

Standard versus uniform binary search and their variants in learned static indexing: The case of the searching on sorted data benchmarking software platform

Domenico Amato | Giosué Lo Bosco  | Raffaele Giancarlo

Dipartimento di Matematica e Informatica, Università degli Studi di Palermo, Palermo, Italy

Correspondence

Giosué Lo Bosco, Dipartimento di Matematica e Informatica, Università degli Studi di Palermo, Via Archirafi 34, 90123 Palermo, Italy.
Email: giosue.lobosco@unipa.it

Funding information

Italian MUR PRIN, Grant/Award Number: 2017WR7SHH; INdAM-GNCS 2022

Abstract

Learned Indexes use a model to restrict the search of a sorted table to a smaller interval. Typically, a final binary search is done using the **lower_bound** routine of the Standard C++ library. Recent studies have shown that on current processors other search approaches (such as k-ary search) can be more efficient in some applications. Using the **SOSD** learned indexing benchmarking software, we extend these results to show that k-ary search is indeed a better choice when using learned indexes. We highlight how such a choice may be dependent on the computer architecture used, for example, Intel I7 or Apple M1, and provide guidelines for the selection of the Search routine within the learned indexing framework.

KEYWORDS

algorithms with prediction, binary search variants, learned index structures, search on sorted data platform

1 | INTRODUCTION

Learned Static Indexes, introduced by Kraska et al.¹ (but see also Reference 2), with follow-up in References 3-9, are a recent approach to search in a sorted table, quite effective with respect to existing procedures and data structures, for example, B-trees,¹⁰ used in important application domains such as Databases¹¹ and Search Engines.¹² As described in Section 2, such a model may be as simple as a straight line or more complex, with a tree-like structure, as the ones mentioned in Section 3.2.2. It is used to make a prediction regarding where a query element may be in the sorted table. Then, the search is limited to the interval so identified and performed via standard binary search. The use of this routine is more of a natural choice rather than a requirement. In fact, the **lower_bound** routine from the standard C++ library is almost exclusively used.

In order to place our contributions on the proper ground, it is useful to recall that two major studies^{12,13} have recently investigated which binary search routines or variants are better suited to take advantage of modern computer

Abbreviations: SOSD, Search on Sorted Data; SA, Standalone Environment; S-BS, Standard Binary Search; U-BS, Uniform Binary Search; S-KS, Standard k-ary Search; U-KS, Uniform k-ary Search; U-EL, Uniform Eytzinger Layout Search; RMI, Recursive Model Index; PGM, Piece-wise Geometric Model; RS, Radix Spline Index; Amzn32, Amazon Dataset with 32 bits integers; Amzn64, Amazon Dataset with 64 bits integers; Face, Facebook UserID Dataset; Osm, Open Street Map Dataset; Wiki, Wikipedia Books ID Dataset.

This is an open access article under the terms of the [Creative Commons Attribution-NonCommercial-NoDerivs](https://creativecommons.org/licenses/by-nc-nd/4.0/) License, which permits use and distribution in any medium, provided the original work is properly cited, the use is non-commercial and no modifications or adaptations are made.

© 2022 The Authors. *Software: Practice and Experience* published by John Wiley & Sons Ltd.

architectures. Those experimental findings hold in the *stand alone* scenario, that is, when no prediction to reduce the search interval is performed, and they provide useful indications on which routine to use in which circumstances. However, to what extent the recommendations coming out of those studies actually hold also for learned indexes has not been investigated. Which version of binary search to use, and when, is unresolved for learned indexes, relying on the natural choice mentioned earlier.

With the use of the **SOSD**¹⁴ benchmarking software platform, we address such a question by experimenting with various binary search routines on both synthetic and real datasets, with executions on two different architectures, that is, Intel I7/9 and Apple M1. Our results further validate and extend the ones provided in References 12 and 13 for the *stand alone* scenario and provide novel indications on when to use **SOSD** for learned indexing executions and with which binary search routine. Our findings are a significant advance with respect to the research performed on learned indexes outlined above. For completeness, we mention that our results hold for the static case of learned indexes, when no insertions or deletions are allowed. For the dynamic case, learned indexing solutions exist.^{5,15} However, how to state a research analogous to ours in that setting is open, to the best of our knowledge.

In order to make our experiments replicable, the software we developed or modified is available at References 16 and 17, while the datasets are available at Reference 18.

2 | A SIMPLE VIEW OF LEARNED SEARCHING IN SORTED SETS

Consider a sorted table A of n keys, taken from a universe U . It is well known that sorted table search can be phrased as the Predecessor Search Problem: for a given query element x , return the $A[j]$ such that $A[j] \leq x < A[j + 1]$. Kraska et al.¹ have proposed an approach that transforms such a problem into a learning-prediction one. With reference to Figure 1, the model learned from the data is used as a predictor of where a query element may be in the table. Binary search is then performed only on the interval returned by the model.

