

# Compressed Indexes for Fast Search of Semantic Data

RAFFAELE PEREGO, ISTI-CNR  
GIULIO ERMANNANO PIBIRI, ISTI-CNR  
ROSSANO VENTURINI, University of Pisa

The sheer increase in volume of RDF data demands efficient solutions for the *triple indexing problem*, that is devising a compressed data structure to compactly represent RDF triples by guaranteeing, at the same time, fast pattern matching operations. This problem lies at the heart of delivering good practical performance for the resolution of complex SPARQL queries on large RDF datasets.

In this work, we propose a trie-based index layout to solve the problem and introduce two novel techniques to reduce its space of representation for improved effectiveness. The extensive experimental analysis conducted over a wide range of publicly available real-world datasets, reveals that our best space/time trade-off configuration substantially outperforms existing solutions at the state-of-the-art, by taking  $30 \div 60\%$  less space and speeding up query execution by a factor of  $2 \div 81\times$ .

## 1 INTRODUCTION

The Resource Description Framework (RDF) is a W3C standard offering a general graph-based model for describing information as a set of (*subject*, *predicate*, *object*) relations, known as triples. Representing data in RDF allows subject and object entities to be unambiguously identified and connected through directed and explainable relationships, thus favoring the integration and reuse of different information sources. Although RDF was initially conceived as a metadata model for the Semantic Web and the Linked Data [5], its generality and flexibility favoured its diffusion in other domains ranging from digital libraries to bioinformatics and business intelligence. Moreover, the success of initiatives such as *schema.org* and *opengraph* made RDF the de-facto standard format for publishing semi-structured information in social networks and Web sites. Such wide popularity encouraged the development of several data management systems able to deal with large RDF datasets and the complexity of querying them via SPARQL, a query language that understands the RDF model and allows to select and join graph-structured data based on both content and patterns. Not surprisingly, the increasing volume of RDF data available on-line and in various repositories pushed researchers to investigate specific solutions enabling users and software agents to store, access and query RDF graph-structured data efficiently. In this direction we can identify three relevant research topics.

- *Compressed string dictionaries*. Each RDF statement has three components: a *subject* (S), an *object* (O), and a *predicate* (P, also called a *property*) that denotes a relationship. Each one of these components is a URI string (or even a literal in the case of an object). Since URI strings can be very long and the same URI generally appears in many RDF statements, the components of triples are commonly mapped to an integer ID to save space, so that each triple in the dataset can be represented with three integers.
- *Triple indexing data structures*. Indexes built over the set of triples should allow fast access to data for processing potentially complex SPARQL queries involving large sequences of *triple selection patterns* over the RDF graph [21, 28].

---

Authors' addresses: Raffaele Peregò, ISTI-CNR, Pisa, Italy; Giulio Ermanno Pibiri, ISTI-CNR, Pisa, Italy; Rossano Venturini, University of Pisa, Pisa, Italy, raffaele.peregò@isti.cnr.it, giulio.ermanno.pibiri@isti.cnr.it, rossano.venturini@unipi.it.

---

- *Query-planning algorithms.* An effective query-planning algorithm has to find a suitable order to the set of atomic selection patterns that are needed to solve a SPARQL query, in order to speed up its execution and optimize expensive join operations [18, 19, 29].

In this work, we focus on the *triple indexing problem* and tackle the issue of reducing the space of representation of the integer triples while attaining to efficient resolution of all possible selection patterns, noting that this is crucial for guaranteeing practical SPARQL query evaluation. Moving from a critical analysis of the state-of-the-art, we note that existing solutions require too much space, because these either: rely on materializing all possible permutations of the S-P-O components [19, 27]; use expensive additional supporting structures [17, 19]; do not use sophisticated data compression techniques to effectively reduce the space for encoding triple identifiers [3, 29]. Furthermore, this additional space overhead does not always pay off in terms of reduced query response time. The aim of this work is addressing these issues by proposing compressed indexes for RDF data that are both compact *and* fast. In particular, our main contributions are the following:

**Our contributions.** Our detailed contributions are as follows.

- (1) We propose the use of a *trie*-based index layout delivering a considerably better efficiency for all triple selection patterns, thanks to the cache-friendly nature of its pattern matching algorithm. Specifically, the index materializes three different permutations of the triples in order to (symmetrically) support all triple selection patterns with one or two wildcard symbols. By leveraging on well-engineered compression techniques, we show that this design is already as compact as the most space-efficient competitor in the literature and  $2 \div 4\times$  faster on average for all selection patterns.
- (2) Starting from the aforementioned index layout, we devise two optimization aimed at reducing the redundancy of the representation. The first technique builds on the observation that the order of the triples given by a permutation can be actually exploited to compress another permutation, hence *cross-compressing* the index. The second technique shows that it is possible to *eliminate* a permutation without affecting (or even improving) triple retrieval efficiency.
- (3) Extensive and thorough experiments aimed to assess the space and time performance of our proposal versus state-of-the-art competitors are conducted on publicly available RDF datasets with a number of triples ranging from 88 millions up to 2 billions and show that our best space/time trade-off configuration substantially outperforms existing solutions at the state-of-the-art, by taking  $30 \div 60\%$  less space *and* speeding up query execution by a factor of  $2 \div 81\times$ .

**Source code.** Our code is publicly available at [https://github.com/jermprdf/rdf\\_indexes](https://github.com/jermprdf/rdf_indexes).

**Organization.** The rest of the paper is structured as follows. Section 2 provides background information and discusses the related work. Section 3 describes our trie-based RDF indexing solution and its features. Then, Section 4 details the experimental settings and reports on the results of our comprehensive evaluation. Finally, we conclude our investigation in Section 5 by summarizing our findings and providing suggestions for future research.

## 2 RELATED WORK

In the last decade many researchers investigated RDF management systems from different perspectives and a complete review of the efforts in this field is out of the scope of this paper. Readers interested in the general topic of RDF data management can refer to two very recent surveys [21, 28]. In the following we focus on the works proposing indexing structures built over the set of triples

to support efficient RDF query processing. These works usually exploit string dictionaries to compactly code the URIs occurring in the triples with unique integer identifiers (IDs) and adopt specific domain-dependent techniques to index various S-P-O permutations in order to enhance locality in the access to all the triples matching a given selection pattern: SPO, OSP, SOP, OPS, POS and PSO. The advantage of such organization is that any selection pattern with wildcards possibly occurring in one or two components among the subject, predicate, or object can be efficiently processed on the most suitable index, e.g., the one where the matching triples are stored contiguously in memory. Moreover, also joins for processing structural SPARQL queries are efficiently supported by fast triple selection over two indexes. However, the clear disadvantage is the additional space occupancy and overhead to store and manage the redundant information.

HexaStore [27] and RDF-3X [18, 19] are among the most notable examples of exhaustive application of the above approach. They both build and materializes six indexes, one for each possible permutation of the three RDF components. RDF-3X introduces also additional indexes over count-aggregated variants for all three two-dimensional and all three one-dimensional projections. The great flexibility of these solutions at querying time is paid with a very large space occupancy. To partially address this issue, inspired by compression methods for inverted lists in text retrieval systems, RDF-3X exploits VByte to encode the delta gaps computed from the increasingly ordered sequences of integers stored in the clustered B+ trees used for the indexes.

