SPARQL Protocol and RDF Query Language (SPARQL) has become the standard [76]. The first version of SPARQL was standardized in 2008, while SPARQL 1.1 was released in 2013 [76]. SPARQL is an expressive language that supports not only joins, but also variants of the broader relational algebra (projection, selection, union, difference, etc.). Various new features were added in SPARQL 1.1, including *property paths* for matching paths in the RDF graph matching a given regular expression. Hundreds of SPARQL query services – called *endpoints* – have emerged on the Web since its standardization [35], with the most popular endpoints receiving millions of queries per day [164, 125]. We refer to engines that support storing, indexing and processing SPARQL (1.1) queries over RDF as *SPARQL engines*. Since SPARQL supports joins, we consider any SPARQL engine to also be an RDF store.

Efficient data storage, indexing and join processing are key to RDF stores (and thus, to SPARQL engines):

- *Storage*. Different engines may store RDF data using a variety of different structures (tables, graphs, etc.), encodings (integer IDs, string compression, etc.) and media (main memory, disk, etc.). Which storage to use may depend on the scale of the data, the types of query features supported, etc.
- *Indexing*. Various indexes are used in RDF stores for fast data lookup and query execution. Different types of indexes can support different operations with differing time-space trade-offs. More indexes generally lead to better query runtimes. However, maintaining such indexes can be costly in terms of space consumption and updates.
- *Join Processing*. At the core of evaluating queries lies efficient methods for processing joins. While traditional pairwise join algorithms – as used in relational databases – are most often employed, recent years have seen the emergence of novel techniques, such as multiway and worst-case optimal joins. Optimizing the order of evaluation of joins can also be important to ensure efficient processing.

Beyond processing joins, SPARQL engines must offer efficient support for more expressive query features:

- **Query Processing**. SPARQL is an expressive language containing a variety of query features beyond joins that need to be supported efficiently, such as filter expressions, optionals, path queries, etc.

RDF stores can further be divided into two categories: (1) local (single-node) stores that manage RDF data on a single machine and (2) distributed stores that partition RDF data among multiple machines. While local stores are more lightweight, the resources available on one machine limit scalability [211, 146, 86]. Various kinds of distributed RDF stores have thus been proposed [72, 86, 170, 171] that typically run on networked clusters of commodity, shared-nothing machines.

In this survey, we describe storage, indexing, join processing and query processing techniques employed by local RDF stores, as well as high-level strategies for partitioning RDF graphs as needed for distributed storage. An appendix attached to this extended version further compares 116 local and distributed RDF engines in terms of the techniques they use, as well as 12 benchmarks in terms of the types of data and queries they contain. The goal of this survey is to crisply introduce the different techniques used by RDF query engines, and also to help users to choose the appropriate engine or benchmark for a given use-case.

The remainder of the paper is structured as follows. Section 2 discusses and contrasts this survey with related studies. Section 3 provides preliminaries relating to RDF and SPARQL. Sections 4, 5, 6 and 7 review

techniques for storage, indexing, join processing and query processing, respectively. Section 8 explains different graph partitioning techniques for distributing storage over multiple machines. Section 9 concludes with research problems and future directions. An appendix attached to this extended version compares specific systems and benchmarks.

2 Literature Review

We first discuss related studies, including peer-reviewed tertiary literature (surveys in journals, short surveys in proceedings, book chapters, surveys with empirical comparisons, etc.) from the last 10 years collating techniques, engines and/or benchmarks for querying RDF.

We summarize the topics covered by these works in Table 1. We use ✓, ~ and blank cells to denote detailed, partial or little/no discussion, respectively, when compared with the current survey (the bottom row). We also include the number of engines and benchmarks included in the work; in some cases the respective publication does not formally list all systems/benchmarks (e.g., as a table), in which case we may write $n+$ as an estimate for the number discussed in the text.

Earlier surveys by Sakr et al. [163], Faye et al. [55], Luo et al. [119] focus on RDF storage for individual machines, particularly using relational storage. Another early survey, provided by Svoboda et al. [187], rather focuses on querying RDF data over the Web. More recent surveys by Ma et al. [122], Öszu [145] and Alaoui [8], Abdelaziz et al. [3], Elzein et al. [52], Yasin et al. [216], Wylot et al. [211], Janke and Staab [93], Zambom Santana and dos Santos Mello [218], and Chawla et al. [40] all focus on distributed, cloud and decentralized settings, with only brief summaries of standard techniques applicable on individual machines. With respect to more general discussions, Pan et al. [146] provide a short survey that serves as a quick overview of the area.

In summary, recent surveys have focused primarily on distributed (e.g., cloud) and decentralized (e.g., Web) environments. However, local RDF stores constitute the bulk of those found in practice [35], where novel storage, indexing and querying techniques continue to emerge. This survey focuses on techniques applicable for individual machines, rather summarizing partitioning schemes for RDF graphs in distributed environments. In the appendix, we present a comprehensive comparison of both local and distributed RDF query engines and benchmarks.

Table 1: Prior tertiary literature on RDF query engines; the abbreviations are: **Sto.**/Storage, **Ind.**/Indexing, **J.Pr.**/Join Processing, **Q.Pr.**/Query Processing, **Dis.**/Distribution, **Eng.**/Engines, **Ben.**/Benchmarks

Study	Year	Type	Techniques				Eng.	Bench.
			Sto.	Ind.	J.Pr.	Q.Pr.		
Chawla et al. [39]	2020	Survey	✓	~	~		46	9
Zambom & dos Santos [218]	2020	Short Survey	~		~		24	
Alaoui [8]	2019	Short Survey	✓				30+	
Janke et al. [91]	2018	Chapter	~	~	~	~	50+	9
Wylot et al. [211]	2018	Survey	✓	~	~	~	24	8
Pan et al. [146]	2018	Short survey	✓	~	~		25+	4
Yasin et al. [216]	2018	Short survey			~	✓	14	
Elzein et al. [52]	2018	Survey	✓		~	~	15+	
Abdelaziz et al. [3]	2017	Comparison	✓	~	~	~	21	4
Ma et al. [122]	2016	Survey	✓	~			17	6
Özsu [145]	2016	Survey	✓		~	~	35+	
Kaoudi et al. [96]	2015	Survey	✓	~	~	~	17	
Faye et al. [55]	2012	Survey	✓	~	~		13	
Luo et al. [119]	2012	Chapter	✓	✓			20+	
Svoboda et al. [187]	2011	Short Survey	~	~			14	6
Sakr et al. [163]	2010	Short Survey	✓	~			10+	
Ali et al.		Survey	✓	✓	✓	✓	116	12

3 Preliminaries

Before beginning the core of the survey, we first introduce some preliminaries regarding RDF and SPARQL.

3.1 RDF

The RDF data model [175] is founded on three pairwise disjoint (and countably infinite) sets of *RDF terms*: the set **I** of *International Resource Identifiers (IRIs)* used to identify resources; the set **L** of *literals* used for (possibly language-tagged) strings and datatype values; and the set **B** of *blank nodes*, interpreted as local existential variables. An *RDF triple* $(s, p, o) \in \mathbf{IB} \times \mathbf{I} \times \mathbf{IBL}$ contains a subject s , a predicate p and an object o .¹ A set of RDF terms is called an *RDF graph* G , where each triple $(s, p, o) \in G$ represents a directed labeled edge $s \xrightarrow{p} o$. We denote by $s(G)$, $p(G)$ and $o(G)$ the set of subjects, predicates and objects in G . We further denote by $so(G) := s(G) \cup o(G)$ the set of *nodes* in G .

An example RDF graph, representing information about two university students, is shown in Figure 1. We include both a graphical representation and a triple-based representation. RDF terms such as `:DB`, `foaf:age`, etc., denote prefixed IRIs, where if we define the prefix `foaf` as `http://xmlns.com/foaf/0.1/`, then the full IRI for `foaf:age` is `http://xmlns.com/foaf/0.1/age`.² Terms such as `"Motor RDF"@es` denote strings with (optional) language tags, and terms such as `"21"^^xsd:int`

denote datatype values. Finally we denote blank nodes with the underscore prefix, where `_:p` refers to the existence of a project shared by Alice and Bob. Terms used in the predicate position (e.g., `foaf:age`, `skos:broader`) are known as *properties*. The RDF standard defines the special property `rdf:type`, which indicates the *class* (e.g., `foaf:Person`, `foaf:Project`) of a resource.

3.2 SPARQL

Various query languages for RDF have been proposed down through the years, such as RQL [98], SeRQL [184], etc. Some of the techniques and systems discussed later were developed for such languages. We focus our discussion on SPARQL [76], which would later become the standard language for query RDF data, referring to the discussion of Haase et al. [71] on its predecessors. Herein we define the core of the SPARQL query language in terms of *basic graph patterns*, *navigational graph patterns*, and *complex graph patterns* [12].

Basic Graph Patterns (BGPs) At the core of SPARQL lie *triple patterns*: RDF triples that also allow variables from the set **V** (disjoint with **IBL**) in any position. A set of triple patterns is called a *basic graph pattern* (BGP). Since blank nodes in BGPs act as variables, we assume they have been replaced with variables. We denote by $\text{vars}(B)$ the set of variables in the BGP B .

Given an RDF graph G , the *evaluation* of a BGP B , denoted $B(G)$, returns a set of *solution mappings*. A solution mapping μ is a partial mapping from the set **V** of variables to the set of RDF terms **IBL**. We denote by

¹ In this paper, we may abbreviate $S \cup T$ as ST .

² We use the blank prefix (e.g., `:DB`) as an arbitrary example. Other prefixes used can be found from `http://prefix.cc/`.

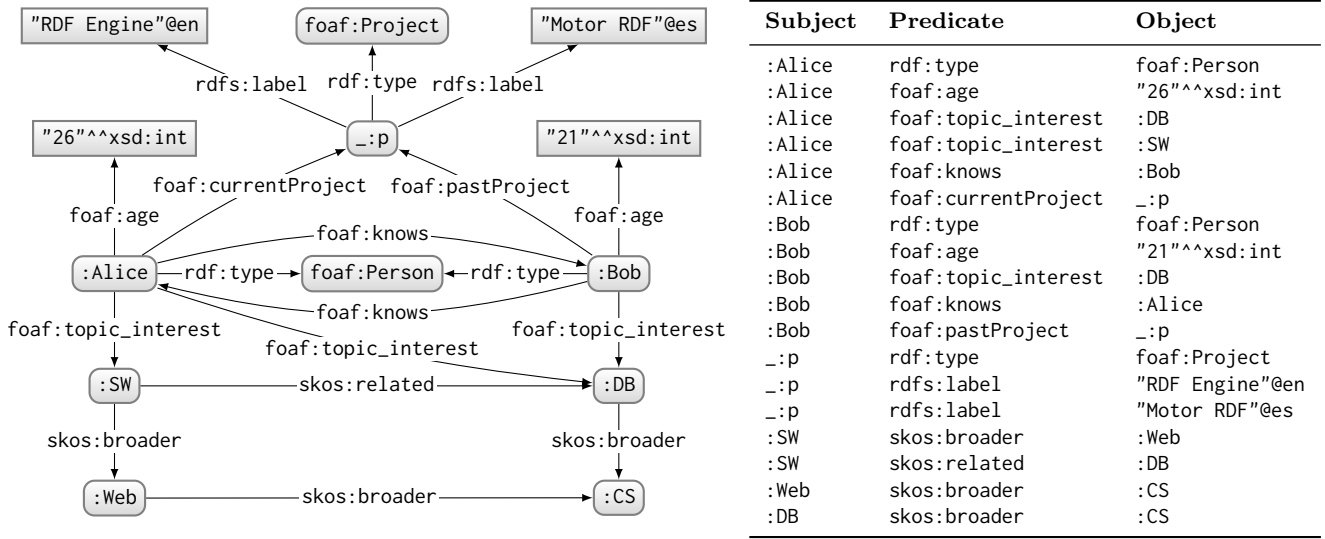


Fig. 1: Graphical (left) and triple-based representation (right) of an example RDF graph

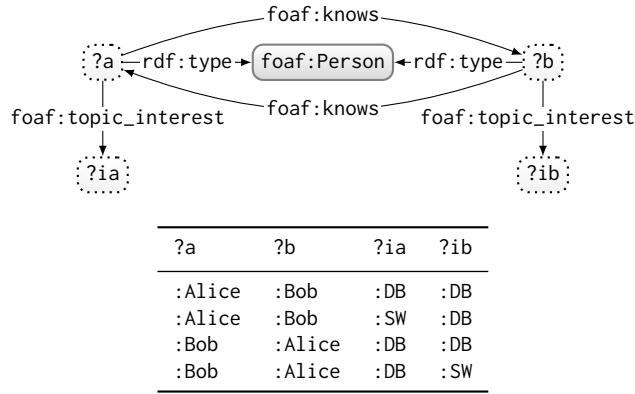


Fig. 2: Example BGP (above) and its evaluation over the graph of Figure 1 (below)

$\text{dm}(\mu)$ the set of variables for which μ is defined. Given a triple pattern t , we denote by $\mu(t)$ the image of t under μ ; i.e., the result of replacing any variable $v \in \text{dm}(\mu)$ appearing in t with $\mu(v)$. We denote by $\mu(B)$ the image of the BGP B under μ ; i.e., $\mu(B) := \{\mu(t) \mid t \in B\}$. The evaluation of the BGP B on G is then given as $B(G) := \{\mu \mid \mu(B) \subseteq G \text{ and } \text{dm}(\mu) = \text{vars}(B)\}$.

We provide an example of a BGP and its evaluation in Figure 2. Each row of the results refers to a solution mapping. Some solutions map different variables to the same term; each such solution is thus a homomorphism from the BGP to the RDF graph. BGPs, when extended with projection, can further be seen as *conjunctive queries* for RDF graphs [12].

Navigational Graph Patterns (NGPs) A key characteristic of graph query languages is the ability to match

paths of arbitrary length [12]. In SPARQL (1.1), this ability is captured by *property paths* [76], which are regular expressions that paths should match. The set \mathbf{E} of *path expressions* can be defined recursively as follows:

- if p is an IRI, then p is a path expression (*property*);
- if e is a path expression, then e (*inverse*), e^* (*zero-or-more*, aka. *Kleene star*), e^+ (*one-or-more*), and $e?$ (*zero-or-one*) are path expressions.
- If e_1, e_2 are path expressions, then e_1/e_2 (*concatenation*) and $e_1|e_2$ (*disjunction*) are path expressions.
- if P is a set of IRIs, then $!P$ and $!^P$ are path expressions (*negated property set*);³

Brackets can also be used to indicate precedence.


The evaluation of these path expressions on an RDF graph G returns pairs of nodes in G connected by paths that match the expression, as defined in Table 2. These path expressions are akin to *2-way regular path queries* (2RPQs) extended with negated property sets [107, 12].

We call a triple pattern (s, e, o) that further allows a path expression as the predicate (i.e., $e \in \mathbf{E}$) a *path pattern*. A *navigational graph pattern* (NGP) is then a set of path patterns. Given a navigational graph pattern N , let $\text{paths}(N) := \text{p}(N) \cap \mathbf{E}$ denote the set of path expressions used in N . Given an RDF graph G and a set of path expressions $E \subseteq \mathbf{E}$, we denote by $G_E := G \cup (\bigcup_{e \in E} \{(s, e, o) \mid (s, o) \in e(G)\})$ the result of materializing all paths matching E in G . The evaluation of the navigational graph pattern N on G is then $N(G) := \{\mu \mid \mu(N) \subseteq G_{\text{paths}(N)} \text{ and } \text{dm}(\mu) = \text{vars}(N)\}$.

³ SPARQL uses the syntax $!(p_1 \dots p_k | p_{k+1} \dots p_n)$ which can be written as $!P|^{P'}$, where $P = \{p_1, \dots, p_k\}$ and $P' = \{p_{k+1}, \dots, p_n\}$ [76, 107].

Table 2: Evaluation of path expressions

$p(G) := \{(s, o) \mid (s, p, o) \in G\}$
$^{\wedge}e(G) := \{(s, o) \mid (o, s) \in e(G)\}$
$e^+(G) := \{(y_1, y_n) \mid \text{for } 1 \leq i < n : \exists(y_i, y_{i+1}) \in e(G)\}$
$e?(G) := e(G) \cup \{(x, x) \mid x \in \text{so}(G)\}$
$e^*(G) := (e^+)(G)$
$e_1/e_2(G) := \{(x, z) \mid \exists y : (x, y) \in e_1(G) \wedge (y, z) \in e_2(G)\}$
$e_1 \mid e_2(G) := e_1(G) \cup e_2(G)$
$!P(G) := \{(s, o) \mid (s, p, o) \in G \wedge p \notin P\}$
$!^{\wedge}P(G) := \{(s, o) \mid (o, p, s) \in G \wedge p \notin P\}$



?a	?b
:Alice	:Alice
:Alice	:Bob
:Bob	:Alice
:Bob	:Bob

Fig. 3: Example NGP (above) and its evaluation over the graph of Figure 1 (below)

We provide an example of a navigational graph pattern and its evaluation in Figure 3.

Complex Graph Patterns (CGPs) Evaluating BGPs and NGPs returns sets of solution mappings. Per Figures 2 and 3, these results can be viewed as *relations* (i.e., tables) where variables are *attributes* (i.e., column names) and *tuples* (i.e., rows) contain the RDF terms bound by each solution mapping. We can then transform and combine the results of BGPs/NGPs with the relational algebra. Formally, let M, M_1, M_2 denote sets of solution mappings and let V denote a set of variables. Let R denote a *filter expression* that maps a solution mapping μ to a true or false value [76]. In Table 3 we define the core relational algebra of SPARQL, including, respectively, *selection* (FILTER), *projection* (SELECT), *natural (equi-)join* (\cdot), *union* (UNION), *semi-join* (FILTER EXISTS), *anti-join* (FILTER NOT EXISTS) and *left-outer-join* (OPTIONAL). By *complex graph patterns* (CGPs) [12] we refer to the closure of BGPs and the relational algebra (akin to *first-order queries on graphs* [12]). By *complex navigational graph patterns* (CNGPs) [12] we refer to the closure of NGPs and the relational algebra.

Named graphs SPARQL allows for querying multiple RDF graphs through the notion of a *SPARQL dataset*

Table 3: Core relational algebra of SPARQL

$\sigma_R(M) := \{ \mu \in M \mid R(\mu) \}$
$\pi_V(M) := \{ \mu' \mid \exists \mu \in M : \mu \sim \mu' \wedge \text{dm}(\mu') = V \cap \text{dm}(\mu) \}$
$M_1 \bowtie M_2 := \{ \mu_1 \cup \mu_2 \mid \mu_1 \in M_1 \wedge \mu_2 \in M_2 \wedge \mu_1 \sim \mu_2 \}$
$M_1 \cup M_2 := \{ \mu \mid \mu \in M_1 \vee \mu \in M_2 \}$
$M_1 \ltimes M_2 := \{ \mu_1 \in M_1 \mid \exists \mu_2 \in M_2 : \mu_1 \sim \mu_2 \}$
$M_1 \triangleright M_2 := \{ \mu_1 \in M_1 \mid \nexists \mu_2 \in M_2 : \mu_1 \sim \mu_2 \}$
$M_1 \bowtie M_2 := (M_1 \bowtie M_2) \cup (M_1 \triangleright M_2)$

$D := \{G, (n_1, G_1), \dots, (n_k, G_k)\}$ where G, G_1, \dots, G_n are RDF graphs; n_1, \dots, n_k are pairwise distinct IRIs; G is known as the *default graph*; and each pair of the form (n_i, G_i) (for $1 \leq i \leq n$) is known as a *named graph*. Letting N', N'' denote sets of IRIs, n', n'' IRIs and v a variable, SPARQL then provides a number of features for querying different graphs:

- FROM N' FROM NAMED N'' : activates a dataset with a default graph composed of the merge of all graphs G' such that $(n', G') \in D$ and $n' \in N'$, and the set of all named graphs $(n'', G'') \in D$ such that $n'' \in N''$;
- GRAPH n' : evaluates a graph pattern on the graph G' if the named graph (n', G') is active;
- GRAPH v : takes the union of the evaluation of a graph pattern over each G' such that (n', G') is active, binding v to n' for each solution generated from G' ;

Without FROM or FROM NAMED, the active dataset is the indexed dataset D . Without GRAPH, graph patterns are evaluated on the active default graph.

SPARQL engines may disallow empty graphs from named graphs, resulting in a *quad store*. In this case, a dataset $D := \{G, (n_1, G_1), \dots, (n_k, G_k)\}$ is seen as a set of quads $D = G \times \{\star\} \cup (\bigcup_{(n_i, G_i) \in D} G_i \times \{n_i\})$, where we use $\star \notin \mathbf{IBL}$ as a special symbol to indicate the default graph. In this case, a quad (s, p, o, n) denotes a triple (s, p, o) in the default graph if $n = \star$, or a triple in the named graph G' such that $(n, G') \in D$ if $n \in \mathbf{I}$. We can define CNGPs involving quad patterns analogously.

Other SPARQL features SPARQL supports features beyond CNGPs such as aggregation, solution modifiers (ordering and slicing solutions), bag semantics (preserving result multiplicity), federation (returning solutions from remote services), entailment and more besides. SPARQL further supports different query types, such as SELECT, which returns a sequence of solution mappings; CONSTRUCT, which returns an RDF graph based on the solution mappings; DESCRIBE, which returns an RDF graph describing indicated RDF terms; and ASK, which returns true if some solution mapping is found, or false otherwise. In general, however, CNGPs form the core of the SPARQL query language and have been the focus of the majority of research on query processing

and optimization. For this reason, and for reasons of conciseness, our preliminaries have focused on CNGBs.

4 Storage

Data storage refers to how data are represented in memory. Different storage mechanisms store different elements of data contiguously in memory, leading to different trade-offs in terms of compression and efficient data access. This section reviews six categories of storage mechanisms used by RDF engines, where the first four leverage relational storage while the latter two depend on native storage and other non-relational stores.

4.1 Triple table

A *triple table* stores an RDF graph G as a single ternary relation, often in a relational database. Figure 1 shows an RDF graph with its triple table on the right-hand side. One complication is that relational engines typically assume a column to be associated with a single type, which may not be true for RDF objects in particular; a workaround is to store a string encoding of the terms, though this may complicate their ordering.

Rather than storing full RDF terms in the triple table, stores may apply dictionary encoding, where RDF terms are mapped one-to-one with numeric object identifiers (OIDs), with OIDs being stored in the table and decoded using the dictionary as needed. Since OIDs consume less memory and are faster to process than strings, such an approach works better for queries that involve many intermediate results but generate few final results; on the other hand, such an approach suffers when queries are simple and return many results (decoding RDF terms will dominate performance), or when selective filters are specified that require the data of the term (e.g., finding names containing “z” may require decoding all names before filtering). To find a better trade-off, some RDF engines (e.g., Jena 2 [203]) only use OIDs for strings with lengths above a threshold.

The most obvious physical storage is to store triples contiguously (row-wise). This allows for quickly retrieving the full triples that match (e.g.) a given triple pattern. However, some RDF engines based on relational storage (e.g., Virtuoso [54]) rather use (or provide an option for) column-wise storage, where the values along a column are stored contiguously, often following a particular order. Such column-wise storage allows for better compression, and thus for caching more data in memory; it offers performance benefits when queries require reading many values from a particular column

rdf:type		foaf:age	
Subject	Object	Subject	Object
:Alice	foaf:Person	:Alice	26
:Bob	foaf:Person	:Bob	21
_:p	foaf:Project		

Fig. 4: A sample of two tables for the vertical partitioning of the RDF graph of Figure 1

(e.g., for aggregations) but may be slower when queries need to match and retrieve entire triples.

Triple tables can be straightforwardly extended to quad tables in order to support SPARQL datasets, as is done in a variety of engines [54, 75].

4.2 Vertical partitioning

Instead of storing an RDF graph G as one ternary relation, the *vertical partitioning* approach [1] uses a binary relation for each property $p \in p(G)$ whose tuples encode subject-object pairs for that property. In Figure 4 we exemplify two such relations. Physical storage can again use OIDs, row-based or column-based storage, etc.

Some benefits of this approach versus triple tables include relations with fewer rows, more specific domains for columns (e.g., the object column for `foaf:age` can be defined as an integer type, which would be particularly beneficial for column-wise storage), and fewer self-joins. A key limitation of this approach is that it becomes complicated to evaluate triple patterns with variable predicates, which require applying a union on all relations in the database. Furthermore, RDF graphs may have upwards of thousands of properties [198], which may lead to a schema with many relations.

Vertical partitioning can be used to store quads by adding a **Graph** column to each table [54, 75].

4.3 Extended vertical partitioning

S2RDF [171] proposes a storage and indexing technique called *extended vertical partitioning* based on *semi-join reductions* (we recall from Table 3 that a semi-join $M_1 \bowtie M_2$, aka. FILTER EXISTS, returns the tuples in M_1 that are “joinable” with M_2). Letting $\mathbf{x}, \mathbf{y}, \mathbf{z}$ denote variables and p, q denote RDF terms, then for each property pair $(p, q) \in p(G) \times p(G)$ such that $p \neq q$, extended vertical partitioning stores three semi-join reductions:

1. $\{(\mathbf{x}, p, \mathbf{y})\}(G) \bowtie \{(\mathbf{y}, q, \mathbf{z})\}(G)$ (O-S),
2. $\{(\mathbf{x}, p, \mathbf{y})\}(G) \bowtie \{(\mathbf{x}, q, \mathbf{z})\}(G)$ (S-S),
3. $\{(\mathbf{x}, p, \mathbf{y})\}(G) \bowtie \{(\mathbf{z}, q, \mathbf{x})\}(G)$ (S-O).

skos:broader \bowtie_{s-s} skos:related		skos:broader \bowtie_{s-o} skos:related	
Subject	Object	Subject	Object
:SW	:Web	:DB	:CS

skos:related \bowtie_{o-s} skos:broader		skos:related \bowtie_{s-s} skos:broader	
Subject	Object	Subject	Object
:SW	:DB	:SW	:DB

Fig. 5: Example semi-join reduction for two properties

The semi-join $\{(x, p, y)\}(G) \bowtie \{(z, q, y)\}(G)$ (O-O) is not stored as most O-O joins have the same predicate, and thus would occur in the same relation. In Figure 5 we give an example of a semi-join reduction for `skos:broader` and `skos:related` over the running example in both directions; empty semi-joins are omitted.

In comparison with vertical partitioning, observing that $(M_1 \bowtie M_2) \bowtie (M_2 \bowtie M_1) \equiv M_1 \bowtie M_2$, we can apply joins over the corresponding semi-join reductions knowing that each tuple read from each side will contribute to the join, thus reducing I/O, which is particularly useful for distributed storage (reducing network I/O). The cost involves storing (and updating) each tuple in up to $3(|p(G)| - 1)$ additional relations; omitting empty semi-joins can help to mitigate this issue in practice [171]. Like vertical partitioning, extended vertical partitioning also presents complications for variable predicates, RDF graphs with many properties, etc.

4.4 Property table

Property tables aim to emulate the n -ary relations typical of relational databases. A property table usually contains one subject column, and n further columns to store objects for the corresponding properties of the given subject. The subject column then forms a primary key for the table. The tables to define can be based on classes, clustering [153], coloring [29], etc., to group subjects with common properties. We provide an example of a property table based on the class `foaf:Person` in Figure 6 for the RDF graph of Figure 1.

Property tables allow for storing and retrieving multiple triples with a given subject as one tuple (e.g., to find people with `age < 30` and `interest = :SW`) without requiring joins; such queries are common in practice [164]. As per vertical partitioning, property tables often store terms of the same type in the same column, enabling better compression and optimizations for column-wise storage. Complications, however, arise for multi-valued (...to-many) or optional (zero-to-...)

foaf:Person					
Subject	age	topic	knows	cProj	pProj
:Alice	26	:DB	:Bob	_:p	NULL
:Bob	21	:DB	:Alice	NULL	_:p

Fig. 6: Example property table for people

properties. In the example of Figure 1, Alice is also interested in SW, which does not fit in the cell. Furthermore, Alice has no past project, and Bob has no current project, leading to nulls. While multi-valued and optional properties in the graph can be normalized into separate relations, effectively using property graphs for a subset of the graph (per, e.g., Jena2 [203]), changes to the graph may require renormalization; for example, even though each person currently has only one value for `knows`, adding that Alice knows another person would require renormalizing the tables. As per vertical partitioning, complications also arise when considering triple patterns with variable predicates, RDF graphs with many properties or classes, etc.

Storing SPARQL datasets using property tables is also quite complicated. One approach might be to add for each column (e.g., `age`) a second column indicating the graph that indicates the age for the given subject (e.g., `age_g`), but this would not support cases where the same triple appears in multiple named graphs. A more general but highly-redundant solution would be to rather duplicate information into, e.g., a quad table.

4.5 Graph-based storage

While the previous three storage mechanisms rely on relational storage, *graph-based storage* is adapted specifically for the graph-based model of RDF and the graph-pattern-based queries of SPARQL. There are a number of key characteristics of such models that can be exploited for storage, including the notion of *adjacency* between connected nodes, the fixed *arity* of graphs, etc.

Unlike in the relational model, graph models have bounded arity (3 for triples, 4 for quads). This can be exploited to build specialized storage. Some engines (e.g., 4store [75], YARS2 [78]) build *native triple/quad tables*, which differ from relational triple/quad tables in that they have fixed arity, fixed attributes (S,P,O(G)), and more general domains (e.g., the O column can contain any RDF term). Such tables may or may not use OID encodings. Another type of storage involves using tensors of bits with order equal to the arity of the relation; a simple variant is to use 2-dimensional adjacency matrices for each property (e.g., BitMat [15], BMatrix [30], QDags [137]), as exemplified in Figure 7, which

skos:broader				
	:CS	:DB	:SW	:Web
:CS	0	0	0	0
:DB	1	0	0	0
:SW	0	0	0	1
:Web	1	0	0	0

Fig. 7: Example bit matrix for skos:broader

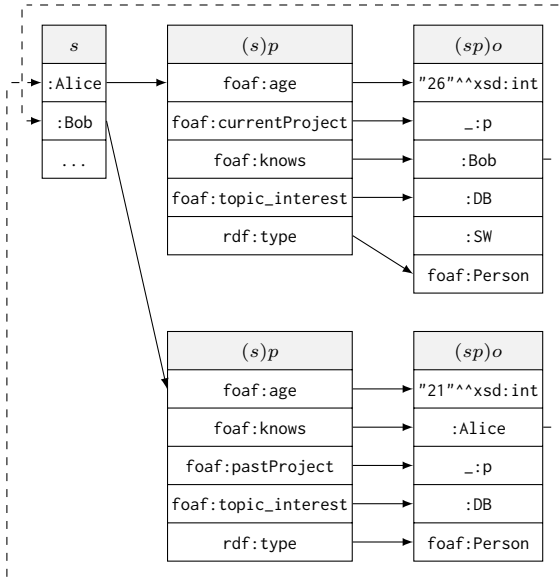


Fig. 8: Example adjacency list for two subjects with dashed links indicating index-free adjacency pointers

constitutes a binary form of vertical partitioning. While such an approach may lead to sparse matrices/tensors, compact data structures offer compressed representations that also support efficient operations [15, 30, 137].