We now outline the basic technique that one can use to build a model for A . It relies on linear regression, with mean square error minimization.²⁰ Consider the mapping of elements in the table to their relative position within the table. Since such a function is reminiscent of the cumulative distribution function over the universe U of elements from which the ones in the table are drawn, as pointed out by Marcus et al.¹⁹ in their benchmarking study on learned indexes, we refer to it as CDF. With reference to the example in Figure 2, and assuming that one wants a linear model, that is, $F(x) = ax + b$, Kraska et al.¹ note that they can fit a straight line to the CDF and then use it to predict where a point x may fall in terms of position and accounting also for approximation errors. In order to perform a query, the model is consulted and an interval in which to search is returned. Then, binary search on that interval is performed. Different models may use different schemes to determine the required range, as outlined in Section 3.2.2. The reader interested in a rigorous presentation of those ideas can consult Marcus et al.¹⁹

For this research, it is important to know how much of the table is discarded once the model makes a prediction on a query element. For instance, binary search, after the first test, discards 50% of the table. Because of the diversity across models to determine the search interval, and in order to place all models on a par, we estimate the reduction factor of a model, that is, the percentage of the table that is no longer considered for searching after a prediction, empirically. That is, with the use of the model and over a batch of queries, we determine the length of the interval to search into for each query.

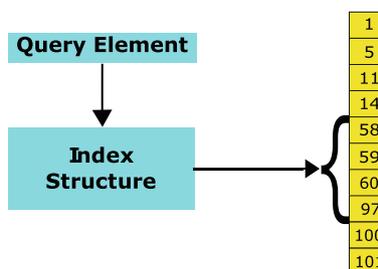


FIGURE 1 A general paradigm of learned searching in a sorted set.¹⁹ The model is trained on the data in the table. Then, given a query element, it is used to predict the interval in the table where to search (included in brackets in the figure).

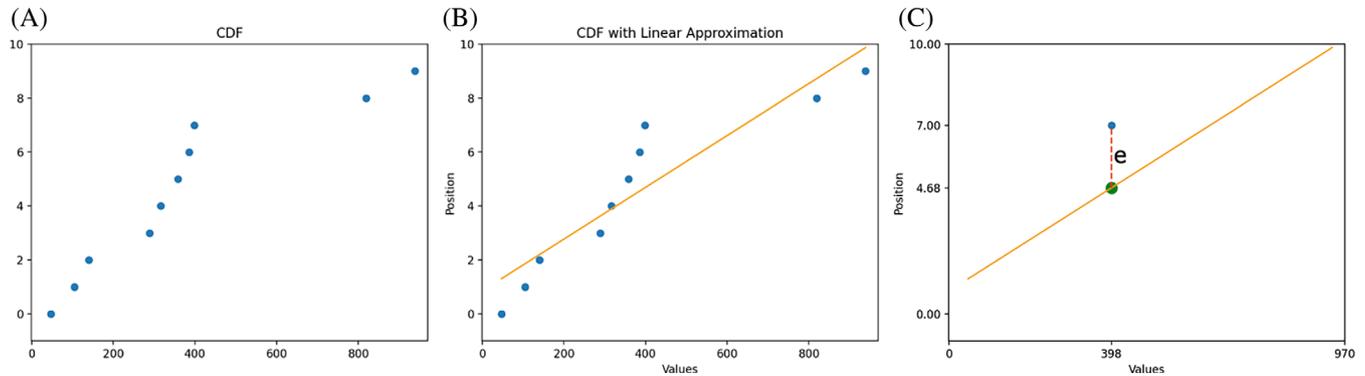


FIGURE 2 The process of learning a simple model via linear regression. Let A be $[47,105, 140,289, 316,358, 386,398, 819,939]$. (A) The CDF of A . In the diagram, the x -axis indicates the value of an element in the table, while the y -axis is its position. (B) The straight line $F(x) = ax + b$ is obtained by determining a and b via linear regression, with mean square error minimization. (C) The maximum error ϵ one can incur in using F is also important. In this case, it is $\epsilon = 3$, that is, accounting for rounding, it is the maximum distance between the position of a point in the table and its position as predicted by F . In this case, the interval to search into, for a given query element x , is given by $[F(x) - \epsilon, F(x) + \epsilon]$.

Based on it, it is possible to compute the reduction factor for that query. Then, we take the average of those reduction factors over the entire set of queries as the reduction factor of the model for the given table.

3 | EXPERIMENTAL METHODOLOGY

Our experimental set-up follows closely the one outlined in the already mentioned benchmarking study by Marcus et al.¹⁹ regarding learned indexes, with some variations. In particular, here we concentrate on the study of how different kinds of Binary and k -ary Searches can affect the performance of learned indexes. Moreover, following Amato et al.,^{3,4} we use datasets of varying sizes in order to understand how the data structures perform on the different levels of the internal memory hierarchy. In addition to that, we also use datasets generated as in Reference 12, in order to establish that the binary search routines we use behave consistently with the findings in the mentioned paper. Details of the entire methodology are provided next.

