

Uplifting the Superpowers of Worst-Case-Optimal Join Algorithms^{*}

Adrián Gómez-Brandón^{1,3}, Aidan Hogan^{2,3}, and Gonzalo Navarro^{2,3}

¹ Universidade da Coruña, CITIC, Facultade de Informática, Spain

² Department of Computer Science, University of Chile, Chile

³ IMFD — Millennium Institute for Foundational Research on Data, Chile

Abstract. Worst-case-optimal (wco) join algorithms have demonstrated their power – in both theory and practice – to efficiently solve complex Basic Graph Patterns (BGPs). Modern graph query languages, such as SPARQL and GQL, have BGPs at their core, but also have a wide range of other features, including filters (aka. selections). Such conditions are typically handled via pre- or post-filtering, before or after processing the BGPs. In this paper we show how to uplift wco join algorithms so as to incorporate such filtering natively, improving efficiency. We demonstrate the superiority of this approach by extending the *Ring* – a compact index that provides wco resolution of BGPs within almost no extra space on top of the graph – so as to handle property graphs using our new techniques while retaining compactness. We implement this extension and experimentally show that it outperforms various baseline systems.

Keywords: worst-case optimal · graph databases · filter · joins · GQL

1 Introduction

There is an ever-growing demand for increasingly-efficient systems for querying knowledge graphs. In the context of open knowledge graphs, RDF and SPARQL dominate [30]. In the context of enterprise knowledge graphs, property graphs, query languages such as Cypher [20] and GQL [17], and commercial engines such as Neo4j [49], are a popular alternative. Despite these seemingly disparate graph models and query languages, the same key primitives – in particular, graph patterns [2] – underlie both RDF/SPARQL and property graph query engines. Techniques for querying one form of knowledge graph model can often be applied for the other, with various systems [12,45,48] natively supporting both.

A promising recent line of research for efficiency evaluating graph queries is to combine compact data structures and worst-case optimal joins [6,8,13,28,4]. Such works have focused on evaluating BGPs over triple-based models such as RDF.

^{*} Funded by ANID – Millennium Science Initiative Program – Code ICN17_002, Chile. AGB is funded by MCIN/AEI/10.13039/501100011033 and EU/ERDF “A way of making Europe” PID2022-141027NB-C21 (EARTHDL). CITIC is funded by the Xunta de Galicia and co-financed by the EU through FEDER Galicia: grant ED431G 2023/01. GN is funded by Fondecyt Grant 1260080, ANID, Chile.

We extend these works in two ways. First, we show how techniques originally proposed for RDF-like models – specifically the Ring [6] – can be adapted to the property graph model. Second, we further add support for selections, aka. *filters*, which are among the most widely-used query operators in practice: Bonifati et al. [14] find, for example, that 40% of SPARQL queries in real-world logs use `FILTER`, second only to `SELECT`. While filters can be evaluated over the results of a BGP, one of the most fundamental query optimizations is to “push down” filters, using them to eliminate intermediate results as soon as possible. We propose techniques to enable the push-down of filters by translating them into efficient operations directly over a custom compact data structure.

Paper structure. Section 2 discusses related works on compact data structures and wco joins for graphs. Section 3 presents preliminaries. Section 4 introduces the core technical abstraction of “leaptors”: iterators enabling wco execution for graph queries with filters. Section 5 proposes the PG-Ring for the concrete setting of property graphs and GQL-like queries, describing the implementation of a static, in-memory database engine using leptors. Section 6 presents experimental results for two benchmarks comparing PG-Ring with internal and external baselines. Section 7 wraps up our contributions and presents future work.

2 Related Work

In the context of efficiently evaluating graph patterns, recent years have seen two promising lines of research emerge.

The first line of research refers to *worst-case optimal* (wco) joins [41], which provide – in the context of querying graphs – theoretical and practical guarantees on the maximum cost of enumerating all solutions for a basic graph pattern (BGP). Such algorithms have been shown empirically to provide notable speed-ups over traditional join algorithms, upholding their theoretical underpinnings for RDF/SPARQL [26,13,48] and similar triple-based graph models [33]. However, the benefits of wco joins in terms of time often come at the cost of space [26,33,48] as they often require storing additional index orders.

The second line of research explores compressed data structures for indexing and querying graphs. In the RDF/SPARQL world, formats such as HDT [18], COTTAS [3], etc., propose compressed RDF file formats that can evaluate triple patterns, which can be used to build more complex query engines [44].

More recent works have combined both lines of research by investigating succinct data structures that support wco joins. These in-memory structures include the Ring (based on wavelet trees) [6], QDags (based on quadtrees) [8], Hypertries (based on bit tensors) [13], RDF-TDAA (based on Directly Addressable Arrays) [28], and RDFCSA (based on compressed suffix arrays) [4]. These works support wco queries while addressing their key limitation: space.

However, to the best of our knowledge, no work has investigated succinct data structures that support wco joins and filters (beyond equality). Furthermore, we are not aware of works applying wco joins over property graphs.

3 Graph Databases and Worst-Case-Optimal Algorithms

3.1 RDF and Basic Graph Patterns

An RDF graph [16] G is a set of triples (s, p, o) , each denoting an edge $s \xrightarrow{p} o$. Node s is called the *subject*, label p the *predicate*, and node o the *object*. A *Basic Graph Pattern (BGP)* Q is a set of *triple patterns* (x, y, z) where any element can be a constant or a variable. A *solution* to a BGP is a function that assigns a constant to each variable of the BGP so that all resulting triples occur in the graph. *Solving* a BGP Q consists of producing the set $Q(G)$ of all its solutions.

Let G have N triples. The *AGM bound* [9] of Q , $Q^* \geq |Q(G)|$, is the maximum size of a solution set of Q over any graph of N triples. An algorithm solving Q is *worst-case-optimal (wco)* [9,26] if it works in time $\tilde{O}(Q^*)$, where $\tilde{O}$ hides factors polylogarithmic on N , or independent of N (such as $|Q|$). The wco concept and wco algorithms have had notable impact on the resolution of complex and cyclic BGPs. Such algorithms outperform every algorithm that enforces the triple pattern conditions *consecutively* (as classic database join plans do) by enforcing all the conditions *simultaneously*. The classical example is the triangle query $R(a, b) \bowtie S(b, c) \bowtie T(c, a)$, which if the tables have N pairs takes $\Omega(N^2)$ worst-case time with any consecutive-enforcing algorithm, whereas wco algorithms solve it in time $\tilde{O}(N^{3/2})$.

3.2 Leapfrog TrieJoin

Leapfrog Triejoin (LTJ) [46] is a seminal wco join algorithm. It precalculates six tries where the graph triples are inserted as strings of length 3 in all different orders, $\mathcal{T} = \{\text{SP0, SOP, POS, PS0, OSP, OPS}\}$. It chooses an order to handle the variables, which is called a *variable elimination order (VEO)*, and then assigns each triple pattern $t \in Q$ to a trie $\tau_t \in \mathcal{T}$, such that (i) the constants of t appear first in the order of τ_t , and (ii) the variables of t appear in τ_t in the same order of the VEO. For each triple pattern $t \in Q$, LTJ first descends by its constants on τ_t , and the node $v_t \in \tau_t$ arrived at is the *locus* of t . Next, LTJ takes one variable x after the other in the VEO. Let $T_x = \{t \in Q, x \text{ appears in } t\}$. For each $t \in T_x$, the possible values of x are the children of v_t . LTJ then *intersects* the children of all the loci $\{v_t, t \in T_x\}$. For each value c in the intersection, it branches with the *binding* $x := c$ and descends by value c from all the involved loci, $v_t \leftarrow \text{child}(v_t, c)$ for all $v_t \in T_x$. By the time all the variables are bound in this way, the set of bindings in that branch forms a new solution to report. Fig. 1 shows a toy graph and its trie POS (ignore property *age* of nodes for now).