Graphs can feature a high degree of local repetition that can be compressed using storage techniques based on (nested) *adjacency lists* (see, e.g., Hexastore [201], gStore [223], SpiderStore [25], Trinity.RDF [219], and GRaSS [120]). These lists are akin to tries, where subject or subject-predicate prefixes are followed by the rest of the triple. Such tries can be stored row-wise, where blocks of triples with a common prefix are stored together; or column-wise, where all subjects are stored contiguously, with each subject pointing to a contiguous predicate block, with pointers in turn to a contiguous object block. *Index-free adjacency* can be used to support efficient navigation, where terms in the suffix directly point to the location on disk of their associated prefix. We refer to Figure 8 for an example. Such structures can be naturally extended to include edges in the other direction (as used by Trinity.RDF [219] and GRaSS [120]). Such approaches can also be seen as storing subgraphs around a central node [120].

L_V	
Node	Attributes
:Alice	{(foaf:age, "26"^^xsd:int)}
:Bob	{(foaf:age, "21"^^xsd:int)}
:DB	{}
...	...
L_E	
Edge	Label
(:Alice, :Bob)	{foaf:knows}
(:Alice, :DB)	{foaf:topic_interest}
(:Bob, :Alice)	{foaf:knows}
...	...

Fig. 9: Example of the multi-graph representation

Another alternative is to decompose an RDF graph into its constituent graph components and store these separately. One such technique is the *multigraph representation* proposed for the AMBER engine [87], where an RDF graph G is decomposed into a set of (non-literal) nodes $V := \text{so}(G) \cap \mathbf{IB}$, a set of (non-literal) edges $E := \{(s, o) \in V \times V \mid \exists p : (s, p, o) \in G\}$, an edge-labelling function of the form $L_E : V \rightarrow 2^{\mathbf{I}}$ such that $L_E(s, o) := \{p \mid (s, p, o) \in G\}$, and an attribute-labelling function of the form $L_V : \mathbf{IB} \rightarrow 2^{\mathbf{I} \times \mathbf{L}}$ such that $L_V(s) := \{(p, o) \mid (s, p, o) \in G \wedge o \in \mathbf{L}\}$. We provide an example in Figure 9 (in practice, AMBER dictionary-encodes nodes, attributes, and edge-labels).

4.6 Miscellaneous storage

Aside from relational-based and graph-based storage, other engines have proposed to leverage other forms of storage as implemented by existing systems. A common example is the use of NoSQL key-value, tabular or document stores for distributed storage [91, 218].

4.7 Discussion

Early works on storing RDF tended to rely on relational storage, which had been subject to decades of developments and optimizations before the advent of RDF (e.g., [203, 1, 54]). Though such an approach still has broad adoption [54], more recent storage techniques aim to exploit the graph-based characteristics of RDF – and SPARQL – in order to develop dedicated storage techniques (e.g., [15, 201, 223]). A more recent trend is to leverage NoSQL storage (e.g., [110, 147, 19]) in order to distribute the management of RDF data; we will discuss distribution in more detail in Section 8.

5 Indexing

Indexing enables efficient lookup operations on RDF graphs (i.e., $O(1)$ or $O(\log |G|)$ time to return the first result or an empty result). The most common such operation is to find all triples that match a given triple pattern. However, indexes can also be used to match non-singleton BGPs (with more than one triple pattern), to match path expressions, etc. We now discuss indexing techniques proposed for RDF graphs.

5.1 Triple indexes

The goal of *triple-pattern indexes* is to efficiently find all triples that match a given triple pattern. Specifically, letting s, p, o denote RDF terms and $\mathbf{s}, \mathbf{p}, \mathbf{o}$ variables, there are $2^3 = 8$ abstract triple patterns: $(\mathbf{s}, \mathbf{p}, \mathbf{o})$, $(\mathbf{s}, \mathbf{p}, o)$, $(\mathbf{s}, p, \mathbf{o})$, $(s, \mathbf{p}, \mathbf{o})$, (s, p, o) , $(\mathbf{s}, \mathbf{p}, o)$, (s, p, \mathbf{o}) and (s, p, o) . Unlike relational databases, where often only the primary key of a relation will be indexed by default and further indexes must be specified by the database administrator, most RDF stores aim to have a complete index by default, covering all eight possible triple patterns. However, depending on the type of storage chosen, this might not always be feasible.

When a storage scheme such as vertical partitioning is used [1], only the five patterns where the predicate is constant can be efficiently supported (variable predicates require a union over all relations, of which there $O(n)$ in the worst case). These five patterns can be supported efficiently by indexing the subject and object columns. In the case that the RDF graph is stored as a (binary) adjacency matrix for each property [15, 137], again only constant-predicate patterns can be efficiently supported. Specialized indexes can be used to quickly evaluate such patterns, where QDags [137] uses *quadrees*: a hierarchical index structure that recursively divides the matrix into four sub-matrices; we provide an example quadtree in Figure 10. A similar structure, namely a k^2 -tree, is used by BMatrix [30].

Otherwise, in triple tables, or similar forms of graph-based storage, all triple patterns can be efficiently supported by using different *triple permutations*. As an example, Figure 8 illustrates a single SPO permutation. A total of $3! = 6$ permutations are possible and suffice to cover all eight abstract triple patterns assuming that the index structure permits prefix lookups; for example, in an SPO permutation we can efficiently support four abstract triple patterns $(\mathbf{s}, \mathbf{p}, \mathbf{o})$, $(s, \mathbf{p}, \mathbf{o})$, $(\mathbf{s}, p, \mathbf{o})$ and (s, p, o) as we require the leftmost terms of the permutation to be filled. Such index permutations can be implemented using standard data structures such as ISAM




Fig. 10: Example of a quadtree index based on the bit matrix of Figure 7; the root represents the full matrix, while children denote four sub-matrices of the parent; a node is colored black if it contains only 1's, white if it contains only 0's, and gray if it contains both; it is only necessary to add children for gray nodes

files [78], B(+)Trees [142], AVL trees [205], tries/adjacency lists [201] (see Figure 8), etc. In fact, with only $\binom{3}{\lfloor 3/2 \rfloor} = 3$ permutations – e.g., SPO, POS and OSP – we can cover all eight abstract triple patterns.⁴

Recent works have proposed methods to reduce the redundancy in the index permutations. In particular, RDFCSA [32] uses a *compact suffix-array (CSA)* such that one permutation suffices to efficiently support all triple patterns. Intuitively speaking, triples can be indexed cyclically in a CSA, such that in an SPO permutation, one can continue from O back to S, thus covering SPO, POS and OSP permutations in one CSA index.

5.2 Entity-based indexes

By *entity-based indexes*, we refer to indexes that support the efficient evaluation of graph patterns that “center on” a particular entity. In particular, BGPs can be reduced to joins over their triple patterns; for example, $\{(\mathbf{x}, p, \mathbf{y}), (\mathbf{y}, q, \mathbf{z})\}(G) = \{(\mathbf{x}, p, \mathbf{y})\}(G) \bowtie \{(\mathbf{y}, q, \mathbf{z})\}(G)$. A common type of join targeted by such indexes is that of *star joins*, typically defined to be a join on a common subject, e.g., $\{(\mathbf{w}, p, \mathbf{x}), (\mathbf{w}, q, \mathbf{y}), (\mathbf{w}, r, \mathbf{z})\}$. Some authors allow star joins that include s-o joins on the common variable, e.g., $\{(\mathbf{w}, p, \mathbf{x}), (\mathbf{w}, q, \mathbf{y}), (\mathbf{z}, r, \mathbf{w})\}$ [120]. Star joins can thus be seen as querying on/for the data surrounding a particular entity (in this case \mathbf{w}). Entity-based indexes permit efficient evaluation of such joins.

Star joins can be efficiently supported through property table storage, so long as the relevant property tables for the star join can be found efficiently, and there are indexes on the relevant columns (e.g., for p, q and/or r). The EAGRE system [221] uses an index for property tables where entities with n properties are encoded in n -dimensional space. A space-filling curve – such as a Z-order curve or a Hilbert curve – is then

⁴ However, with 3 permutations we may lose the ability to iterate over results in order; e.g., given SPO, POS and OSP, and a pattern $(s, \mathbf{p}, \mathbf{o})$, we cannot iterate over the results in order of object, which would be possible with SOP, and which would enable more join algorithms (e.g., merge joins [142]).




Fig. 11: Space-filling indexing with a Hilbert curve

used to index the space. We illustrate the idea in Figure 11, where four entities are indexed (abbreviating :Alice, :Bob, :Carol, :Dave) with respect to two dimensions (say foaf:age for x and integer-encoded values of foaf:knows for y). We show the first-, second- and third-order Hilbert curves from left to right. Letting d denote the number of dimensions, the n^{th} -order Hilbert curve assigns an ordinal to 2^{dn} regions of the space based on the order in which it visits the region; for example, starting with the curve on the bottom left, then :A is in the region of ordinal 2, 7 and 26, respectively, as the order increases. The space-filling curve thus “flattens” multidimensional data into one dimension (the ordinal), which can be indexed sequentially.

Indexes based on property tables encounter complications when presented with multi-valued properties, missing values, etc. A more flexible approach is to index *signatures* of entities, which are bit-vectors encoding the set of property-value pairs of the entity. One such example is the *vertex signature tree* of gStore [223], which first encodes all outgoing (p, o) pairs for a given entity into a bit vector similar to a Bloom filter, and then indexes these bit vectors hierarchically allowing for fast, approximate containment checks that quickly find candidate entities given a subset of such pairs. A similar approach is adopted by the GRaSS engine [120], which optimizes for star subgraphs that include both outgoing and incoming edges on nodes, where a custom *FDD-index* allows for efficient retrieval of the subgraphs containing a triple that matches a triple pattern; each subgraph is further assigned a signature that encodes elements of the subgraph similar to gStore.

5.3 Property-based indexes

Returning to the star join $\{(\mathbf{w}, p, \mathbf{x}), (\mathbf{w}, q, \mathbf{y}), (\mathbf{w}, r, \mathbf{z})\}$, another way to quickly return candidate bindings for the variable \mathbf{w} is to index nodes according to their adjacent properties; then we can find nodes that have at least the adjacent properties p, q, r . Such an approach is used by RDFBroker [179], which defines the signature of a node s as $\Sigma(s) = \{p \mid \exists o : (s, p, o) \in G\}$; for example, the signature of :SW in Figure 1 would be $\Sigma(\text{:SW}) = \{\text{skos:broader}, \text{skos:related}\}$ (corresponding to *char-*




Fig. 12: Lattice of node signatures with abbreviated terms (e.g., s:b denotes skos:broader)

acteristic sets proposed later [139]). A property table is then created for each signature. At query time, property tables whose signatures subsume $\{p, q, r\}$ are found using a lattice of signatures. We provide an example in Figure 12 with respect to the RDF graph of Figure 1, where children subsume the signatures of their parent.

AxonDB [132] uses *extended characteristic sets* such that each triple (s, p, o) in the RDF graph is indexed along with the signatures (i.e., characteristic sets) of its subject and object; i.e., $(\Sigma(s), \Sigma(o))$. Thus for example, the triple $(\text{:SW}, \text{skos:related}, \text{:DB})$ of Figure 1 would be indexed alongside the extended characteristic set $(\{\text{skos:broader}, \text{skos:related}\}, \{\text{skos:broader}\})$. This index then allows for efficiently identifying two star joins that are connected by a given property p .

5.4 Path indexes

A path join involves successive s-o joins between triple patterns; e.g., $\{(\mathbf{w}, p, \mathbf{x}), (\mathbf{x}, q, \mathbf{y}), (\mathbf{y}, r, \mathbf{z})\}$, where the start and end nodes (\mathbf{w}, \mathbf{z}) may be variables or constants. While path joins have fixed length, path queries in navigational graph patterns may further match arbitrary length paths. A number of indexing approaches have been proposed to speed up querying paths.

A path can be seen as a string of arbitrary length; e.g., a path $\{(w, p, x), (x, q, y), (y, r, z)\}$ can be seen as a string $wpxqyrz\$$, where $\$$ indicates the end of the string; alternatively, if intermediate nodes are not of importance, the path could be represented as the string $wpqrz\$$. The Yaanii system [37] builds an index of paths of the form $wpxqyrz\$$ that are clustered according to their template of the form $wpqrz\$$. Paths are then indexed in B+trees, which are partitioned by template. Fletcher et al. [56] also index paths in B+trees, but rather than partition paths, they apply a maximum length of at most k for the paths included. The fact that paths are analogous to strings suggests that text indexing techniques can be applied for paths. Along these lines, Maharjan et al. [124] and the HPRD system [117] both leverage suffix arrays – a common indexing technique for text – to index paths. The downside of such approaches is that they may index an exponential num-




Fig. 13: Preorder and postorder on a `skos:narrower` tree; e.g., `:CS` has preorder 1 and postorder 7




Fig. 14: PLSD index on an example `skos:narrower` hierarchy; terms (e.g., `:CS`) are indexed externally

ber of paths; in the case of HPRD, for example, users are thus expected to specify which paths to index [117].

Other path indexes are inspired by prior works on efficient evaluation of path queries over trees (e.g., for XPath). Bartoń [20] proposes a tree-based index based on preorder and postorder traversal. A preorder traversal starts at the root and traverses children in a depth-first manner from left to right. A postorder traversal starts at the leftmost leaf and traverses all children, from left to right, before moving to the parent. We provide an example preorder and postorder in Figure 13. Given two nodes m and n in the tree, an interesting property is that m is a descendant of n if and only if m is greater than n for preorder and less than n for postorder. Bartoń [20] uses this property to generate an index on ascending preorder so as to linearize the tree and quickly find descendants based on postorder. A disadvantage is that RDF graphs are not trees, where Bartoń thus proposes a decomposition of the graph into a forest of trees that are then indexed [20].

Another type of path index, called *PLSD*, is used in System II [207] for indexing the transitivity of a single property, allowing to efficiently answer path queries of the form (s, p^*, o) , or (s, p^*, o) , etc. For a given property p , each incident (subject or object) node x is assigned a triple of numbers $(i, j, k) \in \mathbb{N}^3$, where i is a unique prime number that identifies the node x , j is the least common multiple of the i -values of x 's parents (i.e., nodes y such that $(y, p, x) \in G$), and k is the least common multiple of the k -values of x 's parents and the i -value of x . We provide an example in Figure 14. PLSD can further handle cycles by multiplying the k -value of all nodes by the i value of all nodes in its strongly-connected component. Given the i -value of a node, the i -values of its parents and ancestors can be retrieved by factorizing j and k/i respectively. The disadvantage is that the least common multiple can reach a large number, where no polynomial time algorithm is known for factorization of binary input numbers.

Gubichev et al. [64] use a path index of directed graphs, called FERRARI [177], for each property in an RDF graph. First, a *condensed graph* is computed that merges each set of nodes that are pairwise reachable from each other (i.e., the nodes of strongly connected

components) into one “supernode”; the result is a directed acyclic graph (DAG) that preserves reachability. If the DAG has multiple nodes without incoming edges, the index adds an artificial root node connected to these nodes. A *spanning tree* – a subgraph that includes all nodes and is a tree – of the DAG is computed and labeled with its postorder. All subtrees thus have contiguous identifiers, where the maximum identifies the root; e.g., in Figure 13, the subtree at `:AI` has the interval $[1, 3]$, where 3 identifies the root. Then there exists a (directed) path from x to y if and only if y is in the subtree interval for x . Nodes in a DAG may, however, be reachable through paths not in the spanning tree. Hence each node is assigned a *set* of intervals for nodes that can be reached from it, where overlapping and adjacent intervals are merged; we must now check that y is in one of the intervals of x . To improve time and space at the cost of precision, approximate intervals are proposed that merge non-overlapping intervals; e.g., $[4, 6], [8, 9]$ is merged to $[4, 9]$, which can reject reachability for nodes with id less than 2 or greater than 9, but has a $\frac{1}{6}$ chance of a false positive for nodes in $[4, 9]$, which must be verified separately. Dynamic programming is used to select an optimal set of approximate intervals based on a budget of intervals per node set by the user. Intervals are marked as approximate or exact and indexed for each node in sorted order, enabling binary search to check reachability for a given node.

5.5 Join indexes

The results of joins can also be indexed. Groppe et al. [62] proposed to construct $6 \times 2^4 = 96$ indexes for joins between two triple patterns, where 6 is the number of non-symmetric types of joins (S-S, S-P, S-O, P-P, P-O, O-O). Hash maps are used to cover the 2^4 permutations of the remaining elements (not considering the join variable). Such an approach enables $O(1)$ join performance for two triple patterns, but at a potentially massive space cost. More commonly, only joins that are frequently encountered in queries are indexed [43, 129].




Fig. 15: Distance-based indexing (GRIN)

5.6 Structural indexes

Another family of indexes – referred to as structural indexes [119] – rather build a high-level summary of the RDF graph. Such summaries may take different forms.

A first type of structural index is based on distance. Udrea et al. propose GRIN [194], which divides the graph hierarchically into regions based on the distance of its nodes to selected centroids. These regions form a tree, where the non-leaf elements indicate a node x and a distance d referring to all nodes at most d steps from x . The root element chooses a node and distance such that all nodes of the graph are covered. Each non-leaf element has two children that capture all nodes of their parent. Each leaf node contains a set of nodes N , which induces a subgraph of triples between the nodes of N ; the leaves can then be seen as partitioning the RDF graph. We provide an example in Figure 15 for the RDF graph of Figure 1, where all nodes are within distance two of :Alice, which are then divided into two regions: one of distance at most two from _:p, and another of distance at most one from :CS. The index can continue dividing the graph into regions, and can then be used to find subgraphs within a particular distance from a given node (e.g., a node given in a BGP).

Another type of structural index relies on some notion of a *quotient graph* [38], where the nodes of a graph $so(G)$ are partitioned into $\{X_1, \dots, X_n\}$, such that $1 \leq i < j \leq n$ implies $X_i \cap X_j = \{\}$ (i.e., the sets are pairwise disjoint) and $\bigcup_{i=1}^n X_i = so(G)$. Then edges of the form (X_i, p, X_j) are added if and only if there exists $(x_i, p, x_j) \in G$ such that $x_i \in X_i$ and $x_j \in X_j$. Intuitively, a quotient graph merges nodes from the input graph into “supernodes” while maintaining the input (labeled) edges between the supernodes. We provide an example of a quotient graph in Figure 16 featuring six supernodes. Any partitioning of nodes can form a quotient graph, ranging from a single supernode with all nodes $so(G)$ and loops for all properties in $p(G)$, to the graph itself replacing each node $x \in so(G)$ with the singleton $\{x\}$. If the input graph yields solutions for a BGP, then the quotient graph will also yield solutions (with variables now matching supernodes). For example, taking the BGP of Figure 2, matching `foaf:Person` to the supernode containing `foaf:Person` in Figure 16, then the variables ?a and ?b will match to the supernode

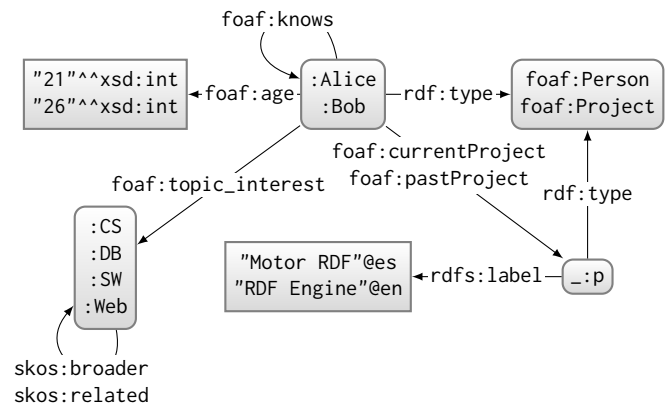


Fig. 16: Quotient graph with six supernodes

node containing :Alice and :Bob, while ?ia and ?ib will match to the supernode containing :CS, :DB, :SW and :Web; while we do not know the exact solutions for the input graph, we know they must correspond to elements of the supernodes matched in the quotient graph.

In summary, quotient graphs provide a structural summary of an RDF graph, where fewer partitions lead to a smaller summary, while more partitions lead to a higher fidelity summary. The DOGMA [33] RDF store partitions an RDF graph into subgraphs, from which a balanced binary tree is computed, where each parent node contains a quotient-like graph of both its children. The (O)SQP approach [191] creates an in-memory index graph, which is a quotient graph whose partition is defined according to various notions of bisimulation. Each supernode indexes (on disk) all triples mentioning any node in the supernode. The query is first evaluated on the index graph, identifying relevant supernodes, where exact solutions are then computed from the triples of those supernodes. The SAINT-DB system [154] adopts a similar approach, but where the supernodes of the indexed quotient graph are rather defined directly as a partition of the triples of the RDF graph, and edges are labeled with the type of join (s-s, p-o, etc.) that exists between them.

5.7 Quad indexes

When indexing quads for SPARQL datasets, most systems follow the triple index scheme [205, 78, 75, 54], extending it to add another element. The number of permutations then grows, where we must now consider $2^4 = 16$ abstract index patterns, $4! = 24$ potential permutations, and $\binom{4}{4/2} = 6$ flat (ISAM/B+Tree/AVL tree/trie) permutations or 2 circular (CSA) permutations to efficiently support all abstract quad patterns. A

practical compromise is to maintain a selection of permutations that cover the most common patterns [54]; for example, a pattern (s, \mathbf{p}, o, g) may be uncommon in practice, and could be supported reasonably well by evaluating (e.g.) $(s, \mathbf{p}, o, \mathbf{g})$ and filtering on $\mathbf{g} = g$.

The RIQ system [100] rather proposes a custom index for quads called a PV-index. The system specifically focuses on the problem of finding (named) graphs that match a BGP. Each graph is indexed using hashing applied for all seven abstract patterns on triples with some constant, generating seven *pattern vectors* for each graph. For example, a triple (s, p, o) in a graph named g will be hashed as (s, p, o) , $(s, p, ?)$, $(s, ?, o)$, $(?, p, o)$, $(s, ?, ?)$, $(?, p, ?)$, $(?, ?, o)$, where $?$ is an arbitrary fixed token, and each result will be added to one of seven pattern vectors for g for that abstract pattern. A triple pattern (s, \mathbf{p}, o) can then likewise be encoded as $(s, ?, o)$, and a BGP can be encoded analogously as pattern vectors. The goal then is to find graphs with similar pattern vectors to the BGP vectors, returning potential matches. Rather than comparing the BGP to each graph, locality sensitive hashing is used to group and retrieve similar pattern vectors. A downside is the potential space requirements when handling many graphs; the PV-index thus uses Bloom filters to encode groups of similar graphs in order to reduce the space used.

5.8 Miscellaneous Indexing

RDF stores may use legacy systems, such as NoSQL stores, for indexing. Since such approaches are not tailored to RDF, and often correspond conceptually to one of the indexing schemes already discussed, we again refer to Janke and Staab [91] and Zambom Santana [218] for further discussion on such approaches. Other stores provide specialized indexes for particular types of values such as spatial or temporal data [197, 109]; we do not discuss such specialized indexes in detail.

5.9 Discussion

While indexing triples or quads is conceptually the most straightforward approach, a number of systems have shown positive results with entity- and property-based indexes that optimize the evaluation of star joins, path indexes that optimize the evaluation of path joins, or structural indexes that allow for isolating regions of the graph that are relevant for a particular graph pattern. The proposal of diverse indexing schemes shows that rather than there being a clear winner, such schemes explore different trade-offs: in general, there exists a time-space trade-off, where more comprehensive indexes en-

able faster queries at the cost of space, and also often at the cost of more costly updates.

6 Join Processing

RDF stores employ diverse query processing strategies, but all require translating logical operators, which represent the query, into “physical operators”, which implement algorithms for efficient evaluation of the operation. The most important such operators – as discussed in this section – are (equi-)joins.

6.1 Pairwise join algorithms

We recall that the evaluation of a BGP $\{t_1, \dots, t_n\}(G)$ can be rewritten as $t_1(G) \bowtie \dots \bowtie t_n(G)$, where the evaluation of each triple pattern t_i ($1 \leq i \leq n$) produces a relation of arity $|\text{vars}(t_i)|$. Thus the evaluation of a BGP B produces a relation of arity $|\text{vars}(B)|$. The relational algebra – including joins – can then be used to combine or transform the results of one or more BGPs, giving rise to CGPs. The core of evaluating graph patterns is thus analogous to processing relational joins. The simplest and most well-known such algorithms perform *pairwise joins*; for example, a pairwise strategy may for computing $\{t_1, \dots, t_n\}(G)$ may evaluate $((t_1(G) \bowtie t_2(G)) \bowtie \dots) \bowtie t_n(G)$.

Without loss of generality, we assume a join of two graph patterns $P_1(G) \bowtie P_2(G)$, where the join variables are denoted by $V = \{v_1, \dots, v_n\} = \text{vars}(P_1) \cap \text{vars}(P_2)$. Well-known algorithms for performing pairwise joins include (*index*) *nested-loop joins*, where $P_1(G) \bowtie P_2(G)$ is reduced to evaluating $\bigcup_{\mu \in P_1(G)} \{\mu\} \bowtie \mu(P_2)(G)$; *hash joins*, where each solution $\mu \in P_1(G)$ is indexed by hashing on the key $(\mu(v_1), \dots, \mu(v_n))$ and thereafter a key is computed likewise for each solution in $P_2(G)$ to probe the index with; and (*sort*-)*merge joins*, where $P_1(G)$ and $P_2(G)$ are (sorted if necessary and) read in the same order with respect to V , allowing the join to be reduced to a merge sort. In general, an index nested-loop join performs well when $|P_1(G)| \ll |P_2(G)|$ (assuming that $\mu(P_2)(G)$ can be reduced to index lookups) since, unlike the other two alternatives, it does not require reading all of $P_2(G)$. Otherwise hash or merge joins perform well, where merge joins are a better choice when solutions can be read directly from the index in the required order [142]. Pairwise join algorithms are used in many RDF stores (e.g., [77, 54, 142]).

A number of techniques are used to complement and optimize traditional pairwise join algorithms. One such technique is that of *sideways information passing* [23], which refers to how information about inter-

mediate results is passed between different parts of the query. Neumann and Weikum [141] propose a strategy called *ubiquitous sideways information passing (U-SIP)* for computing joins over RDF, which allows for globally sharing information on the range of values for a given query variable. The strategy also considers variables that must be bound to the same values (based, for example, on an equality constraint $\text{FILTER}(?a=?b)$), forming equivalence classes of variables. U-SIP is then implemented in different ways for different types of joins. For merge joins, where data are read in order, a maximum value for a given equivalence class of variables can be shared across pairwise joins, allowing individual operators to skip ahead to the current maximum. For hash joins, a global *domain filter* is employed – consisting of a maximum value, a minimum value, and Bloom filters – for each equivalence class of variables, allowing to filter values for those variables elsewhere.

6.2 Multiway joins

Multiway join algorithms exploit the commutativity and associativity of joins to evaluate them over two or more operands at once. For example, in order to compute $\{t_1, \dots, t_n\}(G)$, a multiway join algorithm may evaluate $((t_1(G) \bowtie \dots \bowtie t_k(G)) \bowtie (t_{k+1}(G) \bowtie \dots \bowtie t_n(G)))$ where $k \geq 2$, or even simply $t_1(G) \bowtie \dots \bowtie t_n(G)$.

Some of the previous storage and indexing schemes we have seen lend themselves naturally to processing certain types of multiway joins in an efficient manner. For example, entity-based indexes allow for processing *star joins* efficiently, while path indexes allow for processing path joins efficiently (see Section 5). A BGP can then be decomposed into sub-BGPs that can be evaluated per the corresponding multiway join, with pairwise joins being applied across the sub-BGPs; for example: $\{(\mathbf{w}, p, \mathbf{x}), (\mathbf{w}, q, \mathbf{y}), (\mathbf{w}, r, z), (\mathbf{x}, q, \mathbf{y}), (\mathbf{x}, r, z)\}$ may be divided into the sub-BGP $\{(\mathbf{w}, p, \mathbf{x}), (\mathbf{w}, q, \mathbf{y}), (\mathbf{w}, r, z)\}$ and the sub-BGP $\{(\mathbf{x}, q, \mathbf{y}), (\mathbf{x}, r, z)\}$, which are evaluated individually as multiway joins before being themselves joined. Even in the case of (sorted) triple/quad tables, multiway joins can be applied taking advantage of the locality of processing, where, for example, in an SPO index permutation, triples with the same subject will be grouped together. Similar locality (across sources) is exploited by SMJoin [58], which decomposes BGPs into sub-BGPs corresponding to star joins, evaluating these separately before combining the results.

6.3 Worst case optimal joins

A new family of join algorithms have arisen as a result of the *AGM bound* [16], which puts an upper bound on the number of solutions that can be returned from relational join queries. The result can be adapted straightforwardly to the case of BGPs. Let $B = \{t_1, \dots, t_n\}$ denote a BGP with $\text{vars}(B) = V$. Now define a *fractional edge cover* as a mapping $\lambda : B \rightarrow \mathbb{R}_{[0,1]}$ that assigns a real value in the interval $[0, 1]$ to each triple pattern of B such that for all $v \in V$, it holds that $\sum_{t \in B_v} \lambda(t) \geq 1$, where B_v denotes the set of triple patterns in B that mention v . The AGM bound tells us that if B has the fractional edge cover λ , then for any RDF graph it holds that $|B(G)| \leq \prod_{i=1}^n |\{t_i\}(G)|^{\lambda(t_i)}$. Moreover, if no two triple patterns t_i and t_j in B ($t_i \neq t_j$) can be *unified* – i.e., there does not exist solution mappings μ_i and μ_j such that $\mu_i(t_i) = \mu_j(t_j)$ ⁵ – then the bound is tight, meaning that there exists an RDF graph G and a fractional edge cover λ such that $|B(G)| = \prod_{i=1}^n |\{t_i\}(G)|^{\lambda(t_i)}$. In the case that two triple patterns in B unify (akin to self-joins in the relational setting), the AGM bound only varies by a constant factor.