3.1 | Computer architectures and compilers

All the experiments have been performed on two different architectures, that is, x86 and ARM, using three different CPUs: Intel i7-8700, Intel I9-10850, and Apple M1. In the following, the specifications of the three systems used are reported.

- The i7-8700 works with a 3.2 GHz clock and uses 64 kb of L1 cache per core, 256 kb of L2 cache per core, and 12 Mb of shared L3 cache. The amount of system memory is 32 Gbyte of DDR4. The OS is Ubuntu LTS 20.04.
- The I9-10850 works with a 3.6 GHz clock and uses 64 kb of L1 cache per core, 256 kb of L2 cache per core, and 20 Mb of shared L3 cache. The amount of system memory is 32 Gbyte of DDR4. The OS is Ubuntu LTS 20.04.
- The Apple M1 works with a 3.2 Ghz clock and uses two levels of cache (L1 and L2). The amount of L1 memory cache depends on the cores, which are of two kinds. Namely, high-performance and high-efficiency cores. Moreover, in the Apple M1 architecture, the L1 memory cache is divided into a part for instructions and one for data. In our study, only one core is used, which is of the high-performance kind. As consequence, the L1 cache size is 192 KB for instruction and 128 KB for data. The L2 cache is 12 Mb. The amount of system memory is 8 Gbyte DDR4. The OS is macOS Monterey 12.3.1.

The adopted compiler is the same for all the operating systems we have used, that is, GCC compiler with optimization flag $-O3$. In order to better explain the results obtained, we use a hardware profiler on the Intel architecture, that

is, *Intel Vtune*, which makes possible the extraction of several performance parameters of interest. They are detailed in Section C.1 of the Appendix. The profiler is provided free of charge by Intel. Somewhat unfortunately, as far as the M1 architecture is concerned, we are not aware that an analogous profiler is available. More in general, for the ARM architectures such as Neon, there is a profiler that may give useful information (*Arm Forge Ultimate*), but it is not free of charge. Profilers such as *gperf-tool-profiler* do not give useful information. Based on such a State of the Art, for this research, no profiling for the Apple M1 architecture is performed. In what follows, we report results on the Intel I7 and Apple M1 architectures, while, for conciseness, results on the Intel I9 are omitted when analogous to the I7.

3.2 | Algorithms, code, and software platforms

3.2.1 | Binary search and its variants

We use the standard, that is, textbook, binary search routine,^{21,22} reported in Algorithm 1 and referred to as **S-BS**. We also use another version of the binary search strategy, presented in Reference 21 (see also References 12 and 13) under the name Uniform. The corresponding code is provided in Algorithm 2 and it is referred to as **U-BS**. The **U-BS** routine differentiates itself from the Standard one because there is no test for exit within the main loop. We point out that in all the algorithms we present unless otherwise stated, we include prefetching instructions since their use may be of advantage in terms of execution time.¹² However, whether or not to use them in our research is a fact that needs evaluation. Moreover, depending on the compiler and the architecture, instruction 8 of Algorithm 2 may be translated into a predicated instruction, which is not a branching instruction. Indeed, the main loop of the resulting assembly code reported in Code 1 has no branches. This should be contrasted with the assembly code generated on the I7 for Algorithm 1, which is reported in Code 2 of the Appendix. Such a branch-freeness, as discussed in Reference 12, results in better use of the processor instruction pipeline.

We include in this study also the `lower_bound` routine from the standard C++ library, which is equivalent to **U-BS** in terms of source code (see Algorithm 3 in the Appendix), since, in the **SOSD** software platform, this routine is referred to as `branchy binary search`. More precisely, it is `Uniform and Branchy`.

We also take into consideration `k-ary Search`²³ routines, indicated as **S-KS** and **U-KS** (Algorithms 4 and 5 in the Appendix) in their standard and uniform versions respectively, using $k = 3$ as recommended in the work by.¹³ The reason is that `k-ary Search` is superior to **U-BS** and `lower_bound` on modern computer architectures, according to results in Reference 13.

Finally, we include also the Eytzinger Layout binary search routine (**U-EL** for short, see Algorithm 6 in the Appendix) because, although not directly usable within the Leaned Indexing framework, it is a useful baseline to compare against: it is superior to **U-BS** and **S-BS**.¹² In addition to that, it has not been included by Schulz et al.¹³ in their benchmarking experiments.