A completely different approach relies on representing the triples in memory with bitmaps [3, 29]. BitMat [3] encodes the RDF data with a 3D bit-cube, where the three dimensions correspond to S, P and O, respectively. TripleBit [29] is a recent RDF store that codes triples with a bit matrix. In this matrix each column represents a distinct triple where only two bits are set to 1 in correspondence of the rows associated with the subject and the object of the triple. Symmetrically, each row corresponds to a distinct entity value occurring as subject or object in the triples uniquely identified by the bits set to 1 in the row. Since the resulting bit matrix is very sparse it is compressed at the column level by simply coding the position of the two rows corresponding to the bits set, i.e., the identifiers of the subject and object of the associated triple. The experimental assessment discussed in the paper shows that TripleBit outperforms RDF-3X and BitMat by up to 2 ÷ 4 orders of magnitude on large RDF datasets. Nevertheless, the space occupancy and scalability of such techniques is not very good.

Other authors investigated how variations of well-known data structures like  $k^2$ -trees [7] and the Sadakane’s *compressed suffix array* (CSA) [25] can be exploited to compactly represent RDF datasets. In particular,  $k^2$ -TRIPLES [2] partitions the dataset into disjoint subsets of (*subject*, *object*) pairs, one per predicate, and represents the (highly) sparse bit matrices with  $k^2$ -trees. Another approach called RDFCSA [6] builds an integer CSA index [11] over the sequence of concatenated triple IDs, with the use of truncated Huffman codes on integer gaps and run lengths for optimized performance. The experimental assessment shows that RDFCSA requires roughly twice the space of the  $k^2$ -TRIPLES but it is up to two order of magnitude faster than the former. Although both these solutions outperform existing solutions like RDF-3X in both space usage and query efficiency, no implementation is publicly available to reproduce their results<sup>1</sup>.

The RDF index most similar to our solution is HDT-FoQ (*Focused on Querying*) [12, 17], a trie-based solution that exploits the skewed structure of RDF graphs to reduce space occupancy while supporting fast querying. The HDT-FoQ format includes a *header*, detailing logical and physical metadata, the *dictionary*, encoding all the unique entities occurring in the triples as integers, and the set of *triples* encoded in a single SPO trie data structure. In order to support predicate-based retrieval, the second level of the trie is represented with a *wavelet tree* [15]. Finally, additional

---

<sup>1</sup>Personal communication.

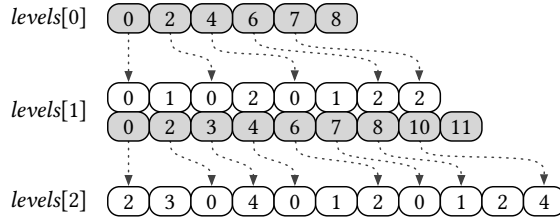


Fig. 1. A trie data structure representing a set of triples. Shaded boxes indicate pointers whereas the others refer to the nodes of the trie. The dashed arrows are just for representational purposes and point to the position written in the corresponding originating box.

inverted lists are maintained for object-based triple retrieval. In particular, for each object  $o$ , an inverted list is built, listing all pairs (*subject, predicate*) of the triples that contain  $o$ . Thus, searches are carried out by accessing each pair and searching for it in the trie. As we are going to detail next, our base solution is similarly based on the trie data structure but instead of maintaining a single trie we exploit efficient and effective compression strategies resulting in a faster solution for all the triple selection patterns and a smaller space occupancy.

### 3 THE PERMUTED TRIE INDEX

In this section we introduce our trie-based index that solves the *triple indexing problem* mentioned in Section 1: compressing a large set of integer triples by granting the efficient resolution of sequences of triple selection patterns. In particular, in Section 3.1 we introduce the base indexing data structure and in Section 3.2 and 3.3 we discuss two variants aimed at reducing redundancies in the representation.

In order to better support our design choices and explain the intuition behind the described ideas, we show in the following the results of some motivating experiments conducted on the DBpedia dataset, dubbed as “the nucleus for a Web of Data” [4], that is the English version of DBpedia (version 3.9) containing more than 351 millions of triples (see also Table 3). In Section 4 we will report on the comprehensive set of experiments conducted to assess the performance in space and time of our implementations versus state-of-the-art competitors on publicly available RDF datasets of varying size and characteristics.

#### 3.1 Data structure

As a high-level overview, our index maintains three different permutations of the triples, with each permutation sorted to allow efficient searches and effective compression as we are going to detail next. The permutations chosen are SPO, POS and OSP in order to (symmetrically) support all the six different triple selection patterns with one or two wildcard symbols: SP? and S?? over SPO; ?PO and ?P? over POS; S?O and ??O over OSP. The two additional patterns with, respectively, all symbols specified or none, can be resolved over any permutation, e.g., over the canonical SPO in order to avoid permuting back each returned triple.

Each permutation of the triples is represented as a 3-level trie data structure, with nodes at the same level concatenated together to form an integer sequence. We keep track of where groups of siblings begin and end in the concatenated sequence of nodes by storing such pointers as absolute positions in the sequence of nodes. Therefore, the pointers are integer sequences as well. Moreover, since the triples are represented by the trie data structure in sorted order, the  $n$  node IDs in the first level of each trie are always complete sequences of integers ranging from 0 to  $n - 1$  and, thus, can

```

1 select(triple)
2   i = triple.first
3   (begin, end) = levels[0].pointers[i]
4   j = begin
5   if triple.second != -1                                ▶ not a wildcard symbol
6     [j = levels[1].nodes.find(begin, end, triple.second)
7     second = iterator(levels[0].pointers.at(i), levels[1].nodes.at(j))
8     i = j
9     (begin, end) = levels[1].pointers[i]
10    j = begin
11    third = iterator(levels[1].pointers.at(i), levels[2].nodes.at(j))
12    return iterator(triple.first, second, third)

```

Fig. 2. Function `select` solving triple selection patterns with one or two wildcard symbols. The input is a *triple* object, assumed to be formed by its *first*, *second* and *third* attributes.

be omitted. We can model each trie data structure with an array `levels[0..2]` of three objects, each one having two integer sequences of *nodes* and *pointers*. An exception is represented by `levels[0]` for which, as above discussed, nodes are missing, and by `levels[2]` for which pointers are missing. Refer to Figure 1 for a pictorial example in which the ordered set of triples  $\{(0,0,2), (0,0,3), (0,1,0), (1,0,4), (1,2,0), (1,2,1), (2,0,2), (2,1,0), (3,2,1), (3,2,2), (4,2,4)\}$  is indexed.

The advantage of this layout is *two-fold*: first, we can effectively compress the integer sequences that constitute the levels of the tries to achieve small storage requirements; second, the triple selection patterns are made *cache-friendly* and, hence, efficient by requiring to simply *scan* ranges of consecutive nodes in the trie levels. We now explore the impact of these two advantages.

**Solving triple selection patterns.** The pseudo code reported in Figure 2 illustrates how triple patterns with one or two wildcard symbols are supported by our index. Given a sequence *S*, function `S.find(i, j, x)` finds the ID *x* in the range `S[i, j)` and returns its absolute position in the sequence. If *x* is not found in the range, a default position (e.g., `-1`) will be returned to signal the event and the number of matches will be 0 (for simplicity, in the pseudo code, we assume the *triple* to be valid, i.e., indexed by the data structure). Function `S.at(i)` instead instantiates an iterator starting at `S[i]`. We assume that accessing the pointers at position *i* fetches a pair of adjacent values (*begin*, *end*), such that `begin = pointers[i]` and `end = pointers[i + 1]` with the value `(end - begin)` giving the number of children of the node at position *i* in the previous level of the trie (line 3 and 9). Furthermore, the select algorithm creates two iterators to scan ranges of the second and third levels of the trie, respectively. These iterators are then used to define a final iterator that combines the iterating capabilities of both objects (line 12).