To illustrate the AGM bound, consider the BGP $B = \{t_1, t_2, t_3\}$ from Figure 17. There exists a fractional edge cover λ of B such that $\lambda(t_1) = \lambda(t_2) = \lambda(t_3) = \frac{1}{2}$; taking $?a$, we have that $B_{?a} = \{t_1, t_3\}$, $\lambda(t_1) + \lambda(t_3) = 1$, and thus $?a$ is “covered”, and we can verify the same for $?b$ and $?c$. Then the AGM bound is given as the inequality $|B(G)| \leq \prod_{i=1}^n |\{t_i\}(G)|^{\lambda(t_i)}$. For G the graph in Figure 17, $|\{t_1\}(G)| = |\{t_2\}(G)| = |\{t_3\}(G)| = 5$, and hence $|B(G)| \leq 5^{\frac{3}{2}}$. In reality, for this graph, $|B(G)| = 5$, thus satisfying the inequality, but we further know that there exists a graph where $B = \prod_{i=1}^n |\{t_i\}(G)|^{\lambda(t_i)}$.

Recently, join algorithms have been proposed that can enumerate the results for a BGP B over a graph G in time $O(\text{agm}(B, G))$, where $\text{agm}(B, G)$ denotes the AGM bound of B over G . Since such an algorithm must at least spend $O(\text{agm}(B, G))$ time writing the results in the worst case, such algorithms can be considered *worst-case optimal (wco)* [143]. Though such algorithms were initially proposed in a relational setting [143, 195], they were later applied for processing joins over RDF graphs [95, 81, 137]. Note that traditional pairwise join algorithms are not wco. If we try to evaluate $\{t_1, t_2\}(G)$ by pairwise join, for example, in order to later join it with $\{t_3\}(G)$, the AGM bound becomes quadratic as $\lambda(t_1) = \lambda(t_2) = 1$, and thus we have the bound $|\{t_1\}(G)| \cdot |\{t_2\}(G)|$, which exceeds the AGM bound for B . The same holds for any pairwise join in B . In practice, note that $\{t_1, t_2\}(G)$ will indeed

⁵ For example, (\mathbf{x}, r, z) can be unified with $(\mathbf{x}, \mathbf{y}, z)$, $(\mathbf{y}, r, \mathbf{x})$, (y, r, z) , etc., but not with (\mathbf{x}, p, z) , (\mathbf{x}, r, q) , etc.

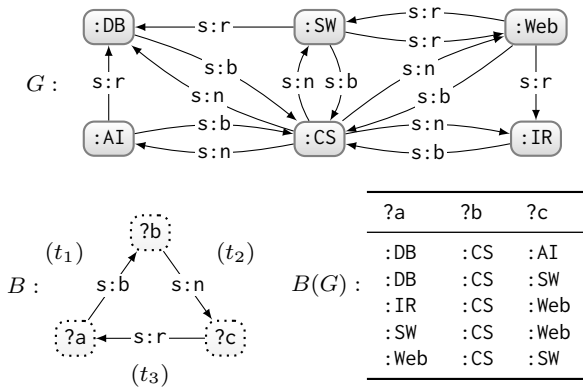


Fig. 17: Example RDF graph G , BGP B and its evaluation $B(G)$; the IRIs $s:b$, $s:n$ and $s:r$ abbreviate $skos:broader$, $skos:narrower$ and $skos:related$, resp.

produce (25) quadratic results, mapping $?a$ to $:CS$ and $?b$ and $?c$ to $\{ :AI, :DB, :IR, :SW, :Web \}^2$.

Wco join algorithms – including *Leapfrog Triejoin* (LTJ) [195] – perform a multiway join that resolves a BGP B variable-by-variable rather than pattern-by-pattern. First an ordered sequence of variables is selected; say $(?a, ?b, ?c)$. Then the set of partial solutions $M_{\{?a\}} = \{ \mu \mid \text{dm}(\mu) = \{?a\} \text{ and } \mu(B_{?a})(G) \neq \emptyset \}$ are computed for the first variable $?a$ such that each image of $B_{?a}$ under $\mu \in M_{\{?a\}}$ has solutions for G ; e.g., $M_{\{?a\}} = \{ \{?a/:DB\}, \{?a/:IR\}, \{?a/:SW\}, \{?a/:Web\} \}$ in Figure 17, since replacing $?a$ in $B_{?a}$ with $:DB$, $:IR$, $:SW$ or $:Web$ yields a BGP with solutions over G . Next we compute $M_{\{?a, ?b\}} = \{ \mu \cup \mu' \mid \mu \in M_{\{?a\}}, \text{dm}(\mu') = \{?b\} \text{ and } \mu'(\mu(B_{?b}))(G) \neq \emptyset \}$, “eliminating” the next variable $?b$. In the example of Figure 17, $M_{\{?a, ?b\}} = \{ \{?a/:DB, ?b/:CS\}, \dots, \{?a/:Web, ?b/:CS\} \}$, where each solution $\mu \in M_{\{?a\}}$ is extended with $\{?b/:CS\}$. Finally, $M_{\{?a, ?b, ?c\}}$ is computed analogously, eliminating the last variable, and yielding the five results seen in Figure 17.

In order to be wco-compliant, the algorithm must always be able to efficiently compute $M_{\{v\}}$, i.e., all solutions μ with $\text{dm}(\mu) = \{v\}$, such that $\mu(B_v)(G) \neq \emptyset$. For instance, to compute $M_{\{?a\}}$ in the running example, we need to efficiently intersect all nodes with an outgoing $s:b$ edge and an incoming $s:r$ edge. This is typically addressed by being able to read the results of a triple pattern, in sorted order, for any variable, which enables efficient intersection by allowing to seek ahead to the maximum current value of all triple patterns involving a given variable. Jena-LTJ [81], which implements an LTJ-style join algorithm for SPARQL, enables this by maintaining all six index permutations over triples.

Wco algorithms have been shown to outperform traditional join algorithms for complex BGPs [95, 81].

6.4 Join reordering

The order of join processing can have a dramatic effect on computational costs. For Figure 17, if we apply pairwise joins in the order $(\{t_1\}(G) \bowtie \{t_2\}(G)) \bowtie \{t_3\}(G)$, the first join $(\{t_1\}(G) \bowtie \{t_2\}(G))$ yields 25 intermediate results, before 5 final results are produced with the second join. On the other hand, if we evaluate the join in the order $(\{t_2\}(G) \bowtie \{t_3\}(G)) \bowtie \{t_1\}(G)$, the first join $(\{t_2\}(G) \bowtie \{t_3\}(G))$ produces only 5 intermediate results, before the second join produces the same 5 final results as before. The second plan should thus be more efficient than the first; if considering a graph at larger scale, the differences may reach orders of magnitude.

A good plan depends not only on the query, but also the graph. Selecting a good plan thus typically requires some assumptions or statistics over the graph. As in relational settings, the most important information relates to *cardinalities*: how many (distinct) solutions a given pattern returns; and/or *selectivity*: what percentage of solutions are kept when fixing variables to constants or adding filters. Statistics can be used not only to select an ordering for joins, but also to decide which join algorithm to apply. For example, given an arbitrary (sub-)BGP $\{t_1, t_2\}$, if we estimate that $|\{t_2\}(G)| \ll |\{t_1\}(G)|$, we may prefer to evaluate $\{t_2\}(G) \bowtie \{t_1\}(G)$ as an index nested loop join, rather than a hash or merge join, to avoid reading $\{t_1\}(G)$ in full.

While cardinality and/or selectivity estimates can be managed in a similar way to relational database optimisers, a number of approaches have proposed custom statistics for RDF graphs. In the context of Jena, Stocker et al. [181] collect statistics relating to the number of triples, the number of unique subjects, and for each predicate, the number of triples and a histogram of associated objects. RDF-3X [142] uses a set of *aggregated indexes*, which store the cardinality of all triple patterns with one or two constants. RDF-3X [142] further stores the exact cardinality of frequently encountered joins, while characteristic sets [139] and extended characteristic sets [132] (discussed in Section 5.3) allow for estimating the cardinality of star joins.

Computing and maintaining such statistics incur costs in terms of space and updates. An alternative is to apply *sampling* while evaluating the query. Vidal et al. [196] estimate the cardinality of star joins by evaluating all solutions for the first pattern of the join, thereafter computing the full solutions of the star pattern for a sample of the initial solutions; the full cardinality of the star pattern is then estimated from the samples. Another alternative is to use *syntactic heuristics* for reordering. Stocker et al. [181] propose heuristics such as assuming that triple patterns with fewer variables

have lower cardinality, and that constant subjects are most selective, followed by constant objects, and constant predicates. Similar syntactic heuristics were explored by Tsialiamanis et al. [193], who further propose to prioritize rarer joins (such as P-S and P-O joins), and to consider literals as more selective than IRIs.

Taking into account the various heuristics and statistics, the simplest strategy to try to find a good join ordering – applied in various RDF stores (e.g., [181,132]) – is to apply a greedy heuristic, starting with the triple pattern t_1 estimated to have the lowest cardinality, and joining it with the triple pattern t_2 with the next lowest cardinality; typically a constraint is added such that t_n ($n > 1$) must have a variable in common with some triple pattern in $\{t_1, \dots, t_{n-1}\}$ to avoid costly Cartesian products. Aside from considering the cardinality of triple patterns, Meimaris and Papastefanatos [131] propose a distance-based planning, where pairs of triple patterns with more overlapping nodes and more similar cardinality estimates have lesser distance between them; the query planner then tries to group and join triple patterns with the smallest distances first in a greedy manner. Greedy strategies may not, however, always provide the best ordering (see, e.g., Figure 17).

More generally, reordering joins based on cost estimates is an optimization problem, where classical methods from the relational literature can be leveraged likewise for BGPs, including dynamic programming [176] (used, e.g., by [78,142,66]) and simulated annealing [88] (used, e.g., by [196]). Other metaheuristics that have been applied for join reordering in BGPs include genetic algorithms [84] and ant colony systems [83,94].

With respect to wco join algorithms, though the wco guarantee does not depend on the ordering of variables (any order will do), runtimes can vary for different orderings. Following the intuitive idea that starting with variables generating fewer results should help lower runtimes, Jena-LTJ [81] uses similar selectivity-based heuristics to Jena [181] in order to estimate the cardinality of individual variables, thereafter ordering variables per a greedy strategy that avoids Cartesian products. While specialized orderings for wco algorithms based on hypertree decompositions have been explored in the relational literature [5], to the best of our knowledge they have not yet been explored for BGPs.

6.5 Caching

Aside from reordering joins, another possible route for optimization – based on the observation that queries in practice may feature overlapping or similar patterns – is to reuse work done previously for other queries. Specifically, we can consider *caching* the results of queries.

In order to increase cache hit rates, we can further try to reuse the results of subqueries, possibly generalizing them to increase usability. Ideally the cache should store solutions for subqueries that (a) have a high potential to reduce the cost of future queries; (b) can reduce costs for many future queries; (c) do not have a high space overhead; and (d) will remain valid for a long time. Some of these aims can be antagonistic; for example, caching solutions for triple patterns satisfy (b) and (c) but not (a), while caching solutions for complex BGPs satisfies (a) but not (b), (c) or (d).

Lampo et al. [111] propose caching of solutions for star joins, which may strike a good balance in terms of reducing costs, being reusable, and not having a high space overhead (as they share a common variable). Other caching techniques try to increase cache hit rates by detecting similar (sub)queries. Stuckenschmidt [183] proposes a similarity measure for caching – based on the edit distance between BGPs – that estimates the amount of computational effort needed to compute the solutions for one query given the solutions to the other. Lorey and Naumann [118] propose a technique for grouping similar queries, which enables a pre-fetching strategy based on predicting what a user might be interested in based on their initial queries. Another direction is to normalize (sub)queries to increase cache hit rates. Wu et al. [208] propose various algebraic normalizations in order to identify common subqueries [118], while Papailiou et al. [149] generalize subqueries by replacing selective constants with variables and thereafter canonically labeling variables (modulo isomorphism) to increase cache hit rates.

A limitation of caching solutions is that an arbitrary BGP can produce an exponential number of results, where an alternative solution is to cache frequently used subgraphs. Along these lines, Zhang et al. [220] rather propose a caching technique based on caching frequently accessed “hot triples” from the RDF graph in memory, which limits space overhead and can improve runtimes for many queries, but still requires all joins to be computed. Addressing dynamic data, Martin et al. [127] propose a cache where results for queries are stored in a relational database but are invalidated when a triple matching a query pattern changes, while Williams and Weaver [204] propose adding last-updated times to the RDF index that indicate whether or not any relevant data have changed for a repeated query.

6.6 Discussion

Techniques for processing BGPs are often based on techniques for processing relational joins. Beyond traditional pairwise evaluation, multiway joins can help

to emulate some of the benefits of property table storage by evaluating star joins more efficiently. Another recent and promising approach is to apply wco join algorithms whose runtime is bounded theoretically by the number of results that the BGP could generate. Aside from specific algorithms, the order in which joins are processed can have a dramatic effect on runtimes. Statistics about the RDF graph help to find a good ordering at the cost of computing and maintaining those statistics; more lightweight alternatives include runtime sampling, or syntactic heuristics that consider only the query. To decide the ordering, options range from simple greedy strategies to complex metaheuristics; while simpler strategies have lower query planning times, more complex strategies may find more efficient plans. Another optimization is to cache results across BGPs, for which a time-space trade-off must be considered.

7 Query Processing

While we have defined RDF stores as engines capable of storing, indexing and processing joins over RDF graphs, SPARQL engines must support additional features beyond joins. We describe some of the proposed techniques for efficiently evaluating such features.

7.1 Property paths

Navigational graph patterns (NGPs) extend BGPs with property paths that allow for matching paths of arbitrary length. Given a property path such as $:p^+$, a direct strategy is to simply apply standard join algorithms recursively, where, in an early work, Stuckenschmidt et al. [184] propose cost models for processing such queries using nested-loop and hash joins. However, specialized algorithms have also been proposed for evaluating property paths in an efficient manner.

One approach is to use graph search algorithms. Though not part of the SPARQL standard, Gubichev and Neumann [65] implement single-source (i.e., a start or end node is given) shortest paths with constraints regarding the length of the path, the elements it contains, etc. Dijkstra's search algorithm is then implemented over B-Trees for all six index permutations of an RDF graph, with optimizations including the ability to find the neighbors of multiple nodes in one lookup, which reduces I/O costs; a dictionary encoding technique that assigns successive IDs to nodes in a breadth-first manner, which achieves better locality of neighborhoods; and estimators for the cardinality of paths, which enables join reordering over NGPs. Baier et al. [17] propose to use the A* search algorithm, where the search

is guided by a heuristic that measures the minimum distance from the current node to completing a full path; though they evaluate property paths over the Web, their techniques can be applied in local settings.

Extending RDF-3X, Gubichev et al. [64] build a FERRARI index [177] (see Section 5.4) for each property $:p$ in the graph that forms one or more directed paths of length at least 2. These indexes are used to evaluate the paths $:p^*$ or $:p^+$. For the purposes of join reordering, the cardinality of paths is estimated based on pre-computing, for each property, the average number of nodes reachable from a given node (i.e., the average number of elements in the intervals of a node in the FERRARI index). A sideways-information passing strategy (see Section 6.1) is further implemented by sharing Bloom filters that encode the FERRARI intervals for particular variables across joins, along with their maximum and minimum values, allowing for globally filtering nodes that fall outside said intervals. However, paths of the form $(:p/:q)^*$, $(:p|:q)^*$, etc., are not directly supported by such methods.

A way to support more expressive property paths is to write them as recursive queries. Yakovets et al. [213] propose a rewriting from property paths to recursive queries in SQL, i.e., queries using the WITH RECURSIVE feature. The evaluation of the property paths can then be delegated to a relational database. Though all features of property paths can be supported, the authors observe that highly nested SQL queries are sometimes generated, and that many popular relational databases cannot (efficiently) detect cycles. Reutter et al. [161], on the other hand, extend SPARQL to support recursion. A temporary RDF graph is built by recursively adding triples produced through CONSTRUCT queries over the base graph and the temporary graph up to a fix-point; a SELECT query can then be evaluated over both graphs. The authors discuss how property paths can then be supported through *linear recursion*, meaning that each new triple only needs to be joined with the base graph, not the temporary graph, to produce further triples, which leads to competitive performance with systems that implement property paths “directly”. Recursive queries are more expressive than property paths, and thus may be more difficult to optimize, but there are known optimizations (e.g., for Datalog), such as linear recursion, magic sets, etc., that can be used.

Yakovet et al. [214] propose Waveguide that first converts the property path into a parse tree, from which plans can be built based on finite automata (FA), or relational algebra with transitive closure (α -RA, where α denotes transitive closure). Figure 18 gives an example of a parse tree and both types of plans. Although there is overlap, FA can express physical plans that α -RA

cannot, and vice versa. For example, in FA we can express non-deterministic transitions (see q_0 in Figure 18), while in α -RA we can materialize (cache) a particular relation in order to apply transitive closure over it. Waveguide then uses hybrid *waveplans*, where breadth first search is guided in a similar manner to FA, but where the results of an FA can be *memoized* (cached) and reused multiple times like in α -RA. We provide an example of a waveplan that uses a memoized waveplan in Figure 18. The authors show that waveplans can express more plans than FA and α -RA combined, and develop algorithms and cost models for them.

Evaluating arbitrarily complex property paths can be costly, but property paths in practice are often relatively simple. Along these lines, Martens and Trautner [126] propose a class of RPQs called *simple transitive expressions* (STEs) that are found to cover 99.99% of the queries found in Wikidata SPARQL logs, and have desirable theoretical properties. Specifically, they call an expression $p_1 \mid \dots \mid p_n$, where p_1, \dots, p_n are IRIs and $n \geq 0$ an *atomic expression*, and an expression of the form $a_1 / \dots / a_k$ or $a_1 ? / \dots / a_k ?$ where a_1, \dots, a_k are atomic expressions, a *bounded expression* of length $k \geq 0$. Then an expression of the form $b_1 / a^* / b_2$, is a *simple transitive expression* (STE), where b_1 and b_2 are bounded expressions, and a is an atomic expression. They then show that simple paths for STEs can be enumerated more efficiently than arbitrary RPQs.

7.2 Relational algebra (beyond joins)

Complex (navigational) graph patterns C(N)GPs introduce additional relational operators beyond joins.

Like in relational databases, *algebraic rewriting rules* can be applied over C(N)GPs in SPARQL in order to derive equivalent but more efficient plans. Schmidt et al. [174] present a comprehensive set of such rules for SPARQL. Key rules involve filters, which allow for selections to be “pushed” closer to the base operators in order to filter intermediate solutions as early as possible. Along these lines, Schmidt et al. [174] derive the following rewriting rules involving filters:

$$\begin{aligned} \sigma_{R_1 \vee R_2}(M) &\equiv \sigma_{R_1}(M) \cup \sigma_{R_2}(M) \\ \sigma_{R_1 \wedge R_2}(M) &\equiv \sigma_{R_1}(\sigma_{R_2}(M)) \\ \sigma_{R_1}(\sigma_{R_2}(M)) &\equiv \sigma_{R_2}(\sigma_{R_1}(M)) \\ \sigma_R(M_1 \cup M_2) &\equiv \sigma_R(M_1) \cup \sigma_R(M_2) \\ \sigma_R(M_1^* \bowtie M_2) &\equiv \sigma_R(M_1^*) \bowtie M_2 \\ \sigma_R(M_1^* \bowtie M_2) &\equiv \sigma_R(M_1^*) \bowtie M_2 \\ \sigma_R(M_1^* \triangleright M_2) &\equiv \sigma_R(M_1^*) \triangleright M_2 \end{aligned}$$

where for each $\mu \in M_1^*$, it holds that $\text{vars}(R) \subseteq \text{dm}(\mu)$. The first two rules allow for simplifying filters, meaning that they can be pushed further down in a query. The third rule allows the order in which filters are applied to be swapped. Finally the latter four rules describe how filters can be pushed “down” inside various operators.

Another feature of importance for querying RDF graphs are optionals (\bowtie), as they facilitate returning partial solutions over incomplete data. Given that an optional can be used to emulate a form of negation (in Table 3 it is defined using an anti-join), it can lead to jumps in computational complexity [152]. Works have thus studied a fragment called *well-designed patterns*, which forbid using a variable on the right of an optional that does not appear on the left but does appear elsewhere in the query; taking an example, the CGP $\{(\mathbf{x}, p, \mathbf{y})\} \text{ OPTIONAL } \{(\mathbf{x}, q, \mathbf{z})\} \cdot \{(\mathbf{x}, r, \mathbf{z})\}$ is not well designed as the variable \mathbf{z} appears on the right of an OPTIONAL and not on the left, but does appear elsewhere in the query. Such variables may or may not be left unbound after the left outer join is evaluated, which leads to complications if they are used outside the optional clause. Most SPARQL queries using optionals in practice are indeed well-designed, where rewriting rules have been proposed to optimize such queries [152, 115].

7.3 Query rewriting

Another approach to processing SPARQL queries is to rewrite them to another language and delegate their optimization and execution to an existing engine that supports the target language. A common strategy (used, e.g., by Virtuoso [54]) is to rewrite the SPARQL query to SQL, which allows for leveraging the existing optimizations of relational database engines, and also allows for querying the same data using SQL or SPARQL; however, care is required to handle SPARQL features that go beyond SQL, such as property paths (perhaps using recursive SQL, as aforementioned [213]). Another approach is to rewrite SPARQL to existing graph query languages that support navigational features, such as Gremlin (e.g., SPARQL-Gremlin [189]); again, however, there are features of SPARQL that typical graph query languages may not support, such as variable predicates, full property path expressions, etc. Other works rather propose to rewrite SPARQL queries to languages executable in a distributed environment, including PigLatin (e.g., PigSPARQL [169], RAPID+ [160]), Hadoop (e.g., Sempala [170]), Spark (e.g., S2RDF [171]), etc. While this allows query execution to be handled by multiple machines, not all SPARQL features are easily supported on existing distributed frameworks.

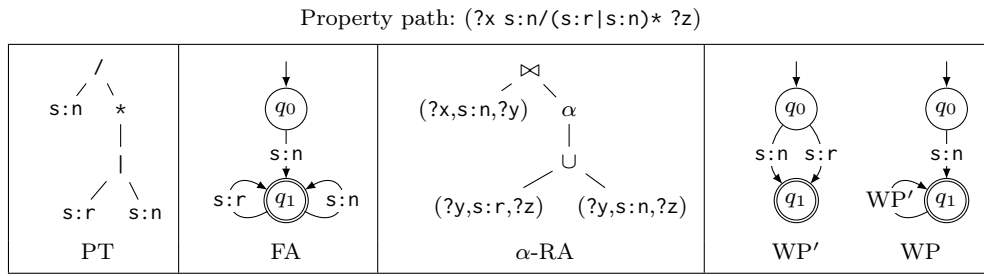


Fig. 18: An example property path with its parse tree (PT) and three plans based on finite automata (FA), relational algebra with transitive closure (α -RA), and a waveplan (WP) that uses a memoized waveplan (WP')

7.4 Discussion

SPARQL supports various features beyond joins that ideally should be implemented in an efficient manner. One option is to rewrite SPARQL queries into a target language and evaluate them using an existing engine for that language. However, it is unlikely that an existing language/engine will support all features of SPARQL in an efficient manner. Better performance for a wider range of features can be achieved with custom implementations and optimizations, where property paths have been the focus of many works. Other features that have been targeted for optimization are filters and optionals, noting that optionals are quite frequently used in the context of querying incomplete RDF data.

8 Partitioning

In distributed RDF stores and SPARQL engines, the data are partitioned over a cluster of machines in order to enable horizontal scale, where additional machines can be allocated to the cluster to handle larger volumes of data. However, horizontal scaling comes at the cost of network communication costs. Thus a key optimization is to choose a partitioning scheme that reduces communication costs by enforcing various forms of locality, principally allowing certain types of (intermediate) joins to be processed on each individual machine [7]. Formally, given an RDF graph G and n machines, an n -partition of G is a tuple of subgraphs (G_1, \dots, G_n) such that $G = \bigcup_{i=1}^n G_i$, with the idea that each subgraph G_i will be stored on machine i .⁶ We now discuss different high-level alternatives for partitioning.

⁶ We relax the typical requirement for a set partition that $G_i \cap G_j = \emptyset$ for all $1 \leq i < j \leq n$ to allow for the possibility of replication or other forms of redundancy.

8.1 Triple/Quad-based Partitioning

A first option is to partition based on individual triples or quads without considering the rest of the graph. For simplicity we will speak about triples as the discussion generalizes straightforwardly to quads. The simplest option is to use *round robin* or *random partitioning*, which effectively places triples on an arbitrary machine. This allows for an equal number of triples to be assigned to each machine, but does not support any locality of processing, and does not allow for finding the particular machine storing triples that match a given pattern.

An alternative is to partition according to a deterministic function over a given key; for example, a partition key of S considers only the subject, while a partition key of PO considers both the predicate and object. Later given a triple pattern that covers the partition key (e.g., with a constant subject if the key is S), we can find the machine(s) storing all triples that match that pattern. We show some examples using different functions and partition keys in Figure 19 considering four machines. *Range-based partitioning* assigns a range over the partition key to each function, where the example of Figure 19 splits s into $[a:1, a:3]$, $[a:4, a:6]$, $[b:1, b:3]$, $[c:1, d:1]$. This approach allows for range-based queries to be pushed to one machine, but requires maintaining a mapping of ranges to machines, and can be complicated to keep balanced. An alternative is *hash-based partitioning* where we compute the hash of the partition key modulo the number of machines, where the second example of Figure 19 splits P by hash. This does not require storing any mapping, and techniques such as consistent hashing can be used to rebalance load when a machine enters or leaves; however, if partition keys are skewed (e.g., one predicate is very common), it may lead to an unbalanced partition. A third option is to apply a *hierarchical-based partition* based on prefixes, where the third example of Figure 19 partitions O by their namespace. This may lead to increased locality of data with the same prefix [93], where different levels of

prefix can be chosen to enable balancing, but choosing prefixes that offer balanced partitions is non-trivial.

Any such partitioning function will send any triple with the same partition key to the same machine, which ensures that (equi-)joins on partition keys can be pushed to individual machines. Hash-based partitioning is perhaps the most popular among distributed RDF stores (e.g., YARS2 [78], SHARD [162], etc.). Often triples will be hashed according to multiple partition keys in order to support different index permutations, triple patterns, and indeed joins across different types of partition keys (e.g., with *s* and *o* as two partition keys, we can push *s-s*, *o-o* and *s-o* (equi-)joins to each machine). When the partition key *P* is chosen, the scheme is analogous to distributed vertical partitioning⁷ as discussed in Section 4.2. Care must be taken to avoid imbalances caused by frequent terms, such as the *rdf:type* predicate, or frequent objects such as classes, countries, etc. Omitting partitioning on highly-skewed partition keys may be advantageous for balancing purposes [78].

8.2 Graph-based Partitioning

Graph-based partitioning takes into consideration the entire graph when computing a partition. A common strategy is to apply a *k-way partition* of the RDF graph *G* [99]. Formally, letting $V = \text{so}(G)$ denote the nodes of *G*, the goal is to compute a node partition V_1, \dots, V_n such that $V = \bigcup_{i=1}^k V_n$, $V_i \cap V_j = \emptyset$ for all $1 \leq i < j \leq k$, $\lfloor \frac{|V|}{k} \rfloor \leq |V_i| \leq \lceil \frac{|V|}{k} \rceil$ for all $1 \leq i \leq k$, and the number of triples $(s, p, o) \in G$ such that *s* and *o* are in different node partitions is minimised. In Figure 20, we show the optimal 4-way partitioning of the graph seen previously, where each partition has 3 nodes, there are 10 edges between partitions (shown dashed), and no other such partition leads to fewer edges (<10) between partitions. Edges between partitions may be replicated in the partitions they connect. Another alternative is to *k-way partition* the *line graph* of the RDF graph: an undirected graph where each triple is a node, and triples sharing a subject or object have an edge between them.

⁷ The term “vertical partitioning” in the context of RDF (per Section 4.2) may be misleading. Vertical partitioning in a relational context typically refers to splitting a table with many columns into several tables with fewer columns sharing the same primary key (a natural join over the partitions should recreate the original table). Horizontal partitioning refers to splitting a table with many rows into several tables with fewer rows sharing the same columns (a union over the partitions should recreate the original table). “Vertical partitioning” in an RDF context is actually horizontal partitioning in a traditional sense as partitioning is done by row: we can still read entire triples from each partition.

Finding an optimal *k*-way partition is intractable⁸, where approximations are thus necessary for large-scale graphs, including *spectral methods*, which use the eigenvectors of the graph’s Laplacian matrix to partition it; *recursive bisection*, which recursively partitions the graph in two; *multilevel partitioning*, which “coarsens” the graph by computes a hierarchical graph summary (similar to a multilevel quotient graph, per Figure 16), then partitions the smaller graph summary (using, e.g., spectral methods), and finally “uncoarsens” by expanding back out to the original graph maintaining the partitions; etc. We refer for more details to Buluç et al. [36], who argue that multilevel partitioning is “*clearly the most successful heuristic for partitioning large graphs*”. Such techniques have been used by H-RDF-3x [86], EA-GRE [221], Koral [92], and more besides.

8.3 Query-based Partitioning

While the previous partitioning schemes only consider the data, other partitioning methods are (also) based on queries. Workload-based partitioning strategies have been broadly explored, where the idea is to identify common joins in query logs that can be used to partition or replicate parts of the graph in order to ensure that high-demand joins can be pushed to individual machines. Partitioning can then be *a priori*, for example, based on a query log; or *dynamic* (aka. adaptive), where the partitions change as queries are received. Such strategies are used by WARP [85], Partout [57], WORQ [123], chameleon-DB [10], AdPart [73], etc.