A way to implement LTJ intersections is to take the smallest child value c of one locus $v_t \in T_x$, and search the children of the next locus $v_{t'} \in T_x$ for the smallest child value $c' \geq c$; this operation is called *leap*, that is, $c' \leftarrow \text{leap}(c)$. We then assign $c \leftarrow c'$ and go on with the third locus $v_{t''} \in T_x$, and so on. We continue cycling over the loci until a whole pass over T_x retains the same c value, so c is the next element in the intersection. After branching on $x := c$, we restart the process with the next child value in some loci. This algorithm is optimal




Fig. 1: On the left, a graph database of where people live and area they work in; person nodes have also property age. On the right, the trie POS for the graph.

in the sense that its cost is $\tilde{O}(alt)$, where alt is the *alternation complexity* [11] of the sets to intersect. This is the least number of switches from one sorted sequence of values to another needed to cover them all; for example intersecting $\{2i, 1 \leq i \leq n\} \cap \{2i-1, 1 \leq i \leq n\} = \emptyset$ is hard because we must switch at every new value ($alt = 2n-1$), whereas intersecting $\{i, 1 \leq i \leq n\} \cap \{i, n < i \leq 2n\} = \emptyset$ is easy ($alt = 1$; it suffices to compare two elements to certify the output). Other intersection techniques used in some LTJ implementations, like probing the elements of the smallest set in the others, do not yield this guarantee.

3.3 Beyond RDF: Filtering by properties and labels

Most graph query languages, like SPARQL [24], GQL [19], and SQL/PGQ [17], have BGPs at their core, but also include filters. GQL and SQL/PGQ assume property graphs [1], which extend triples to allow nodes and edges to have *properties* with *values*. Queries may then specify that, to bind some variable x to a node or edge in the triple pattern, that node or edge must contain a property π and that its corresponding value $x.\pi$ must be within a range, say $a \leq x.\pi \leq b$ (where a and b can be constants or other property values, say $x.\pi \leq y.\pi'$). It is also customary to assign *labels* to nodes and edges, so that $x.\ell$ requires that variable x is assigned to an element having label ℓ . While SPARQL filters are applied directly on nodes and predicates in triples, we will allow zero-to-many values per property and zero-to-many labels, and hence our scheme is also flexible enough to capture features of RDF, as we will discuss in Section 4.3.

While the BGP core of the query can be solved efficiently with a wco algorithm, conditions on properties and labels have been typically handled outside of this wco algorithm by pre- or post-filtering: we can either filter all nodes and edge values v to leave for consideration only those holding $a \leq v.\pi \leq b$ (or where label $v.\ell$ exists), or first solve the BGP and filter every output, removing bindings of v that do not satisfy the conditions. If we use LTJ, a better solution is to adapt it so as to post-filter the values v as soon as they come out of the intersection of loci children, thereby pruning branches earlier.

However, as with non-wco algorithms, these approaches may force us to examine many more solutions than necessary. Consider, for the graph of Fig. 1, the

BGP $\{(x, \textit{lives}, \textit{Europe}), (x, \textit{works}, \textit{CS})\}$ complemented with the condition $x.\textit{age} \geq 50$, aimed at finding senior people living in Europe and working in CS. LTJ will choose two loci, say in POS (it could also be OPS), by descending through $(\textit{lives}, \textit{Europe})$ and $(\textit{works}, \textit{CS})$, respectively. The intersection of the loci children yields all people x in Europe working in CS. If we pre-filter by age on a large real graph of this form, we will generate a lot of senior people, very few of which will end up passing the BGP filter. If we post-filter, we will generate a lot of people in Europe working in CS, very few of which will pass the seniority filter. None of those methods meet the alternation complexity bound, which is never asymptotically larger than the minimum of both sets, and can be much smaller.

4 Leaptors: Incorporating Filters into LTJ

We generalize the logic of LTJ so as to handle not only triple patterns mapped to trie loci, but in general *leaptors* (these extend “iterators”, which can only deliver the elements sequentially). A leaptor represents a set S with a total order and supports operation $\textit{leap}(c) = \min\{c' \in S \cup \{+\infty\}, c' \geq c\}$, in time polylogarithmic on $|S|$. An example leaptor is the algorithm that finds the child of a trie node (the locus) with the smallest value $c' \geq c$; this can be implemented, say, via exponential search on the increasing sequence of child values.

4.1 Properties and ranges of values

We now define a new leaptor to incorporate filters on property values. Let $\mathcal{U}$ be the set of node identifiers in the graph (we defer properties and labels on edges to Section 5, though they could be supported similarly), and $\mathcal{A}_\pi$ be the set of values a property π takes in the graph. Assume that both sets are integers; if not, we map them to integer ranges $[1, |\mathcal{U}|]$ and $[1, |\mathcal{A}_\pi|]$ while preserving order. For example, by storing an array of all the node values, and another of all the values in $\mathcal{A}_\pi$, we map actual values to integers with a binary search (in $O(\log N)$ time, only once per query), and can map back any result in constant time.

We now define a *discrete grid* M_π of $[1, |\mathcal{U}|] \times [1, |\mathcal{A}_\pi|]$, and set a point at (v, a) in M_π for every node v having property π with value a . This supports v having zero-to-many values for a property π . Query $\textit{leap}(c)$ on the leaptor for $a \leq x.\pi \leq b$ is implemented on M_π via the orthogonal range query $\textit{leftmost}([c, +\infty], [a^+, b^-])$, which finds a point with minimum x -value $c' \geq c$ whose y -value is within $[a^+, b^-]$ ($\textit{leap}(c)$ returns $+\infty$ if this range is empty). Here, a^+ is the smallest value $a^+ \geq a$ in $\mathcal{A}_\pi$ and b^- is the largest value $b^- \leq b$ in $\mathcal{A}_\pi$; both are found in $O(\log |\mathcal{A}_\pi|) \subseteq O(\log N)$ time via binary searches (done only once per query). The geometric query itself is called an *orthogonal range successor query*, and can be solved in time $O(\log^\epsilon N)$ for any constant $\epsilon > 0$, using space linear in the number of points in M_π [38]. This space is proportional to the database size, which includes one unit of space per value each property takes for each node. The leaptor time is then always in $O(\log N)$. See Fig. 2.




Fig. 2: The grid M_{age} for the graph of Fig. 1 (only showing person nodes). The shadowed area corresponds to the query $\text{leftmost}([Carla, +\infty] \times [50, +\infty])$.