Apart from the cost of the find operation called for patterns with one wildcard symbol and the time needed to access the pointer sequences, the time spent per matched triple is  $O(1)$ , that is optimal. In the following, we discuss how the find operation and the random access can be efficiently supported directly over the compressed node sequences.

**Representation.** A key characteristic of the trie data structure is to conceptually replace runs of the same ID *x* in a sequence with the pair  $(x, \textit{pointer})$ , where the *pointer* information indicates the run length and where the run is located in the sequence. This already produces significant space

		SPO			
		bits/triple	access	find	scan
Level 2	Compact	4.74 (4.41%)	1.6	22	3
	EF	2.15 (2.33%)	24.4	74	4
	PEF	2.01 (2.51%)	65.9	83	5
	VByte+SIMD	3.79 (4.07%)	68.2	77	2
Level 3	Compact	27.00 (25.12%)	1.6	31	4
	EF	26.90 (29.19%)	31.1	97	6
	PEF	26.36 (32.95%)	77.0	115	6
	VByte+SIMD	26.57 (28.52%)	418.8	421	6
		POS			
		bits/triple	access	find	scan
Level 2	Compact	10.42 (9.69%)	1.4	285	2
	EF	4.54 (4.93%)	18.6	533	3
	PEF	1.34 (1.67%)	73.1	200	4
	VByte+SIMD	3.42 (3.67%)	154.3	13067	1
Level 3	Compact	25.00 (23.26%)	2.6	69	4
	EF	24.17 (26.23%)	41.6	185	5
	PEF	20.31 (25.39%)	90.0	150	6
	VByte+SIMD	22.95 (24.64%)	431.6	910	5
		OSP			
		bits/triple	access	find	scan
Level 2	Compact	22.15 (20.61%)	2.2	63	4
	EF	21.37 (23.19%)	34.0	168	5
	PEF	18.06 (22.58%)	79.8	135	6
	VByte+SIMD	20.40 (21.90%)	429.3	905	5
Level 3	Compact	11.00 (10.24%)	2.2	8	5
	EF	5.86 (6.36%)	30.4	60	6
	PEF	4.76 (5.95%)	67.5	97	7
	VByte+SIMD	8.87 (9.52%)	71.8	75	5

Table 1. Performance of various compressors applied to the sequences of *nodes* of all tries and levels for the DBpedia dataset. Performance is reported in bits/triple and in nanoseconds per integer when performing access, find and scan operations. In parentheses we report the percentage of space occupancy taken out by the specific sequence with respect to the whole index.

savings when triples share many repeated IDs, as it holds for large RDF datasets. Again, refer to Figure 1 for a basic graphical evidence of this fact. We note that the sequences of *nodes* and *pointers* on each trie level can be effectively compressed using a wealth of available techniques, developed for compactly represent (long) integer sequences with excellent search capabilities, e.g., the ones appearing in inverted indexes [23, 30]. A review and discussion of such techniques is out of scope for the contents of this paper: we point the interested reader to the more general survey by Zobel and Moffat [30] and the one by Pibiri and Venturini [23] discussing the most recent developments.

Therefore, the problem we face now is the one of compactly representing the levels of each trie by achieving, at the same time, an efficient execution of select algorithm that, besides the find function, demands fast random access to the pointer sequences (see Figure 2). Below we discuss some different options to achieve this goal, chosen as representative of the many available.

A first (obvious) option would be to minimize the number of bits needed to encode a single integer, taking  $\lceil \log_2 \max \rceil$  bits per value with  $\max$  being the maximum one in the sequence. The advantage of this technique, that we indicate as Compact in the following, is simplicity and efficiency, given that a few bit-wise operations are needed to implement the random access operation and, consequently, the  $\text{find}(i, j, x)$  function that can be supported by binary searching the interval  $[i, j)$ .

Instead, an elegant integer encoder fulfilling the select requirements is *Elias-Fano* (EF) [9, 10] that combines a fast, namely constant-time, random access algorithm with compressed space close to the information theoretic minimum. Again, see [23] for a detailed description of such encoding as well as of its *partitioned* variant PEF [20]. We test this encoder as (one of the) representative of the best space/time trade-off in the literature. Not surprisingly, the adoption of this encoder to model the levels of a trie has been recently advocated in [22] and shown to reduce the space of massive  $n$ -gram datasets by a factor of 2 over the most compact trie representation and to provide state-of-the-art time performance.

Other approaches focus on representing a sequence of  $d$ -gaps, i.e., differences between successive integers [30], with the requirement of computing a prefix-sum when decoding the  $d$ -gapped stream. A common way of speeding up the operations on the compressed sequence, is to divide the sequence into fixed-size blocks, e.g., 128 integers. As a meaningful representative of such technique, we use VByte [26] that achieves the highest decompression speed in the literature, especially when combined with the SIMD-ized decoding algorithm devised by Plaisance et al. [24].

**Performance.** Table 1 reports the performance of the mentioned techniques when applied to represent the node ID sequences of the tries of DBpedia. In order to operate in a meaningful benchmarking setting for the select algorithm, instead of accessing random positions in the sequences or resolving a find in random intervals, we use a set of 5000 actual triples randomly extracted from the indexed DBpedia dataset. Given a triple  $(s, p, o)$  belonging to the trie SPO, we report in the table the time required to access the second level at the position of  $p$  (pre-calculated); similarly, we measure the time needed for finding the position of  $p$  among the children of  $s$  (find). The time reported for scan indicates the time spent per node, when decoding the level sequentially. The same hold for the third level of SPO, as well as for the other tries POS and OSP. The details of the experimental setting are reported at the beginning of Section 4.

By looking at the results reported in Table 1 we can observe the following.

- All the tested techniques require about the same space for representing the sequences of subjects and objects, with PEF being the most space-efficient.
- There is only a little difference between the space of PEF and Compact when representing the objects of SPO (third level). The reason for this behavior is the *monotonicity* of the sequences required by the Elias-Fano encoder. In fact, while the sequences of pointers are monotone by construction and, thus, are immediately encodable, the same does not hold for node ID sequences where only sub-sequences of sibling nodes are ordered. Node ID sequences are thus formed by *sorted sub-sequences* with the last element of a sub-sequence being not necessarily smaller than the first one of the next one. In order to encode such sequences with a compressor like Elias-Fano we need to apply a simple transformation adding to the value of each node ID the value of the prefix-sum of the previously coded sub-sequence. This transformation makes Elias-Fano codes to be expensive whenever the prefix-sum is updated too frequently as it happens in the presence of many small ranges, because the universe of representation grows

(very) quickly, thus enlarging the space of representation. As we will show later, most levels of the tries are indeed populated by such tiny ranges, meaning that each node has only few children on average.

- PEF is roughly 2× smaller than the other options on the sequences of predicates.
- The Compact representation is, in general, the fastest for all operations, especially when performing random access that requires 1.4 ÷ 2.6 nanoseconds. However, it is the largest too.
- The PEF variant imposes a penalty of roughly 2× over its un-partitioned counterpart when performing random access, but results to be faster on average for the find operation because it operates inside a partition that is much smaller than the whole sequence. As expected, VByte codes spend more time for find because of the (expensive) decompression of blocks that happens to be competitive only when decoding the sequences of predicates.

For the reasons discussed above, we adopt the following design choices. We use PEF to represent all the sequences of node IDs, except for the last of the trie SPO where we adopt the Compact representation. The pointer sequences are represented using plain EF (not partitioned) given its superior random access efficiency.

From now on, we refer to this solution as the 3T index.

**Space breakdowns.** Lastly in this subsection, we discuss the space breakdowns reported in parentheses in Table 1. These values indicate the percentage of space out of the whole index required to encode with a given compressor the specific sequence of IDs (the space for pointers is excluded from the count).