8.4 Replication

Rather than partitioning data, another option is to replicate data across partitions. This may vary from replicating the full graph on each machine, such that queries can be answered in full by any machine increasing query throughput (used, e.g., by DREAM [72]), to replicating partitions that are in high-demand (e.g., containing schema data, central nodes, etc.) so that more queries can be evaluated on individual machines and/or machines have equal workloads that avoid hot-spots (used, e.g., by Blazegraph [190] and Virtuoso [54]).

8.5 Discussion

Triple/quad-based partitioning is the simplest to compute and maintain, requiring only the data present in

⁸ Given a graph, deciding if there is a *k*-way partition with fewer than *n* edges between partitions is NP-complete.

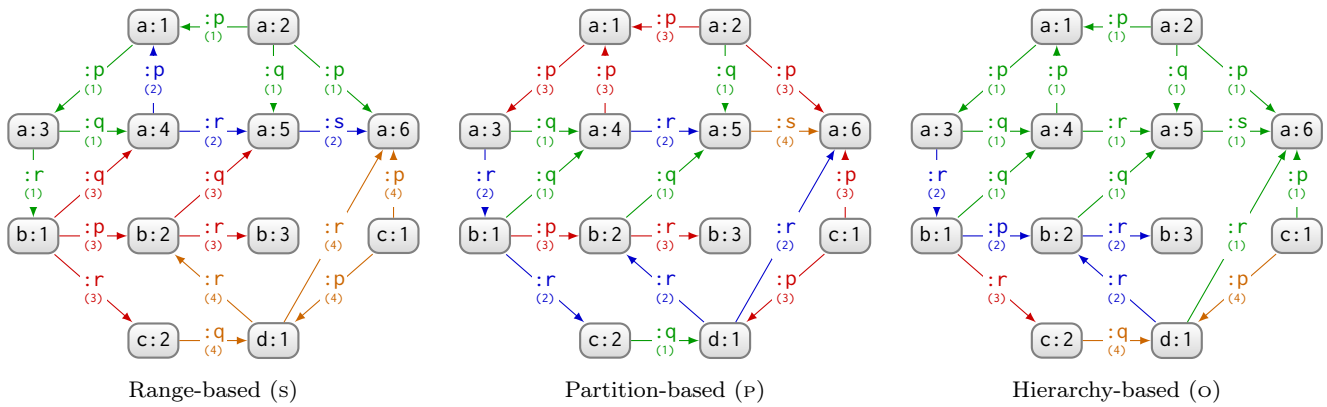



Fig. 19: Examples of triple-based partitioning schemes

Fig. 20: Example of optimal k -way partitioning ($k = 4$)

an individual tuple, allowing joins on the same partition key to be pushed to individual machines. Graph-based partitions can allow for evaluating more complex graph patterns on individual machines, but are more costly to compute and more difficult to maintain (considering, e.g., dynamic data). Information about queries, where available, can also be used for the purposes of workload-based partitioning, which partitions or replicates data in order to enable locality for common sub-patterns. Replication can further improve load balancing, locality and fault-tolerance at the cost of redundant storage.

9 Conclusion

We have introduced the key techniques underlying state-of-the-art RDF stores and SPARQL engines, specifically for storing and indexing RDF graphs, as well as processing joins and more complex forms of graph queries. We have further summarized techniques for partitioning and managing large RDF graphs over multiple machines. We now conclude this survey by introducing additional material that can be found in the

appendix complementing the current paper, as well as general research challenges in the area.

9.1 Systems and Benchmarks

We have conducted a comprehensive survey of over one hundred individual RDF stores and SPARQL query engines – both distributed and local – in terms of which of the techniques discussed herein they use. We have further surveyed the synthetic and real-world benchmarks available for evaluating these systems under a variety of criteria. For discussion of individual systems and benchmarks, we refer to the appendix of this extended version.

9.2 Research challenges

Though major advances have been made in terms of the *scale* and *efficiency* of RDF stores in recent years, these will remain central challenges as the scale of RDF graphs and demand for querying them in more complex ways increases. Other challenges have only been occasionally or partially addressed, such as support for *dynamicity*, where many of the techniques surveyed assume static data, and would not be able to handle updates gracefully. Another challenge relates to *query expressivity*, where most works have focused on optimizing joins, with some works further exploring optimizations for property paths; however, SPARQL contains a rich set of query features that could benefit from further optimizations, and the expressivity of the language is sure to grow (e.g., in the context of SPARQL 1.2). Yet another emerging challenge is that of *query volume*, where leading SPARQL endpoints process millions of queries per day; though related to efficiency, this challenge specifically motivates further research on

workload-aware or caching strategies that leverage frequent patterns. With the diversity of techniques and systems that have been proposed, and the RDF graphs and queries that can be considered, *evaluation* remains a key challenge in order to understand the complex trade-offs that underlie such systems. In the broader context of the Semantic Web, challenges also arise from providing support for *federated querying*, *entailment*, *validation*, etc., which raise many interesting questions.

Acknowledgments Hogan was supported by Fondecyt Grant No. 1181896. and by ANID – Millennium Science Initiative Program – Code ICN17_002.

References

1. D. J. Abadi, A. Marcus, S. Madden, and K. Hollenbach. SW-Store: a vertically partitioned DBMS for Semantic Web data management. *VLDB J.*, 18(2):385–406, 2009.
2. D. J. Abadi, A. Marcus, S. Madden, and K. J. Hollenbach. Scalable Semantic Web Data Management Using Vertical Partitioning. In *VLDB*, pages 411–422. ACM, 2007.
3. I. Abdelaziz, R. Harbi, Z. Khayyat, and P. Kalnis. A Survey and Experimental Comparison of Distributed SPARQL Engines for Very Large RDF Data. *PVLDB*, 10(13):2049–2060, 2017.
4. I. Abdelaziz, R. Harbi, S. Salihoglu, and P. Kalnis. Combining Vertex-Centric Graph Processing with SPARQL for Large-Scale RDF Data Analytics. *IEEE Trans. Parallel Distributed Syst.*, 28(12):3374–3388, 2017.
5. C. R. Aberger, A. Lamb, S. Tu, A. Nötzli, K. Olukotun, and C. Ré. EmptyHeaded: A Relational Engine for Graph Processing. *ACM Trans. Database Syst.*, 42(4):20:1–20:44, 2017.
6. F. Abiri, M. Kahani, and F. Zarinkalam. An entity based RDF indexing schema using Hadoop and HBase. In *International Conference on Computer and Knowledge Engineering (ICCKE)*, pages 68–73, 2014.
7. A. Akhter, A.-C. N. Ngonga, and M. Saleem. An empirical evaluation of RDF graph partitioning techniques. In *European Knowledge Acquisition Workshop*, pages 3–18. Springer, 2018.
8. K. Alaoui. A Categorization of RDF Triplestores. In *International Conference on Smart City Applications (SCA)*, pages 1–7. ACM, 2019.
9. G. Aluç, O. Hartig, M. T. Özsu, and K. Daudjee. Diversified stress testing of RDF data management systems. In *International Semantic Web Conference (ISWC)*, pages 197–212. Springer, 2014.
10. G. Aluç, M. T. Özsu, K. Daudjee, and O. Hartig. chameleon-db: a Workload-Aware Robust RDF Data Management System, 2013.
11. S. Álvarez-García, N. R. Brisaboa, J. D. Fernández, M. A. Martínez-Prieto, and G. Navarro. Compressed vertical partitioning for efficient RDF management. *Knowl. Inf. Syst.*, 44(2):439–474, 2015.
12. R. Angles, M. Arenas, P. Barceló, A. Hogan, J. L. Reutter, and D. Vrgoc. Foundations of Modern Query Languages for Graph Databases. *ACM CSUR*, 50(5):68:1–68:40, 2017.
13. A. Aranda-Andújar, F. Bugiotti, J. Camacho-Rodríguez, D. Colazzo, F. Goasdoué, Z. Kaoudi, and I. Manolescu. AMADA: Web Data Repositories in the Amazon Cloud. In *Conference on Information and Knowledge Management (CIKM)*, page 2749–2751. ACM, 2012.
14. M. Atre, V. Chaoji, M. J. Zaki, and J. A. Hendler. Matrix "Bit" loaded: a scalable lightweight join query processor for RDF data. In *International Conference on World Wide Web (WWW)*, pages 41–50. ACM, 2010.
15. M. Atre and J. A. Hendler. BitMat: A Main Memory Bit-matrix of RDF Triples. In *Workshop on Scalable Semantic Web Knowledge Base Systems (SSWS)*, page 33, 2009.
16. A. Atserias, M. Grohe, and D. Marx. Size Bounds and Query Plans for Relational Joins. *SIAM J. Comput.*, 42(4):1737–1767, 2013.
17. J. Baier, D. Daroch, J. L. Reutter, and D. Vrgoc. Evaluating Navigational RDF Queries over the Web. In *ACM Conference on Hypertext and Social Media (HT)*, pages 165–174. ACM, 2017.
18. S. Bail, S. Alkiviadous, B. Parsia, D. Workman, M. van Harmelen, R. S. Gonçalves, and C. Garilao. FishMark: A Linked Data Application Benchmark. In *Workshop on Scalable and High-Performance Semantic Web Systems (SSWS+HPCSW)*, pages 1–15. CEUR, 2012.
19. M. Banane. RDFMongo: A MongoDB Distributed and Scalable RDF management system based on Meta-model. *Int. J. of Adv. Trends in Comp. Sci. and Eng.*, 8:734–741, 2019.
20. S. Bartoň. Designing Indexing Structure for Discovering Relationships in RDF Graphs. In *International Workshop on Databases, Texts, Specifications and Objects (DATESO)*, pages 7–17. CEUR, 2004.
21. B. R. Bebee, D. Choi, A. Gupta, A. Gutmans, A. Khandelwal, Y. Kiran, S. Mallidi, B. McGaughy, M. Personick, K. Rajan, S. Rondelli, A. Ryazanov, M. Schmidt, K. Sengupta, B. B. Thompson, D. Vaidya, and S. Wang. Amazon Neptune: Graph Data Management in the Cloud. In *ISWC Posters & Demonstrations, Industry and Blue Sky Ideas Tracks*. CEUR-WS.org, 2018.
22. D. Beckett. The design and implementation of the redland RDF application framework. In *World Wide Web Conference (WWW)*, pages 449–456. ACM, 2001.
23. C. Beeri and R. Ramakrishnan. On the Power of Magic. *J. Log. Program.*, 10(3&4):255–299, 1991.
24. A. Bigerl, F. Conrads, C. Behning, M. Sherif, M. Saleem, and A.-C. N. Ngomo. Tentris: A Tensor-Based Triple Store. In *International Semantic Web Conference (ISWC)*, pages 56–73. Springer, 2020.
25. R. Binna, W. Gassler, E. Zangerle, D. Pacher, and G. Specht. SpiderStore: A Native Main Memory Approach for Graph Storage. In *Grundlagen von Datenbanken (GI-Workshop)*, pages 91–96. CEUR, 2011.
26. B. Bishop, A. Kiryakov, D. Ognyanov, I. Peikov, Z. Tashchev, and R. Velkov. OWLIM: A family of scalable semantic repositories. *Semantic Web*, 2(1):33–42, 2011.
27. C. Bizer, R. Meusel, and A. Pimpl. Web Data Commons – Microdata, RDFa, JSON-LD, and Microformat Data Sets, 2020.
28. C. Bizer and A. Schultz. The Berlin SPARQL Benchmark. *Int. J. Sem. Web Inf. Syst.*, 5(2):1–24, 2009.
29. M. A. Bornea, J. Dolby, A. Kementsietsidis, K. Srinivas, P. Dantressangle, O. Udrea, and B. Bhattacharjee. Building an Efficient RDF Store over a Relational Database. In *International Conference on Management of Data (SIGMOD)*, page 121–132. ACM, 2013.

30. N. R. Brisaboa, A. Cerdeira-Pena, G. de Bernardo, and A. Fariña. Revisiting Compact RDF Stores Based on k2-Trees. In *Data Compression Conference (DCC)*, pages 123–132. IEEE, 2020.
31. N. R. Brisaboa, A. Cerdeira-Pena, G. de Bernardo, A. Fariña, and G. Navarro. Space/time-efficient RDF stores based on circular suffix sorting. *CoRR*, abs/2009.10045, 2020.
32. N. R. Brisaboa, A. Cerdeira-Pena, A. Fariña, and G. Navarro. A Compact RDF Store Using Suffix Arrays. In *String Processing and Information Retrieval (SPIRE)*, pages 103–115. Springer, 2015.
33. M. Bröcheler, A. Pugliese, and V. S. Subrahmanian. DOGMA: A Disk-Oriented Graph Matching Algorithm for RDF Databases. In *International Semantic Web Conference (ISWC)*, pages 97–113. Springer, 2009.
34. J. Broekstra, A. Kampman, and F. Van Harmelen. Sesame: A generic architecture for storing and querying RDF and RDF Schema. In *International Semantic Web Conference (ISWC)*, pages 54–68. Springer, 2002.
35. C. Buil-Aranda, A. Hogan, J. Umbrich, and P.-Y. Vandenbussche. SPARQL Web-Querying Infrastructure: Ready for Action? In *International Semantic Web Conference (ISWC)*, pages 277–293. Springer, 2013.
36. A. Buluç, H. Meyerhenke, I. Safro, P. Sanders, and C. Schulz. Recent Advances in Graph Partitioning. In *Algorithm Engineering - Selected Results and Surveys*, pages 117–158. Springer, 2016.
37. P. Cappellari, R. D. Virgilio, and M. Roantree. Path-oriented keyword search over graph-modeled Web data. *World Wide Web Conference (WWW)*, 15(5-6):631–661, 2012.
38. S. Cebiric, F. Goasdoué, H. Kondylakis, D. Kotzinos, I. Manolescu, G. Troullinou, and M. Zneika. Summarizing semantic graphs: a survey. *VLDBJ*, 28(3):295–327, 2019.
39. T. Chawla, G. Singh, E. S. Pilli, and M. Govil. Storage, partitioning, indexing and retrieval in big rdf frameworks: A survey. *Computer Science Review*, 38:100309, 2020.
40. T. Chawla, G. Singh, E. S. Pilli, and M. C. Govil. Storage, partitioning, indexing and retrieval in Big RDF frameworks: A survey. *Comput. Sci. Rev.*, 38:100309, 2020.
41. X. Chen, H. Chen, N. Zhang, and S. Zhang. SparkRDF: Elastic Discreted RDF Graph Processing Engine with Distributed Memory. In *ISWC Posters & Demos (ISWC-PD)*, pages 261–264. CEUR, 2014.
42. L. Cheng and S. Kotoulas. Scale-out processing of large RDF datasets. *IEEE Trans. Big Data*, 1(4):138–150, 2015.
43. E. I. Chong, S. Das, G. Eadon, and J. Srinivasan. An Efficient SQL-Based RDF Querying Scheme. In *International Conference on Very Large Databases (VLDB)*, page 1216–1227. VLDB End., 2005.
44. F. Conrads, J. Lehmann, M. Saleem, M. Morsey, and A. N. Ngomo. IGUANA: A Generic Framework for Benchmarking the Read-Write Performance of Triple Stores. In *International Semantic Web Conference (ISWC)*, pages 48–65. Springer, 2017.
45. O. Corby, R. Dieng-Kuntz, and C. Faron-Zucker. Querying the semantic web with corese search engine. In R. L. de Mántaras and L. Saitta, editors, *European Conference on Artificial Intelligence (ECAI)*, pages 705–709. IOS Press, 2004.
46. O. Corby and C. Faron-Zucker. Implementation of SPARQL Query Language Based on Graph Homomorphism. In *International Conference on Conceptual Structures (ICCS)*, pages 472–475. Springer, 2007.
47. M. Cossu, M. Färber, and G. Lausen. PRoST: Distributed Execution of SPARQL Queries Using Mixed Partitioning Strategies. In *International Conference on Extending Database Technology (EDBT)*, pages 469–472. OpenProceedings, 2018.
48. O. Curé, G. Blin, D. Revuz, and D. C. Faye. Waterfowl: A compact, self-indexed and inference-enabled immutable RDF store. In *Extended Semantic Web Conference (ESWC)*, pages 302–316. Springer, 2014.
49. G. Demartini, I. Enchev, M. Wylot, J. Gapany, and P. Cudré-Mauroux. BowlognaBench - Benchmarking RDF Analytics. In *Symposium on Data-Driven Process Discovery and Analysis (SIMPDA)*, pages 82–102. Springer, 2011.
50. F. Du, H. Bian, Y. Chen, and X. Du. Efficient SPARQL Query Evaluation in a Database Cluster. In *BigData Congress*, pages 165–172. IEEE, 2013.
51. S. Duan, A. Kementsietsidis, K. Srinivas, and O. Udrea. Apples and oranges: A comparison of RDF benchmarks and real RDF datasets. In *International Conference on Management of Data (SIGMOD)*, pages 145–156. ACM, 2011.
52. N. M. Elzein, M. A. Majid, I. A. T. Hashem, I. Yaqoob, F. A. Alaba, and M. Imran. Managing big RDF data in clouds: Challenges, opportunities, and solutions. *Sustainable Cities and Society*, 39:375–386, 2018.
53. O. Erling, A. Averbuch, J. Larriba-Pey, H. Chafi, A. Gubichev, A. Prat-Pérez, M. Pham, and P. A. Boncz. The LDBC Social Network Benchmark: Interactive Workload. In *International Conference on Data Management (SIGMOD)*, pages 619–630. ACM, 2015.
54. O. Erling and I. Mikhailov. *Virtuoso: RDF Support in a Native RDBMS*, pages 501–519. Springer, 2010.
55. D. C. Faye, O. Curé, and G. Blin. A survey of RDF storage approaches. In *Revue Africaine de la Recherche en Informatique et Mathématiques Appliquées*, page pp. 25, 2012.
56. G. H. L. Fletcher, J. Peters, and A. Poulouvasilis. Efficient regular path query evaluation using path indexes. In *Extending Database Technology (EDBT)*, pages 636–639. OpenProceedings.org, 2016.
57. L. Galárraga, K. Hose, and R. Schenkel. Partout: A distributed engine for efficient RDF processing. In *WWW Companion*, pages 267–268. ACM, 2014.
58. M. Galkin, K. M. Endris, M. Acosta, D. Collarana, M. Vidal, and S. Auer. SMJoin: A Multi-way Join Operator for SPARQL Queries. In *International Conference on Semantic Systems (SEMANTICS)*, pages 104–111. ACM, 2017.
59. F. Goasdoué, Z. Kaoudi, I. Manolescu, J. Quiané-Ruiz, and S. Zampetakis. CliqueSquare: Flat plans for massively parallel RDF queries. In *International Conference on Data Engineering (ICDE)*, pages 771–782. IEEE, 2015.
60. D. Graux, L. Jachiet, P. Genevès, and N. Layaïda. SPARQLGX: Efficient distributed evaluation of SPARQL with Apache Spark. In *International Semantic Web Conference (ISWC)*, pages 80–87. Springer, 2016.
61. J. Groppe, S. Groppe, A. Schleifer, and V. Linnemann. LuposDate: a semantic web database system. In *ACM Conference on Information and Knowledge Management (CIKM)*, pages 2083–2084. ACM, 2009.

62. S. Groppe, J. Groppe, and V. Linnemann. Using an index of precomputed joins in order to speed up SPARQL processing. In *International Conference on Enterprise Information Systems (ICEIS)*, pages 13–20, 2007.
63. S. Groppe, T. Kiencke, S. Werner, D. Heinrich, M. Stelzner, and L. Gruenwald. P-LUOSDATE: Using Precomputed Bloom Filters to Speed Up SPARQL Processing in the Cloud. *Open J. Semantic Web*, 1(2):25–55, 2014.
64. A. Gubichev, S. J. Bedathur, and S. Seufert. Sparqling kleene: fast property paths in RDF-3X. In *Workshop on Graph Data Management Experiences & Systems (GRADES)*, page 14. CWI/ACM, 2013.
65. A. Gubichev and T. Neumann. Path Query Processing on Very Large RDF Graphs. In *International Workshop on the Web and Databases (WebDB)*, 2011.
66. A. Gubichev and T. Neumann. Exploiting the query structure for efficient join ordering in SPARQL queries. In *International Conference on Extending Database Technology (EDBT)*, pages 439–450. OpenProceedings.org, 2014.
67. X. Guo, H. Gao, and Z. Zou. Leon: A Distributed RDF Engine for Multi-query Processing. In *International Conference on Database Systems for Advanced Applications (DASFAA)*, pages 742–759. Springer, 2019.
68. X. Guo, H. Gao, and Z. Zou. WISE: Workload-Aware Partitioning for RDF Systems. *Big Data Research*, 22:100161, 2020.
69. Y. Guo, Z. Pan, and J. Heflin. LUBM: A benchmark for OWL knowledge base systems. *J. Web Sem.*, 3(2-3):158–182, 2005.
70. S. Gurajada, S. Seufert, I. Miliaraki, and M. Theobald. TriAD: A distributed shared-nothing RDF engine based on asynchronous message passing. In *International Conference on Management of Data (SIGMOD)*, pages 289–300. ACM, 2014.
71. P. Haase, J. Broekstra, A. Eberhart, and R. Volz. A Comparison of RDF Query Languages. In *International Semantic Web Conference (ISWC)*, pages 502–517. Springer, 2004.
72. M. Hammoud, D. A. Rabbou, R. Nouri, S. M. R. Beheshti, and S. Sakr. DREAM: Distributed RDF engine with adaptive query planner and minimal communication. *PVLDB*, 8(6):654–665, 2015.
73. R. Harbi, I. Abdelaziz, P. Kalnis, N. Mamoulis, Y. Ebrahim, and M. Sahli. Accelerating SPARQL queries by exploiting hash-based locality and adaptive partitioning. *VLDBJ*, 25(3):355–380, 2016.
74. S. Harris and N. Gibbins. 3store: Efficient Bulk RDF Storage. In *International Workshop on Practical and Scalable Semantic Systems (PSSS)*, pages 1–20, 2003.
75. S. Harris, N. Lamb, and N. Shadbolt. 4store: The design and implementation of a clustered RDF store. In *International Workshop on Scalable Semantic Web Systems (SSWS)*, pages 94–109, 2009.
76. S. Harris, A. Seaborne, and E. Prud’hommeaux. SPARQL 1.1 Query Language. W3C Recommendation, 2013. <http://www.w3.org/TR/sparql11-query/>.
77. A. Harth and S. Decker. Optimized index structures for querying RDF from the Web. In *Latin American Web Congress (LA-WEB)*, pages 71–80. IEEE, 2005.
78. A. Harth, J. Umbrich, A. Hogan, and S. Decker. YARS2: A Federated Repository for Querying Graph Structured Data from the Web. In *International Semantic Web Conference (ISWC)*, pages 211–224. Springer, 2007.
79. M. Hassan and S. K. Bansal. Data Partitioning Scheme for Efficient Distributed RDF Querying Using Apache Spark. In *International Conference on Semantic Computing (ICSC)*, pages 24–31. IEEE, 2019.
80. L. He, B. Shao, Y. Li, H. Xia, Y. Xiao, E. Chen, and L. J. Chen. Stylus: A Strongly-Typed Store for Serving Massive RDF Data. *PVLDB*, 11(2):203–216, 2017.
81. A. Hogan, C. Riveros, C. Rojas, and A. Soto. A worst-case optimal join algorithm for SPARQL. In *International Semantic Web Conference (ISWC)*, pages 258–275. Springer, 2019.
82. A. Hogan, C. Riveros, C. Rojas, and A. Soto. Wiki-data Graph Pattern Benchmark (WGPB) for RDF/SPARQL, 2019.
83. A. Hogenboom, F. Frasincar, and U. Kaymak. Ant colony optimization for RDF chain queries for decision support. *Expert Syst. Appl.*, 40(5):1555–1563, 2013.
84. A. Hogenboom, V. Milea, F. Frasincar, and U. Kaymak. RCQ-GA: RDF Chain Query Optimization Using Genetic Algorithms. In *E-Commerce and Web Technologies (EC-Web)*, pages 181–192. Springer, 2009.
85. K. Hose and R. Schenkel. WARP: Workload-aware replication and partitioning for RDF. In *ICDE Workshops*, pages 1–6, 2013.
86. J. Huang, D. J. Abadi, and K. Ren. Scalable SPARQL querying of large RDF graphs. *PVLDB*, 4(11):1123–1134, 2011.
87. V. Ingalalli, D. Ienco, and P. Poncelet. *Querying RDF Data: a Multigraph-based Approach*, chapter 5, pages 135–165. John Wiley & Sons, 2018.
88. Y. E. Ioannidis and E. Wong. Query Optimization by Simulated Annealing. In *International Conference on Management of Data (SIGMOD)*, pages 9–22. ACM, 1987.
89. H. Jabeen, E. Haziiev, G. Sejdiu, and J. Lehmann. DISE: A Distributed in-Memory SPARQL Processing Engine over Tensor Data. In *International Conference on Semantic Computing (ICSC)*, pages 400–407. ACM, 2020.
90. M. Janik and K. Kochut. BRAHMS: a workbench RDF store and high performance memory system for semantic association discovery. In *International Semantic Web Conference (ISWC)*, pages 431–445. Springer, 2005.
91. D. Janke and S. Staab. Storing and Querying Semantic Data in the Cloud. In *Reasoning Web Summer School*, pages 173–222. Springer, 2018.
92. D. Janke, S. Staab, and M. Thimm. Koral: A Glass Box Profiling System for Individual Components of Distributed RDF Stores. In *Workshop on Benchmarking Linked Data (BLINK)*. CEUR, 2017.
93. D. Janke, S. Staab, and M. Thimm. On Data Placement Strategies in Distributed RDF Stores. In *International Workshop on Semantic Big Data (SBD)*, pages 1–6. ACM, 2017.
94. E. G. Kalayci, T. E. Kalayci, and D. Birant. An ant colony optimisation approach for optimising SPARQL queries by reordering triple patterns. *Inf. Syst.*, 50:51–68, 2015.
95. O. Kalinsky, O. Mishali, A. Hogan, Y. Etsion, and B. Kimelfeld. Efficiently Charting RDF. *CoRR*, abs/1811.10955, 2018.
96. Z. Kaoudi and I. Manolescu. RDF in the clouds: a survey. *VLDB Journal*, 24(1):67–91, 2015.
97. G. Karvounarakis, S. Alexaki, V. Christophides, D. Plexousakis, and M. Scholl. RQL: a declarative query language for RDF. In *International World Wide Web Conference (WWW)*, pages 592–603. ACM, 2002.

98. G. Karvounarakis, A. Magkanaraki, S. Alexaki, V. Christophides, D. Plexousakis, M. Scholl, and K. Tolle. Querying the Semantic Web with RQL. *Comput. Net.*, 42(5):617–640, 2003.
99. G. Karypis and V. Kumar. A Fast and High Quality Multilevel Scheme for Partitioning Irregular Graphs. *SIAM J. Sci. Comput.*, 20(1):359–392, 1998.
100. A. Katib, V. Slavov, and P. Rao. RIQ: Fast processing of SPARQL queries on RDF quadruples. *Journal of Web Sem.*, 37:90–111, 2016.
101. V. Khadilkar, M. Kantarcioglu, B. Thuraisingham, and P. Castagna. Jena-HBase: A Distributed, Scalable and Efficient RDF Triple Store. In *ISWC Posters and Demos (ISWC-PD)*, page 85–88. CEUR, 2012.
102. J. Kim, H. Shin, W.-S. Han, S. Hong, and H. Chafi. Taming Subgraph Isomorphism for RDF Query Processing. *PVLDB*, 8(11), 2015.
103. K. Kim, B. Moon, and H. Kim. Rg-index: An RDF graph index for efficient SPARQL query processing. *Expert Syst. Appl.*, 41(10):4596–4607, 2014.
104. K. Kim, B. Moon, and H. Kim. R3F: RDF triple filtering method for efficient SPARQL query processing. *WWW J.*, 18(2):317–357, 2015.
105. A. Kiryakov, D. Ognyanov, and D. Manov. OWLIM - A pragmatic semantic repository for OWL. In *WISE Workshops*, pages 182–192. Springer, 2005.
106. D. Kolas, I. Emmons, and M. Dean. Efficient linked-list RDF indexing in Parliament. In *International Workshop on Scalable Semantic Web Knowledge Base Systems (SSWS)*, pages 17–32. CEUR, 2009.
107. E. V. Kostylev, J. L. Reutter, M. Romero, and D. Vrgoc. SPARQL with Property Paths. In *International Semantic Web Conference (ISWC)*, pages 3–18. Springer, 2015.
108. M. Koubarakis and K. Kyzirakos. Modeling and Querying Metadata in the Semantic Sensor Web: The Model StRDF and the Query Language StSPARQL. In *Extended Semantic Web Conference (ESWC)*, pages 425–439. Springer, 2010.
109. K. Kyzirakos, M. Karpathiotakis, and M. Koubarakis. Strabon: A Semantic Geospatial DBMS. In *International Semantic Web Conference (ISWC)*, pages 295–311. Springer, 2012.
110. G. Ladwig and A. Harth. CumulusRDF: Linked Data Management on Nested Key-Value Stores, 2011.
111. T. Lampo, M. Vidal, J. Danilow, and E. Ruckhaus. To Cache or Not To Cache: The Effects of Warming Cache in Complex SPARQL Queries. In *On the Move to Meaningful Internet Systems (OTM)*, pages 716–733. Springer, 2011.
112. K. Lee and L. Liu. Scaling queries over big RDF graphs with semantic hash partitioning. *PVLDB*, 6(14):1894–1905, 2013.
113. J. Leeka and S. Bedathur. RQ-RDF-3X: going beyond triplestores. In *ICDE Workshops*, pages 263–268. IEEE, 2014.
114. J. Lehmann, G. Sejdin, L. Böhmann, P. Westphal, C. Stadler, I. Ermilov, S. Bin, N. Chakraborty, M. Saleem, A.-C. Ngonga Ngomo, and H. Jabeen. Distributed Semantic Analytics Using the SANSA Stack. In *International Semantic Web Conference (ISWC)*, pages 147–155. Springer, 2017.
115. A. Letelier, J. Pérez, R. Pichler, and S. Skritek. Static analysis and optimization of semantic web queries. *ACM TODS*, 38(4):25:1–25:45, 2013.
116. W. Li, B. Zhang, G. Rao, R. Chen, and Z. Feng. Hash Tree Indexing for Fast SPARQL Query in Large Scale RDF Data Management Systems. In *ISWC Posters & Demonstrations and Industry Tracks*. CEUR-WS.org, 2017.
117. B. Liu and B. Hu. HPRD: a high performance RDF database. *Int. J. Parallel Emergent Distr. Syst.*, 25(2):123–133, 2010.
118. J. Lorey and F. Naumann. Caching and Prefetching Strategies for SPARQL Queries. In *ESWC Satellite Events*, pages 46–65. Springer, 2013.
119. Y. Luo, F. Picalausa, G. H. L. Fletcher, J. Hidders, and S. Vansummen. Storing and Indexing Massive RDF Datasets. In *Semantic Search over the Web*, pages 31–60. Springer, 2012.
120. X. Lyu, X. Wang, Y. Li, Z. Feng, and J. Wang. GraSS: An Efficient Method for RDF Subgraph Matching. In *Web Information Systems Engineering Conference (WISE)*, pages 108–122. Springer, 2015.
121. L. Ma, Z. Su, Y. Pan, L. Zhang, and T. Liu. RStar: an RDF storage and query system for enterprise resource management. In *International Conference on Information and Knowledge Management (CIKM)*, pages 484–491, 2004.
122. Z. Ma, M. A. Capretz, and L. Yan. Storing massive Resource Description Framework (RDF) data: A survey. *Knowl. Eng. Rev.*, 31(4):391–413, 2016.
123. A. Madkour, A. M. Aly, and W. G. Aref. WORQ: Workload-Driven RDF Query Processing. In *International Semantic Web Conference (ISWC)*, pages 583–599. Springer, 2018.
124. R. Maharjan, Y. Lee, and S. Lee. Exploiting Path Indexes to Answer Complex Queries in Ontology Repository. In *International Conference on Computational Science and Its Applications (ICCSA)*, pages 56–61. IEEE, 2009.
125. S. Malyshev, M. Krötzsch, L. González, J. Gonsior, and A. Bielefeldt. Getting the Most Out of Wikidata: Semantic Technology Usage in Wikipedia’s Knowledge Graph. In *International Semantic Web Conference (ISWC)*, pages 376–394. Springer, 2018.
126. W. Martens and T. Trautner. Evaluation and Enumeration Problems for Regular Path Queries. In *International Conference on Database Theory (ICDT)*, pages 19:1–19:21, 2018.
127. M. Martin, J. Unbehauen, and S. Auer. Improving the Performance of Semantic Web Applications with SPARQL Query Caching. In *Extended Semantic Web Conference (ESWC)*, pages 304–318. Springer, 2010.
128. B. McBride. Jena: A semantic web toolkit. *IEEE Internet Computing*, 6(6):55–58, 2002.
129. J. P. McGlothlin and L. R. Khan. RDFJoin: A Scalable Data Model for Persistence and Efficient Querying of RDF Datasets. Tech. Rep. UTDCS-08-09, Univ. of Texas at Dallas, 2009.
130. J. P. McGlothlin and L. R. Khan. RDFKB: Efficient support for RDF inference queries and knowledge management. In *International Database Engineering & Applications Symposium (IDEAS)*, pages 259–266. ACM, 2009.
131. M. Meimaris and G. Papastefanatos. Distance-based triple reordering for SPARQL query optimization. In *International Conference on Data Engineering (ICDE)*, pages 1559–1562. IEEE Computer Society, 2017.
132. M. Meimaris, G. Papastefanatos, N. Mamoulis, and I. Anagnostopoulos. Extended Characteristic Sets: Graph Indexing for SPARQL Query Optimization. In *International Conference on Data Engineering (ICDE)*, pages 497–508. IEEE, 2017.