Algorithm 1. C++ implementation of standard binary search with prefetching. The version without prefetching is obtained by deleting lines 4–5

```

1: int StandardBinarySearch(int *A, int x, int left, int right){
2:   while (left < right) {
3:     int m = (left + right) / 2
4:     __builtin_prefetch(&(data[lo + m / 2]), 0, 0);
5:     __builtin_prefetch(&(data[m + m / 2]), 0, 0);
6:     if(x < A[m]) right = m;
7:     else if( x > A[m]) left = m+1;
8:     else return m;
9:   }
10:  return right;
11: }
```

Algorithm 2. C++ Implementation of Uniform binary search with prefetching. The code is as in Reference 12 (see also References 13 and 21). The version without prefetching is obtained by deleting lines 6–7

```

1: int prefetchUniformBinarySearch(int *A, int x, int left, int right){
2:   const int *base = A;
3:   int n = right;
4:   while (n > 1) {
5:     const int half = n / 2;
6:     __builtin_prefetch(base + half/2, 0, 0);
7:     __builtin_prefetch(base + half + half/2, 0, 0);
8:     base = (base[half] < x) ? &base[half] : base;
9:     n -= half;
10:  }
11:  return (*base < x) + base - A;
12: }

```

TABLE 1 Branchy and branch-free assembly code production

	Intel I7	Intel I9	M1
S-BS	Branchy	Branchy	Branchy
U-BS	Branch-free	Branch-free	Branch-free
lower_bound	Branchy	Branchy	Branch-free

Note: The first row indicates the computer architecture, while the first column the routines. For each entry, we report the kind of code produced by the compiler, that is, Branchy or Branch-free code.

3.2.2 | Index model classes in SOSD

From the many models available in **SOSD**, we choose the ones that have been the most successful among the ones benchmarked in Reference 19. That is, the Recursive Model Index (**RMI**, for short),¹ the Radix Spline (**RS**, for short)⁹ and the Piecewise Geometric Model (**PGM**, for short).^{5,7} For the convenience of the reader, a brief outline of each of those indexes is provided in Section A.2 of the Appendix. We point out that we have modified the **SOSD** library so that an implementation of a learned index can use one of **U-BS**, **S-BS**, **U-KS**, or **S-KS**, for the final search stage. It is to be noted that each of those models can err in making a prediction. However, each of them has a mechanism to correct for such a mistake in order to return a valid interval in which to search into. The interested reader is referred to the original papers for a description of those mechanisms, which are somewhat more complex than the one we have considered in Figure 2. It is also worth recalling that those models can only use binary search procedures with a sorted table layout, that is, **U-EL** cannot be used with those models.

3.3 | Computer architecture and compilers: The production of branch-free code

Given the binary search routines and their variants described in Section 3.2.1, it is not clear that branch-free assembly code is actually produced by the compiler. Therefore, we have inspected the assembly code generated by the compiler in each of the used architectures. The results are summarized in Table 1. For conciseness, we report only the branch-free assembly code in Section A.3 of the Appendix, in addition to Code 2. The remaining code regarding Branchy binary search, rather lengthy, is available upon request. Furthermore, it is to be noted that the table does not report the case for the k-ary Searches, because the program to extract the assembly code (Linux obj-dump) does not provide specific details capable to determine the presence of predicated instructions.

Interestingly, the **lower_bound** routine is translated into branchy code on the Intel architectures and in branch-free code on the Apple M1. Such a difference may be explained as follows. Although **lower_bound** is a Uniform routine, its inner loop makes explicit use of if-then-else (lines 7-11 of Algorithm 3 in the Appendix) rather than a conditional operator

Code 1. Assembly code of uniform binary search on the Intel I7 (only main loop). The predicated instruction is line 1485 (in bold). No prefetching is used

```
0000000000001460 <_Z21Uniform_Binary_SearchPmmmm>:
.
.
1473: jbe 1492 <_Z21Uniform_Binary_SearchPmmmm+0x32>
1475: nopl (%rax)
1478: mov %rcx,%rdx
147b: shr %rdx
147e: lea (%rax,%rdx,8),%r8
1482: cmp (%r8),%rsi
1485: cmovae %r8,%rax
1489: sub %rdx,%rcx
148c: cmp $0x1,%rcx
1490: ja 1478 <_Z21Uniform_Binary_SearchPmmmm+0x18>
.
.
```

(line 8 in Algorithm 2). This accounts for the difference between the two assembly codes on the Intel architectures. As for the Apple M1, being an ARM architecture, it has an extensive set of predicated instructions,¹² which apparently the compiler is able to use even in the presence of simple if-then-else constructs in the high-level code.

3.4 | Datasets and index model training

We have used two kinds of datasets. The first kind was generated as described in Reference 12, that is, we generated 24 synthetic datasets with an increasing number n of elements from 2^4 to 2^{28} , containing only odd 64-bit integers in $[1, 2n + 1]$. For each of these, we generated a 2 million element query file consisting of 1 million odd elements which are present in the dataset and 1 million even elements that are missing.