As a general consideration, we observe that each permutation takes roughly 1/3 of the space of the whole index, with the SPO trie being a little larger than the other two. For example, by considering the values for the PEF encoder, we have that the levels of SPO require 2.51% + 32.95% = 35.46% of the whole index, whereas the ones of POS 27.06%. Summing up all the values for PEF, we get a total percentage of 91.05%, thus the space for the pointer sequences takes less than 9% of the whole index.

The levels of the index that contribute the most to the overall index space are: the third levels of SPO and POS; the second level of OSP. Specifically, we observe that such levels take practically the *whole* space of the single tries which they belong to because the other two levels take just a small fraction of the space (e.g., between 5% and 7.5% for PEF). The first levels take almost no space because they comprise only the pointers. The second levels of SPO and POS take little space for different reasons: for the former, because the number of bits needed to represent a predicate is small compared to the number of bits necessary to represent subjects and objects (see Table 3); for the latter, because the high associativity of predicates partitions the objects' space into a few but very long sorted sequences that are highly compressible. For example, only 1.34 bits/triple are needed for the objects to be represented in the second level of the trie POS. The former reason also causes the third level of OSP to require a small space.

These considerations suggest that the efforts in trying to reduce the space of our RDF index should be mainly targeted to reduce the space for encoding level three of SPO and POS and level two of OSP. This will be the objective of the next subsection.

### 3.2 Cross compression

The index layout that we have introduced represents the triples three times in different (cyclic) permutations in order to optimally solve all triple selection patterns. However, the current description does not take advantage of the fact that in this way the same set of triples is represented multiple times and that, consequently, the index has an abundance of redundant information. In this subsection we develop a novel compression framework that does not compress each trie



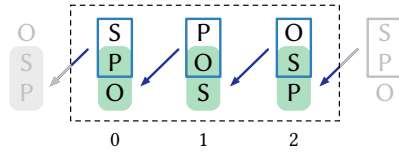


Fig. 3. Graphical representation of our cross-compression technique, in which an arrow  $(X)_i \rightarrow (Y)_j$  indicates that level  $X$  of trie  $i$  is used to cross-compress the level  $Y$  of trie  $j = (i + 2) \bmod 3$ , for  $i = 0, 1, 2$ .

independently but employs levels of the tries to compress levels of other tries, thus holistically *cross-compressing* the different permutations.

The compression framework is graphically illustrated in Figure 3. In the picture, we depict the three different permutations and highlight that the levels enclosed in the squared boxes are used to cross-compress the (lower) levels enclosed by the shaded boxes as pointed to by the arrows. Cross-compression works by noting this crucial property: *the nodes belonging to a sub-tree rooted in the second level of trie  $j$  are a subset of the nodes belonging to a sub-tree rooted in the first level of trie  $i$ , with  $j = (i + 2) \bmod 3$ , for  $i = 0, 1, 2$* . The correctness of this property follows automatically by taking into account that the triples indexed by each permutation are the same. Therefore, the children of  $x$  in the second level of trie  $j$  can be re-written as the *positions* they take in the (larger, enclosing) set of children of  $x$  in the first level of trie  $i$ . Re-writing the node IDs as positions relative to the set of children of a sub-tree yields a clear space optimization because the number of children of a given node is much smaller (on average) than the number of distinct subjects or objects (see also Table 3 for the precise statistics).

Trie	Level	Average	Maximum
SPO	1	5.54	52
	2	2.32	8489
POS	1	91,578.32	21,219,244
	2	2.59	10,141,311
OSP	1	2.70	10,141,327
	2	1.13	10

Table 2. Number of children of the trie nodes for DBpedia.

We claim that the average number of children for the nodes in the first levels of the tries is the key statistic affecting the effectiveness of the described cross-compression technique because smaller numbers require a smaller number of bits to be represented. We report this statistic in Table 2 for the DBpedia dataset. The data reported in the table is self-explanatory: except for the first level of POS, the average number of children is *very* low, being actually less than 3 for OSP, across the different permutations. As already mentioned, the high associativity of the predicates causes each predicate to have many children, i.e., several orders of magnitude more than the children of subjects and objects. We also show the maximum number of children, i.e., a maximum of  $m$  indicates that there is a least one node that has  $m$  children. While this value can be very distant from the average, for some levels like the second of SPO and OSP is actually similar, meaning that the selection patterns  $SP?$  and  $S?O$  have a robust worst-case guarantee.

```

1 map(parent, child)
2   (begin, end) = levels[0].pointers[parent]
3   return levels[1].nodes.find(begin, end, child) – begin

1 unmap(parent, child)
2   (begin, end) = levels[0].pointers[parent]
3   return levels[1].nodes[child + begin]

```

Fig. 4. Functions used to map and unmap a child ID *conditionally* to its sibling IDs.

Figure 4 shows a compact pseudo code illustrating how a *child* ID can be rewritten conditionally to its *parent* ID in the described cross-compression framework. Conversely, the function *unmap* shows what actions are needed to recover the original ID before returning the result to the user. In particular, this latter function needs to be applied to the third component of *every* triple matched by a selection pattern, hence partially eroding the retrieval efficiency. Note that fast random access is needed by the *unmap* function, thus it could be convenient to model *levels[1].nodes* with a Compact representation (see also Table 1 for benchmark numbers).

**Discussion.** We have seen that the nodes in the third level of the POS trie can be cross-compressed using the sub-trees rooted in the first level of the trie OSP, whose average number of children is actually less than 3 on real-world datasets. Therefore, what we should expect is to have the third level of POS mostly populated with values in  $\{0, 1, 2\}$  that requires just two bits to be written instead of more than 20 bits to represent each subject. Therefore, we argue that a significant space saving can be achieved by applying this technique, introducing a slowdown for two (out of eight) selection patterns only, i.e., ?PO and ?P? that are solved on the trie POS. As already motivated, the second level of OSP is represented with Compact.

We also argue that cross-compressing the other two permutations, i.e., SPO and OSP, yields only modest advantages. Again, refer to Table 2. To cross-compress the third level of SPO we use the children branching out from the predicates to the objects. But, as already noted, a predicate has many children, thus dwarfing the corresponding reduction in the average number of bits needed to represent a mapped ID. For the permutation OSP, instead, we have to shrink the node ID of the predicates. But these are already encoded compactly, e.g., in less than 4.8 bits/triple on average for DBpedia (see Table 1). Thus, given that we can not expect great space savings in these two cases, we consider the cross-compressed index to be the one with cross-compression on the POS triples and that we indicate with CC in the following.

### 3.3 Eliminating a permutation

The low number of predicates exhibited by RDF data in combination to the corresponding very few children of the nodes in the second level of the trie OSP, leads us to consider a different select algorithm for the resolution of the query pattern S?O, able to take advantage of such skewed distribution of the predicates. In fact, recall from Table 2 that on the DBpedia dataset, for a given (*subject*, *object*) pair, just 1.13 predicates are returned on average and 10 at most. Also recall from Section 3.1 that the pattern S?O is solved by the select algorithm in Figure 2 by performing a find operation followed by a scan on the OSP trie.