133. M. Morsey, J. Lehmann, S. Auer, and A. N. Ngomo. DBpedia SPARQL Benchmark - Performance Assessment with Real Queries on Real Data. In *International Semantic Web Conference (ISWC)*, pages 454–469. Springer, 2011.
134. R. Mutharaju, S. Sakr, A. Sala, and P. Hitzler. D-SPARQ: Distributed, Scalable and Efficient RDF Query Engine. In *ISWC Posters & Demos (ISWC-PC)*, page 261–264. CEUR, 2013.
135. A. Muys. Building an Enterprise-Scale Database for RDF Data. Netymon technical paper, 2006.
136. H. Naacke and O. Curé. On distributed SPARQL query processing using triangles of RDF triples. *Open J. Semantic Web*, 7(1):17–32, 2020.
137. G. Navarro, J. L. Reutter, and J. Rojas-Ledesma. Optimal Joins Using Compact Data Structures. In *International Conference on Database Theory (ICDT)*, pages 21:1–21:21. Schloss Dagstuhl, 2020.
138. Y. Nenov, R. Piro, B. Motik, I. Horrocks, Z. Wu, and J. Banerjee. RDFox: A Highly-Scalable RDF Store. In *International Semantic Web Conference (ISWC)*, pages 3–20. Springer, 2015.
139. T. Neumann and G. Moerkotte. Characteristic sets: Accurate cardinality estimation for RDF queries with multiple joins. In *International Conference on Data Engineering (ICDE)*, pages 984–994. IEEE, 2011.
140. T. Neumann and G. Weikum. RDF-3X: a RISC-style engine for RDF. *Proc. VLDB Endow.*, 1(1):647–659, 2008.
141. T. Neumann and G. Weikum. Scalable join processing on very large RDF graphs. In *International Conference on Management of Data (SIGMOD)*, pages 627–640. ACM, 2009.
142. T. Neumann and G. Weikum. The RDF-3X engine for scalable management of RDF data. *VLDBJ*, 19(1):91–113, 2010.
143. H. Q. Ngo, E. Porat, C. Ré, and A. Rudra. Worst-case Optimal Join Algorithms. *J. ACM*, 65(3):16:1–16:40, 2018.
144. A. Owens, A. Seaborne, N. Gibbins, and m. schraefel. Clustered TDB: A clustered triple store for Jena. EPrints Server, 2008.
145. M. T. Özsu. A survey of RDF data management systems. *Frontiers of Comp. Sci.*, 10(3):418–432, 2016.
146. Z. Pan, T. Zhu, H. Liu, and H. Ning. A survey of RDF management technologies and benchmark datasets. *Journal of Ambient Intelligence and Humanized Computing*, 9(5):1693–1704, 2018.
147. N. Papailiou, I. Konstantinou, D. Tsoumakos, P. Karras, and N. Koziris. H2RDF+: High-performance distributed joins over large-scale RDF graphs. *International Conference on Big Data (Big Data)*, pages 255–263, 2013.
148. N. Papailiou, I. Konstantinou, D. Tsoumakos, and N. Koziris. H2RDF: Adaptive Query Processing on RDF Data in the Cloud. In *World Wide Web Conference (WWW)*, page 397–400. ACM, 2012.
149. N. Papailiou, D. Tsoumakos, P. Karras, and N. Koziris. Graph-Aware, Workload-Adaptive SPARQL Query Caching. In *International Conference on Management of Data (SIGMOD)*, pages 1777–1792. ACM, 2015.
150. P. Peng, L. Zou, L. Chen, and D. Zhao. Adaptive Distributed RDF Graph Fragmentation and Allocation based on Query Workload. *IEEE TKDE*, 31(4):670–685, 2019.
151. P. Peng, L. Zou, M. T. Özsu, L. Chen, and D. Zhao. Processing SPARQL Queries over Distributed RDF Graphs. *VLDBJ*, 25(2):243–268, 2016.
152. J. Pérez, M. Arenas, and C. Gutierrez. Semantics and Complexity of SPARQL. *ACM TODS*, 34(3), 2009.
153. M. Pham and P. A. Boncz. Exploiting Emergent Schemas to Make RDF Systems More Efficient. In *International Semantic Web Conference (ISWC)*, pages 463–479, 2016.
154. F. Picalausa, Y. Luo, G. H. L. Fletcher, J. Hidders, and S. Vansummeren. A Structural Approach to Indexing Triples. In *Extended Semantic Web Conference (ESWC)*, pages 406–421. Springer, 2012.
155. A. Potocki, A. Polukhin, G. Drobyazko, D. Hladky, V. P. Klintsov, and J. Unbehauen. OntoQuad: Native High-Speed RDF DBMS for Semantic Web. In *Knowledge Engineering and the Semantic Web (KESW)*, volume 394 of *Communications in Computer and Information Science*, pages 117–131. Springer, 2013.
156. A. Potter, B. Motik, Y. Nenov, and I. Horrocks. Dynamic Data Exchange in Distributed RDF Stores. *IEEE TKDE*, 30(12):2312–2325, 2018.
157. R. Punnoose, A. Crainiceanu, and D. Rapp. Rya: A Scalable RDF Triple Store for the Clouds. In *International Workshop on Cloud Intelligence (Cloud-I)*. ACM, 2012.
158. S. Qiao and Z. M. Özsoyoglu. RBench: Application-Specific RDF Benchmarking. In *International Conference on Management of Data (SIGMOD)*, pages 1825–1838. ACM, 2015.
159. A. Rajith, S. Nishimura, and H. Yokota. JARS: join-aware distributed RDF storage. In *International Database Engineering & Applications Symposium (IDEAS)*, pages 264–271. ACM, 2016.
160. P. Ravindra, H. Kim, and K. Anyanwu. An Intermediate Algebra for Optimizing RDF Graph Pattern Matching on MapReduce. In *Extended Semantic Web Conference (ESWC)*, pages 46–61. Springer, 2011.
161. J. L. Reutter, A. Soto, and D. Vrgoc. Recursion in SPARQL. In *International Semantic Web Conference (ISWC)*, pages 19–35. Springer, 2015.
162. K. Rohloff and R. E. Schantz. High-Performance, Massively Scalable Distributed Systems Using the MapReduce Software Framework: The SHARD Triple-Store. In *Programming Support Innovations for Emerging Distributed Applications (PSI Eta)*. ACM, 2010.
163. S. Sakr and G. Al-Naymat. Relational processing of RDF queries: a survey. *SIGMOD Record*, 38(4):23–28, 2010.
164. M. Saleem, M. I. Ali, A. Hogan, Q. Mehmood, and A. N. Ngomo. LSQ: The Linked SPARQL Queries Dataset. In *International Semantic Web Conference (ISWC)*, pages 261–269. Springer, 2015.
165. M. Saleem, A. Hasnain, and A. N. Ngomo. LargeRDF-Bench: A billion triples benchmark for SPARQL endpoint federation. *J. Web Sem.*, 48:85–125, 2018.
166. M. Saleem, Q. Mehmood, and A. N. Ngomo. FEASIBLE: A Feature-Based SPARQL Benchmark Generation Framework. In *International Semantic Web Conference (ISWC)*, pages 52–69. Springer, 2015.
167. M. Saleem, G. Szárnyas, F. Conrads, S. A. C. Bukhari, Q. Mehmood, and A.-C. Ngonga Ngomo. How Representative Is a SPARQL Benchmark? An Analysis of RDF Triplestore Benchmarks. In *World Wide Web Conference (WWW)*, page 1623–1633. ACM, 2019.
168. A. Schätzle, M. Przyjaciół-Zablocki, T. Berberich, and G. Lausen. S2X: Graph-Parallel Querying of RDF with GraphX. In *Workshop on Biomedical Data Management and Graph Online Querying (Big-O(Q))*, pages 155–168. Springer, 2016.

169. A. Schätzle, M. Przyjaciół-Zablocki, and G. Lausen. PigSPARQL: Mapping SPARQL to Pig Latin. In *International Workshop on Semantic Web Information Management (SWIM)*, 2011.
170. A. Schätzle, M. Przyjaciół-Zablocki, A. Neu, and G. Lausen. Sempala: Interactive SPARQL query processing on Hadoop. In *International Semantic Web Conference (ISWC)*, pages 164–179, 2014.
171. A. Schätzle, M. Przyjaciół-Zablocki, S. Skilevic, and G. Lausen. S2RDF: RDF querying with SPARQL on Spark. *PVLDB*, 9(10):804–815, 2016.
172. M. Schmidt, O. Görlitz, P. Haase, G. Ladwig, A. Schwarte, and T. Tran. FedBench: A Benchmark Suite for Federated Semantic Data Query Processing. In *International Semantic Web Conference (ISWC)*, pages 585–600. Springer, 2011.
173. M. Schmidt, T. Hornung, G. Lausen, and C. Pinkel. SP2Bench: A SPARQL Performance Benchmark. In *International Conference on Data Engineering ICDE*, pages 222–233. IEEE, 2009.
174. M. Schmidt, M. Meier, and G. Lausen. Foundations of SPARQL query optimization. In *International Conference on Database Theory (ICDT)*, pages 4–33, 2010.
175. G. Schreiber and Y. Raimond. RDF 1.1 Primer. W3C Working Group Note, 2014. <http://www.w3.org/TR/rdf11-primer/>.
176. P. G. Selinger, M. M. Astrahan, D. D. Chamberlin, R. A. Lorie, and T. G. Price. Access Path Selection in a Relational Database Management System. In *International Conference on Management of Data (SIGMOD)*, pages 23–34. ACM, 1979.
177. S. Seufert, A. Anand, S. J. Bedathur, and G. Weikum. FERRARI: Flexible and efficient reachability range assignment for graph indexing. In *International Conference on Data Engineering (ICDE)*, pages 1009–1020. IEEE, 2013.
178. J. Shi, Y. Yao, R. Chen, H. Chen, and F. Li. Fast and Concurrent RDF Queries with RDMA-Based Distributed Graph Exploration. In *Conference on Operating Systems Design and Implementation (OSDI)*, pages 317–332. USENIX, 2016.
179. M. Sintek and M. Kiesel. RDFBroker: A Signature-Based High-Performance RDF Store. In *European Semantic Web Conference (ESWC)*, pages 363–377. Springer, 2006.
180. C. Stadler, G. Sejdić, D. Graux, and J. Lehmann. Sparkliff: A Scalable Software Component for Efficient Evaluation of SPARQL Queries over Distributed RDF Datasets. In *International Semantic Web Conference (ISWC)*, pages 293–308. Springer, 2019.
181. M. Stocker, A. Seaborne, A. Bernstein, C. Kiefer, and D. Reynolds. SPARQL basic graph pattern optimization using selectivity estimation. In *World Wide Web Conference (WWW)*, pages 595–604. ACM, 2008.
182. M. Stonebraker, D. J. Abadi, A. Batkin, X. Chen, M. Cherniack, M. Ferreira, E. Lau, A. Lin, S. Madden, E. O’Neil, P. O’Neil, A. Rasin, N. Tran, and S. Zdonik. C-Store: A Column-Oriented DBMS. In *VLDB*, pages 553–564. VLDB End., 2005.
183. H. Stuckenschmidt. Similarity-Based Query Caching. In *International Conference on Flexible Query Answering Systems (FQAS)*, pages 295–306. Springer, 2004.
184. H. Stuckenschmidt, R. Vdovjak, J. Broekstra, and G. Houben. Towards distributed processing of RDF path queries. *Int. J. Web Eng. Tech.*, 2(2/3):207–230, 2005.
185. P. Stutz, A. Bernstein, and W. Cohen. Signal/Collect: Graph Algorithms for the (Semantic) Web. In *International Semantic Web Conference (ISWC)*, pages 764–780. Springer, 2010.
186. P. Stutz, M. Verman, L. Fischer, and A. Bernstein. TripleRush: A Fast and Scalable Triple Store. In *Workshop on Scalable Semantic Web Knowledge Base Systems (SSWS)*, page 50–65. CEUR, 2013.
187. M. Svoboda and I. Mlýnková. Linked Data indexing methods: A survey. In *OTM Confederated International Conferences (OTM)*, pages 474–483. Springer, 2011.
188. G. Szárnyas, B. Izsó, I. Ráth, and D. Varró. The Train Benchmark: Cross-technology performance evaluation of continuous model queries. *Softw. Syst. Model.*, 17(4):1365–1393, 2018.
189. H. Thakkar, R. Angles, M. Rodriguez, S. Mallette, and J. Lehmann. Let’s build Bridges, not Walls: SPARQL Querying of TinkerPop Graph Databases with Sparql-Gremlin. In *International Conference on Semantic Computing (ICSC)*, pages 408–415. IEEE, 2020.
190. B. B. Thompson, M. Personick, and M. Cutcher. The bigdata® rdf graph database. In *Linked Data Management*, pages 193–237. CRC Press, 2014.
191. T. Tran, G. Ladwig, and S. Rudolph. Managing Structured and Semistructured RDF Data Using Structure Indexes. *IEEE TKDE*, 25(9):2076–2089, 2013.
192. C. Tschanz. Query-Driven Index Partitioning for TripleRush. Bachelor Thesis, 2014.
193. P. Tsialiamanis, L. Sidirourgos, I. Fundulaki, V. Christophides, and P. A. Boncz. Heuristics-based query optimisation for SPARQL. In *International Conference on Extending Database Technology (EDBT)*, pages 324–335. ACM, 2012.
194. O. Udrea, A. Pugliese, and V. S. Subrahmanian. GRIN: A Graph Based RDF Index. In *Conference on Artificial Intelligence (AAAI)*, pages 1465–1470. AAAI, 2007.
195. T. L. Veldhuizen. Triejoin: A Simple, Worst-Case Optimal Join Algorithm. In *International Conference on Database Theory (ICDT)*, pages 96–106. OpenProceedings.org, 2014.
196. M. Vidal, E. Ruckhaus, T. Lampo, A. Martínez, J. Sierra, and A. Polleres. Efficiently Joining Group Patterns in SPARQL Queries. In *Extended Semantic Web Conference (ESWC)*, pages 228–242. Springer, 2010.
197. A. Vlachou, C. Doukeridis, A. Glenis, G. M. Santipantakis, and G. A. Vouros. Efficient Spatio-Temporal RDF Query Processing in Large Dynamic Knowledge Bases. In *Symposium on Applied Computing (SAC)*, page 439–447. ACM, 2019.
198. D. Vrandečić and M. Krötzsch. Wikidata: a free collaborative knowledgebase. *CACM*, 57(10):78–85, 2014.
199. X. Wang, L. Chai, Q. Xu, Y. Yang, J. Li, J. Wang, and Y. Chai. Efficient Subgraph Matching on Large RDF Graphs Using MapReduce. *Data Sci. Eng.*, 4(1):24–43, 2019.
200. X. Wang, S. Wang, Y. Xin, Y. Yang, J. Li, and X. Wang. Distributed Pregel-based provenance-aware regular path query processing on RDF knowledge graphs. *World Wide Web*, 23(3):1465–1496, 2020.
201. C. Weiss, P. Karras, and A. Bernstein. Hexastore: Sextuple indexing for semantic web data management. *PVLDB*, 1(1):1008–1019, 2008.
202. R. T. Whitman, B. G. Marsh, M. B. Park, and E. G. Hoel. Distributed Spatial and Spatio-Temporal Join on Apache Spark. *ACM TSAS*, 5(1), 2019.
203. K. Wilkinson, C. Sayers, H. Kuno, and D. Reynolds. Efficient RDF storage and retrieval in Jena2. In *International Conference on Semantic Web and Databases (SWDB)*, pages 120–139. CEUR, 2003.

204. G. T. Williams and J. Weaver. Enabling Fine-Grained HTTP Caching of SPARQL Query Results. In *International Semantic Web Conference (ISWC)*, pages 762–777. Springer, 2011.
205. D. Wood, P. Gearon, and T. Adams. Kowari: A platform for semantic web storage and analysis. In *XTech Conference*, pages 1–7, 2005.
206. B. Wu, Y. Zhou, P. Yuan, H. Jin, and L. Liu. Sem-Store: A Semantic-Preserving Distributed RDF Triple Store. In *Conference on Information and Knowledge Management (CIKM)*, pages 509–518. ACM, 2014.
207. G. Wu, J. Li, J. Hu, and K. Wang. System Pi: A Native RDF Repository Based on the Hypergraph Representation for RDF Data Model. *J. Comput. Sci. Technol.*, 24(4):652–664, 2009.
208. G. Wu and M. Yang. Improving SPARQL query performance with algebraic expression tree based caching and entity caching. *J. Zhejiang Univ. Sci. C*, 13(4):281–294, 2012.
209. H. Wu, T. Fujiwara, Y. Yamamoto, J. Bolleman, and A. Yamaguchi. BioBenchmark Toyama 2012: An evaluation of the performance of triple stores on biological data. *J. Bio. Sem.*, 5, 2014.
210. M. Wylot and P. Cudré-Mauroux. DiploCloud: Efficient and Scalable Management of RDF Data in the Cloud. *IEEE TKDE*, 28(3):659–674, 2016.
211. M. Wylot, M. Hauswirth, P. Cudré-Mauroux, and S. Sakr. RDF Data Storage and Query Processing Schemes: A Survey. *ACM CSUR*, 51(4):84:1–84:36, 2018.
212. M. Wylot, J. Pont, M. Wisniewski, and P. Cudré-Mauroux. DipLODocus[RDF]: Short and Long-Tail RDF Analytics for Massive Webs of Data. In *International Semantic Web Conference (ISWC)*, pages 778–793. Springer, 2011.
213. N. Yakovets, P. Godfrey, and J. Gryz. Evaluation of SPARQL Property Paths via Recursive SQL. In *Alberto Mendelzon International Workshop on Foundations of Data Management (AMW)*. CEUR, 2013.
214. N. Yakovets, P. Godfrey, and J. Gryz. Query Planning for Evaluating SPARQL Property Paths. In *International Conference on Management of Data (SIGMOD)*, pages 1875–1889. ACM, 2016.
215. S. Yang, X. Yan, B. Zong, and A. Khan. Towards effective partition management for large graphs. In *International Conference on Management of Data (SIGMOD)*, pages 517–528. ACM, 2012.
216. M. Q. Yasin, X. Zhang, R. Haq, Z. Feng, and S. Yitagesu. A Comprehensive Study for Essentiality of Graph Based Distributed SPARQL Query Processing. In *International Conference on Database Systems for Advanced Applications (DASFAA)*, pages 156–170. Springer, 2018.
217. P. Yuan, P. Liu, B. Wu, H. Jin, W. Zhang, and L. Liu. TripleBit: A Fast and Compact System for Large Scale RDF Data. *PVLDB*, 6(7):517–528, 2013.
218. L. H. Zambom Santana and R. dos Santos Mello. An Analysis of Mapping Strategies for Storing RDF Data into NoSQL Databases. In *Symposium on Applied Computing (SAC)*, pages 386–392. ACM, 2020.
219. K. Zeng, J. Yang, H. Wang, B. Shao, and Z. Wang. A distributed graph engine for web scale RDF data. In *PVLDB*, pages 265–276, 2013.
220. W. E. Zhang, Q. Z. Sheng, K. Taylor, and Y. Qin. Identifying and Caching Hot Triples for Efficient RDF Query Processing. In *Database Systems for Advanced Applications (DASFAA)*, pages 259–274. Springer, 2015.
221. X. Zhang, L. Chen, Y. Tong, and M. Wang. EAGRE: Towards scalable I/O efficient SPARQL query evaluation on the cloud. In *International Conference on Data Engineering (ICDE)*, pages 565–576, 2013.
222. W. Zheng, L. Zou, X. Lian, H. Zhang, W. Wang, and D. Zhao. SQBC: an efficient subgraph matching method over large and dense graphs. *Inf. Sci.*, 261:116–131, 2014.
223. L. Zou, J. Mo, L. Chen, M. Tamer Özsu, and D. Zhao. gStore: Answering SPARQL queries via subgraph matching. *PVLDB*, 4(8):482–493, 2011.

A Survey of RDF Stores

We now present a survey of local and distributed RDF stores, and how they use the aforementioned techniques. At the end of this section we will discuss some general trends for RDF stores. We include here systems for which we could find technical details regarding (at least) the storage, indexing and processing of joins over RDF graphs. In the case of distributed RDF stores, we expect similar technical details, along with the type of partitioning and/or replication used. We include systems with associated publications, as well as systems that are unpublished but widely known in practice. Both local and distributed systems are presented in approximate chronological order, based on the year of publication, or an approximate year in which the system was released. For unpublished local stores, we include the year where RDF was first supported. For unpublished distributed stores, we include the approximate year when distributed features were added. Some stores that are often deployed in local environments also support distribution; they are included in both sections. Some systems are unnamed; if it is a distributed store that extends an existing local store, we append the suffix “-D” or “-D2” to the local store’s name; otherwise we use an abbreviation based on authors and year. Where systems change name, we prefer the more modern name. The papers sometimes use different terminology to refer to similar concepts; we often map the original terminology to that used in the body of the survey in order to increase coherence and improve readability.

A.1 Local RDF Stores

The local RDF stores we include, and the techniques they use, are summarized in Table 4.

Redland [22] (2001) is a set of RDF libraries for native storage of RDF that has seen various developments over the years. The original paper describes triple-table like storage based on creating three hash maps – $SP \rightarrow O$, $PO \rightarrow S$, $SO \rightarrow P$ – which, given two elements of an RDF triple, allow for finding the third element; for example, using $PO \rightarrow S$, we can find the subjects of triples with a given predicate and object. The hash maps can be stored either in-memory or on a persistent storage. Support for the RDQL and SPARQL query languages were later enabled through the Rasqal RDF query library.

Jena [128] (2002) uses relational databases to store RDF graphs as triple tables, with entries for subject, predicate, object IRIs, and object literals. IRIs and literals are encoded as IDs and two separate dictionaries are

created for both. Indexing is delegated to an underlying relational DBMS (e.g., Postgresql, MySQL, Oracle, etc.). RDQL is used as a query language and is translated into SQL and run against the underlying relational DBMS. The Jena store would later be extended in various directions, with SDB referring to the use of relational-style storage (per the original system), and TDB referring to the use of native storage.

RDF4J [34] (2002), known originally as Sesame, provides persistent storage and querying of RDF data. RDF4J provides storage-independent solutions and can thus be deployed on top of a variety of storage engines such as RDBMSs and object-oriented databases. Graph queries can be expressed with the RQL language. The storage, indexing, and query processing techniques depend on the underlying storage engine used by RDF4J. Recent versions of RDF4J features improved functionalities such as both in-memory and persistent data storage, SerQL and SPARQL support etc.

RSSDB [97,98] (2002) stores an RDF graph using a vertical partitioning approach with Postgres as the underlying database. Two variants are considered for class instances: creating a unary table per class (named after the class, with rows indicating instances), or creating one binary table called instances (with rows containing both the instance and the class) in order to reduce the number of tables. Four tables are also added to model RDFS definitions (classes, properties with their domain and range, sub-classes and sub-properties). The system supports queries in RQL (proposed in the same paper [98]), which are translated to SQL by an RQL interpreter and evaluated over Postgres.

3store [74] (2003) uses MySQL as its back-end, storing RDF graphs in four tables, namely a triple table, a models table, a resource table, and a literal table. The triple table stores RDF triples (one per row) with additional information: (1) the model the triple belongs to, (2) a boolean value to indicate if the object is a literal, and (3) a boolean value to indicate if this triple is inferred. The models, resource, and literal tables are two-column tables that dictionary encode models, resources, and literals, respectively. Queries expressed in RDQL are rewritten to SQL for execution over MySQL.

AllegroGraph ⁹ (2003) is a general purpose store for semi-structured data that can be used to query documents (e.g., JSON) and graph data (e.g., RDF). RDF data is stored and indexed in six permutations as quads,

⁹ <https://franz.com/agraph/allegrograph/>

Table 4: Categorization of local RDF Engines.

Storage: T = Triple Table, Q = Quad Table, V = Vertical Partitioning, P = Property table, G = Graph-based, M = Miscellaneous**Indexing:** T = Triple, Q = Quad, E = Entity, P = Property, N = Path/Navigational, J = Join, S = Structural, M = Miscellaneous**Join Proc.:** P = Pairwise, M = Multiway, W = Worst case optimal**Query Proc.:** N = Paths/Navigational, R = Relational, Q = Query rewriting

Engine	Year	Storage						Indexing								Join Pr.			Query Pr.		
		T	Q	V	P	G	M	T	Q	E	P	N	J	S	M	P	M	W	N	R	Q
RedLand [22]	2001	✓						✓													
Jena [128]	2002	✓													✓	✓					✓
RDF4J [34]	2002	✓		✓											✓	✓					✓
RSSDB [97,98]	2002	✓													✓	✓					✓
3store [74]	2003	✓													✓	✓					✓
AllegroGraph	2003		✓						✓							✓			✓	✓	
Jena2 [203]	2003	✓			✓										✓	✓	✓				✓
CORESE [45,46]	2004	✓						✓								✓			✓	✓	
Jena TDB	2004	✓	✓					✓	✓							✓			✓	✓	
RStar [121]	2004	✓													✓	✓					✓
BRAHMS [90]	2005	✓						✓											✓		
GraphDB [105,26]	2005	✓	✓					✓	✓							✓			✓	✓	
Kowari [205]	2005		✓						✓							✓			✓	✓	
RDF_MATCH [43]	2005		✓		✓				✓	✓			✓		✓	✓	✓		✓	✓	
YARS [77]	2005		✓						✓							✓			✓	✓	
RDFBroker [179]	2006				✓										✓	✓	✓		✓	✓	
Virtuoso [54]	2006	✓	✓												✓	✓			✓	✓	
GRIN [194]	2007					✓								✓		✓					
SW-Store [1]	2007			✓																	
Blazegraph [190]	2008	✓	✓					✓	✓							✓	✓		✓	✓	
Hexastore [201]	2008					✓		✓								✓				✓	
RDF-3X [142]	2008	✓						✓								✓	✓			✓	
BitMat [15]	2009					✓		✓		✓							✓				
DOGMA [33]	2009					✓				✓				✓			✓				
LuposDate [61]	2009	✓						✓								✓					
Parliament [106]	2009	✓	✓					✓	✓												
RDFJoin	2009	✓						✓					✓			✓	✓			✓	✓
System II [207]	2009					✓		✓				✓	✓			✓				✓	
HPRD [117]	2010	✓	✓					✓	✓		✓	✓				✓	✓				
Stardog	2010		✓					✓	✓							✓			✓	✓	
dipLODocus [212]	2011					✓		✓		✓						✓				✓	
gStore [223]	2011					✓		✓		✓							✓				
SpiderStore [138]	2011					✓		✓		✓											
SAINT-DB [154]	2012					✓								✓		✓					
Strabon [109]	2012			✓			✓	✓							✓	✓				✓	
DB2RDF [29]	2013				✓					✓					✓	✓				✓	✓
OntoQuad [155]	2013		✓						✓							✓	✓			✓	
OSQP [191]	2013					✓								✓		✓				✓	
Triplebit [217]	2013			✓		✓		✓								✓	✓				
R3F [103,104]	2014	✓				✓		✓			✓	✓				✓					
RQ-RDF-3X [113]	2014		✓						✓							✓				✓	
SQBC [222]	2014					✓				✓							✓				
WaterFowl [48]	2014	✓						✓								✓				✓	
GraSS [120]	2015					✓		✓		✓							✓				
k ² -triples [11]	2015			✓				✓								✓					
RDFCSA [32,31]	2015	✓						✓								✓					
RDFox [138]	2015	✓						✓								✓			✓	✓	
TurboHOM++ [102]	2015					✓				✓							✓				
RIQ [100]	2016		✓			✓		✓						✓			✓				
axonDB [131]	2017	✓						✓			✓					✓	✓				
HTStore [116]	2017					✓				✓											
AMBER [87]	2018					✓															
Jena-LTJ [81]	2019	✓						✓										✓			
BMatrix [30]	2020	✓						✓													
Tentris [24]	2020	✓						✓										✓		✓	

which are additionally associated with a triple identifier. SPARQL queries are supported, where the most recent version provides an option for two query engines: SBQE, which is optimized for SPARQL 1.0-style queries, and MJQE, which features merge joins and caching techniques optimized for property paths.