Let us return to our example of Section 3.3 to illustrate the power of our method. Apart from the two leaptors implemented on tries for the triple patterns, we have the one implemented with grids for the condition $x.\text{age} \geq 50$. Once a candidate c passes the filters of the two loci, we find the leftmost point (c', a) in $[c, +\infty] \times [a^+, b^-]$. This will yield the next senior person $c' \geq c$, and then the next person in Europe working in CS will be sought after c' , not after c . In the example, the first candidate coming out of the intersection of children of trie loci (*lives*, Europe) and (*works*, CS) is Bob. The grid leaptor also returns Bob for the query $\text{leap}(\text{Bob}) = \text{leftmost}([\text{Bob}, +\infty] \times [50, +\infty])$, and thus Bob is an output of the query. If we now start from the grid to get the next candidate, its leaptor returns Diego (see the shaded area in Fig. 2), which lets us discard Carla already.

This process achieves the alternation complexity, visiting fewer candidates than pre- or post-filtering, and potentially not fully traversing either set (those passing the BGP and those passing the age filter). For example, if the person ids happen to be clustered by country, occupation, or age, the alternation complexity can be much lower than the number of candidates processed by standard filtering.

4.2 Labels

We also define leaptors for labels. While labels can be handled as a special property (giving them value 0 when the label exists, and asking for $\text{leftmost}([c, +\infty] \times [0, 0])$ on their grid), a faster leaptor uses a *successor* data structure [42] storing the set D_ℓ of node ids having label ℓ . The structure can find the smallest integer $\geq c$ in D_ℓ (or return $+\infty$ if there is none) in time $O(\log \log |\mathcal{U}|)$, using $O(|D_\ell|)$ space. The leaptor for label ℓ then simply implements $\text{leap}(c) = \text{successor}(D_\ell, c)$.

It is also generally possible to ask for nodes having a property π , independently of its value. The corresponding leaptor can, again, be simulated with query $\text{leftmost}([c, +\infty] \times [-\infty, +\infty])$ on M_π , but within linear space we can also store a successor data structure on the set D_π of the node ids having some value for property π , and run $\text{leap}(c)$ in time $O(\log \log |\mathcal{U}|)$.

SP0 order			OSP order			POS order		
		L_{SP0}			L_{OSP}			L_{POS}
Alice	<i>lives</i>	Europe	America	Diego	<i>lives</i>	<i>lives</i>	America	Diego
Alice	<i>works</i>	Finance	America	Emily	<i>lives</i>	<i>lives</i>	America	Emily
Bob	<i>lives</i>	Europe	CS	Bob	<i>works</i>	<i>lives</i>	Europe	Alice
Bob	<i>works</i>	CS	CS	Carla	<i>works</i>	<i>lives</i>	Europe	Bob
Carla	<i>lives</i>	Europe	CS	Diego	<i>works</i>	<i>lives</i>	Europe	Carla
Carla	<i>works</i>	CS	Europe	Alice	<i>lives</i>	<i>works</i>	CS	Bob
Diego	<i>lives</i>	America	Europe	Bob	<i>lives</i>	<i>works</i>	CS	Carla
Diego	<i>works</i>	CS	Europe	Carla	<i>lives</i>	<i>works</i>	CS	Diego
Emily	<i>lives</i>	America	Finance	Alice	<i>works</i>	<i>works</i>	Finance	Alice
Emily	<i>works</i>	Finance	Finance	Emily	<i>works</i>	<i>works</i>	Finance	Emily

Fig. 3: A Ring structure on the graph of Fig. 1, assuming the node id order is the lexicographic order of the strings. The Ring only stores the three shaded sequences. The arrows show how the triple **(Alice, lives, Europe)** is tracked across the three orders. The rectangle on POS, corresponding to the child of the trie root by **works**, descends by **Carla** to the rectangle on SP0.

4.3 ERing: A succinct wco index handling properties and labels

The Ring [6] is a *succinct* index that handles BGPs in wco time on RDF graphs. An index is said to be succinct if, including the data, it uses $\mathcal{S} + o(\mathcal{S})$ space, where $\mathcal{S}$ is the size of a plain representation of the data. The Ring represents the set of triples (s, p, o) with three sequences, L_{SP0} , L_{OSP} , and L_{POS} . Sequence L_{SP0} stores the objects of the N triples sorted by the order SP0, that is, triples (s, p, o) are sorted lexicographically and L_{SP0} collects the sequence of their o components. Analogously, sequence L_{OSP} collects the N predicates p in the order OSP and sequence L_{POS} collects the N subjects s in the order POS. Overall, letting $\mathcal{U}$ be the universe of nodes and $\mathcal{L}$ that of labels (or predicates), the Ring uses $(2N \log_2 |\mathcal{U}| + N \log_2 |\mathcal{L}|)(1 + o(1))$ bits, which is the space needed to store the N triples (s, p, o) in plain form, plus a sublinear redundant space. See Fig. 3.

The sequences $L_*[1, N]$ are represented with a data structure called a *wavelet tree* [23,35], which on a universe Σ (where Σ is $\mathcal{U}$ or $\mathcal{L}$) uses $N \log_2 |\Sigma|(1 + o(1))$ bits and in particular implements the operation $rank_c(L_*, i)$, the number of times c occurs in $L_*[1, i]$, in time $O(\log |\Sigma|)$. We also represent arrays A_0 , A_P , and A_S , where $A_*[1, |\Sigma|]$ holds in $A_*[e]$ the number of positions in L_* with values less than e . Those arrays can be represented with bitvectors $B_*[1, N]$ where $B_*[A_*[e]] = 1$ for all e and 0 otherwise, so $A_*[e] = select_1(B_*, e)$, where $select_b(B_*, e)$ is the position of the e th bit b in B_* . Bitvector representations solving $select_b$ in constant time [15,34] require $N + o(N)$ bits of space, which fits within the sublinear extra space of the Ring.

With those structures we can *track* triples across the arrays L_* : the triple at $L_{SP0}[i]$ has object $o = L_{SP0}[i]$, and it corresponds to $L_{OSP}[i']$, where $i' = A_0[o] + rank_o(L_{SP0}, i)$. The triple then has predicate $p = L_{OSP}[i']$, and it corresponds to $L_{POS}[i'']$, where $i'' = A_P[p] + rank_p(L_{OSP}, i')$. The subject of the triple is then $s = L_{POS}[i'']$ and if we compute $A_S[s] + rank_s(L_{POS}, i'')$ we return to $L_{SP0}[i]$.

Trie nodes correspond to ranges in some sequence L_* . A locus in SOP or SP0, reached by descending by s , corresponds to the range $L_{SP0}[i, j] = L_{SP0}[A_s[s] + 1, A_s[s + 1]]$ of the triples $(s, *, *)$. If we now descend by o in SOP, the resulting locus corresponds to range $(o, s, *)$ in the order SOP, which we do not represent, but also to range $(o, s, *)$ in the order OSP, which we do represent in L_{OSP} . This is $L_{OSP}[i', j']$, where $i' = A_0[o] + rank_o(L_{SP0}, i - 1) + 1$ and $j' = A_0[o] + rank_o(L_{SP0}, j)$. If, instead, we descend by p after s in SP0, we stay in the same sequence L_{SP0} , which represents that trie, now restricting the range to the triples $(s, p, *)$. The new range is computed by descending from the root of POS with p , and from the resulting range $L_{POS}[i'', j'']$ descending by s to $L_{SP0}[i''', j''']$, with $i''' = A_s[s] + rank_s(L_{POS}, i'' - 1) + 1$ and $j''' = A_s[s] + rank_s(L_{POS}, j'')$. We will still call those operations “descending” by a value, even if they are mapped to ranges in L_* . See again Fig. 3 for examples of tracking and descending.