**An alternative approach for solving S?O.** The idea is to pattern match S?O directly over the SPO permutation with the algorithm illustrated in Figure 5. For a given subject *s* and object *o*, in

```

1 enumerate(s, o)
2   (begin, end) = levels[0].pointers[s]
3   for i = begin; i != end; i + 1
4     p = levels[1].nodes[i]
5     (j, k) = levels[1].pointers[i]
6     position = levels[2].nodes.find(j, k, o)
7     if position != -1                                     ▶ o is found in [j, k]
8     | (s, p, o) is a matching triple

```

Fig. 5. The algorithm specialized to pattern match S?O.

short, we operate as follows. We consider the set of all the predicates that are children of  $s$ . For each predicate  $p_i$  in the set, we determine if the object  $o$  is a child of  $p_i$  with the find function: if it is, then  $(s, p_i, o)$  is a triple to return. We refer to such strategy as the enumerate algorithm, whose correctness is immediate.

As partially motivated, we argue that the efficiency of the outlined algorithm is due to several distinct facts.

- The small number of predicates as children of a given subject implies that the for loop in Figure 5 performs few iterations: while these children are few per se, e.g., *at most* 52 for DBpedia, the iterations will be on average far less. For example, less than 6 iterations are needed on average for DBpedia (see detailed statistics in Table 2).
- For a given (subject, predicate) pair, we have a very limited number of children to be searched for  $o$ , thus making the find function run *faster* by *short scans* rather than via binary search. On DBpedia, an average of 2.32 children have to be scanned per pair.
- Furthermore, the enumerate algorithm operates on the SPO permutation whereas the select algorithm on OSP: but if we consider the rows for the PEF encoder in Table 1, we see that a find operation issued on the second level of OSP costs  $1.6\times$  more the amount of nanoseconds spent on SPO. Thus, avoiding percolating the trie OSP produces further savings.

**Discussion.** In the light of devising a competitive algorithm to pattern match S?O on SPO, we consider another index layout. Combining the design introduced in Section 3.1 with the present considerations, five of out the eight different selection patterns can be solved efficiently by the trie SPO, i.e.: SPO, SP?, S??, S?O and ????. In order to support other two selection patterns, we can either chose to: (1) materialize the permutation POS for *predicate-based* retrieval (query patterns ?PO and ?P?); (2) materialize the permutation OPS for *object-based* retrieval (query patterns ?PO and ??O). The choice of which permutation to maintain depends on the statistics of the selection patterns that have to be supported. We stress that the introduced enumerate algorithm allows us to actually *save the space* for a third permutation that, as we have already pointed out in Section 3.1, costs roughly 1/3 of the whole space of the index. We call this solution the 2T index; with two concrete instantiations 2Tp (predicate-based) and 2To (object-based).

The only selection pattern that is not immediately supported is ?P? for 2To (symmetrically, ??O for 2Tp). However, we know that predicates are very associative, thus the list of subject nodes having a predicate parent is, on average, very long and highly compressible. Therefore, we can afford to maintain a two-level structure PS that, for every predicate  $p$ , maintains the list of all subjects that appear in triples having predicate  $p$ . Recall from our discussion concerning the space breakdowns in Section 3.1 that the similar structure PO costs just 1.34 bits/triple for DBpedia, thus

Dataset	Triples	Subjects (S)	Predicates (P)	Objects (O)	SP pairs	PO pairs	OS pairs
DBLP	88,150,324	5,125,936	27	36,413,780	58,476,283	46,468,249	70,234,083
Geonames	123,020,821	8,345,450	26	42,728,317	118,410,418	45,096,877	112,961,698
DBpedia	351,592,624	27,318,781	1480	115,872,941	151,464,424	135,673,814	311,567,728
WatDiv	1,092,155,948	52,120,385	86	92,220,397	230,085,646	111,561,465	1,092,137,931
LUBM	1,334,681,190	217,006,852	17	161,413,040	1,060,824,925	195,085,216	1,334,459,593
Freebase	2,067,068,154	102,001,451	770,415	438,832,462	878,472,435	722,280,094	1,765,877,943

Table 3. Datasets statistics.

PS should be only a little more costly given that the SP pairs are more in number than the PO pairs (see Table 3). Therefore, we proceed as follows.

(1) Access the list  $p \rightarrow [s_1, \dots, s_n]$ .

(2) For  $i = 1, \dots, n$ , consider the pair  $(s_i, p)$  and pattern match  $s_i p$ ? against the permutation SPO.

Similarly for the index 2Tp, if we consider a query pattern ??O with specified object  $o$ , |P| find operations are needed to locate all the occurrences of  $o$  in the second level of the trie POS.

(1) For every predicate  $p$ , consider its children.

(2) Determine if  $o$  is among them with the find function. If yes, then return all triples  $(s, p, o)$  with  $s$  being a child of the pair  $(p, o)$ .

Again, the correctness of these two algorithms is immediate. In order to distinguish them from select and enumerate, in the following analysis we refer to both approaches as inverted.

## 4 EXPERIMENTAL ANALYSIS

In this section we first compare the performance of the three solutions that we have introduced in the previous section, i.e., 3T (Section 3.1), 3T with *cross compression*, indicated as CC (Section 3.2) and 2T (Section 3.3). Then, we compare against the competitive approaches at the state-of-the-art that we have discussed in Section 2.

**Datasets.** We perform our experimental analysis on the following standard datasets, whose statistics are summarized in Table 3.

- DBLP [8] is the dump of the DBLP Computer Science Bibliography collected during 2017;
- Geonames [13] is the official 2012 dump collected by geonames.org;
- DBpedia [4] is the English version of DBpedia (version 3.9), dubbed as “the nucleus for a Web of Data”;
- WatDiv [1] is the Waterloo SPARQL Diversity Test Suite largest dataset;
- LUBM [16] is a synthetic dataset generated using the methodology described by Lehigh University Benchmark, that is an established and widely-used benchmark for semantic data;
- Freebase [14] is the last available data dump of Google Freebase collected in December 2013.

**Experimental setting and methodology.** All the experiments are performed on a server machine with 4 Intel i7-7700 cores (@3.6 GHz), with 64 GB of RAM DDR3 (@2.133 GHz), running Linux 4.4.0, 64 bits. We implemented the indexes in C++14 and compiled the code with gcc 7.3.0 using the highest optimization setting, i.e., with compilation flags `-O3` and `-march=native`.

The indexes are saved to the disk after construction and loaded in internal memory to be queried. In order to avoid disk caching effects, we clear the cache of the disk before measuring query timings. To measure the query processing speed, we use a set of 5000 triples drawn at random from the

Index		DBLP	Geonames	DBpedia	Freebase
		bits/triple	bits/triple	bits/triple	bits/triple
	3T	75.24(+31%)	71.59 (+32%)	80.64(+33%)	74.20(+30%)
	CC	63.54(+18%)	67.04 (+27%)	66.91(+19%)	70.46(+26%)
	2To	56.46 (+8%)	53.23 (+8%)	57.51 (+6%)	55.72 (+6%)
	2Tp	<b>51.99</b>	<b>48.98</b>	<b>54.14</b>	<b>52.17</b>
		ns/triple	ns/triple	ns/triple	ns/triple
SPO	<i>all</i>	203	221	353	521
SP?	<i>all</i>	197	347	11	3
S??	<i>all</i>	28	40	10	3
???	<i>all</i>	11	13	9	9
S?O	3T,CC	2490 (5.6×)	3767 (7.7×)	1833 (2.6×)	6547 (1.8×)
	2To,2Tp	<b>445</b>	<b>490</b>	<b>692</b>	<b>3736</b>
?PO	3T,2To,2Tp	<b>5</b>	<b>5</b>	<b>5</b>	<b>5</b>
	CC	12 (2.4×)	15 (3.0×)	16 (3.2×)	14 (2.8×)
??O	3T,CC	12 (2.4×)	12 (2.4×)	12 (2.4×)	10 (2.0×)
	2To	<b>5</b>	<b>5</b>	<b>5</b>	<b>5</b>
	2Tp	5 (1.0×)	5 (1.0×)	6 (1.2×)	10 (2.0×)
?P?	3T,2Tp	<b>9</b>	<b>8</b>	<b>6</b>	<b>6</b>
	CC	21 (2.3×)	36 (4.5×)	30 (5.0×)	29 (4.8×)
	2To	81 (9.0×)	138(17.2×)	22 (3.7×)	18 (3.0×)

Table 4. Comparison between the performance of 3T, CC and 2T indexes, expressed as the total space in bits/triple and in average ns/triple for all the different selection patterns.

datasets and set 0, 1 or 2 wildcard symbols. The reported results are averages over five runs of queries in order to smooth fluctuations during measurements. All queries run on a single CPU core.