Jena2 [203] (2003) is a revised version of the original Jena database schema, with support for both triple and property tables. Unlike the original version, IRIs and literals are stored directly in the tables, unless they exceed a certain length, in which case they are dictionary encoded by two separate tables; this allows filter operations to be directly performed on the triple and property tables, thus reducing dictionary lookups, but increasing storage sizes as string values are stored multiple times. Indexing is handled by an underlying relational database, and graph queries in RDQL are rewritten to SQL queries evaluated over the database.

CORESE [45,46] (2004) began as a search engine with path-finding functionality and inference support for ontologies [45], but was extended to support SPARQL query features [46]. *CORESE* models RDF graphs as conceptual graphs; for simplicity we discuss their methods in terms of the RDF model. RDF graphs are indexed according to the terms, enabling the efficient evaluation of triple patterns. Given a basic graph pattern, the triple patterns are reordered based on heuristics – such as the number of constants or filters associated with the triple pattern, or the number of variables bound by previous triple patterns in the order – as well as cardinality estimates. A nested-loop style algorithm is then applied to perform joins. Filters are also evaluated as soon as possible to reduce intermediate results.

Jena TDB ¹⁰ (2004) is a native RDF store that has seen continuous development in the past decades. A TDB instance consists of three tables: a node table (a dictionary, allowing to encode/decode RDF terms to/from 8-byte identifiers), a triple/quad table (with dictionary-encoded terms), and a prefixes table (used to store common prefixes used for abbreviations). Storage is based on custom B+ trees, which are used to build indexes for various triple/quad permutations. Join processing uses pairwise (nested-loop) joins, with a variety of statistic- and heuristic-based methods available for join reordering. SPARQL 1.1 query processing is implemented in the custom Jena ARQ query processor. Jena TDB has become the recommended RDF store for Jena, with older relational-based storage (later named Jena SDB) having been deprecated.

RStar [121] (2004) is designed to store (RDFS-style) ontology information and instance data. The storage is based on multiple relations, stored in the IBM DB2 RDBMS. Five two-column tables are used to store ontological data (property dictionary, sub-property relations, class dictionary, sub-class relations, and domain and range relations). Another five two-column tables are used to store instance-related data (literal dictionary, IRI dictionary, triples, class instances, namespace dictionary). *RStar* pushes indexing and other tasks to the underlying database. The *RStar* Query Language (RSQL) is used to express queries, where *RStar* translates RSQL into the SQL of an underlying database.

BRAHMS [90] (2005) is an in-memory RDF store. The RDF graph is stored and indexed in three hash tables – $S \rightarrow PO$, $O \rightarrow SP$, $P \rightarrow SO$ – which allow for finding triples that use a particular constant. The main motivating use-case of *BRAHMS* is to find semantic associations – i.e., paths between two subject/object nodes – in large RDF graphs. This path-search functionality was implemented in *BRAHMS* using depth-first search and breadth-first search algorithms.

GraphDB [105,26] (2005) (formerly known as OWLIM) stores RDF graphs using a mix of triple and quad tables. In the most recent version, indexes are built for two triple permutations (POS and PSO) as well as a quad permutation (GPSO). Predicate lists (SP and OP) can also be indexed in order to quickly find the predicates associated with a given subject or object. Terms are dictionary encoded. Joins are reordered according to cardinality estimations. SPARQL 1.1 is supported, along with a wide range of other features, including spatial features, full-text indexing, inference, semantic similarity, integration with MongoDB, and more besides.

Mulgara [205,135] (2005), a fork of an earlier RDF store known as Kowari, implements native RDF storage in the form of quads tables using AVL trees. Dictionary encoding based on 64-bit longs is used. Support for transactions is provided using immutable arrays that store quads on disk in compressed form, with skiplists enabling fast search; insertions and deletions lead to a new immutable array being generated on disk. Indexing is based on six permutations of quads (which is sufficient to efficiently evaluate all sixteen possible quad patterns). Joins are evaluated pairwise and reordered (possibly on-the-fly) based on cardinality estimations. Queries are expressed in the iTQL language, where SPARQL support was added later.

¹⁰ <https://jena.apache.org/documentation/tdb/architecture.html>

RDF_MATCH [43] (2005) is an RDF store that uses the Oracle RDBMS as an underlying database. It stores RDF data in two different tables: a dictionary table, and a quads table. Indexes are defined based on B-trees. Queries are evaluated as self-joins over the quads table, which is further joined with the dictionary. Materialized views can further be used to index S-S, S-P, S-O, P-P, P-O and O-O joins on-demand, as well as selected property tables. Support for Datalog-style rules is also provided. Queries are expressed in an SQL-style syntax, with functions used to express graph patterns, which in turn are interpreted and evaluated by Oracle.

YARS [77] (2005) is a native RDF store that indexes quad tables in B+ trees using dictionary encoded terms. It uses four types of indexes: dictionary indexes, keyword indexes, quad indexes (with six permutations), and cardinality indexes that count occurrences of quad patterns. YARS implements pairwise (nested-loop) joins that are reordered based on cardinalities. Basic graph patterns in Notation3 (N3) syntax are supported.

RDFBroker [179] (2006) is an RDF store that follows a property table approach. For each subject in the graph, its signature (equivalent to the notion of characteristic sets that would come later) is extracted, with a property table defined for each signature, including a column for the subject, and a column for each property in the signature. Support for RDFS reasoning is also described. An index over signatures is proposed based on a lattice that models set containment between signatures. Given a signature extracted from the query, the lattice can be used to find tables corresponding to signatures that subsume that of the query. A prototype based on in-memory storage is described, implementing typical relational query optimizations such as join reordering.

Virtuoso [54] (2006) stores RDF data as a quad table, where in the most recent version, by default, the quad table includes five indexes: PSOG (for P, PS, PSO and PSOG prefixes), POGS (for PO and POG), SP (for S), OP (for O), and GS (for G). The table is sorted according to the primary key PSOG. If only the subject or object are known, the SP/OP index can be used to identify predicates, allowing the PSOG/POGS index to be subsequently used. In the case that only the graph is known, the GS index is joined with SP and then with PSOG. Subjects, predicates and graph names are dictionary encoded; objects are stored in raw form (for fast filtering). By default, column-wise storage is used. Query execution is based on translating SPARQL queries into SQL to be executed on a custom underlying database.

GRIN [194] (2007) is an RDF store based on a structural index (see Section 5.6). This index is a binary tree, where the root refers to all the nodes of the graph, and both children divide the nodes of its parents based on a given distance from a given node in the graph. The leaves can then be seen as forming a partition of the triples in the graph induced by the nodes in its division. The structural index is used to find small subgraphs that may generate results for a query, over which an existing subgraph matching algorithm is applied.

SW-Store [2, 1] (2007) is an RDF store based on vertical partitioning. SW-Store relies on a column-oriented DBMS called C-store [182], which is shown to pair well with vertical partitioning in terms of performance (e.g., the object column of a `foaf:age` table with have integers in an interval $[0, 150]$, which are highly compressible). Each table is indexed by both subject and object. An “overflow” triple table is used for inserts alongside the compressed, vertically partitioned tables. Jena ARQ is used to translate SPARQL queries into SQL for evaluation over C-Store. Pairwise joins are used, preferring merge joins when data are sorted appropriately, otherwise using index-nested loop joins. Materialization of S-O joins is also discussed.

Blazegraph [190] (2008), formerly known as *BigData*, is a native RDF store supporting SPARQL 1.1. Blazegraph allows for indexing triples for RDF graphs, or quads for SPARQL datasets. Three index permutations are generated for triples, and six permutations are generated for quads; indexes are based on B+ trees. Both row and column data storage models are supported, which can be saved both in-memory or on-disk. Dictionary encoding with 64-bit integers is used for compressed representation of RDF triples. Two query optimization strategies are available: the default approach uses static analysis and cardinality estimation; the second approach uses runtime sampling of join graphs. Supported joins include hash joins, index nested-loop joins, merge joins, and multiway star joins.

Hexastore [201] (2008) is an in-memory RDF store based on adjacency lists similar to Figure 8. Six indexes are built for all $3! = 6$ possible orderings of the elements of a triple. For example, in the SPO index, each subject s is associated with an ordered vector of predicates, wherein each p in turn points to an ordered vector of objects. In the PSO index, each p points to a vector of subjects, wherein each s points to the same vector of objects as used for sp in the SPO index. Terms are dictionary encoded. Having all six index orders allows for (pairwise) merge joins to be used extensively.

RDF-3X [140,141,142] (2008) stores an RDF graph as triple tables in compressed indexes, which are based on clustered B+trees whose values are delta-encoded to reduce space. Triples are indexed in all six possible ways. RDF-3X also indexes counts for all triple patterns with one or two (distinct) constants; thus, it can find the exact number of triples matching, for example, (s,p,o) , (s,p,o) , etc. Counts are also maintained for frequent path and star joins. Join reordering is based on a dynamic programming procedure, which uses the aforementioned counts as cardinality statistics, and leverages the complete indexing of all six permutations to enable merge joins, further evaluating multiway joins in the case of star patterns. A later extension adds sideways information passing strategies [141] in order to filter intermediate results based on global heuristics.

BitMat [15,14] (2009) considers a one-hot encoding of an RDF graph, i.e., a 3-dimensional bit array (or matrix/tensor) of dimension $|S| \times |P| \times |O|$, where S , P , O , indicate the set of subjects, predicates and objects of the graph; and where index s,p,o contains a 1 if the corresponding (dictionary encoded) triple (s,p,o) exists, or 0 otherwise. The system stores slices of this array (called BitMats), where for each predicate, SO and OS BitMats are stored; for each subject, a PO BitMat is stored; and for each object, a PS BitMat is stored. Each BitMat is a 2-dimensional bit array; e.g., the OS matrix of a given predicate p enables finding all objects matching o in (s,p,o) or all subjects matching s in (s,p,o) . Though OP and SP BitMats could be indexed for subjects and objects, resp., the authors argue they would be rarely used. BitMats also store the count of 1's (triples) they contain, a row vector indicating which columns contain a 1, and a column vector indicating which rows contain a 1. In total, $2|P||S||O|$ BitMats are generated, gap compressed, stored on disk, and loaded into memory as needed. Bitwise AND/OR/NOT operators are used, in memory, to process multiway joins.

DOGMA [33] (2009) is a graph-based RDF store, where an RDF graph is first decomposed into subgraphs using a graph partitioning algorithm. These subgraphs are indexed as the leaves of a balanced binary tree stored on disk. Each non-leaf node in this tree encodes the k -merge of its two children, which is a graph with k nodes that is isomorphic to a quotient graph (see Section 5.6) of both children. DOGMA then proposes a variety of algorithms for evaluating basic graph patterns with constant predicates. The basic algorithm generates a set of candidate results for each individual variable node based on its incoming and outgoing edges; starting with the node with the fewest candidates, the algorithm then proceeds to check the edges between

them in a depth-first manner (similar to wco joins). Further algorithms prune the sets of candidates based on their distance from the candidates of other nodes in the query, leveraging the distance within and between nodes in the subgraphs in the leaves of the binary tree.

LuposDate [61] (2009) stores data in a triple table. A total of seven hash indexes are added for S , P , O , SP , SO , PO , SPO , enabling efficient evaluation of all eight triple patterns. Triples are also annotated with their rank (position in the order) with respect to the six permutations of the triple; for example, the OPS rank indicates the position of the triple when the graph is sorted in OPS order. These ranks are used for fast sorting of intermediate results when applying sort-merge joins.

Parliament [106] (2009) stores RDF graphs in three types of tables: a resource table encoding details of individual terms, a statement table encoding triples, and a dictionary table. These tables are stored as linked lists. For each RDF term, the resource table stores the first triple containing the term in the statement table; the number of triples that use it in the subject, predicate and object position; a pointer to its entry in the dictionary; and a bitvector encoding metadata about the term. The statement table contains eight components: a statement identifier; three identifiers for the subject, predicate, and object of the triple; three statement identifiers pointing to the next triple with the same subject, predicate, and object; and a bitvector for encoding metadata of the statement. This storage scheme avoids the need for multiple orders and enables fast lookups when triple patterns have one constant; for triple patterns with multiple constants, however, the most selective constant is looked up, with filters run to check the other constants. SPARQL query processing is enabled through existing libraries, such as Jena ARQ.

RDFJoin [129] (2009) stores RDF graphs using three types of tables. Two dictionary tables are used to encode and decode subjects/objects and predicates. Three triple tables are used, where each has two positions of the triple as the primary key, and the third position is encoded as a bit vector; for example, in the PO table, predicate and object are used as a primary key, and for each predicate-object pair, a bit vector of length $|so(G)|$ is given, with a 1 at index k encoding a triple with the subject identifier k for the given predicate and object. Join tables store the results of S-S, O-O, and S-O joins, encoded with the two predicates of both triples as primary key (joins using the same predicate twice are excluded), and a bit vector to encode the join terms for that predicate pair (the subject/object that

matches the join variable). MonetDB and LucidDB are proposed as underlying databases. SPARQL queries are supported, where joins are evaluated using the join indexes and pairwise algorithms. Inference support would later be added in the extended RDFKB system [130].

System II [207] (2009) is a graph-based RDF store. Nodes of the graph are indexed with an identifier, value and type, while edges are indexed as triples, with their subject, predicate and object identifiers. Nodes are then linked to their inward and outward edges, which enables lookups for triples with a given subject or object. A more compressed version where outward edges only store the predicate and object, or inward edges only store the subject and predicate, is also proposed. Edges are then grouped by their vertex (outward edges by subject, inward edges by object). The compressed variant is thus similar to an adjacency list. Indexes are built for three triple permutations, as well as for reachability queries on a single predicate (PLSD – see Section 5.4). SPARQL (1.0) queries are supported over the proposed indexes using pairwise joins, with the PLSD index used to support entailment over transitive properties.

HPRD [117] (2010) is an RDF store based on three types of index over the dictionary-encoded graph. Triple indexes are based on B+trees and cover three triple permutations: SPO, PO, OS. A path index is built using a suffix arrays, and is used only to cache paths that are commonly accessed in queries; the paths are indexed by their predicates. Context indexes are based on B+Trees and are used to support temporal data, versioning, or named graphs; six permutations are covered, namely GSPO, SPO, POC, OCS, CP and SO. Cardinality statistics are further stored for triple patterns, and used for greedy join reordering. RDQL queries supported.

dipLODocus [212] (2011) is an RDF store based on the notion of a “molecule”, which is a subgraph surrounding a particular “root” node. The root nodes are defined based on matching triple patterns provided by the administrator. The molecule of a root node is then the subgraph formed by expanding outward in the graph until another root node is encountered. Dictionary encoding is used. Indexes are further built mapping nodes and the values of properties indicated by the administrator to individual molecules. SPARQL is supported through the Rasqal query library, with joins pushed within individual molecules where possible; otherwise hash joins are used. Aggregate queries are further pushed to the indexes on values of individual properties (which offer benefits similar to column-wise storage).

*Stardog*¹¹ (2010) is a commercial RDF store based on RocksDB: a persistent key-value store for fast storage. Stardog indexes quads in various permutations using RocksDB. Different types of pairwise joins – such as hash join, bind join, merge join etc. – are used. Stardog supports SPARQL 1.1, full-text search through Lucene, ACID transactions, versioning, and a variety of other features, including support for property graphs.

gStore [223] (2011) is a graph-based RDF store. The RDF graph is stored using adjacency lists (see Section 4.5) where each node is associated with a bit vector – which serves as a vertex signature (see Section 5.2) – that encodes the triples where the given node is the subject. gStore then indexes these signatures in a vertex signature tree (VS-tree) that enables multi-way joins. The leaves of the VS-tree encode signatures of nodes, and non-leaf nodes encode the bitwise OR of their children; the leaves are further connected with labelled edges corresponding to edges between their corresponding nodes in the graph. Basic graph patterns can then be encoded in a similar manner to the graph, where gStore then evaluates the pattern by matching its signature with that of the indexed graph.

SpiderStore [138] (2011) is an in-memory graph store based on adjacency lists. Specifically, for each node in the RDF graph, an adjacency list for incoming and outgoing edges is stored. Likewise, for each predicate, a list of subject nodes is stored. Rather than storing the constants directly in these lists, pointers are stored to the location of the term (with the adjacency lists for the node or the subjects of the predicate). Alongside these pointers, cardinality metadata are stored. (Though SPARQL queries with basic graph patterns and filters are evaluated in the experiments, the types of join algorithms used are not described.)

SAINT-DB [154] (2012) is an RDF store with a structural index that organizes triples in the graph according to the type of join that exists between them (S-S, P-O, etc.). The index itself is then a directed edge-labeled graph whose nodes represent a set of triples from the graph, edges indicate that some pair of triples in both nodes are joinable, and edge labels indicate the type of join that exists (which makes the graph directed as S-O differs from O-S). The nodes of the index then form a partition of the graph: no triple appears in more than one node, and their union yields the graph. This index can range from a single node with all triples in the graph (with loops for each type of join present), to singleton nodes each with one triple of the graph. A

¹¹ <https://docs.stardog.com/>

condition based on semi-joins is used to strike a balance, minimizing the intermediate results generated for individual triple patterns. Given a basic graph pattern, each triple pattern is then mapped to nodes in the structural index, where the triple patterns it joins with must match some triple in a neighbor on an edge whose label corresponds to the type of join.

Strabon [109] (2012) is an RDF store that supports custom features for storing and querying geospatial data (specifically in the form of stRDF [108] data). Strabon is built upon Sesame/RDF4J, which is chosen as an open-source solution that can easily integrate with PostGIS: a DBMS with spatial features. Strabon then stores RDF using a vertical partitioning scheme with dictionary encoding; an identifier for each triple is also included. B+tree indexes are built for the three columns of each table (subject, predicate, identifier). Strabon supports an extension of SPARQL, called stSPARQL [108], for querying stRDF based datasets, with spatial features supported through PostGIS.

DB2RDF [29] (2013) uses a relational schema similar to property tables to store RDF data. However, rather than having a column for each property/predicate associated with a given subject, DB2RDF uses a “primary hash” table with columns $S, P_1, O_1, \dots, P_k, O_k$, where each P_i, O_i pair of columns indicates the i^{th} predicate-object pair associated with the subject listed in the S column. A binary “spill” column is added, with a 1 indicating that a subject has more than k triples, in which case it will occupy more than one row of the table. Rows for subjects with fewer than k triples are completed with NULLs. A second table is used to deal with multi-valued properties, where if a subject s has multiple values for the property p – say o_1, \dots, o_n – then a single fresh value v is used in the primary hash table, which is mapped to o_1, \dots, o_n (as a binary relation) in the second table. Two orders are indexed: in the second order, subjects and objects are reversed. IBM DB2 is used as the underlying database for experiments, with SPARQL queries being optimized and translated to SQL.

OntoQuad [155] (2013) is an RDF store that extends the triple-based representation of Hexastore to additionally support quads. A structure similar to a trie is used, where the top layer is a vector of values for S, P, O, G ; the second level encodes SP, \dots, GO , etc., with three children for each parent in the top layer (e.g., SP, SO, SG for S); the third layer has two children for each parent in the second layer encoding SPO, \dots, GOP ; the fourth layer has one child for each parent in the third layer, completing the quad permutation. B-trees

are then used for indexing. Both pairwise and multiway joins are supported using zig-zag joins that seek forward to the maximum compatible join value across the triple patterns. Cardinality estimates and query rewriting rules are used to optimize SPARQL query plans.

OSQP [191] (2013) is an RDF store based on a structural index using various notions of bisimulation, where two nodes in the graph are bisimilar if they cannot be distinguished from their paths. The nodes of the graph are then partitioned into sets that are pairwise bisimilar. The index is then based on a quotient graph, where supernodes correspond to a set of bisimilar nodes in the input graph. In order to reduce index sizes, a parameter corresponding to path lengths is added, such that bisimulation only considers paths within a bounded region of the graph rather than the entire graph. A basic graph pattern is then matched over the quotient graph (kept in-memory), where the triples corresponding to each matched node are retrieved (from the disk) and used to compute the final results. Custom optimizations are considered for triples with unprojected variables, whose triple patterns can be definitively “satisfied” and thus pruned based on the index; and selective triple patterns, which are evaluated directly over the RDF graph.

Triplebit [217] (2013) represents a dictionary-encoded RDF graph as a compressed 2-dimensional bit matrix. Each column of the matrix represents a triple, and each row represents a subject/object node. The subject and object rows are assigned 1 for the corresponding column of the triple. Columns are sorted by predicate, where a range of columns corresponds to the triples for that predicate. The columns for triples are sparse (at most two 1’s) and thus the two identifiers for subjects and objects are used, rather than storing 1’s; two orders are maintained for SO and OS (thus effectively covering PSO and POS orders). Two auxiliary indexes are used in TripleBit. Given a subject or object node and a predicate node, the first index (called *ID-Chunk*) supports lookups for finding the range for the unspecified object or subject. Given a subject or object node alone, the second index (called *ID-predicate*) finds predicates associated with that subject or object. Basic graph patterns are evaluated using multiway merge-joins for star joins, with semi-joins used to reduce the number of intermediate results across star joins. Join ordering uses a greedy strategy based on selectivity estimates.

R3F [103,104] (2014) is an extension of RDF-3X with path-based indexes and novel join processing techniques. The first addition is the “RP-index”, which indexes all nodes with a given incoming path expression up to a

certain length; for example, the incoming path expression pqr (of length 3) indexes all nodes z such that there exists w, x, y , such that $(w, p, x), (x, q, y), (y, r, z)$ are all triples of the graph. The RP-index is structured as a trie indexing the prefixes of the incoming path expressions, whose leaves are the list of nodes (which are dictionary encoded, sorted and delta encoded). Virtual inverse predicates are added to the RDF graph to support paths in both directions. The second extension is a modification to the sideways information passing strategy of RDF-3X to incorporate information about paths for filtering additional intermediate results.

RQ-RDF-3X [113] (2014) is an extension of RDF-3X – which supports triples – towards support for quads. The extension follows the same principles and techniques for RDF-3X, but the extension to quads requires covering additional permutations. Indexes are built for all $4! = 24$ quad permutations, similar to how RDF-3X indexes all $3! = 6$ triple permutations; having all permutations enables reading the results for any variable of any triple pattern in sorted order, which in turn enables merge joins. The delta encoding used by RDF-3X is extended to the fourth element. Like in RDF-3X, counts are indexed for all quad patterns with 1, 2, or 3 constants, requiring 4, 12 and 24 indexes, respectively (40 in total). Join and query processing use RDF-3X's techniques. RQ-RDF-3X then offers optimized support for reification using named graphs/triple identifiers.

SQBC [222] (2014) is a graph store – with support for RDF graphs – inspired by existing subgraph matching techniques for efficiently finding subgraph isomorphisms.¹² In order to index the graph, codes are extracted for each node that capture structural information about it, including its label, the largest clique containing it, the degrees of its neighbours, etc. Given a basic graph pattern, candidates are identified and filtered for variable nodes. If the basic graph pattern has no cliques, degree information is used; otherwise clique sizes can be used to filter candidate matches.

WaterFowl [48] (2014) is a compact RDF store based on succinct data structures. The RDF graph is dictionary encoded and sorted in SPO order, and represented as a trie: the first layer denotes subjects, connected to their predicates in a second layer, connected to their objects in the third layer. This trie structure is encoded

¹² The evaluation of basic graph patterns in SPARQL is defined in terms of homomorphisms, rather than subgraph isomorphisms as supported by SQBC, with the difference being that two or more variables in a basic graph pattern can match one node in the RDF graph.

in a compact representation using a combination of bit strings that indicate the number of children for a parent (e.g., for predicates, $100101\dots$ tells us that the first subject has three children (unique predicates) and the second has two); and wavelet trees that encode the sequence of terms themselves (e.g., the sequence of predicates). Pairwise joins are evaluated in terms of left-deep plans, with further support for SPARQL (1.0) features. RDFS inference is also supported.

GraSS [120] (2015) is an RDF store that is based on decomposing basic graph patterns into subgraph patterns forming star joins (considering S–S, S–O, or O–O joins). An “FFD-index” for star joins is proposed, where for each node, a bit-string signature is computed that encodes its incoming and outgoing edges, i.e., the triples in which it appears as subject or object. A neighbourhood table is constructed: each row denotes a node, which is associated with its signature and edges. Five triple permutations are further indexed (covering SP^* , OP^* , S^* , P^* , O^*), where in the SP^* permutation, for example, (s, p) pairs are mapped to a list of objects and their degrees. A basic graph pattern is then decomposed into sub-patterns forming star joins, which are evaluated using the available indexes.

k^2 -triples [11] (2015) is a compact in-memory RDF store based on k^2 trees. The RDF graph is first dictionary encoded. For each predicate, a k^2 tree is used to index its subjects and objects. In order to support variable predicates in triple patterns, SP and OP indexes are used to map subjects and objects, respectively, to their associated predicates; these indexes are encoded using compressed predicate lists. For processing basic graph patterns, S–S, S–O, O–S and O–O index nested-loop joins and merge joins are supported. A sideways information passing optimization is supported where two k^2 trees involved in a join can be descended in a coordinated fashion to filter intermediate results.

RDFox [138] (2015) is an in-memory RDF engine that supports materialization-based Datalog reasoning. The RDF graph is stored as a triple table implemented as a linked list, which stores identifiers for subject, predicate and object, as well as three pointers in the list to the next triple with the same subject, predicate and object (similar to Parliament [106]). Four indexes are built: a hash table for three constants, and three for individual constants; the indexes for individual constants offer pointers to the first triple in the list with that constant, where patterns with two constants can be implemented by filtering over this list, or (optionally) by using various orderings of the triple list to avoid filtering (e.g., a

triple list ordered by SPO can be used to evaluate patterns with constant subject and predicate without filtering). These in-memory indexes support efficient parallel updates, which are key for fast materialization. According to the implementation, (index) nested-loop joins are supported; optionally join plans can be generated based on tree decompositions. SPARQL 1.1 is further supported over the engine.

RDFCSA [32,31] (2015) is a compact in-memory RDF store based on text indexes. Specifically, triples of the RDF graph are dictionary encoded and considered to be strings of length 3. The graph is thus sorted and encoded as a string of length $3n$, where n is the number of triples. This string is indexed in a compressed suffix array (CSA): a compact data structure commonly used for indexing text. The CSA is modified by shifting elements so that instead of indexing a string of $3n$ elements, triples cycle back on themselves, giving n circular strings of length 3. Thus in an SPO permutation, after reading the object of a triple, the next integer will refer to the subject of the same triple rather than the next one in the order. With cyclical strings, one triple permutation is sufficient to support all triple patterns; SPO is in fact equivalent to POS and OSP. Merge joins, sort-merge joins and a variant of index nested-loop joins (called “chain joins”) are supported.

TurboHOM++ [102] (2015) is an in-memory, graph-based RDF store. The RDF graph is stored as the combination of adjacency lists for incoming and outgoing triples (see Section 4.5), and an index that allows for finding nodes of a particular type (based on `rdf:type`). Evaluation of basic graph patterns is then conducted by generating candidates for an initial node of the query graph based on local information (intersecting adjacency lists and type information in order to match all triple patterns that the node appears in), where the neighbors of the candidates are explored recursively in the graph guided by the graph pattern, generating candidates for further query nodes (in a manner akin to DOGMA [33]). A number of optimizations are included, including a multiway join that can check if a new candidate is connected to the candidates of multiple query nodes in one operation in a way that satisfies the query.

RIQ [100] (2016) provides a layer on top of an existing RDF store that indexes similar named graphs in a SPARQL dataset. A bit vector – called a “pattern vector” – is computed for each named graph in the dataset. The pattern vector consists of seven vectors for S, P, O, SP, SO, PO and SPO, where, e.g., the SP vector hashes all

subject–predicate pairs in the named graph. An index over the pattern vectors (PV-index) is constructed by connecting similar pattern vectors (based on locality-sensitive hashing) into a graph; each connected component of the graph forms a group of similar graphs. The union of the graphs in each group is further encoded into Bloom filters. In order to evaluate a basic graph pattern, a pattern vector is computed combining the triple patterns (e.g., a triple pattern (s, p, o) will generate a single SP sub-vector). The PV-index is then used to optimize an input query by narrowing down the candidate (named) graphs that match particular basic graph patterns (if any) before evaluating the optimized query over the underlying SPARQL store.

axonDB [131] (2017) uses two dictionary-encoded triple tables to store RDF graphs. In the first table, each triple is additionally associated with the characteristic set (CS) of its subject (see Section 5.3). The CS is assigned a unique identifier and one-hot encoded, i.e., represented by a bit vector with an index for each property that carries a 1 if the property is part of the CS, or a 0 otherwise. Triples are then sorted by their CS, grouping subjects with the same CS together. A second triple table stores each triple, along with the corresponding extended characteristic set (ECS; again see Section 5.3). The ECS is encoded with a unique identifier, and the identifiers for the subject and object CSs. The triple table is sorted by ECS. When evaluating a basic graph pattern, its analogous CSs and ECSs are extracted, along with the paths that connect them. The CSs and ECSs are matched with those of the graph, enabling multiway joins; binary hash joins are used to join the results of multiple CSs/ECSs.