Descending operations allow the Ring to simulate the LTJ algorithm by mapping loci to ranges in L_* . The Ring leaptor is implemented as follows. Say we are again in $L_{SP0}[i, j]$ after descending by s in trie SOP. To implement $leap(c)$ on the values of o , we must find the smallest value in $L_{SP0}[i, j]$ that is $\geq c$. This operation, called *range next value*, is supported in $O(\log |\Sigma|)$ time by the wavelet tree [21]. If, instead, we are in trie SP0 and want to find the smallest value of p that is $\geq c$, we start at range $L_{POS}[A_p[c] + 1, N]$ (the range of all predicates $\geq p$ in order POS), and compute $L_{SP0}[i'] = A_s[s] + rank_s(L_{POS}, A_p[c]) + 1$. Tracking $L_{SP0}[i']$ to L_{OSP} yields the desired value of p . The Ring then implements the leaptors in time $O(\log |\Sigma|)$, yielding a total query time complexity of $O(Q^* \cdot |Q| \log |\Sigma|) \subseteq \tilde{O}(Q^*)$.

Incorporating properties and labels. We add to the Ring succinct representations of grids M_π that use $N_\pi \log_2 |\mathcal{A}_\pi| (1 + o(1))$ bits, where N_π is the sum of the number of values property π has over all nodes. This is again the size of the data plus a sublinear term. Such a succinct representation is the wavelet tree of the sequence $S_\pi[1, N_\pi]$ of values property π takes, sorted by increasing identifiers of the nodes. We also represent a bitvector $B_\pi[|\mathcal{U}| + N_\pi]$ where, for each node identifier v , we append a 1 and then as many 0s as values v has for property π ; note that positions in S_π align with 0s in B_π . The grid of Fig. 2, for example, would be represented with the sequence $S_{age} = \langle 55, 45, 60, 25 \rangle$ and the bitvector $B_{age} = 110101010$ (considering only the node ids that appear in the figure).

The bitvector representation of B_π [15,34] uses $(|\mathcal{U}| + N_\pi)(1 + o(1))$ bits of space (thus fitting within the sublinear extra space of the index) and supports queries $select_b$ in constant time. We thus compute with $c' \leftarrow select_1(B_\pi, c) - c + 1$ the first position in S_π of node identifiers $\geq c$. The wavelet tree then solves the orthogonal range successor query from $S_\pi[c'..]$, in time $O(\log |\mathcal{U}|)$ [10]. The answer c'' is finally converted back to an identifier with $select_0(B_\pi, c'') - c''$. In our example, to compute $leftmost([Carla, +\infty] \times [50, +\infty])$ on M_{age} , where $Carla = 3$, we start with $c' \leftarrow select_1(B_{age}, 3) - 3 + 1 = 2$, so we find the first value ≥ 50 on $S_{age}[2..]$. This is at $S_{age}[3] = 60$. We finally convert that position, $c'' = 3$, back to the identifier $select_0(B_{age}, 3) - 3 = 4 = Diego$.

Queries on labels were handled with generic successor data structures in Section 4.2. While using linear space, this representation is not succinct: If a label ℓ is

assigned to m elements out of n , a succinct representation uses space close to the bits required to specify those m elements, that is, $\log_2 \binom{n}{m} = m \log \frac{n}{m} + O(m)$ bits. To match that space, we store one bitvector B_ℓ for each label ℓ , with one bit per node, marking with 1 the nodes having label ℓ . A *sparse* bitvector representation [36, Sec. 4.4], if m out of n bits are 1, uses $m \log \frac{n}{m} + 2m$ bits, and supports $rank_b$ and $select_0$ in time $O(\log m)$ and $select_1$ in constant time. Operation $leap(c)$ for the condition $x.\ell$ is then implemented as $select_1(B_\ell, rank_1(B_\ell, c - 1) + 1)$ with this leaptor, in time $O(\log m)$ (which is logarithmic on the data size, and thus $\tilde{O}(1)$). This representation also provides a leaptor for the next element *not* having a label ℓ , with $leap(c) = select_0(B_\ell, rank_0(B_\ell, c - 1) + 1)$, also in time $O(\log m)$.

The ERing. The resulting structure, which we call *Extended Ring (ERing)*, is succinct on the triples, properties and labels it represents. It implements LTJ and supports leaptors for filtering by the existence of node properties and labels, and by restricting the values of node properties to ranges. The range limits can be constants or other property values. When a range involves several variables, as in $x.age \leq y.age$, the leaptor is enabled only when binding the last variable (x or y in our example); by then, the condition has become a range with constants.

4.4 Worst-case optimality

The concept of worst-case optimality for BGPs (i.e., intersections of triple patterns) is well-defined for the RDF model. Node properties π can be expressed in RDF by creating a node for each literal value $a \in \mathcal{A} = \cup_\pi \mathcal{A}_\pi$, and represent $v.\pi = a$ with a triple (v, π, a) . For labels we can use a reserved property *lab* and represent $v.\ell$ with $(v, \textit{lab}, \ell)$. Conditions like $x.\pi = a$ or $x.\ell$ are then expressed as triple patterns (x, π, a) or $(x, \textit{lab}, \ell)$. Range conditions $a \leq x.\pi \leq b$ can be translated into $\{(x, \pi, z), R(a, z), R(z, b)\}$ with variables x and z . The (virtual) table $R(i, j)$ contains all the pairs $i, j \in \mathcal{A}$ such that $i \leq j$. The range condition is then simulated by making node ids respect the order of the literals in $\mathcal{A}$ and using a simple leaptor for z : $leap(c) = \max(a, c)$ if $c \leq b$ and $+\infty$ otherwise.

Mapping properties and labels into plain RDF and using the standard Ring [6] is then wco, but not succinct: it creates one triple per node label or property value. We show that our ERing, instead, is a succinct-space wco solution.

Definition 1. *A solution to BGPs extended with clauses for the existence of properties and labels, and comparing property values, is worst-case optimal if its time complexity is within the AGM bound of the corresponding query mapped onto the RDF model, as described above.*

Theorem 1. *The ERing is worst-case-optimal when solving extended BGPs.*

Proof. It suffices to observe that the ERing’s time complexity is not worse than the standard Ring’s [6] on the corresponding RDF graph, with a certain VEO: for each range condition $a \leq x.\pi \leq b$, translated into $\{(x, \pi, z), R(a, z), R(z, b)\}$, we bind z immediately after x . In the ERing, the potential values of x and z are filtered simultaneously, which per the alternation complexity is never worse than

filtering one after the other. With the standard Ring, each potential value of x will be instantiated and only then filtered by z . The conditions on the existence of properties $x.\pi$ are particular cases, $-\infty \leq x.\pi \leq +\infty$.

Simpler conditions like (x, π, a) or $(x, \textit{lab}, \ell)$ are filtered optimally by the Ring, by finding the locus of the path (π, a) or $(\textit{lab}, \ell)$ in the trie POS: later, when binding x , this locus filters the corresponding values. ERing achieves the same by adding, at the time of binding x , the leaptor $\textit{leftmost}([c, +\infty], [a, a])$ on S_π for properties or the leaptor $\textit{select}_1(B_\ell, \textit{rank}_1(B_\ell, c - 1) + 1)$ for labels.