The second kind of datasets have origin from the carefully chosen ones in Reference 19 (and therein referred to as **amzn32**, **amzn64**, **face**, **osm**, **wiki**). They have been derived from them in References 3 and 4, in order to fit well each level of the main memory hierarchy with respect to the Intel I7 architecture. The essential point of the derivation is that, for each of the generated datasets, the CDF of the corresponding original dataset is well approximated. The details are as follows, where n is the number of elements in a table.

- **Fitting in L1 cache: cache size 64 Kb.** Therefore, we choose $n = 3.7$ K, that is, about 30 Kb. For each dataset, the table corresponding to this type is denoted with the suffix **L1**, for example, **amzn32-L1**, when needed.
- **Fitting in L2 cache: cache size 256 Kb.** Therefore, we choose $n = 31.5$ K, that is, about 252 Kb. For each dataset, the table corresponding to this type is denoted with the suffix **L2**, when needed.
- **Fitting in L3 cache: cache size 8 Mb.** Therefore, we choose $n = 750$ K, that is, about 6 Mb. For each dataset, the table corresponding to this type is denoted with the suffix **L3**, when needed.
- **Fitting in PC Main Memory: memory size 32Gb.** Therefore, we choose $n = 200$ M, that is, the entire dataset of about 1.6 Gb and about 800 Mb in the case of 64 bit and 32 bit respectively. For each dataset, the table corresponding to this type is denoted with the suffix **L4**.

As for query dataset generation, for each of the tables built as described above, we extract uniformly and at random (with replacement) from the Universe U a total of two million elements, 50% of which are present and 50% absent, in each table.

For the case of the Apple M1 architecture, we take into consideration only datasets fitting at most into the L3 cache of the Intel I7, because our system is equipped with 8 Gbyte of main memory. Indeed, when the code used in this research is executed on the Apple M1 using as input each of the full datasets, the performance degrades due to substantial swapping. As for model training, **SOSD** has ten predefined Models each for the **PGM** and **RS**. For the **RMI** Model family, following the Literature, we use **CDFShop**,¹⁹ which returns up to ten versions of an **RMI** for a given table. The selection process is heuristic and tries to choose good models in terms of query time that use little space. The interested reader can find details in Reference 19

All the experiments involving the mentioned datasets are reported in full in the Appendix and in part here. The query time that we report is an average taken on a batch of two million queries executed by a search routine or a learned index. This is essential for learned indexes: a measure of a single query performance would be unreliable,²⁴ while the method we choose is compliant with the Literature.¹⁹ Such a limitation makes it unreliable to measure some relevant performance parameters of a learned index, as for instance, for each query, the amount of time spent for prediction and the amount of time spent for searching. In fact, to the best of our knowledge, none of the papers reporting on learned indexing provides such a breakdown. Rather they concentrate on the accuracy of a prediction. In terms of theoretic worst-case analysis, the prediction for the **RMIs** used here takes $O(1)$ time and $O(\log n)$ time for the **PGM** and the **RS**.

4 | EXPERIMENTS: SEARCHING IN CONSTANT ADDITIONAL SPACE, WITH OR WITHOUT SOSD

The aim of this Section is to shed light on the consistency of our experimental setting with the current Literature, that is, Khuong and Morin¹² and Schultz et al..¹³ However, the results reported here provide also useful indications regarding the use of **SOSD** with binary search routines only. This scenario is meaningful since those routines require only constant additional space with respect to the table to be searched into, while the learned indexes may require additional space that depends on the parameters of the model. It may be a small percentage or a really large one. The interested reader can find a study in References 3 and 4. The scenario we consider is the one in which one can only use constant additional space with respect to the input table. To the best of our knowledge, it is not clear whether **SOSD** is worth using and with which routine. As a further result, we get also indications on how to set up the search routines in **SOSD**, when learned indexing is to be used.

4.1 | Replication of the experiments by Khuong and Morin

According to a study by Khuong and Morin,¹² modern processor architectures are best used with branch-free binary search code. In order to assess to what extent those findings hold in our experimental set-up, we have experimented with all the routines mentioned in Section 3.2.1, that have been executed as *stand-alone* C++ code (as in Reference 12) and as code included into the highly engineered **SOSD** platform. We have considered all the architectures mentioned in Section 3.1, the synthetic datasets described in Section 3.4, with the inclusion also of the **osm** dataset for completeness. No prefetching is used here since the advantage of its use is discussed separately within this Section.

In regard to the **lower_bound** routine, which is the standard within **SOSD** and the associated benchmarking of learned indexes, we have compared it with all the other routines, finding that it is inferior to all of them. The full set of experiments is available upon request and, for brevity, we only report some interesting observations regarding this important routine as compared with **S-BS** and **U-BS**. On the Intel processors, despite being a Uniform binary search, its assembly code is branchy (code omitted for brevity and available upon request). The experiments reported in Figure C1 of the Appendix point to the fact that, on those architectures, there is very little difference with **S-BS**. On the Apple M1 architecture, its assembly code is branch-free (Code 5 in the Appendix) but, as shown in Figure C2 of that File, both **S-BS** and **U-BS** are better. For those reasons, we no longer consider **lower_bound** in this study.