HTStore [116] (2017) uses hash-based indexes to build an RDF store. Specifically, the RDF graph is indexed in a hash tree whose top layer forms a hash table over the nodes of the graph. The hash tree is based on a sequence of prime numbers. When hashing a node, the first prime number is used, and if no collision is detected, the node is inserted in the first layer. Otherwise the second prime number is used, and if no collision is detected, it is inserted in that layer as a child of the bucket of the first layer that caused the collision. Otherwise the third prime number is used, and so forth. Nodes in the hash tree then point to their adjacency lists in the graph. SPARQL queries are supported. When querying, constant nodes in the query are hashed in the same manner in order to retrieve the data for the node. SPARQL queries are supported, though details about join and query processing are omitted.

AMBER [87] (2018) stores RDF graphs using a “multi-graph” representation, where IRIs form nodes, whereas predicate–literal pairs form “attributes” on nodes. All nodes, predicates and attributes are dictionary encoded. *AMBER* then generates three indexes: the first stores the set of nodes for each attribute, the second stores vertex signatures that encode meta-data about the triples where a given node is subject or object, and the third stores adjacency lists. Basic graph patterns are evaluated by classifying query nodes with degree greater than one as core nodes, and other nodes as satellite nodes. Core nodes are processed first, where candidates are produced for each query node based on the available indexes, recursively producing candidates for neighbors; the algorithm starts with the core query node with the most satellite nodes attached, or the highest degree. For each solution over the core nodes, each satellite node is then evaluated separately as they become disconnected once the core nodes are bound to constants.

Jena-LTJ [81] (2019) extends the Jena TDB RDF store with the ability to perform worst-case optimal (wco) joins (see Section 6.3). Specifically, Jena TDB is extended with an algorithm inspired by Leapfrog TrieJoin (LTJ), which is adapted from a relational setting for the RDF/SPARQL settings. The algorithm evaluates basic graph patterns variable-by-variable in a manner that ensures that the overall cost of enumerating all of the results is proportional to the number of results that it can return in the worst case. In order to reach wco guarantees, the three-order index of Jena TDB – based on B+ trees – is extended to include all six orders. This ensures that for any triple pattern, the results for any individual variable can be read in sorted order directly from the index, which in turn enables efficient intersection of the results for individual variables across triple patterns. Thus Jena-LTJ uses twice the space of Jena TDB, but offers better query performance, particularly for basic graph patterns with cycles.

BMatrix [30] (2020) is a compact in-memory RDF store, where the RDF graph is first dictionary encoded and sorted by predicate. Two binary matrices are created: an $s \times n$ matrix called ST and an $o \times n$ matrix called OT, where s , o and n are the number of unique subjects, objects and triples respectively. The ST/OT matrix contains a 1 in index i, j if the subject/object of the j^{th} triple corresponds to the i^{th} term (or a 0 otherwise). Both matrices are indexed with k^2 -trees, while a bit string of length n encodes the predicate boundaries with a 1, i.e., in which columns of the matrix (denoting triples sorted or grouped by predicate) the predicate changes. These indexes are sufficient to cover all

eight possible triple patterns. Further compression techniques can be applied to the leaf matrices of the k^2 .tree in order to trade space for time. The authors mention that joins can be supported in a similar fashion as used for RDFCSA and k^2 -triples.

Tentris [24] (2020) is an in-memory RDF store wherein an RDF graph is viewed as a one-hot encoded 3-order tensor (equivalent to the 3-dimensional array used in BitMat [14]), which in turn is viewed as a trie of three levels for S, P and O. However, rather than storing tries for all permutations, a *hypertrie* is used with three levels. The leaves in the third level correspond to all possible combinations of two constants in a triple: for each triple (s, p, o) , there exists a leaf that maps $(s, p, ?)$ to the set of all objects that replace $?$ in the graph (including o), with analogous leaves for $(?, p, o)$ and $(s, ?, o)$. The second level refers to single constants, where three nodes are present for $(s, ?, ?)$, $(?, p, ?)$, $(?, ?, o)$ such that $(s, ?, ?)$ will map to the leaves for $(s, p, ?)$ and $(s, ?, o)$, and likewise for $(?, p, ?)$ and $(?, ?, o)$. Finally, the top level – the root, representing zero constants – simply maps to all the second level elements. Basic graph patterns (with projection) are translated into tensor operations that can be evaluated on the hypertrie using a worst-case optimal join algorithm.

A.2 Distributed RDF Engines

We now survey distributed RDF stores. Table 5 summarizes the surveyed systems and the techniques they use. We further indicate the type of underlying storage used, where italicized entries refer to local stores. Some systems that appear in the following may have appeared before in the local discussion if they are commonly deployed in both settings.

YARS2 [78] (2007) is an RDF store based on similar principles to YARS (see local stores) but for a distributed environment. The index manager in YARS2 uses three indexes namely a quad index, keyword index, and a join index for evaluating queries. The quad indexes cover six permutations of quads. The keyword index is used for keyword lookups. The join indexes help speeding up query execution for common joins. The core index on quads is based on hashing the first element of the permutation, except in the case of predicates (e.g., for a POGS permutation), where hashing creates skew and leads to imbalance, and where random distribution is thus used. Indexed nested loop joins are used, with triple patterns being evaluated on one machine where possible (based on hashing), or otherwise on all machines in parallel (e.g., for constant predicates

Table 5: Categorization of distributed RDF Engines.

Storage: T = Triple Table, Q = Quad Table, V = Vertical Partitioning, P = Property table, G = Graph-based, M = Miscellaneous
Indexing: T = Triple, Q = Quad, E = Entity, P = Property, N = Path/Navigational, J = Join, S = Structural, M = Miscellaneous
Join P.: P = Pairwise, M = Multiway, W = Worst case optimal
Query P.: N = Paths/Navigational, R = Relational, Q = Query rewriting
Partitioning: S = Statement (Triple/Quad)-based, G = Graph-based, Q = Query-based, R = Replication

Engine	Year	Storage						Indexing								Join P.			Query P.			Partitioning				
		T	Q	V	P	G	M	T	Q	E	P	N	J	S	M	P	M	W	N	R	Q	S	G	Q	R	Store
YARS2 [78]	2007	✓	✓					✓								✓			✓			✓				Custom
Clustered TDB [144]	2008	✓						✓								✓						✓				Jena TDB
Virtuoso EE [54]	2008	✓	✓					✓							✓				✓	✓	✓				✓	Custom
4store [75]	2009	✓	✓					✓								✓			✓	✓	✓	✓				Custom
Blazegraph	2009	✓	✓					✓	✓							✓	✓		✓	✓	✓	✓			✓	Custom
SHARD [162]	2009					✓										✓	✓		✓	✓	✓	✓				HDFS
Allegrograph	2010		✓					✓											✓	✓	✓	✓			✓	Custom
GraphDB [105, 26]	2010	✓	✓					✓	✓							✓			✓	✓	✓	✓			✓	Custom
AnzoGraph	2011		✓													✓			✓	✓	✓	✓				Custom
CumulusRDF [110]	2011	✓	✓					✓	✓						✓	✓			✓		✓	✓				Cassandra
H-RDF-3X [86]	2011	✓						✓								✓	✓		✓	✓	✓	✓			✓	RDF-3X
PigSPARQL [169]	2011	✓														✓	✓		✓	✓	✓	✓				HDFS
Rapid+ [160]	2011			✓												✓	✓		✓	✓	✓	✓				HDFS
AMADA [13]	2012	✓													✓	✓			✓		✓	✓				SimpleDB
H2RDF(+) [148]	2012	✓						✓								✓	✓		✓		✓	✓				HBase
Jena-HBase [101]	2012	✓		✓				✓								✓	✓		✓		✓	✓				HBase
Rya [157]	2012	✓						✓								✓	✓		✓		✓	✓				Accumulo
Sedge [215]	2012					✓										✓						✓	✓	✓	✓	Pregel
chameleon-db [10]	2013					✓		✓				✓				✓			✓			✓	✓	✓	✓	Custom
D-SPARQ [134]	2013	✓						✓							✓	✓						✓			✓	MongoDB
EAGRE [221]	2013	✓				✓		✓		✓						✓	✓		✓		✓	✓				HDFS
MR-RDF [50]	2013	✓						✓								✓	✓		✓		✓	✓				RDF-3X
SHAPE [112]	2013	✓						✓								✓	✓		✓		✓	✓				RDF-3X
Trinity.RDF [219]	2013	✓						✓		✓						✓	✓		✓		✓	✓				Trinity
TripleRush [186]	2013	✓						✓								✓	✓				✓	✓				Signal/Collect
WARP [85]	2013	✓						✓								✓	✓							✓	✓	RDF-3X
AKZ14 [6]	2014	✓			✓			✓							✓	✓	✓		✓	✓	✓	✓				HBase
Partout [57]	2014	✓						✓								✓	✓		✓					✓		RDF-3X
P-LUOSDATE [63]	2014	✓						✓								✓	✓		✓	✓	✓	✓				HBase
Sempala [170]	2014				✓										✓	✓	✓		✓		✓	✓				HDFS
SemStore [206]	2014			✓		✓		✓								✓	✓					✓				TripleBit
SparkRDF [41]	2014			✓				✓				✓				✓	✓		✓		✓	✓				HDFS/Spark
TriAD [70]	2014	✓						✓					✓			✓	✓					✓	✓			Custom
CK15 [42]	2015			✓				✓				✓				✓	✓					✓		✓		Custom
CliqueSquare [59]	2015	✓						✓								✓	✓				✓	✓			✓	HDFS
DREAM [72]	2015	✓						✓								✓			✓						✓	RDF-3X
AdPart [73]	2016	✓						✓								✓					✓	✓			✓	Custom
DiploCloud [210]	2016					✓		✓		✓						✓			✓			✓	✓			Custom
gStore-D [151]	2016					✓		✓		✓						✓	✓		✓		✓	✓				gStore
JARS [159]	2016	✓						✓								✓	✓		✓		✓	✓				Custom
S2RDF [171]	2016			✓											✓	✓			✓	✓	✓	✓				HDFS
S2X [168]	2016	✓														✓	✓		✓	✓	✓	✓				Spark
SPARQLGX [60]	2016			✓												✓	✓		✓		✓	✓				HDFS
Wukong [178]	2016	✓						✓							✓	✓										DrTM-KV
Koral [92]	2017	✓						✓								✓					✓	✓				Custom
SANSA [114]	2017	✓		✓				✓								✓				✓		✓				HDFS
Spartex [4]	2017					✓		✓								✓			✓		✓	✓				GSP/Custom
Stylus [80]	2017					✓		✓		✓	✓					✓			✓	✓	✓	✓			✓	Trinity
Neptune	2018		✓					✓		✓						✓			✓	✓					✓	Custom
PRoST [47]	2018			✓	✓										✓	✓	✓			✓	✓	✓				HDFS
RDFox-D [156]	2018	✓						✓								✓							✓			RDFox
WORQ [123]	2018			✓								✓				✓	✓									Spark
DiStRDF [202]	2019	✓			✓										✓	✓			✓		✓	✓				HDFS
gStore-D2 [150]	2019					✓						✓				✓	✓					✓	✓			Custom
Leon [67]	2019	✓						✓								✓	✓					✓	✓			Custom
SPT+VP [79]	2019			✓	✓											✓	✓		✓	✓						Spark
StarMR [199]	2019					✓										✓	✓				✓	✓				HDFS
DISE [89]	2020	✓						✓								✓					✓	✓				Spark
DP2RPQ [200]	2020					✓													✓		✓	✓				Spark
Triag [136]	2020					✓						✓				✓	✓				✓	✓		✓		Spark
WISE [68]	2020	✓						✓			✓					✓	✓				✓	✓	✓	✓		Leon

or keyword searches). Dynamic programming is used for join reordering in order to optimize the query.

Clustered TDB [144] (2008) is a distributed RDF store based on Jena TDB storage (a local system). The system is based on a master-slave architecture where the master receives and processes queries, and slaves index parts of the graph and can perform joins. Hash-based partitioning is used to allocate dictionary-encoded triples to individual slaves based on each position of the triple; more specifically, distributed SPO, POS and OSP index permutations are partitioned based on S, P and O, respectively. An exception list is used for very frequent predicates, which are partitioned by PO instead of P. Index-nested loop joins are supported and used to evaluate SPARQL basic graph patterns.

Virtuoso EE [54] (2008) is a local RDF store whose enterprise edition also offers support for indexing over a cluster of machines. Recalling that Virtuoso stores RDF graphs as a quads table in a custom relational database, the most recent version of Virtuoso offers three options for each table: partitioned, replicated or local. Partitioning is based on partition columns specified by the administrator, which are used for hash-based partitioning; partitions can also be replicated, if specified. Replication copies the full table to each machine, which can be used for query-based partitioning, or to store a global schema that is frequently accessed by queries. Local tables are only accessible to the individual machine, and are typically used for local configuration.

4store [75] (2009) stores quads over a cluster of machines, where subject-based hash partitioning is used. Three types of indexes are used in 4Store namely R, M, and P indexes. The R index is a hash table that dictionary encodes and stores meta-data about individual RDF terms (called “resources”). The M index is a hash table that maps graph names (called “models”) to the corresponding triples in the named graph. The P Indexes consist of radix tries, with two for each predicate (similar to vertical partitioning): one for SOG order and another for OSG order. Joins are pushed, where possible, to individual machines. Join reordering uses cardinality estimations. SPARQL queries are supported.

Blazegraph [190] (2009), discussed previously as a local store, also features partitioning. In particular, key-range shards are supported that allow for partitioning B+tree indexes, potentially across multiple machines. An alternative replication cluster is supported that indexes the full RDF graph or SPARQL dataset on each

machine, allowing queries to be evaluated entirely on each machine without network communication.

SHARD [162] (2009) is a distributed, Hadoop-based RDF store. It stores an RDF graph in flat files on HDFS such that each line presents all the triples associated with a given subject resource of the RDF triple, which can be seen as an adjacency list. The graph is hash partitioned, so that every partition contains a distinct set of triples. As the focus is on batch processing of joins, rather than evaluating queries in real-time, there is no specific indexing employed in SHARD. Query execution is performed through MapReduce iterations: first, it collects the results for the subqueries, which are joined and finally filtered according to bound variables and to remove redundant (duplicate) results.

AllegroGraph (2010), discussed previously as a local store, also features a distributed version, where data are horizontally partitioned into shards, which are indexed locally on each machine per the local version. Alongside these shards, “knowledge bases” can be stored, consisting of triples that are often accessed by all shards (e.g., schema or other high level data), such that queries can be evaluated (in a federated manner) over one shard, and potentially several knowledge bases.

GraphDB [105,26] (2010), included as a local store, also offers the ability to store RDF graphs on a cluster of machines. Specifically, the enterprise edition offers a master-slave architecture, where each cluster has at least one master node that manages one or more worker nodes, each replicating the full database copy, thus allowing for queries to be evaluated in full on any machine. Updates are coordinated through the master.

*AnzoGraph*¹³ (2011) is an in-memory, massively parallel processing (MPP) RDF store based on a master-slave architecture. The system indexes named graphs, where partitioning and replication are also organized by named graphs. By default, all triples involving a particular term are added into a named graph for that term. A dictionary is provided to map terms to named graphs. The query is issued at a master node, which features a query planner that decides the type of join (hash or merge joins are supported) or aggregation needed. Individual operations are then processed over the slaves in parallel, generating a stream of intermediate results that are combined on the master.

¹³ <https://docs.cambridgesemantics.com/anzograph/userdoc/features.htm>

CumulusRDF [110] (2011) works on top of Apache Cassandra: a distributed key-value store with support for tabular data. Three triple permutations – SPO, POS, OSP – and one quad permutation – GSPO – are considered. A natural idea would be to index the first element as the row key (e.g., S for SPO), the second (e.g., P) as the column key, and the third (e.g., O) as the cell value, but this would not work in multi-valued cases as columns are unique per row. Two other data storage layouts are thus proposed. Taking SPO, the “hierarchical layout” stores S as the row key (hashed and used for partitioning), P as the supercolumn key (sorted), O as the column key (sorted), with the cell left blank. An alternative that outperforms the hierarchical layout is a “flat layout”, where for SPO, S remains the row key, but PO is concatenated as the column key, and the cell is left blank. In the POS permutation, the P row key may create a massive row; hence PO is rather used as the row key, with P being indexed separately. Join and query processing is enabled through Sesame.

H-RDF-3X [86] (2011) is a Hadoop-based distributed RDF store that uses RDF-3X on a cluster of machines. A graph-based partitioning (using the METIS software package) is used to distribute triples among multiple worker nodes. It also employs a k -hop guarantee, which involves replicating nodes and edges that are k hops away from a given partition, thus increasing the locality of processing possible, and reducing communication costs. Local joins are optimized and evaluated on individual machines by RDF-3X, while joins across machines are evaluated using Hadoop. The use of Hadoop – which involves expensive coordination across machines, and heavy use of the disk – is minimized by leveraging the k -hop guarantee and other heuristics.

PigSPARQL [169] (2011) is a Hadoop-based distributed RDF store that uses a vertical partitioning strategy. Data are stored on HDFS without indexes, and thus the focus is on batch processing. SPARQL queries are translated into PigLatin: an SQL-inspired scripting language that can be compiled into Hadoop tasks by the Pig framework. Jena ARQ is used to parse SPARQL queries into an algebra tree, where optimizations for filters and selectivity-based join reordering are applied. The tree is traversed in a bottom-up manner to generate PigLatin expressions for every SPARQL algebra operator. The resulting PigLatin script is then translated to – and run as – MapReduce jobs on Hadoop.

Rapid+ [160] (2011) is a Hadoop-based system that uses a vertical partitioning strategy for storing RDF data. Without indexing, the system targets batch processing. Specifically, Pig is used to generate and access

tables under a vertical partitioning strategy. In order to translate SPARQL queries into PigLatin scripts, user-defined-functions are implemented that allow for optimizing common operations, such as loading and filtering in one step. Other optimizations include support for star joins using grouping, and a look-ahead heuristic that reduces and prepares intermediate results for operations that follow; both aim to reduce the number of Hadoop tasks needed to evaluate a query.

AMADA [13] (2012) is an RDF store designed for use on the Amazon Web Services (AWS) cloud infrastructure. Indexes for the RDF graph are built using Amazon SimpleDB: a key-value storage solution that supports a subset of SQL. SimpleDB offers several indexing strategies, where “attribute indexing” can be used to create three indexes for the three elements of a triple. In AMADA, a query is submitted to a query processing module running on EC2, which in turn evaluates triple patterns using the SimpleDB-based indexes.

H2RDF(+) [148,147] (2012) stores RDF graphs in the HBase distributed tabular NoSQL store. Three triple permutations (SPO, POS, and OSP) are created over the HBase tables in the form of key-value pairs. A join executor module creates the query plan, which decides between the execution of joins in a centralized (local) and distributed (Hadoop-based) manner. It further reorders joins according to selectivity statistics. Multiway (sort-)merge joins are implemented in Hadoop.

Jena-HBase [101] (2012) is a distributed RDF store using HBase as its back-end. Jena-HBase supports three basic storage layouts for RDF graphs in HBase namely “simple”: three triple tables, the first indexed and partitioned by S, the second by P, the third by O; “vertical partitioning”: two tables for each predicate, one indexed by S, the other by O; “indexed”: six triple tables covering all permutations of a triple. Hybrid layouts are also proposed that combine the basic layouts, and are shown to offer better query times at the cost of additional space. Jena is used to process joins and queries.

Rya [157] (2012) is a distributed RDF store that employs Accumulo – a key-value and tabular store – as its back-end. However, it can also use other NoSQL stores as its storage component. Rya stores three index permutations namely SPO, POS, and OSP. Query processing is based on RDF4J, with index-nested loop joins being evaluated in a MapReduce fashion. The count of the distinct subjects, predicates, and objects is maintained and used during join reordering and query optimization.

Sedge [215] (2012) is an RDF store based on Pregel: a distributed (vertex-centric) graph processing framework. Pregel typically assumes a strict partition of the nodes in a graph, where *Sedge* relaxes this assumption to permit nodes to coexist in multiple partitions. A complementary graph partitioning approach is proposed involving two graph partitionings, where the cross-partition edges of one are contained within a partition of the other, reducing cross-partition joins. Workload-aware query-based partitioning is also proposed, where commonly accessed partitions and frequently-queried cross-partition “hotspots” are replicated. The store is implemented over Pregel, where indexes are built to map partitions to their workloads and to their replicas, and to map nodes to their primary partitions.

chameleon-db [10] (2013) is a distributed RDF store using custom graph-based storage. Partitioning is graph-based and is informed by the queries processed, which may lead to dynamic repartitioning to optimize for the workload being observed. An incremental indexing technique – using a decision tree – is used to keep track of partitions relevant to queries. It also uses a hash-table to index the nodes in each partition, and a range-index to keep track of the minimum and maximum values for literals of each distinct predicate in each partition. The evaluation of basic graph patterns is delegated to a sub-graph matching algorithm over individual partitions, whose results are then combined in a query processor per the standard relational algebra. Optimizations involve rewriting rules that preserve the equivalence of the query but reduce intermediate results.

D-SPARQ [134] (2013) is a distributed RDF store built on MongoDB: a NoSQL store for JSON-like documents. *D-SPARQ* partitions the RDF graph by subject. Partial data replication is used whereby selected triples are replicated across partitions to increase parallelism when executing (sub-)queries. Indexes are built for SP and PO permutations. *D-SPARQ* optimizes multiway s-s (star) joins, taking advantage of the s-based partitioning; selectivity estimates are used to reorder joins.

EAGRE [221] (2013) stores RDF data on HDFS, where data are pre-processed using Hadoop to extract entities and their classes, thereafter applying graph-based data partitioning based on METIS. For each entity class, *EAGRE* adopts a space-filling curve technique (see Section 5.2): an in-memory index structure that is used to index high-dimensional data, and more specifically in this case, to decide where the data for a given entity should be stored. Joins are pushed to individual nodes where possible, with multiway joins between nodes being evaluated using Hadoop. A strategy similar to a distributed form of sideways-information-passing is used

to reduce network traffic, where nodes share information about the possible ranges of constants matching individual variables, filtering intermediate results outside those ranges before they are sent over the network.

MR-RDF [50] (2013) is a distributed RDF store that uses RDF-3X for local storage, and Hadoop for join processing. A partition is generated for each characteristic set, where a triple is added to the partition for the characteristic set of its subject. Given that this may give rise to a large number of partitions, similar characteristic sets are clustered together to form larger partitions corresponding to the number of machines available. The larger partition is then described by the union of the characteristic sets it contains, which can be used for matching star joins (with constant predicates) to partitions. Star joins are evaluated locally by RDF-3X, and their results are joined over Hadoop.

SHAPE [112] (2013) uses RDF-3X to store and index RDF triples on a distributed cluster of machines. Triples are partitioned using a semantic hash partitioning scheme that is based on the IRI prefix hierarchy: triples with the same subject or object prefixes are identified and are placed in the same partition. The intuition is that such triples are more likely to be queried together. A distributed query execution planner coordinates the intermediate results from different nodes, which are joined using Hadoop.

Trinity.RDF [219] (2013) is an RDF store implemented on top of Trinity: a distributed memory-based key-value storage system. A graph-based storage scheme is used, where an inward and outward adjacency list is indexed for each node. Hash-based partitioning is then applied on each node such that the adjacency lists for a given node can be retrieved from a single machine; however, nodes with a number of triples/edges exceeding a threshold may have their adjacency lists further partitioned. Aside from sorting adjacency lists, a global predicate index is also generated, covering the POS and PSO triple permutations. Queries are processed through graph exploration, with dynamic programming over cardinality estimates used to choose a query plan.

TripleRush [186] (2013) is based on the Signal/Collect framework [185]: a distributed graph processing framework. Within this framework, *TripleRush* considers an in-memory graph with three types of nodes. Triple nodes embed an RDF triple with its subject, predicate, and object. Index nodes embed a triple pattern. Query nodes

coordinate the query execution. The index graph is formed by index and triple nodes, which are linked based on matches. A query execution is initialized when a query node is added to a TripleRush graph. The query vertex emits a query particle (a message) which is routed by the Signal/Collect framework to index nodes for matching. Partitioning of triples and triple patterns is based on the order S, O, P, where the first constant in this order is used for hash-based partitioning. Later work explored workload-aware query-based partitioning methods [192].

WARP [85] (2013) uses RDF-3X to store triples in partitions among a cluster of machines. Like H-RDF-3X, graph-based partitioning is applied along with a replication strategy for k -hop guarantees. Unlike H-RDF-3X, WARP proposes a query-based, workload-aware partitioning, whereby the value of k is kept low, and selective replication is used to provide guarantees specifically with respect to the queries of the workload, reducing storage overheads. Sub-queries that can be evaluated on one node are identified and evaluated locally, with custom merge joins (rather than Hadoop, as in the case of H-RDF-3X) used across nodes. Joins are reordered to minimize the number of single-node subqueries.

AKZ14 [6] (2014) is a distributed RDF store based on the HBase tabular store. A property table storage scheme is implemented over HBase, which is built based on clustering entities with similar properties. A secondary triple table is used for multi-valued properties and (infrequent) properties that do not appear in the clusters. Property tables are used to solve subject-based star joins, with other joins being evaluated over Hadoop by translating SPARQL queries to Hive (an SQL-like language for Hadoop). Metadata for the HBase tables are stored in a relational database (MySQL).

Partout [57] (2014) is a distributed RDF store that uses RDF-3X for underlying storage on each machine. The RDF graph is partitioned using a workload-aware query-based partitioning technique, aiming to group together triples that are likely to be queried together. Each partition is indexed using standard RDF-3X indexing. The SPARQL query is issued to a query processing master, which uses RDF-3X to generate a suitable query plan according to a global statistics file. The local execution plan of RDF-3X is transformed into a distributed plan, which is then refined by a distributed cost model that assigns subqueries to partitions. This query plan is executed by slave machines in parallel, whose results are combined in the master.

P-LUPOSDATE [63] (2014) is a distributed RDF store based on HBase. Triples are distributed according to six triple permutations – partitioning on S, P, O, SP, SO, PO – enabling lookups for any triple pattern. In order to reduce network communication, Bloom filters are pre-computed for each individual variable of each triple pattern with at least one constant and one variable that produces some result; e.g., for SP, a Bloom filter is generated encoding the objects of each subject–predicate pair; for S, a Bloom filter is generated for each subject encoding its predicates, and optionally, another Bloom filter is generated for its objects. These Bloom filters are sent over the network in order to compute approximate semi-join reductions, i.e., to filter incompatible results before they are sent over the network. SPARQL (1.0) queries are evaluated by translating them to PigLatin, which are compiled into Hadoop jobs by Pig.

Sempala [170] (2014) stores RDF triples in a distributed setting, using the columnar Parquet format for HDFS that supports queries for specific columns of a given row (without having to read the full row). In this sense, Parquet is designed for supporting a single, wide (potentially sparse) table and thus Sempala uses a single “unified property table” for storing RDF triples with their original string values; multi-valued properties are stored using additional rows that correspond to a Cartesian product of all values for the properties of the entity. SPARQL queries are translated into SQL, which is executed over the unified property table using Apache Impala: a massively parallel processing (MPP) SQL engine that runs over data stored in HDFS.

SemStore [206] (2014) is a distributed RDF store with a master-slave architecture. A custom form of graph partitioning is used to localize the evaluation of subqueries of particular patterns – star, chain, tree, or cycle – that form the most frequent elements of basic graph patterns. A k -means partitioning algorithm is used to assign related instances of patterns to a particular machine, further increasing locality. The master creates a global bitmap index over the partitions and collects global cardinality-based statistics. Slave nodes use the TripleBit local RDF engine for storage, indexing and query processing. The master node then generates the query plan using dynamic programming and global cardinality statistics, pushing joins (subqueries) to individual slave nodes where possible.

SparkRDF [41] (2014) is a Spark-based RDF engine that distributes the graph into subgraphs using vertical partitioning, adding tables for classes as well as properties. SparkRDF then creates indexes over the class and property tables, and further indexes class–property, property–class, and class–property–class joins. These indexes are loaded into an in-memory data structure in Spark (a specialized RDD) that implements query processing functionalities such as joins, filters, etc. Class information is used to filter possible results for individual variables, where a greedy selectivity-based strategy is used for reordering joins. Joins themselves are evaluated in a MapReduce fashion.

TriAD [70] (2014) implements a distributed RDF store in main-memory based on a master–slave architecture. The master maintains a dictionary of terms, a graph summary that allows for pruning intermediate results, as well as global cardinality-based statistics that allow for query planning. The graph summary is a quotient graph based on graph partitioning (using METIS): each partition forms a supernode, and labeled edges between supernodes denote triples that connect nodes in different partitions; the graph summary is indexed in two permutations: PSO and POS. The triples for each partition are stored on a slave; triples connecting two partitions are stored on both slaves. Each slave indexes their subgraph in all six triple permutations. Given a basic graph pattern, the graph summary is used to identify relevant partitions, which are shared with the slaves and used to prune results; dynamic programming uses the global statistics to optimize the query plan. Alongside distributed hash and merge joins, an asynchronous join algorithm using message passing is implemented.

CK15 [42] (2015) is a distributed in-memory RDF store that combines two types of partitioning: triple-based partitioning and query-based partitioning. The graph is initially divided over the machines into equal-size chunks and dictionary-encoded in a distributed manner (using hash-based partitioning of terms). The encoded triples on each machine are then stored using a vertical partitioning scheme, where each table is indexed by hashing on subject, and on object, providing $P \rightarrow SO$, $PS \rightarrow O$ and $PO \rightarrow S$ lookups. Parallel hash joins are proposed. Secondary indexes are then used to cache intermediate results received from other machines while processing queries, such that they can be reused for future queries. These secondary indexes can also be used for computing semi-join reductions on individual machines, thus reducing network traffic.

CliqueSquare [59] (2015) is a Hadoop-based RDF engine used to store and process massive RDF graphs. It

stores RDF data in a vertical partitioning scheme using semantic hash partitioning, with the objective of enabling co-located or partitioned joins that can be evaluated in the map phase of the MapReduce paradigm. CliqueSquare also maintains three replicas for fast query processing and increased data locality. In order to evaluate SPARQL queries, CliqueSquare uses a clique-based algorithm, which works in an iterative way to identify cliques in a query–variable graphs and to collapse them by evaluating joins on the common variables of each clique. The process terminates when the query–variable graph consists of only one node.