Conditions that compare properties, like $x.\pi \leq y.\pi'$, are translated into $\{(x, \pi, z), (y, \pi', z'), R(z, z')\}$. In this case, the ERing eliminates (say) x before y , and only during the elimination of y it enforces $z \leq z'$ (where $z = x.\pi$ is by then a constant). This corresponds to, and is never worse than, eliminating variables in the order x, z, y, z' with the standard Ring.

Since the ERing solution is not worse than the standard Ring solution with some VEO, and it implements all leaps in time $\tilde{O}(1)$, it is wco. $\square$

The ERing can be better than the (standard) Ring (mapped to RDF): it can filter simultaneously what the Ring can only filter sequentially, which is equivalent to pre- or post-filtering. On the other hand, the Ring can use VEOs that are not available to the ERing: consider conveniently translating the clause $x.\pi = y.\pi$ to $\{(x, \pi, z), (y, \pi, z)\}$, where binding z first might be advantageous.

Conversely, one could represent RDF triples using properties and labels to enable simultaneous filtering over RDF: triples (u, p, v) where v is a literal could map to properties on u ($u.p = v$), while triples $(u, \textit{type}, v)$ could map to labels on u ($u.\ell = v$). Our model already supports zero-or-more labels and property values per node, but would require some amendments to cover RDF properties mixing datatype and node values, joins on predicates, and other such details.

5 Property Graphs

In the previous section we described techniques for efficiently handling filters on properties and labels. We now apply these techniques for a concrete data model and query language, namely property graphs and a fragment of GQL. Some details of this concrete setting (e.g., edge identifiers, Boolean label expressions) require additional care, while others (e.g., one edge label, at most one value for a property on a given node or edge) provide opportunities to optimize further.

5.1 Data model and query language

A *property graph* is a directed labeled multigraph where any node or edge may feature properties and labels. Let $\mathcal{G}(\mathcal{N}, \mathcal{E})$ be a property graph with $\mathcal{N} \cap \mathcal{E} = \emptyset$. Let $\mathcal{P}$ be a set of properties, and $\mathcal{A}$ a set of atomic values properties can take (e.g., integers, dates, strings). Each node or edge has zero or more properties, but a given property has at most one value for a given node or edge. Nodes have zero or more labels in $\mathcal{L}_N$ and edges have exactly one label in $\mathcal{L}_E$ ($\mathcal{L}_N \cap \mathcal{L}_E = \emptyset$).

Our graph of Fig. 1 is a property graph, with node property **age** and edge labels *lives* and *works*. Possible edge properties could be, for example, a date **since** (for edges with label *lives*) or an integer **salary** (for edges with label *works*). Possible labels for nodes (representing persons) could be *person* and *has-phd*.

Two main standards exist to query property graphs: SQL/PGQ and GQL [19]; their common graph pattern matching language is called GPML [17]. A GPML query is of the form **MATCH patterns**, where **patterns** specifies the subgraph to match as a set of *edge patterns*, optionally augmented with a **WHERE** clause to filter the results. We now describe the fragment we support [17,1,19]; the formal syntax and semantics are given in Appendix A [22].

The most basic edge patterns are of the form $(s) \rightarrow (o)$, where **s** and **o** can be constants in $\mathcal{N}$ or node variables from a set $\mathcal{V}_N$. Those are analogous to triple patterns but do not specify any label. Conditions on labels are associated with node variables, for example we can add a clause $(x:person)$ to specify that $x \in \mathcal{V}_N$ can match only nodes having the label *person* $\in \mathcal{L}_N$. Conditions with no variable names, like $(:person) \rightarrow (y)$, are also permitted. We can include these clauses in edge patterns, for example in $(x:person) \rightarrow (y)$. Label conditions may consist of expressions including conjunctions (**&**), disjunctions (**|**), negations (**!**), and groups of labels. For example, a clause $(x:person\&has-phd)$ admits only the nodes with both labels, *person* and *has-phd*.

Label conditions can also be given on edge patterns, as in $(s) - [:works] \rightarrow (o)$, which allows matching only edges whose label is *works* $\in \mathcal{L}_E$; expressions on labels can also be used. It is also possible to give variable names, from a set $\mathcal{V}_E$ disjoint from $\mathcal{V}_N$, to edges: $(s) - [e] \rightarrow (o)$ gives variable name $e \in \mathcal{V}_E$ to the edge; we can also incorporate expressions as in $(s) - [e : works] \rightarrow (o)$.

The **WHERE** clause filters the results by restricting the existence or the values of some of their properties or by comparing the identifiers of edges/nodes (yet we cannot compare edge variables with node variables). Conditions can be combined with **AND**, **OR**, and **NOT**. Comparisons are of the form $(A \text{ op } B)$, where **A** and **B** are constants or variable property values $x.\pi$ of the same type, and **op** $\in \{=, \neq, <, \leq, \geq, >\}$. We can directly compare node or edge identifiers as if they were properties, as in $(x <= y)$ **AND** $(y <= z)$, which avoids reporting the same tuples (x, y, z) in every possible order. One can also state that a variable must have, or not have, some property using $(x.since \text{ IS NOT NULL})$ or $(x.since \text{ IS NULL})$.

A final **RETURN** clause projects the variable or property values to return. We use *bag semantics*, so repetitions caused by projection are not removed.

The following query, for example, looks for names and ages of senior people in Europe working in CS that started in their job before 2010 or have no phd:

```
MATCH (x:person)-[:lives]->(Europe), (x)-[e:works]->(CS)
WHERE (x.age >= 50) AND ((e.since < 01/01/2010) OR (!x.has-phd))
RETURN x.name, x.age
```

5.2 PG-Ring: A succinct index for property graphs

We build on the ERing described in Section 4.3 to define PG-Ring, which indexes property graphs within succinct space.

For compatibility with the ERing, we create new variables where edge patterns omits them; those are naturally projected out by the **RETURN** clauses.

Since the edges of property graphs have exactly one label in $\mathcal{L}_E$, we treat edges $e = s \rightarrow o$ with label ℓ as triples (s, ℓ, o) in order to create the sequences L_* . The symbols of L_{POS} and L_{SPQ} then range over $\mathcal{N}$ and those of L_{OSP} range over $\mathcal{L}_E$. The wavelet tree of L_{OSP} uses $N \log_2 |\mathcal{L}_E| (1 + o(1))$ bits, which is succinct.

Storing the additional edge identifiers of property graphs would not be succinct, so we rather exploit the fact that such identifiers are internal by defining the identifier of an edge as its index in the order **POS**. We choose this index order since all edge identifiers with some label ℓ are precisely in the range $[A_P[\ell] + 1, A_P[\ell + 1]]$, and edge patterns of the form $(s) - [: \ell] \rightarrow (o)$ are treated as triple patterns with a constant predicate, instead of using a leaptor.

Label expressions L on edges, $(s) - [: L] \rightarrow (o)$, are converted into a set of $O(L)$ disjoint ranges: the complement of ℓ becomes $\{[1, A_P[\ell]], [A_P[\ell + 1] + 1, N]\}$, disjunction becomes union of ranges, and conjunction becomes intersection. We then maintain, in general, a set of disjoint ranges during query processing, not just one. For simplicity we consider a single range in the discussion that follows.