In all tables, percentages and speed up factors are taken with respect to the values highlighted in bold font.

#### 4.1 The permuted trie index

To assess the performance of our solutions we chose to show the results for the *real-world* datasets of DBLP, Geonames, DBpedia and Freebase, given that similar trends and conclusions were observed on the others. We will use the other two (synthetic) datasets, WatDiv and LUBM, to test the speed of selection patterns occurring in the corresponding SPARQL query logs in Section 4.2.

**Compression effectiveness.** We first comment on the compression effectiveness achieved by our different solutions. Refer to the upper part of Table 4, where we report the average number of bits spent per represented triple on the different datasets. We sorted the rows in order of increasing effectiveness, thus we can immediately conclude that the most compact index is 2Tp, with CC in between the values of 3T and 2T. The 2Tp variant is even smaller than 2To because it materializes the permutation POS that requires less space than OPS as maintained by 2To. In particular, the compression level exhibited by 3T lies in the range of  $72 \div 81$  bits/triple and as expected, the

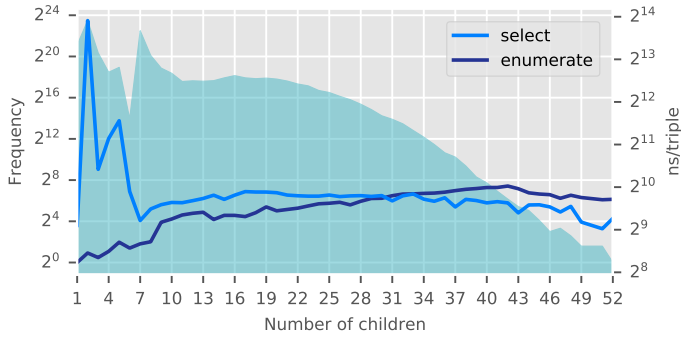


Fig. 6. Comparison between the select and enumerate algorithm when pattern matching S?O on DBpedia, for queries having subjects whose number of children (predicates) is between 1 and 52. On the background, we show the distribution of the number of children.

refinements devised in Section 3.2 and 3.3 pay off. The CC index brings a further saving of 11% on average, because the space of representation of the third level of POS is almost dissolved, being actually reduced by roughly 4× on the different datasets (but the Compact compressor on the second level of OSP takes more space than PEF, as used by 3T). Instead, the 2T indexes reduces the space by 30% on average since they basically eliminate one permutation of the triples, hence optimizing by roughly 1/3.

**Triple selection patterns.** Now we examine the speed of all the different triple selection patterns by commenting on the lower part of Table 4, where we report the timings in nanoseconds spent per returned triple. In particular, patterns are grouped together when solved by the same trie (or have exactly the same performance). Let us proceed in order, by discussing each group.

- The index component represented by the trie SPO is common by the three solutions, thus the results for SPO, SP?, S??, ??? are the same. We can see that the indexes solve a lookup operation, i.e., the query pattern SPO with all symbols specified, in a fraction of a microsecond, specifically  $1/5 \div 1/2 \mu\text{sec}$ . This is basically the time needed for two find operation, respectively on the second and third level of the trie. Observe how, instead, the “fixed” cost of the find operation on the second level is amortized by the number of matching triples for the pattern SP? on the larger DBpedia and Freebase. Similarly, the time spent per triple by the patterns S?? and ??? is just a handful of nanoseconds.
- We now discuss the pattern S?O that is solved by 3T and CC with the select algorithm (Figure 2) on the trie OSP, but by 2T with the enumerate algorithm (Figure 5) on the trie SPO. In Section 3.3, we have discussed about the efficiency of the enumerate algorithm and motivated its advantages over the select algorithm. In fact, we observe from Table 4 that enumerate is faster than select by 4.4× on average (a minimum of 1.8× on Freebase, and up to 7.7× on Geonames). In particular, we have claimed that the crucial statistic affecting the performance of the enumerate algorithm is the very low associativity of the subjects, i.e., the limited number of predicates as children of a given subject. Therefore, in order to validate our hypothesis, we inspect the full spectrum of possibilities in Figure 6, where we illustrate how the time for S?O changes for queries whose subjects have different number of children. We show the result on DBpedia for space constraints (similar behaviors were observed for the other datasets anyway). For DBpedia the number of children of a subject, henceforth indicated with  $C$ , varies between 1 and 52 (see also Table 2).

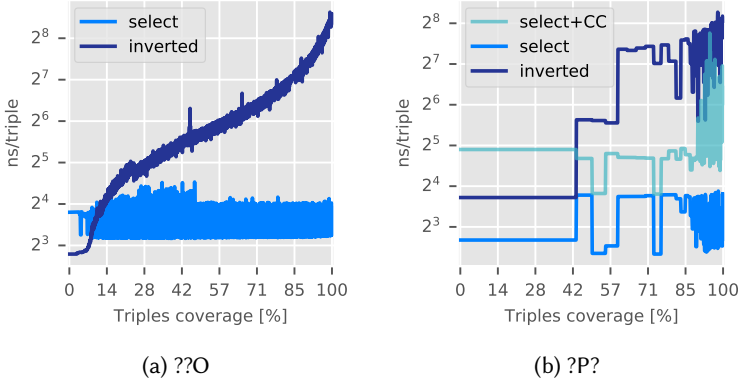


Fig. 7. Average ns/triple by *decreasing* number of matches, for the query pattern ??O and ??P?

As it is clear from the plot, whenever  $C < 30$  the enumerate algorithm is faster and loses only for higher values, because of the cache-friendly behaviour of select due to its scan-based nature. In particular, on the background of the plot we also show the distribution of  $C$ , i.e., how many subjects have  $C = c$  children, for  $c = 1..52$ . The distribution explains that the majority of subjects have only few children, and in correspondence of such values, the enumerate algorithm is *much* faster than select (e.g., up to  $43\times$  for  $C = 2$ ), thus explaining why on average is faster as well.

- The pattern ?PO illustrates how the cross-compression affects the retrieval of the subjects on the third level of the POS permutation. For every match, the unmap function must be executed, hence requiring a random access to be performed on the second level of OSP (a potential cache-miss). This results in a slowdown of  $\approx 3\times$ .
- The number of nanoseconds per triple for the pattern ??O is controlled by the average selectivity of the queries: while most of the queries have only a few matches, the presence of low-selective queries maintains the average response time very low. Also, notice that it is solved  $2.4\times$  faster by 2To because it traverses the trie OPS and not OSP as it is done by the 3T and CC indexes, thus employing find operations on the sequence of predicates rather than subjects that are faster (see details in Table 1). It is also interesting to note that the pattern is solved efficiently by the 2Tp variant on average, given that the cost of the  $|P|$  find operations is well amortized by the number of returned matches. However, when  $|P|$  grows significantly and the results per query are very few, the response time per occurrence is much higher. This *output-sensitive* behaviour is a well-known characteristic of RDF stores. Therefore, for completeness, we also show in Figure 7a how the time per triple changes by *decreasing* number of matches, until we cover the whole set of triples. The experiment shows that while the select is less output-sensitive, indeed for a reasonably good fraction of the triples, e.g., 25%, the *average* query time of inverted is close to the one of select and even faster for  $\approx 10\%$  of the triples.
- Similar considerations hold for the ??P? query pattern, whose corresponding stress behaviour is shown in Figure 7b: the select offers the best worst-case guarantee across all different output sizes, with the cross-compression technique (select+CC) and the inverted algorithm both imposing overheads due to additional cache-misses. Observe how inverted is actually faster than select+CC for more than 40% of the triples, with the latter being in trade-off position between the former and select.