For the remaining routines, the results are reported in Figure 3 on synthetic datasets and in Figure C4 of the Appendix for the **osm** dataset. We also anticipate that, from the experiments reported and discussed in this Section, there is no substantial difference in performance between **S-KS** and **U-KS**. For this reason, **U-KS** is not considered in the following Sections.

4.1.1 | Stand-alone

In this setting, on both the Intel and the ARM architectures, the results in Reference 12 are confirmed (see Figure 3B,D for synthetic datasets and Figure C4B,D of the Appendix for the **osm** datasets). That is, Uniform branch-free is better than Standard branchy for datasets fitting in the cache memory. It is useful to recall that, for the ARM architecture, we use only datasets that fit in the cache memory. Moreover, there is also confirmation of the findings in Reference 13, stating that k-ary Search is better than both of those routines for datasets fitting in main memory. A new finding, overlooked both in Reference 12 and in Reference 13, is that **U-EL** is always better than k-ary Search, on both architectures we have considered.

4.1.2 | SOSD

In the Intel architectures setting, we find the same results as in the *stand-alone* case, with some notable differences: (a) for tables fitting in the cache memory, for example, of size in the interval $[1, 2^{22}]$, the gap between the performance of **U-BS** and **S-BS** is reduced and virtually unnoticeable for small tables, as shown by the two curves in Figure 3A,B;