Leaptors for edge variables. Variables in $\mathcal{V}_E$ can be assigned to edge identifiers. Those will be handled like any other variable in LTJ, with the particularity that their identifiers are not explicit in any sequence L_* , but implicit as indices in L_{POS} (i.e., on order **POS**). The leaptor for a variable $z \in \mathcal{V}_E$ – which is mentioned on an edge pattern t represented by a trie $\tau_t \in T_z$ – will thus implement $\text{leap}(c)$ differently depending on which sequence L_* holds the range that represents the locus v_t . See Appendix B [22] for details.

Leaptors for node properties and labels. Some simplifications with respect to the treatment in Section 4.3 are possible (see Appendix B [22]). For label expressions on nodes, we can handle the complement of a single label ℓ with the leaptor for not having a label we described in Section 4.3. A leaptor for a label expression L builds a syntax tree where negations are pushed down to the leaves and $\text{leap}(c)$ is solved by recursively calling $\text{leap}(c)$ on both children of internal nodes, taking the minimum returned value c' on disjunctions. For conjunctions, we create multiary nodes containing all the subexpressions that descend by conjunction nodes, and intersect all children as a set, just as for our intersection algorithm. This matches the alternation complexity of the problem with an extra penalty factor of $O(|L| \log N) \subseteq \tilde{O}(1)$. Bitvectors B_π of properties are used in the same way as bitvectors B_ℓ of labels, to filter nodes having or not having property π .

Boolean leaptors. We use a similar syntax tree, with multiary conjunction nodes, to process together the **MATCH** edge patterns and the **WHERE** clauses. This generalizes our set intersections by allowing disjunction nodes (to represent **OR** in **WHERE** clauses), which leap by taking the least of their children’s leap values. Negations are pushed to the leaves and applied to the atomic condition, both for properties (e.g., **NOT** $(x < y)$ becomes $(x \geq y)$) and for labels (e.g. **NOT** $(x.\text{person})$ becomes $(!x.\text{person})$). See Fig. 4, and Appendix B [22] for details.




Fig. 4: The syntax tree of our example of Section 5.1. Each internal node implements a Boolean leaptor; leaf nodes implement trie, grid, and bitvector leaptors.

6 Experiments

We implement PG-Ring in C++11 on top of the succinct data structures library, *SDSL*, as described in more detail in Appendix C [22]. We compare it with various state-of-the-art alternatives that support property graphs and GQL-like languages for two different benchmarks: LSQB and Wikidata. We run a set of queries with a time limit of 600 seconds. Following the LSQB benchmark [32], for each query, we report the average runtime over 5 runs. That measurement is computed after a first warm-up to mitigate the advantage of in-memory systems. Disk activity was monitored using *iostat* at one-second intervals. In LSQB, a single pass over the queries was enough to avoid disk reads, whereas in Wikidata, additional warm-up queries (e.g., reporting all the nodes) were required beforehand. For both benchmarks, the effect on time performance of a *cold cache* scenario is reported in Appendix G [22]. All experiments were conducted on an AMD Ryzen 9 9900X at 4.4GHz, with 24 cores, 32 MB cache, and 249 GB RAM.

6.1 Labelled Subgraph Query Benchmark

We used the Labelled Subgraph Query Benchmark (LSQB) [32] to generate a graph with 35.5 million nodes and 219.4 million edges. Each node can be given one of 14 possible labels, whereas each edge can be given one of 15 possible labels. In this graph, nodes and edges do not have properties. We use the first six queries from the benchmark; the others are discarded because they require unsupported optional edges and negation. See Appendix D [22] for details.

We compare systems supporting GQL (details in Appendix E [22]): Umbra [40], DuckDB [43], Kùzu [27], Neo4j [39], Memgraph [31], MillDB [48], and our PG-Ring. Our figures show the space required by their data and indexes, not the extra needed for query resolution. This extra space is negligible for PG-Ring but not for others that produce intermediate results (we discuss an extreme case soon). We run the six queries on all the systems, limiting results to 1, 10, 100, 1,000, and 10,000. Appendix D [22] shows all results; Fig. 5 summarizes some.

Fig. 5 (left) shows the space and average time of each system with the number of results limited to 10. We can see that PG-Ring achieves a good balance between space and time. In space, PG-Ring is dominated only by DuckDB, which




Fig. 5: The space–time trade-off of the evaluated systems on six queries of LSQB. The first plot is limited to 10 results, and the second to 1,000.

is 1.6 times smaller due to its columnar compression, but its average times are 12–2,700 times slower (see Appendix D [22]). The next competitor in space is Umbra, which uses 1.8 times more space than PG-Ring. In time performance, the PG-Ring wins in three cyclic queries (Q2, Q3, and Q6, see Appendix D [22]), showing the effect of using the LTJ algorithm to solve these queries. In the other queries, the clear winner is MillDB. In fact, MillDB is very competitive in all queries, but it requires around 5.4 times more space than PG-Ring.

Fig. 5 (right) shows the same evaluation when restricting the results to 1,000. In this scenario, PG-Ring performs competitively on cyclic queries, with the exception of Q2, where Neo4j performs best. In contrast, Neo4j shows weaker performance in Q3, where it exceeds the time limit, and Q6, where it is 11 times slower. In Q1, Memgraph stands out by being 1.8 times faster than MillDB, but it is the most space demanding structure and times out in Q3. For all remaining queries, MillDB is the top performer, taking less than 4 milliseconds per query.

The plot also shows that DuckDB and Umbra fare better for obtaining many results. These systems use hash-joins, which require processing the left-side of the join completely and indexing it for lookup before processing the join’s right-side (looking up each such value). Even with an output limit, all joins except those on the right extreme of the plan are computed completely, where the last, rightmost table-scan (and the joins above it) can halt early once the limit is reached. Thus the number of intermediate results they process varies little with the limit. The consequence is that they may generate huge intermediate results, preventing their use at larger scale. For example, DuckDB, although using less space than PG-Ring in the static index, generated one billion intermediate results on the cyclic query Q3, which has only 15 million final results. Even on the relatively small LSQB dataset, it ran out of memory on Q1 and Q6 without limit.

In conclusion, PG-Ring offers a good space-time trade-off, it is highly competitive for cyclic queries, highlighting the impact of the wco join algorithm; and ranks second in space usage thanks to the use of compact data structures.




Fig. 6: On the left, the space–time trade-off of the systems evaluated on Wikidata. On the right, boxplots represent the distribution of the query execution times. The numbers on top indicate the median time for each system.

6.2 Wikidata Benchmark

We use the Wikidata knowledge graph [47] to evaluate our performance in a large scale graph with properties (via qualifiers). The Wikidata graph is represented using the RDF model; therefore, we transformed it into a property graph. The resulting graph consists of 109 million nodes, which can take 103 thousand distinct labels and have 10,569 different properties. In total, there are 719.5 million edges, with 1,640 possible edge labels and 603 distinct properties. From the Wikidata SPARQL query logs, we extracted a set of 19.6 thousand queries and converted them to GQL syntax. See Appendix F [22] for details, which includes discussions on the limitations of property graphs for representing Wikidata [25], in particular for its multi-valued and node-valued qualifiers/properties.