Index	DBLP	Geonames	DBpedia	Freebase
	bits/triple	bits/triple	bits/triple	bits/triple
2Tp	<b>51.99</b>	<b>48.98</b>	<b>54.14</b>	<b>52.17</b>
HDT-FoQ	76.89 (+32%)	88.73 (+45%)	76.66 (+29%)	83.11 (+37%)
TripleBit	125.10 (+58%)	120.03 (+59%)	130.07 (+58%)	—
	ns/triple	ns/triple	ns/triple	ns/triple
2Tp	5	5	5	5
?PO	HDT-FoQ 12 (2.4×)	13 (2.6×)	14 (2.8×)	13 (2.6×)
	TripleBit 15 (3.0×)	13 (2.6×)	14 (2.8×)	—
	ns/triple	ns/triple	ns/triple	ns/triple
2Tp	<b>445</b>	<b>490</b>	<b>692</b>	<b>3736</b>
S?O	HDT-FoQ 1789 (4.0×)	2097 (4.3×)	3010 (4.3×)	$0.7 \times 10^7$ (2057×)
	TripleBit 11872 (26.7×)	13008 (26.5×)	18023 (26.0×)	—
	ns/triple	ns/triple	ns/triple	ns/triple
2Tp	<b>197</b>	<b>347</b>	<b>11</b>	<b>3</b>
SP?	HDT-FoQ 640 (3.2×)	897 (2.6×)	30 (2.7×)	9 (3.0×)
	TripleBit 1222 (6.2×)	927 (2.7×)	42 (3.8×)	—
	ns/triple	ns/triple	ns/triple	ns/triple
2Tp	<b>28</b>	<b>40</b>	<b>10</b>	<b>3</b>
S??	HDT-FoQ 110 (3.9×)	154 (3.9×)	29 (2.9×)	9 (3.0×)
	TripleBit 2275 (81.2×)	3261 (81.5×)	490 (49.0×)	—
	ns/triple	ns/triple	ns/triple	ns/triple
2Tp	<b>9</b>	<b>8</b>	<b>6</b>	<b>4</b>
?P?	HDT-FoQ 108 (12.0×)	173 (21.6×)	32 (5.3×)	41 (6.8×)
	TripleBit 28 (3.1×)	28 (3.5×)	40 (6.7×)	—
	ns/triple	ns/triple	ns/triple	ns/triple
2Tp	<b>5</b>	<b>5</b>	<b>6</b>	<b>10</b>
??O	HDT-FoQ 17 (3.4×)	17 (3.4×)	18 (3.0×)	18 (1.8×)
	TripleBit 24 (4.8×)	60 (12.0×)	24 (4.0×)	—

Table 5. Comparison between the performance of different indexes, expressed as the total space in bits/triple and in average ns/triple.

**Discussion.** From the above considerations concerning the performance of our solutions, we conclude that: (1) the 3T index is the one delivering best *worst-case* performance guarantee for all triple selection patterns; (2) the 2T variants reduce its space of representation by 25 ÷ 33% *without* affecting or even improving the retrieval efficiency on most triple selection patterns (*only one* out of the eight possible has a lower query throughput in the worst-case); (3) the cross-compression technique is outperformed by the 2T index layouts for space usage but offers a better worst-case performance guarantee than 2Tp for the pattern ?P?. Therefore, as a reasonable trade-off between space and time, we elect 2Tp as the solution to compare against the state-of-the-art alternatives.

## 4.2 Overall comparison

In this section, we compare the performance of our selected solution 2Tp against the competitive approaches HDT-FoQ [17] and TripleBit [29] described in Section 2, since these both outperform RDF-3X [19] and BitMat [3]. In particular, the experimental analysis provided by the authors of HDT-FoQ [17] reports that RDF-3X is larger than their own index by 3.8×, 4.6× and 3× on DBLP, Geonames and DBpedia respectively, and also slower by a factor of 4 ÷ 8× on most selection patterns



Index	WatDiv		LUBM	
	bits/triple	sec/query	bits/triple	sec/query
2Tp	<b>54.16</b>	<b>0.08</b>	<b>50.01</b>	<b>0.73</b>
HDT-FoQ	68.79(+21%)	19.67(238.2×)	92.41(+46%)	17.26(23.7×)
TripleBit	129.49(+58%)	0.34 (4.1×)	111.93(+55%)	4.01 (5.5×)

Table 6. Performance in bits/triple and in average seconds spent by the different solutions to execute the sequence of triple selection patterns generated for the SPARQL queries in the WatDiv and LUBM logs.

(HDT-FoQ scores worse only for the pattern ?P? because of the penalty induced by representing predicate sequences with wavelet trees, as we are going to confirm next). This is not surprising given that all the six triple permutations, plus aggregate indexes, are materialized by RDF-3X with a severe query processing overhead due to expensive I/O transfers. Similar results are obtained by the authors of TripleBit [29], whose experiments on LUBM show that RDF-3X is 2× larger and 2 ÷ 14× slower. In the same paper BitMat is shown to be even larger than RDF-3X and up to several orders of magnitude slower<sup>2</sup>.

The figures reported for our solution, HDT-FoQ and TripleBit do not consider the (common) space needed for the string dictionaries, as well as the time needed for dictionary access/lookup. We use the C++ libraries as provided by the corresponding authors and available at <https://github.com/rdfhdt/hdt-cpp> and <https://github.com/nitingupta910/TripleBit>, respectively.

Table 5 reports the space of the indexes and the timings for the different selection patterns, but excluding (due to page limit) the ones for SPO and ???: our approach is anyway faster for both by at least a factor of 3× (TripleBit does not support the query pattern SPO). Concerning the space, we see that the 2Tp index is significantly more compact, specifically by 30% and almost 60% compared to HDT-FoQ and TripleBit respectively, on average across all different datasets (TripleBit fails in building the index on Freebase). Concerning the speed of triple selection patterns, most factors of improved efficiency range in the interval 2 ÷ 5× and, depending on the pattern examined, we report peaks of 26×, 49×, 81× or even 2057×.

To further confirm the results, we execute all triple selection patterns needed to solve the SPARQL queries from the publicly available query logs of WatDiv<sup>3</sup> and LUBM<sup>4</sup>, that are both well-established benchmarks for RDF data. We used the query planning algorithm of TripleBit to obtain a serial decomposition of the SPARQL queries into atomic selection patterns. We argue that this methodology has the advantage of fairly testing the different indexes over the same set of patterns as executed in the same order, thus really testing the raw speed of the underlying indexing data structure that is our goal in this work. The results of the benchmark, illustrated in Table 6, closely match the ones already discussed in Table 5. In particular, our index optimizes the space of representation by 21 ÷ 58% and is remarkably faster than TripleBit by 4.5× on average and by up to 238× than HDT-FoQ on WatDiv. Most selection patterns for both query logs have ?P? and ?PO forms. In fact, observe how our solution is 3 ÷ 6.7× faster than TripleBit on these patterns, with HDT-FoQ being the worst because of the cache-misses introduced by representing the predicate sequences with (potentially, tall) wavelet trees.