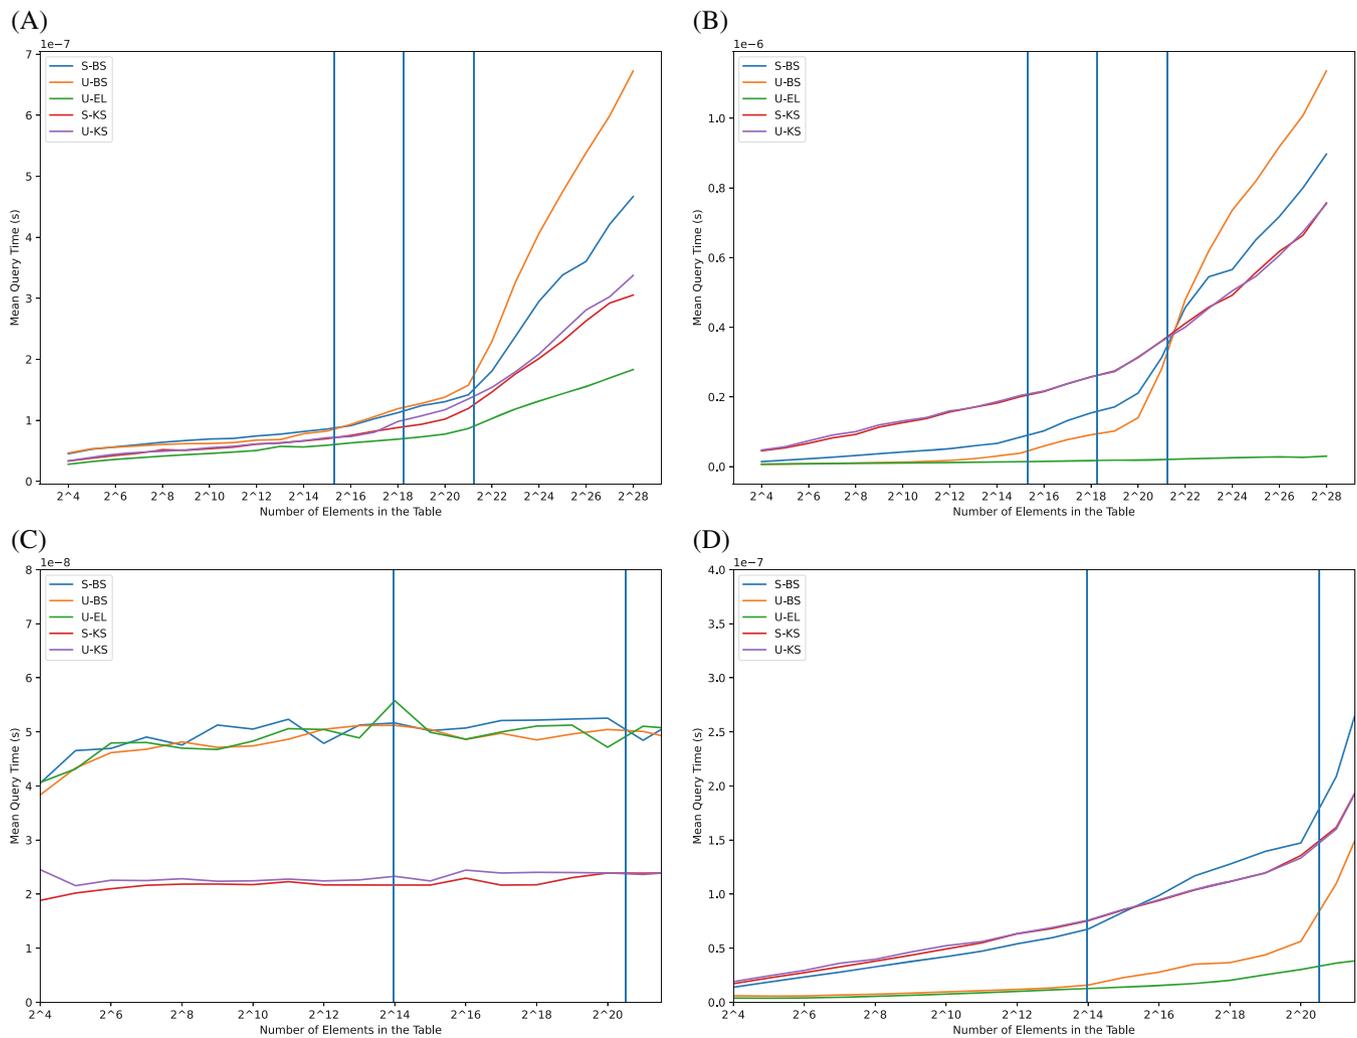


FIGURE 3 Mean query times of search methods on synthetic data on the Intel I7 and the Apple M1. (A,C) report results using **SOSD** on the Intel I7 and the Apple M1, respectively. (B,D) reports *Stand-Alone* implementation for the same architectures. The x-axis reports number of elements in the input Table, while the y-axis the mean query time in seconds. The vertical lines indicate the size of each cache memory level.

(b) **S-KS** and **U-KS** are always better than **S-BS** and **U-BS**. In regard to the **osm** datasets, the same results are confirmed (see Figure C4A,C of the Appendix). Concerning the Apple M1 architecture and at least for sizes that fit in L3, and in reference to Figure 3C, we find that **S-BS**, **U-BS**, and **U-EL** perform analogously, while **S-KS** and **U-KS** are always better.

4.1.3 | *Stand-alone* versus SOSD

It is also of interest to assess whether there are differences in terms of query execution time between the *stand-alone* and the **SOSD** case. With reference to Figure C3 of the Appendix for the Intel I7, it is to be noted that the **S-BS** and **U-BS** routines perform better in their *stand-alone* version with datasets fitting the memory cache. Moreover, in regard to **U-EL**, the *stand-alone* version is always better than the **SOSD** counterpart. Finally, the **S-KS** and **U-KS** perform always better in the **SOSD** case.

4.1.4 | Profiler analysis

In order to gain insights into the differences outlined above, we make use of the *Intel Vtune* profiler on two synthetic datasets, that is, a small table (size 2^{10} elements) and a large one (size 2^{24} elements). The results are reported in Tables C1 and C2 in the Appendix. Although the profiling does not give a definite indication regarding the superiority of one setting over the other, there are a few facts that are worth noting.

Concerning the **U-EL**, all the profiler parameters get worse within **SOSD**, except for the Front-end one. In particular, the Bad Speculation parameter (which indicates wrong predictions such as in the case of branches), gets much worse in going from *stand-alone* to **SOSD**. Interestingly, the same thing happens with **U-BS**. On the other hand, such a parameter decreases for **S-BS**. Although no definitive conclusion can be drawn, those facts indicate that the optimizations within **SOSD** seem to be oriented towards branchy code.

4.2 | Pros and cons of prefetching

As pointed out in Reference 12, explicit prefetching can improve performance by loading blocks of data into the cache memory before they are accessed, avoiding processor stalls. However, this operation is expensive. Therefore, the advantage of using it must be carefully evaluated. To this end, using the same set-up of Section 4.1, we have studied whether explicit prefetching can improve the performance of the search routines included in this research.

The results are summarized in Figure 4 for the Intel I7 architecture and Figure 5 for the Apple M1. We report only results for the execution within **SOSD** because in the *stand-alone* case they are analogous.

The findings of Reference 12 for the use of explicit prefetching are confirmed on the Intel I7 architectures, that is, it is never useful for **S-BS** (Figure 4A), it is useful for **U-BF** only for the on big-sized datasets (Figure 4B), and it is always useful for **U-EL** (Figure 4C). In regard to the Apple M1 architecture, prefetching is not useful across any search method. Presumably, this is due to the fact that we only use datasets fitting the cache memory.

As for the Intel architecture, in order to get insights into those findings, we have again profiled the code, as in the previous Section. The results are reported in Section C.2 of the Appendix. With reference to them, it is to be noted that the profiler parameters confirm, for datasets larger than the cache memory, that the number of stalls due to data cache misses decreases with the use of explicit prefetching.