We exclude Umbra [40], DuckDB [43], and Kuzu [27] because Wikidata nodes can have multiple labels (classes), but these systems assume a single type/table per entity. Memgraph [31] was also excluded due to its high memory usage and the impracticality of loading the Wikidata graph. In Neo4j [39], we create indexes on the entity identifiers and on the properties used in the queries.

To study the effect of incorporating the filters directly to the LTJ algorithm, we add two more variants of PG-Ring: The *pre-filter* baseline first applies the property and label filters, and only then evaluates the edge patterns. The *post-filter* baseline evaluates the edge patterns first and subsequently applies filters.

Fig. 6 shows the performance in space and time of different systems on the selected queries, limiting the output to 1,000,000 results. PG-Ring achieves the most favorable space-time trade-off: it requires only 32% of the storage used by MillDB, which is on average just 32% faster than PG-Ring. MillDB shows, however, more stability across all queries, with its third quartile being at around 5 milliseconds, while the other systems exceed 10 milliseconds. Conversely, Neo4j is slower than PG-Ring and consumes significantly more space, exceeding 100 GB: in a real-world scenario Neo4j will consume much more space because we are restricting the indexes to just those properties involved in these specific queries.

The experiment also shows that pushing the filters directly into the LTJ algorithm leads to superior performance. On average, PG-Ring achieves a 16% improvement over the *post-filter* baseline and is 15 times faster than *pre-filter*. The boxplots show a more marked difference in distributions, for example PG-Ring is almost 60% faster than *post-filter* in the median, and much more stable in general. As an example to explain this difference, consider one of the queries:

```
MATCH (v:Q5)-[e:P21]->(u)
WHERE (v.P569 >= 1554-01-01) AND (v.P569 < 1555-01-01)
RETURN v, e, u
```

which gets *the gender (P21) of humans (Q5) born in (P569) 1554*, for which PG-Ring achieves a speedup over 5,000 times compared to the *post-filter* baseline, which needs to check for every human if they were born in 1554. Conversely, the *pre-filter* is twice as slow as PG-Ring because it first retrieves all entities born in 1554 and only then verifies whether they are humans. Note that Wikidata can also store birth dates for non-human entities (e.g., fictional characters, animals).

7 Conclusions

We proposed succinct data structures that support wco joins and filters. We show the advantages of pushing the filtering within the wco processing of the joins, thereby uplifting the superpowers of those join processing algorithms, and outperforming the classic pre- and post-filtering methods that handle the wco joins as a black box. We adapt and optimize these structures for the concrete setting of evaluating the corresponding fragment of GQL queries over property graphs, giving rise to PG-Ring. Experiments on two property graphs show that although existing systems sometimes outperform PG-Ring in terms of time (MillDB), or space (DuckDB), PG-Ring provides a strong time–space trade-off in this setting.

For future work, we plan to implement PG-Ring for RDF/SPARQL, optimize it specifically for this setting, and compare it with leading systems. Though the Ring [6] supports this setting, an interesting question is when it is preferable to represent (e.g.) datatype values as property values versus graph nodes.

We only support fragments of languages like GQL and SPARQL, missing typical relational operators (UNION, OPTIONAL), as well as path queries. It would be of interest to explore pushing further query operators from GQL, SPARQL, etc., into low-level operations on the succinct data structures; for example, existing methods for evaluating paths on the Ring could be deployed [7].

A limitation of PG-Ring is that it is currently read-only, but it can be adapted to support updates. The original Ring article [6, Sec. 8 & App. E] discusses in detail how to handle changes on nodes and edges. The PG-Ring represents property values and labels via wavelet trees and bitvectors, for which dynamic versions exist [29] while retaining succinctness and enabling updates in time $O(\log^2 N)$, but at the cost of an $O(\log N)$ time penalty factor on queries. This can be alleviated with adaptive dynamic bitvectors [37], a recent development that supports dynamism with a much smaller penalty when the update-to-query ratio is very low, something proven to work well for Wikidata-style workloads [5].

Supplemental Material Statement: Source code and query sets are available from Github at <https://github.com/adriangbrandon/PG-Ring>. The datasets are available from Zenodo at <https://zenodo.org/records/19818050>. Appendices are published on arXiv, at <https://arxiv.org/abs/2608.03840> [22].

Declaration of use of Generative AI: Generative AI was used only as a support tool for code completion, for assisting with the parsing and preprocessing of some data, and to produce Figs. 1–3 in tikz from textual descriptions. All suggestions were reviewed, tested, and adapted before being included in the final work.

References

1. R. Angles. The Property Graph Database Model. In *Proc. Alberto Mendelzon Workshop (AMW)*, 2018.
2. R. Angles, M. Arenas, P. Barceló, A. Hogan, J. L. Reutter, and D. Vrgoc. Foundations of modern query languages for graph databases. *ACM Computing Surveys*, 50(5):68:1–68:40, 2017.
3. J. Arenas-Guerrero and S. Ferrada. COTTAS: Columnar triple table storage for efficient and compressed RDF management. In *Proc. 24th International Semantic Web Conference (ISWC), part II*, pages 313–331, 2025.
4. D. Arroyuelo, F. Barisione, A. Fariña, A. Gómez-Brandón, and G. Navarro. New compressed indices for multijoins on graph databases. *Information Systems*, 137:article 102647, 2026.
5. D. Arroyuelo, D. Campos, A. Gómez-Brandón, Y. Linker, G. Navarro, C. Rojas, and D. Vrgoc. CompactLTJ: Space & time efficient Leapfrog Triejoin on graph databases. *The Very Large Databases Journal*, 34:article 67, 2025.
6. D. Arroyuelo, A. Gómez-Brandón, A. Hogan, G. Navarro, J. Reutter, J. Rojas-Ledesma, and A. Soto. The Ring: Worst-case optimal joins in graph databases using (almost) no extra space. *ACM Transactions on Database Systems*, 49(2):1–45, 2024.
7. D. Arroyuelo, A. Gómez-Brandón, A. Hogan, G. Navarro, and J. Rojas-Ledesma. Optimizing RPQs over a compact graph representation. *The VLDB Journal*, 33(2):349–374, 2024.
8. D. Arroyuelo, G. Navarro, J. L. Reutter, and J. Rojas-Ledesma. Optimal joins using compressed quadrees. *ACM Transactions on Database Systems*, 47(2):8:1–8:53, 2022.
9. A. Atserias, M. Grohe, and D. Marx. Size bounds and query plans for relational joins. *SIAM Journal on Computing*, 42(4):1737–1767, 2013.
10. J. Barbay, F. Claude, and G. Navarro. Compact binary relation representations with rich functionality. *Information and Computation*, 232:19–37, 2013.
11. J. Barbay and C. Kenyon. Deterministic algorithm for the t-threshold set problem. In *Proc. 14th ISAAC*, pages 575–584, 2003.
12. 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 *Proc. ISWC 2018 Posters & Demonstrations, Industry and Blue Sky Ideas Tracks*, 2018.