DREAM [72] (2015) is a distributed RDF store that uses RDF-3X for its underlying storage and indexing. The entire RDF graph is replicated on every machine, with standard RDF-3X indexing and query processing being applied locally. However, only dictionary-encoded terms are communicated within the cluster. In the query execution phase, the SPARQL query is initially represented as a directed graph, which is divided into multiple subqueries to be evaluated by different machines. The results of subqueries are combined using hash joins.

AdPart [73] (2016) is a distributed in-memory RDF store following a master–slave architecture. The master initially performs a hash-based partitioning based on the subjects of triples. The slave stores the corresponding triples using an in-memory data structure. Within each slave, AdPart indexes triples by predicate, predicate–subject, and predicate–object. Each slave machine also maintains a replica index, which incrementally replicates data that are accessed by many queries; details of this replication are further indexed by the master machine. Query planning then tries to push joins locally to slaves (hash joins are used locally), falling back to distributed semi-joins where that is not possible. Join reordering then takes communication costs and cardinalities into account.

DiploCloud [210] (2016) is a distributed version of the local RDF store dipLODocus. The store follows a master–slave architecture, where slaves store “molecules” (see the previous discussion on dipLODocus). The master provides indexes for a dictionary, for the class hierarchy (used for inference), as well as an index that maps the individual values of properties selected by the administrator to their molecule. Each slave stores the molecule subgraphs, along with an index mapping nodes to molecules, and classes to nodes. Query processing pushes joins where possible to individual slaves; if intermediate results are few, the master combines results, or otherwise a distributed hash join is employed.

Molecules can be defined as a k -hop subgraph around the root node, based on input from an administrator, or based on a given workload of queries.

gStore-D [151] (2016) is a distributed RDF store that uses a variation of the local gStore RDF engine for local storage and indexing (it uses adjacency lists and vertex signatures, as discussed previously). A graph partitioning algorithm is applied over the RDF graph, with the subgraphs induced by each partition being assigned to individual machines. Edges that connect distinct partitions are indexed within the subgraphs of both partitions. Where possible, joins are then pushed to the individual subgraphs. Joins across subgraphs can be evaluated in a central (i.e., by a single master) or distributed (i.e., by several slaves) manner. Support for the relational features of SPARQL (1.0) is described.

JARS [159] (2016) is a distributed RDF store that combines triple-based and query-based partitioning. The graph is partitioned by hashing on subject, and hashing on object, constructing two distributed triple tables. The subject-hashed table is indexed on the POS, PSO, OSP and SPO permutations, while the object-hashed table is indexed on POS, PSO, SOP and OPS. Specifically, by hashing each triple on subject and object, the data for S-S, O-O and S-O are on one machine; the permutations then allow for such joins to be supported as merge joins on each machine. Basic graph patterns are then decomposed into subqueries answerable on a single machine, with a distributed hash join applied over the results. Jena ARQ is used to support SPARQL.

S2RDF [171] (2016) is a distributed RDF store based on HDFS (using the Parquet format). The specific storage scheme is based on an extended version of vertical partitioning that uses semi-join reductions (see Section 4.3). This scheme has a high space overhead, but ensures that only data useful for a particular (pairwise) join will be communicated over the network. In order to reduce the overhead, semi-join tables are not stored in cases where the selectivity of the join is high; in other words, semi-join tables are stored only when many triples are filtered by the semi-join (the authors propose a threshold of 0.25, meaning that at least 75% of the triples must be filtered by the semi-join for the table to be included). SPARQL queries are optimized with cardinality-based join reordering, and then translated into SQL and evaluated using Spark.

S2X [168] (2016) processes SPARQL queries over RDF graphs using GraphX: a distributed graph processing framework built on the top of Spark. The triples are

stored in-memory on different slave machines with Spark (RDDs), applying a hash-based partitioning on subject and objects (per GraphX's default partitioner). S2X does not maintain any custom indexing. For SPARQL query processing, graph pattern matching is combined with relational operators (implemented in the Spark API) to produce solution mappings.

SPARQLGX [60] (2016) stores RDF data on HDFS using a vertical partitioning scheme. Specifically, a separate file is created for each unique predicate in the RDF graph, with each file containing the subjects and objects of that triple. No indexes are provided, and thus the system is intended for running joins in batch-mode. SPARQL queries are first optimized by applying a greedy join reordering based on cardinality and selectivity statistics; the query plan is then translated into Scala code, which is then directly executed by Spark.

Wukong [178] (2016) stores RDF graphs in DrTM-KV: a distributed key-value store using "remote direct memory access" (RDMA), which enables machines to access the main memory of another machine in the same cluster while by-passing the remote CPU and OS kernel. Within this store, Wukong maintains three kinds of indexes: a node index that maps subjects or (non-class) objects to their corresponding triples; a predicate index, which returns all subjects and objects of triples with a given predicate; and a type index, which returns the class(es) to which a node belongs. Hash-based partitioning is used for the node index, while predicate and type indexes are split and replicated to improve balancing. A graph-traversal mechanism is used to evaluate basic graph patterns, where solutions are incrementally extended or pruned. For queries involving fewer data, the data are fetched from each machine on the cluster and joined centrally; for queries involving more data, subqueries are pushed in parallel to the individual machines. A work-stealing mechanism is employed to provide better load balancing while processing queries.

Koral [92] (2017) is a distributed RDF store based on a modular master-slave architecture that supports various options for each component of the system. Among these alternatives, various triple-based and graph-based partitioning schemes are supported. In order to evaluate basic graph patterns, joins are processed in an analogous way to TriAD, using asynchronous execution, which makes the join processing strategy independent of the partitioning chosen. The overall focus of the system is to be able to quickly evaluate different alternatives for individual components – particularly partitioning strategies – in a distributed RDF store.

SANSA [114] (2017) is a Spark-based distributed RDF store. RDF data are stored on HDFS, where triple-based partitioning – such as predicate-based vertical partitioning – is employed. Queries are then transformed into Spark (or Flink) programs, using Sparklify [180]: a query engine for SPARQL-to-SQL translations, which can, in turn, be run over Spark. SANSA is part of a larger stack that supports RDF-based inferencing and machine learning in a distributed environment.

Spartex [4] (2017) is a distributed RDF store with analytical capabilities. An extension of SPARQL queries is proposed with user-defined procedures for analytics (e.g., PageRank), among other features. The system is built on top of GPS: an open-source implementation of Pregel's distributed, vertex-centric graph processing framework. A master-slave architecture is employed. The master is responsible for query planning and manages global statistics. The RDF graph is partitioned among its slaves; namely each (subject/object) node and its incident edges (triples) is assigned to a slave. Each slave stores and indexes its subgraph in-memory using PS and PO permutations. Basic graph patterns are then evaluated using graph traversals in GSP, with nodes (vertexes) sharing intermediate results as messages, which are joined with local data. Optimizations are based on minimizing duplicate traversals involving (non-Eulerian) cycles, and cardinality estimates.

Stylus [80] (2017) is a distributed RDF store that uses Trinity: a graph engine based on an in-memory key-value store. Terms of the RDF graph are dictionary encoded. Each subject and object node is associated with a dictionary identifier and its characteristic set. A sorted adjacency list (for inward and outward edges) is then stored for each node that also encodes an identifier for the characteristic set of the node. Schema-level indexes for characteristic sets are replicated on each machine. Hash-based partitioning is employed on the data level. Indexes are used to efficiently find characteristic sets that contain a given set of properties, as well as to evaluate common triple patterns. Given a basic graph pattern, the characteristic sets are used to prune intermediate results on star joins, where candidates are kept for each variable node in the query. Cardinality-based join reordering is applied. Relational features are supported for SPARQL (1.0) queries.

Neptune [21] (2018) is an RDF store that is hosted as a service on Amazon's S3 cloud storage infrastructure. Neptune stores SPARQL datasets in the form of quads with three index permutations: SPOG, POGS and GSPO; this is sufficient to cover 9 out of 16 possible quad patterns. Neptune makes use of cardinality estimations and

static analysis to rewrite queries. Partitioning is not supported, where Neptune rather offers up to 16 replicas of the full graph to increase query throughput; a primary replica is nominated to receive and coordinate updates. Graphs in Neptune can be queried (and processed) through the SPARQL 1.1, Apache TinkerPop and Gremlin languages.

ProST [47] (2018) is a distributed RDF store based on HDFS storage and Spark query processing. The storage scheme uses a combination of vertical partitioning and property tables that aims to leverage the strengths and minimize the weaknesses of both schemes. Like Sem-pala, the property table is stored in the column-wise Parquet format; multi-valued properties are supported by adding lists of values. Star joins on a common subject variable are evaluated on the property table, while other patterns and joins are addressed with the vertical partitioning tables. Selectivity-based heuristics are used to reorder joins. Queries are then rewritten into SQL for execution with Spark.

RDFox-D [156] (2018) is a distributed in-memory RDF store based on RDFox that uses index nested loops joins in the distributed setting. A global index is built mapping the terms of the RDF graph to the partitions it appears in. The graph is partitioned by a weighted graph-based partitioning scheme, where nodes are weighted by the number of triples they appear in as subject. The partitioning minimizes cross-partition edges while balancing the sum of the node weights in each partition. Triples with a subject in the same partition are sent to the same machine; the weights used for partitioning then help to ensure more even balancing. Joins are evaluated in a pairwise manner, where each machine extends solutions asynchronously, without central coordination, based on its partition of the graph; it then sends the extended partial solution to the machines that can potentially extend it further (based on the global index). Termination occurs when all partial solutions have been forwarded. Various optimizations are discussed. Joins are reordered based on cardinalities.

WORQ [123] (2018) is a distributed RDF store that uses a workload-aware approach to partition data. In order to reduce the number of intermediate results, Bloom filters are used to index the constants matching the variable of a given triple pattern, which are shared and used to filter results for that variable elsewhere. Bloom filters provide an approximate membership function (i.e., they may yield false positives), and thus a distributed join algorithm must be applied over

the (reduced) intermediate results. Further Bloom filters can be computed for multiway joins, analogous to an approximate form of semi-join reduction (as used by S2RDF). These reductions can be cached for later re-use, where they are partitioned across machines based on the join element. WORQ is implemented over Spark.

DiStRDF [202] (2019) is a massively parallel processing (MPP) RDF store based on Spark with support for spatio-temporal queries. A special dictionary-encoding mechanism is used where the identifier concatenates a bit-string for spatial information, a bit-string for temporal information, and a final bit-string to ensure that the overall identifier is unique. Thus spatial and temporal processing can be applied directly over the identifiers. Storage based on both on a triple table and property tables is supported, where range-based partitioning is applied to the triples (based on the spatio-temporal information). Data is stored on HDFS in CSV or Parquet formats. Query processing is implemented in Spark. Distributed hash joins and sort-merge joins are supported; selections and projections are also supported. Three types of query plans are proposed that apply RDF-based selections, spatio-temporal selections and joins in different orders.

gStore-D2 [150] (2019) is a distributed RDF store that proposes various forms of workload-aware query-based graph partitioning methods. Frequently accessed (sub-graph) patterns are mined from the workload, where all subjects and objects are mapped to variables. Sub-graphs that instantiate these patterns are assigned DFS codes that are indexed as a tree, and associated with various meta-data, including identifiers for queries that use the pattern, cardinality estimations, partition identifiers, etc. Three partitioning methods are based on these patterns, with partitions stored locally in gStore. “Vertical partitioning” indexes all instances of a given pattern on the same machine. “Horizontal partitioning” distributes instances of the same pattern across various machines based on its constants. “Mixed partitioning” combines the two. Basic graph patterns are then decomposed into frequent sub-patterns, over which various types of joins can be applied; the join order and algorithms are selected to reduce communication costs.

Leon [67] (2019) is an in-memory distributed RDF store based on a master-slave architecture. Triples are partitioned based on the characteristic set of their subject; the characteristic sets are ordered in terms of the number of triples they induce, and assigned to machines with the goal of keeping a good balance. Indexes (similar to those of Stylus [80]) are built, including a bidirectional index between subjects and their characteristic

sets, an index to find characteristic sets that contain a given set of properties, and indexes to evaluate certain triple patterns. A multi-query optimization technique is implemented where, given a workload (a set) of queries, the method searches for an effective way to evaluate and share the results for common subqueries – in this case, based on characteristic sets – across queries.

SPT+VP [79] (2019) is a distributed RDF store based on the principle of combining two partitioning techniques. First, a modified property table scheme is used for storage, where one table is maintained with a column for subject and all properties in the RDF graph; instead of storing multi-valued properties in multiple rows, as in Sempala’s unified property table, such values are stored as lists nested in the given row. The property table is then split (vertically) into multiple tables, similar to a clustering-based definition of a property table, but where a subject may appear in multiple tables. This “subset property table” approach is combined, secondly, with vertical partitioning storage. Given a SPARQL query, joins are reordered based on global statistics, with the property tables used for s-s joins and vertical partitioning used for other joins. The query is then translated into Spark SQL for execution.

StarMR [199] (2019) is a distributed RDF store that centers around optimizations for star joins. A graph-based storage scheme is employed, where for each node in the graph, its outward edges are represented in an adjacency list; this then supports efficient evaluation for s-s star joins. No indexing is provided, where the system targets batch-based (e.g., analytical) processing. A basic graph pattern is then decomposed into (star-shaped) sub-patterns, which are evaluated and joined. Hadoop is then used to join the results of these individual sub-patterns. Optimizations include the use of characteristic sets to help filter results, and the postponement of Cartesian products, which are used to produce the partial solutions for star joins including the non-join variables; these partial solutions are not needed if the corresponding join value is filtered elsewhere.

DISE [89] (2020) is an in-memory, distributed RDF store that conceptualizes an RDF graph as a 3-dimensional binary tensor, similar to local approaches such as Bit-Mat; however, physical representation and storage is based on dictionary encoded triples. Partitioning is based on slicing the tensor, which is equivalent to a triple-based partitioning. Joins are evaluated starting with the triple pattern with the fewest variables. SPARQL queries are supported through the Jena (ARQ) query library and evaluated using Spark.

DP2RPQ [200] (2020) is an RDF store based on a distributed graph processing framework with support for regular path queries (RPQs), which form the core of SPARQL's property paths. Unlike the standard RPQ semantics, the evaluation returns the "provenance" of the path, defined to be the subgraph induced by matching paths. Automata are used to represent the states and the potential transitions of paths while evaluating the RPQ, and are thus used to guide a navigation-based evaluation of the RPQ implemented by passing messages between nodes in the framework. Optimizations include methods to filter nodes and edges that cannot participate in the solutions to the RPQ, compression techniques on messages, as well as techniques to combine multiple messages into one. The system is implemented on top of Spark's GraphX framework.

Triag [136] (2020) is a distributed RDF store that optimizes for triangle-based (sub)-patterns in queries. Two types of triangular RDF subgraphs are extracted using Spark: cyclic ones (e.g., $(a, p, b), (b, q, c), (c, r, a)$) and (directed) acyclic ones (e.g., $(a, p, b), (b, q, c), (a, r, c)$). The predicates of these triangular subgraphs are extracted, ordered, hashed, and indexed in a distributed hash table using the predicate-based hash as key and the three nodes (e.g., a, b, c) as value. An encoding is used to ensure that the ordering of predicates is canonical for the pattern (assuming that nodes are variables) and that the subgraph can be reconstructed from the node ordering, even when some predicates are the same. Parallel versions of hash joins and nested loop joins are supported, where triangular subqueries can be pushed to the custom index. Queries are executed over Spark. Support for inferencing is also described.

WISE [68] (2020) is a distributed RDF store that features workload-aware query-based partitioning. The system follows a master-slave architecture. As queries arrive at the master, aside from being processed, they are analyzed in terms of workload: common sub-patterns are extracted from a generalized version of the queries where constant subject and object nodes are first converted to variables. Query-based partitioning is then applied so that common sub-patterns can be pushed to individual machines. Partitioning is dynamic, and may change as queries are received. A cost model is thus defined for the dynamic partitioning, taking into account the benefits of the change in partitioning, the cost of migrating data, and potential load imbalances caused by partition sizes; a greedy algorithm is then used to decide on which migrations to apply. The system uses Leon – an in-memory distributed RDF store discussed previously – for underlying storage and indexing.

A.3 Trends

We remark on some general trends based on the previous survey of local and distributed systems.

In terms of local systems, earlier approaches were based on underlying relational stores given that their implementations were already mature when interest began to coalesce around developing RDF stores. Thus, many of these earlier stores could be differentiated in terms of the relational schema (triple table, vertical partitioning, property tables, etc.) used to represent and encode RDF graphs. Systems that came later tended to rather build custom native storage solutions, optimizing for specific characteristics of RDF in terms of its graph structure, its fixed arity, etc.; relating to the fixed arity, for example, native stores began to develop complete indexes, by default, that would allow efficient lookups for any triple pattern possible. Also, many engines began to optimize for star-joins, which are often used to reconstruct n -ary relations from RDF graphs. Engines would soon start to explore graph-inspired storage and indexing techniques, including structural indexes, compressed adjacency lists, etc. A more recent trend – likely following developments in terms of hardware – has been an increased focus on in-memory stores using compact representations. Another recent development has been the application of worst-case optimal join algorithms for evaluating basic graph patterns, replacing traditional pairwise join techniques from the database literature.

With respect to distributed RDF stores, in line with an increased demand for managing RDF graphs at very large scale, proposals began to emerge around 2007 regarding effective ways to store, index and query RDF over a cluster of machines.¹⁴ Initial proposals were based on existing native stores, which were extended with triple/quad-based partitioning and distributed join processing techniques to exploit a cluster of machines. A second trend began to leverage the maturation and popularity of "Big Data" platforms, including distributed processing frameworks like Hadoop and later Spark, and distributed NoSQL stores like Cassandra, HBase, MongoDB, etc., in order to build distributed RDF stores. During this time, graph-based and later query-based partitioning methods began to emerge. Like in the local case – and following a more general trend in Big Data towards exploiting the main memory of machines, including Spark, Trinity, etc. – more and more in-memory distributed RDF stores began to emerge. Another trend was to explore the use of distributed graph processing

¹⁴ We highlight that decentralized proposals for managing RDF graphs existed before this, including federated systems, P2P systems, etc., but are not considered in-scope here.

frameworks – that offer a vertex-based computation and messaging paradigm – for evaluating queries over RDF.

While proposed solutions have clearly been maturing down through the years, and much attention has been given to evaluating basic graph patterns over RDF, some aspects of SPARQL query processing have not gained much attention. Most stores surveyed manage triples rather than quads, meaning that named graphs are often overlooked. A key feature of SPARQL – and of graph query languages in general – is the ability to query paths of arbitrary length, where optimizing property paths in SPARQL has not received much attention, particularly in the distributed setting. Many works also focus on a WORM (write once, read many) scenario, with relatively little attention paid (with some exceptions) to managing dynamic RDF graphs.

A final aspect that is perhaps not well-understood is the trade-off that exist between different proposals, what precisely are their differences on a technical level (e.g., between relational- and graph-based conceptualizations), and which techniques perform better or worse in which types of settings. In this regard, a number of benchmarks have emerged to try to compare RDF stores in terms of performance; we discuss these next.

B SPARQL Benchmarks for RDF Stores

We now discuss a variety of SPARQL benchmarks for RDF stores. We speak specifically of SPARQL benchmarks since benchmarks for querying RDF either came after the standardization of SPARQL (and thus were formulated in terms of SPARQL), or they were later converted to SPARQL for modern use. The discussion herein follows that of Saleem et al. [167], who analyze different benchmarks from different perspectives. We first discuss the general design principles for benchmarks, and then survey specific benchmarks.

B.1 SPARQL Benchmark Design

SPARQL query benchmarks comprise of three elements: RDF graphs (or datasets), SPARQL queries, and performance measures. We first discuss some design considerations regarding each of these elements.

Datasets The RDF graphs and datasets proposed for use in SPARQL benchmarks are of two types: *real-world* and *synthetic*. Both have strengths and weaknesses.

Real-world graphs reflect the types of graphs that one wishes to query in practice. Graphs such as DBpedia, Wikidata, YAGO, etc., tend to be highly complex and diverse; for example, they can contain hun-

dreds, thousands or tens of thousands of properties and classes. Presenting query performance over real-world graphs is thus a relevant test of how a store will perform over RDF graphs found in practice. Certain benchmarks may also include a number of real-world graphs for the purposes of distributed, federated or even decentralized (web-based) querying [172].

Synthetic graphs are produced using specific generators that are typically parameterized, such that graphs can be produced at different scales, or with different graph-theoretic properties. Thus synthetic graphs can be used to test performance at scales exceeding real-world graphs, or to understand how particular graph-theoretic properties (e.g., number of properties, distributions of degrees, cyclicity, etc.) affect performance. Synthetic graphs can also be constructed in order to emulate certain properties of real-world graphs [51].

A number of measures have been proposed in order to understand different properties of benchmark graphs. Obvious ones include basic statistics, such as number of nodes, number of triples, number of properties and classes, node degrees, etc. [165,51]. Other (less obvious) proposals of measures include *structuredness* [51], which measures the degree to which entities of the same class tend to have similar characteristic sets; *relationship specialty* [158], which indicates the degree to which the multiplicity of individual properties varies for different nodes, etc. Observations indicate that the real-world and synthetic graphs that have been used in benchmarks tend to vary on such measures, with more uniformity seen in synthetic graphs [51,158,167]. This may affect performance in different ways; e.g., property tables will work better over graphs with higher structuredness and (arguably) lower relationship specialty.

SPARQL Queries The second important element of the benchmark is the queries proposed. There are four ways in which the queries for a benchmark may be defined:

- *Manually-generated*: The benchmark designer may manually craft queries against the RDF graph, trying to balance certain criteria such as query features, complexity, diversity, number of results, etc.
- *Induced from the graph*: The queries may be induced from the RDF graph by extracting sub-graphs (e.g., using some variation on random walks), with constants in the sub-graphs replaced by variables to generate basic graph patterns.
- *Extracted from logs*: The queries to be used may be extracted from real-world SPARQL logs reflecting realistic workloads; since logs may contain millions of queries, a selection process is often needed to identify an interesting subset of queries in the log.

Aside from concrete queries, benchmarks may also define query templates, which are typically queries where a subset of variables are marked as placeholders. These placeholders are replaced by constants in the data, typically so that the resulting partially-evaluated query still returns results over the RDF graph. In this way, each template may yield multiple concrete queries for use in the benchmark, thus smoothing variance for performance that may occur for individual queries.

Queries can vary in terms of the language considered (SPARQL 1.0 vs. SPARQL 1.1) and the algebraic features used (e.g., projection, filters, paths, distinct, etc.), but also in terms of various measures of the complexity and diversity of the queries – and in particular, the basic graph patterns – considered. Some basic measures to characterize the complexity and diversity of queries in a benchmark include the number of queries using different features, measures for the complexity of the graph patterns considered (e.g., number of triple patterns, number of variables, number of joins variables, number of cyclic queries, mean degree of variables, etc.), etc. Calculating such measures across the queries of the benchmark, a high-level *diversity score* can be computed for a set of queries [167], based on the average coefficient of variation (dividing the mean by the standard deviation) across the measures.

Performance Measures The third element of a benchmark is the performance measures used. Some benchmarks may be provided without a recommended set of measures, but at the moment in which a benchmark is run, the measures to be used must be selected. Such measures can be divided into four categories [167]:

- *Query Processing Related*: The most important dimension relating to query processing relates to runtimes. Since a benchmark usually contains many queries, reporting the runtime for each and every query is often too fine-grained. Combined results can rather be presented using measures like Query Mix per Hour (QMpH), Queries per Second (QpS), or measures over the distributions of runtimes (max, mean, percentile values, standard deviation, etc.). Other statistics like the number of intermediate results generated, disk/memory reads, resource usage, etc., can be used to understand lower-level performance issues during query processing load [173].
- *Data Storage Related*: This category includes measures like data loading time, storage space, index sizes, etc. Often there is a space–time trade-off inherent in different approaches, where more aggressive indexing can help to improve query runtimes but at the cost of space and more expensive updates. Hence these measures help to contextualize query-processing related measures.

- *Result Related*: Some systems may produce partial results for a query based on fixed thresholds or timeouts. Thus an important consideration for a fair comparison between two RDF engines relates to the results produced in terms of correctness and completeness. This can often be approximately captured in terms of the number of results returned, the number of queries returning empty results (due to timeouts), the recall of individual queries, etc.
- *Update Related*: In real-world scenarios, queries must often be executed while the underlying data are being updated in parallel. While the previous categories consider a read-only scenario, benchmarks may also record measures while updates are being executed [53,44]. Measures may include the number of insertions or deletions per second, the number of read/write transactions processed, etc.

Often a mix of complementary measures will be presented in order to summarize different aspects of the performance of the tested systems.

B.2 Synthetic Benchmarks

We now briefly survey the SPARQL benchmarks that have been proposed and used in the literature, and that are available for download and use. We start with benchmarks based on synthetic data.

LUBM (Lehigh) [69] (2005) generates synthetic RDF graphs that describe universities, including students, courses, professors, etc. The number of universities described by the graph is a parameter that can be changed to increase scale. The benchmark includes 14 hand-crafted queries. LUBM further includes an OWL ontology to benchmark reasoning capabilities, though often the benchmark is run without reasoning.

BSBM (Berlin) [28] (2009) is based on an e-commerce use-case describing entities in eight classes relating to products. The number of products can be varied to produce RDF graphs of different scales. A total of 12 query templates are defined with a mix of SPARQL features. The benchmark is also given in SQL format, allowing to compare RDF stores with RDBMS engines.

SP²Bench [173] (2009) creates synthetic RDF graphs that emulate an RDF version of the DBLP bibliographic database. Various distributions and parameters from the DBLP data are extracted and defined in the generator. A total of 17 queries are then defined for the benchmark in both SPARQL and SQL formats.

BowlognaBench [49] (2012) generates synthetic RDF graphs inspired by the Bologna process of reform for European universities. The dataset describes entities such as students, professors, theses, degrees, etc. A total of 13 queries are defined that are useful to derive analytics for the reform process.

WatDiv [9] (2014) provides a data generator that produces synthetic RDF graphs with an adjustable value of structuredness, and a query template generator that generates a specified number of query templates according to specified constraints. The overall goal is to be able to generate diverse graphs and queries.

LDBC-SNB [53] (2015) is a benchmark based on synthetically generated social networking graphs. Three workloads are defined: *interactive* considers both queries and updates in parallel; *business intelligence* considers analytics that touch a large percentage of the graph; *algorithms* considers the application of graph algorithms.

TrainBench [188] (2018) is a synthetic benchmark inspired by the use-case of validating a railway network model. The graph describes entities such as trains, switches, routes, sensors, and their relations. Six queries are defined that reflect validation constraints. TrainBench is expressed in a number of data models and query languages, including RDF/SPARQL and SQL.

B.3 Real-World Benchmarks

Next we survey benchmarks that are based on real-world datasets and/or queries from real-world logs.

DBPSB (DBpedia) [133] (2011) clusters queries from the DBpedia logs, generating 25 query templates representative of the common queries found within. These queries can then be evaluated over DBpedia, where a dataset of 153 million triples is used for testing, though smaller samples are also provided.

FishMark [18] (2012) is based on the FishBase dataset and is provided in RDF and SQL formats. The full RDF graph uses 1.38 billion triples, but a smaller graph of 20 million triples is used for testing. In total, 22 queries from a log of real-world (SQL) queries are converted to SPARQL for testing.

BioBenchmark [209] (2014) is based on queries over five real-world RDF graphs relating to bioinformatics – Allie, Cell, DDBJ, PDBJ and UniProt – with the largest dataset (DDBJ) containing 8 billion triples. A total of 48 queries are defined for the five datasets based on queries generated by real-world applications.

FEASIBLE [166] (2015) is a framework for generating SPARQL benchmarks from real-world query logs based on clustering and feature selection techniques. The framework is applied to DBpedia and Semantic Web Dog Food (SWDF) query logs and used to extract 15–175 benchmark queries from each log. The DBpedia and SWDF datasets used contain 232 million and 295 thousand triples, respectively.

WGPPB [82] (2019) is a benchmark of basic graph patterns over Wikidata. The queries are based on 17 abstract patterns, corresponding to binary joins, paths, stars, triangles, squares, etc. The benchmark contains 850 queries, with 50 instances of each abstract pattern mined from Wikidata using guided random walks. Two Wikidata graphs are given: a smaller one with 81 million triples, and a larger one with 958 million triples.

B.4 Benchmark Comparison and Results

For a quantitative comparison of (most of) the benchmarks mentioned here, we refer to the work by Saleem et al. [167], which provides a breakdown of various measures for SPARQL benchmarks. For benchmarks with results comparing different RDF stores, we refer to the discussion for (italicizing non-RDF/SPARQL engines):

- BSBM [28] (2009) including results for Jena, RDF4J, Virtuoso and *MySQL*;
- SP²Bench [173] (2009) including results for Kowari, Jena, RDF4J, Redland and Virtuoso.
- DBPSB [133] (2011) including results for GraphDB, Jena, RDF4J and Virtuoso;
- BowlognaBench [49] (2012) including results for 4store, dipLODocus, RDF-3X and Virtuoso.
- FishMark [18] (2012) including results for Virtuoso, *MySQL* and *Quest*;
- BioBench [209] (2014) including results for 4store, Blazegraph, GraphDB, Kowari and Virtuoso.
- WatDiv [9] (2014) including results for 4store, gStore, RDF-3X, Virtuoso and *MonetDB*;
- FEASIBLE [166] (2015) including results for GraphDB, Jena, RDF4J and Virtuoso;
- LDBC-SB [53] (2015), including results for SparkSee and Virtuoso;
- TrainBench [188] (2018) including results for Jena, RDF4J, *Neo4j* and *SQLite*, among others.

For a performance comparison of eleven distributed RDF stores (SHARD, H2RDF+, CliqueSquare, S2X, S2RDF, AdPart, TriAD, H-RDF-3x, SHAPE, gStore-D and DREAM) and two local RDF stores (gStore and RDF-3X) over various benchmarks (including LUBM and WatDiv), we refer to Abdelaziz et al. [3].