<sup>2</sup>To confirm this result, we built the BitMat index on DBpedia, with flag COMPRESS\_FOLDED\_ARRAY as suggested by the author Medha Atre, and measured a space occupancy of 483.72 bits/triple.

<sup>3</sup><http://swat.cse.lehigh.edu/projects/lubm/queries-sparql.txt>

<sup>4</sup><http://grid.hust.edu.cn/triplebit/watdiv.txt>

## 5 CONCLUSIONS AND FUTURE WORK

In this work we have proposed compressed indexes for the storage and search of large RDF datasets, delivering a remarkably improved effectiveness/efficiency trade-off against existing solutions. In particular, the extensive experimentation provided has shown that our best trade-off configuration reduces storage requirements by 30 ÷ 60% and provides 2 ÷ 81× better efficiency, on real-world RDF datasets with up to 2 billions of triples.

Future work will target the two related problems mentioned in Section 1, that is: (1) providing an efficient *string dictionary* data structure and (2) devising a novel *query planning* algorithm.

## ACKNOWLEDGMENTS

This work was partially supported by the BIGDATAGRAPES project (grant agreement #780751, European Union’s Horizon 2020 research and innovation programme) and MIUR-PRIN 2017 “Algorithms, Data Structures and Combinatorics for Machine Learning”.

## REFERENCES

- [1] Güneş Aluç, Olaf Hartig, M Tamer Özsu, and Khuzaima Daudjee. 2014. Diversified stress testing of RDF data management systems. In *International Semantic Web Conference*. Springer, 197–212.
- [2] Sandra Álvarez-García, Nieves Brisaboa, Javier D Fernández, Miguel A Martínez-Prieto, and Gonzalo Navarro. 2015. Compressed vertical partitioning for efficient RDF management. *Knowledge and Information Systems* 44, 2 (2015), 439–474.
- [3] Medha Atre, Vineet Chaoji, Mohammed J Zaki, and James A Hendler. 2010. Matrix Bit loaded: a scalable lightweight join query processor for RDF data. In *Proceedings of the 19th international conference on World wide web*. ACM, 41–50.
- [4] Sören Auer, Christian Bizer, Georgi Kobilarov, Jens Lehmann, Richard Cyganiak, and Zachary Ives. 2007. Dbpedia: A nucleus for a web of open data. (2007), 722–735.
- [5] Tim Berners-Lee, James Hendler, and Ora Lassila. 2001. The Semantic Web. *Scientific American* 284, 5 (2001), 34–43.
- [6] Nieves R Brisaboa, Ana Cerdeira-Pena, Antonio Farina, and Gonzalo Navarro. 2015. A compact RDF store using suffix arrays. In *International Symposium on String Processing and Information Retrieval*. Springer, 103–115.
- [7] Nieves R Brisaboa, Susana Ladra, and Gonzalo Navarro. 2009.  $k^2$ -trees for compact web graph representation. In *International Symposium on String Processing and Information Retrieval*. Springer, 18–30.
- [8] DBLP. 2017. Computer Science Bibliography. <http://www.rdfhdt.org/datasets>. (2017).
- [9] Peter Elias. 1974. Efficient Storage and Retrieval by Content and Address of Static Files. *Journal of the ACM (JACM)* 21, 2 (1974), 246–260.
- [10] Robert Mario Fano. 1971. On the number of bits required to implement an associative memory. *Memorandum 61, Computer Structures Group, MIT* (1971).
- [11] Antonio Fariña, Nieves R Brisaboa, Gonzalo Navarro, Francisco Claude, Ángeles S Places, and Eduardo Rodríguez. 2012. Word-based self-indexes for natural language text. *ACM Transactions on Information Systems (TOIS)* 30, 1 (2012), 1.
- [12] Javier D Fernández, Miguel A Martínez-Prieto, and Claudio Gutierrez. 2010. Compact representation of large RDF data sets for publishing and exchange. In *International Semantic Web Conference*. Springer, 193–208.
- [13] geonames.org. 2012. Geonames. <http://www.rdfhdt.org/datasets>. (2012).
- [14] Google. 2013. Freebase Data Dumps. <https://developers.google.com/freebase>. (2013).
- [15] Roberto Grossi, Ankur Gupta, and Jeffrey Scott Vitter. 2003. High-order entropy-compressed text indexes. In *Proceedings of the fourteenth annual ACM-SIAM symposium on Discrete algorithms*. 841–850.
- [16] LUBM. 2019. Lehigh University Benchmark. <http://swat.cse.lehigh.edu/projects/lubm>. (2019).
- [17] Miguel A Martínez-Prieto, Mario Arias Gallego, and Javier D Fernández. 2012. Exchange and consumption of huge RDF data. In *Extended Semantic Web Conference*. Springer, 437–452.
- [18] Thomas Neumann and Gerhard Weikum. 2010. The RDF-3X engine for scalable management of RDF data. *The VLDB Journal—The International Journal on Very Large Data Bases* 19, 1 (2010), 91–113.
- [19] Thomas Neumann and Gerhard Weikum. 2010. x-RDF-3X: fast querying, high update rates, and consistency for RDF databases. *Proceedings of the VLDB Endowment* 3, 1-2 (2010), 256–263.

- [20] Giuseppe Ottaviano and Rossano Venturini. 2014. Partitioned elias-fano indexes. In *Proceedings of the 37th international ACM SIGIR conference on Research & development in information retrieval*. ACM, 273–282.
- [21] M. Tamer Özsu. 2016. A Survey of RDF Data Management Systems. *Front. Comput. Sci.* 10, 3 (June 2016), 418–432.
- [22] Giulio Ermanno Pibiri and Rossano Venturini. 2017. Efficient data structures for massive n-gram datasets. In *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval*. ACM, 615–624.
- [23] Giulio Ermanno Pibiri and Rossano Venturini. 2018. Inverted index compression. *Encyclopedia of Big Data Technologies* (2018), 1–8.
- [24] Jeff Plaisance, Nathan Kurz, and Daniel Lemire. 2015. Vectorized VByte Decoding. In *International Symposium on Web Algorithms (iSWAG)*.
- [25] Kunihiko Sadakane. 2003. New text indexing functionalities of the compressed suffix arrays. *Journal of Algorithms* 48, 2 (2003), 294–313.
- [26] Larry H Thiel and HS Heaps. 1972. Program design for retrospective searches on large data bases. *Information Storage and Retrieval (ISR)* 8, 1 (1972), 1–20.
- [27] Cathrin Weiss, Panagiotis Karras, and Abraham Bernstein. 2008. Hexastore: sextuple indexing for semantic web data management. *Proceedings of the VLDB Endowment* 1, 1 (2008), 1008–1019.
- [28] Marcin Wylot, Manfred Hauswirth, Philippe Cudré-Mauroux, and Sherif Sakr. 2018. RDF Data Storage and Query Processing Schemes: A Survey. *ACM Computing Surveys (CSUR)* 51, 4 (2018), 84.
- [29] Pingpeng Yuan, Pu Liu, Buwen Wu, Hai Jin, Wenya Zhang, and Ling Liu. 2013. TripleBit: a fast and compact system for large scale RDF data. *Proceedings of the VLDB Endowment* 6, 7 (2013), 517–528.
- [30] Justin Zobel and Alistair Moffat. 2006. Inverted files for text search engines. *ACM computing surveys (CSUR)* 38, 2 (2006), 6.