13. A. Bigerl, F. Conrads, C. Behning, M. Ahmed Sherif, M. Saleem, and A.-C. Ngonga Ngomo. Tentris - a tensor-based triple store. In *Proc. 19th International Semantic Web Conference (ISWC), part I*, pages 56–73, 2020.
14. A. Bonifati, W. Martens, and T. Timm. An analytical study of large SPARQL query logs. *The VLDB Journal*, 29(2-3):655–679, 2020.
15. D. R. Clark. *Compact PAT Trees*. PhD thesis, University of Waterloo, Canada, 1996.
16. R. Cyganiak, D. Wood, and M. Lanthaler. RDF 1.1 Concepts and Abstract Syntax. W3C Recommendation, February 2014. <https://www.w3.org/TR/rdf11-concepts/>.
17. A. Deutsch, N. Francis, A. Green, K. W. Hare, B. Li, L. Libkin, T. Lindaaker, V. Marsault, W. Martens, J. Michels, F. Murlak, S. Plantikow, P. Selmer, O. van Rest, H. Voigt, D. Vrgoc, M. Wu, and F. Zemke. Graph pattern matching in GQL and SQL/PGQ. In *Proc. International Conference on Management of Data (SIGMOD)*, pages 2246–2258, 2022.
18. Javier D. Fernández, Miguel A. Martínez-Prieto, Claudio Gutierrez, Axel Polleres, and Mario Arias. Binary RDF representation for publication and exchange (HDT). *Journal of Web Semantics*, 19:22–41, 2013.
19. N. Francis, A. Gheerbrant, P. Guagliardo, L. Libkin, V. Marsault, W. Martens, F. Murlak, L. Peterfreund, A. Rogova, and D. Vrgoc. A Researcher’s Digest of GQL. In *Proc. 26th International Conference on Database Theory (ICDT)*, pages 1:1–1:22, 2023.
20. N. Francis, A. Green, P. Guagliardo, L. Libkin, T. Lindaaker, V. Marsault, S. Plantikow, M. Rydberg, P. Selmer, and A. Taylor. Cypher: An evolving query language for property graphs. In *Proc. International Conference on Management of Data (SIGMOD)*, pages 1433–1445, 2018.
21. T. Gagie, G. Navarro, and S.J. Puglisi. New algorithms on wavelet trees and applications to information retrieval. *Theoretical Computer Science*, 426-427:25–41, 2012.
22. A. Gómez-Brandón, A. Hogan, and G. Navarro. Uplifting the superpowers of worst-case-optimal join algorithms. arXiv:2608.03840, 2026. <https://arxiv.org/abs/2608.03840>.
23. R. Grossi, A. Gupta, and J. S. Vitter. High-order entropy-compressed text indexes. In *Proc. 14th Symposium on Discrete Algorithms (SODA)*, pages 841–850, 2003.
24. S. Harris, A. Seaborne, and E. Prud’hommeaux. SPARQL 1.1 Query Language. W3C Recommendation, March 2013. <https://www.w3.org/TR/sparql11-query/>.
25. D. Hernández, A. Hogan, C. Riveros, C. Rojas, and E. Zerega. Querying Wikidata: Comparing SPARQL, relational and graph databases. In *Proc. 15th International Semantic Web Conference (ISWC), Part II*, pages 88–103, 2016.
26. A. Hogan, C. Riveros, C. Rojas, and A. Soto. A worst-case optimal join algorithm for SPARQL. In *Proc. 18th International Semantic Web Conference (ISWC)*, pages 258–275, 2019.
27. G. Jin, X. Feng, Z. Chen, C. Liu, and S. Salihoglu. Kùzu graph database management system. In *Proc. 13th Conference on Innovative Data Systems Research (CIDR)*, 2023.
28. Y. Liu, Y. Lin, X. Peng, Y. Li, and J. Zhang. RDF-TDAA: Optimizing RDF indexing and querying with a trie based on directly addressable arrays and a path-based strategy. *Expert Systems and Applications*, 279:article 127384, 2025.
29. V. Mäkinen and G. Navarro. Dynamic entropy-compressed sequences and full-text indexes. *ACM Transactions on Algorithms*, 4(3):article 32, 2008.

30. 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 *Proc. 17th International Semantic Web Conference (ISWC), part II*, pages 376–394, 2018.
31. Memgraph Ltd. Memgraph: An in-memory graph database, 2024.
32. A. Mhedhbi, M. Lissandrini, L. Kuiper, J. Waudby, and G. Szárnyas. LSQB: A large-scale subgraph query benchmark. In *Proc. 4th ACM SIGMOD Joint International Workshop on Graph Data Management Experiences & Systems (GRADES) and Network Data Analytics (NDA)*, pages 1–11, 2021.
33. A. Mhedhbi and S. Salihoglu. Optimizing subgraph queries by combining binary and worst-case optimal joins. *Proc. VLDB Endowment*, 12(11):1692–1704, 2019.
34. J. I. Munro. Tables. In *Proc. 16th Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS)*, pages 37–42, 1996.
35. G. Navarro. Wavelet trees for all. *Journal of Discrete Algorithms*, 25:2–20, 2014.
36. G. Navarro. *Compact Data Structures – A practical approach*. Cambridge University Press, 2016.
37. G. Navarro. Practical adaptive dynamic bitvectors. *Software Practice and Experience*, 55(9):1539–1559, 2025.
38. Y. Nekrich and G. Navarro. Sorted range reporting. In *Proc. 13th Scandinavian Symposium on Algorithmic Theory (SWAT)*, LNCS 7357, pages 271–282, 2012.
39. Neo4j, Inc. Neo4j graph database management system, 2024.
40. T. Neumann and M. J. Freitag. Umbra: A disk-based system with in-memory performance. In *Proc. 10th Conference on Innovative Data Systems Research (CIDR)*, 2020.
41. H. Q. Ngo, E. Porat, C. Ré, and A. Rudra. Worst-case optimal join algorithms. *Journal of the ACM*, 65(3):16:1–16:40, 2018.
42. M. Pătraşcu and M. Thorup. Time-space trade-offs for predecessor search. In *Proc. 38th Annual ACM Symposium on Theory of Computing (STOC)*, pages 232–240, 2006.
43. M. Raasveldt and H. Mühleisen. DuckDB: An embeddable analytical database. In *Proc. ACM International Conference on Management of Data (SIGMOD)*, pages 1981–1984, 2019.
44. R. Taelman, J. Van Herwegen, M. Vander Sande, and R. Verborgh. Comunica: A modular SPARQL query engine for the web. In *Proc. 17th International Semantic Web Conference (ISWC), part II*, pages 239–255, 2018.
45. W. van Leeuwen, T. Mulder, B. Van De Wall, G. Fletcher, and N. Yakovets. Avant-graph query processing engine. *Proc. VLDB Endowment*, 15(12):3698–3701, 2022.
46. T. L. Veldhuizen. Triejoin: A simple, worst-case optimal join algorithm. In *Proc. 17th International Conference on Database Theory (ICDT)*, pages 96–106, 2014.
47. D. Vrandečić and M. Krötzsch. Wikidata: A free collaborative knowledgebase. *Communications of the ACM*, 57(10):78–85, 2014.
48. D. Vrgoc, C. Rojas, R. Angles, M. Arenas, D. Arroyuelo, C. Buil-Aranda, A. Hogan, G. Navarro, C. Riveros, and J. Romero. MillenniumDB: An open-source graph database system. *Data Intelligence*, 5(3):560–610, 2023.
49. J. Webber. A programmatic introduction to Neo4j. In *Proc. ACM Conference on Systems, Programming, and Applications: Software for Humanity (SPLASH)*, pages 217–218. ACM, 2012.