4.3 | A summary of useful indications

Based on the experiments reported so far, we get some useful indications regarding the use of binary search routines that account also for the **SOSD** platform.

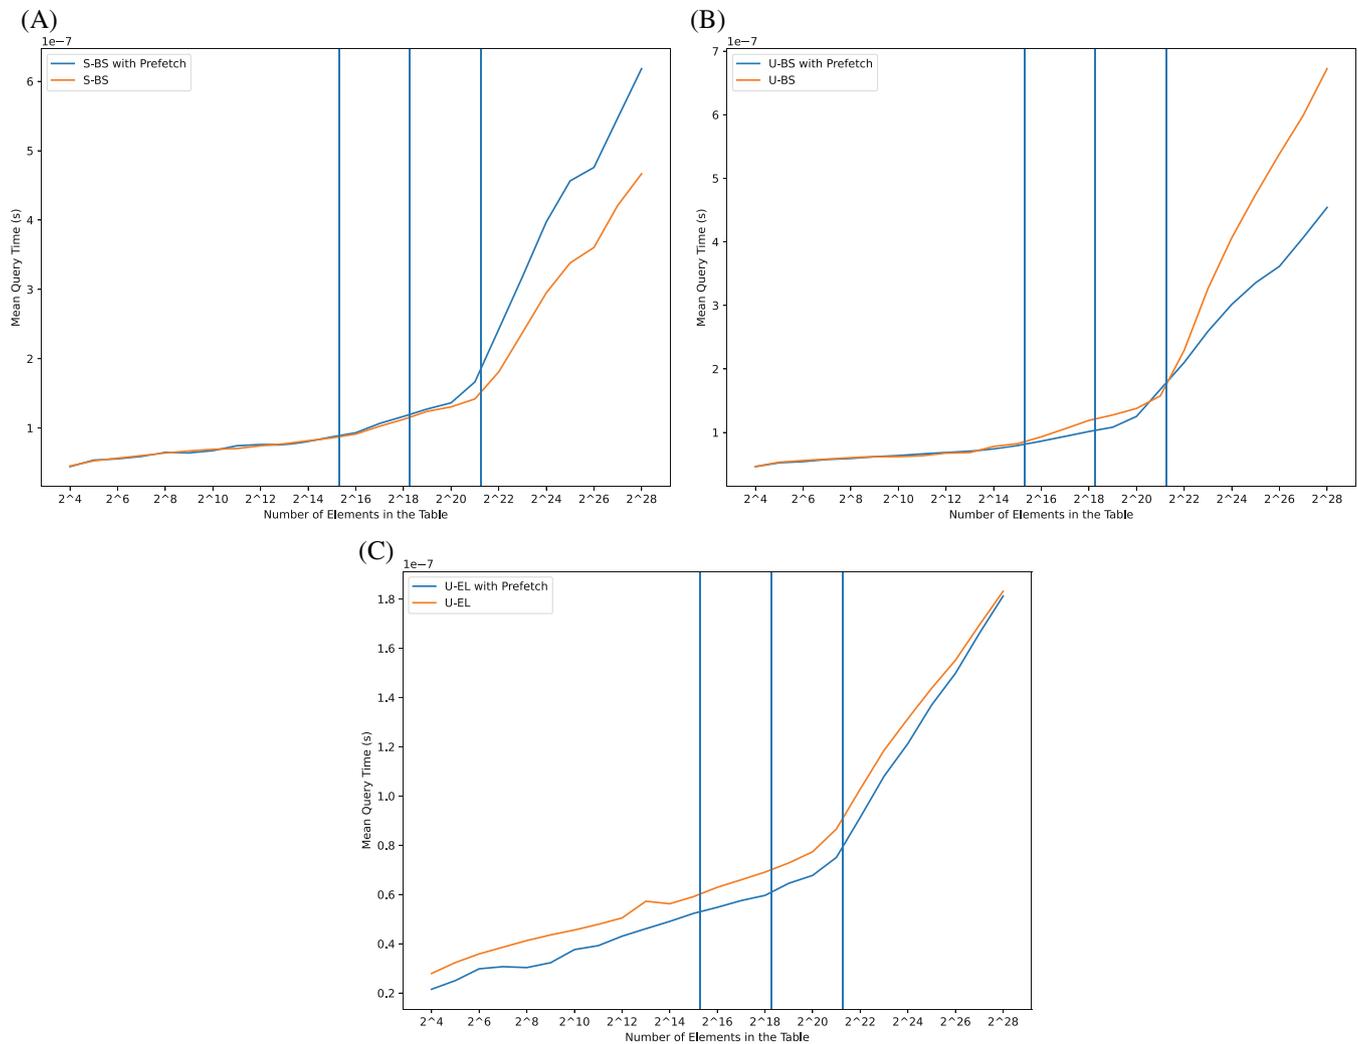


FIGURE 4 Comparison between routines with and without explicit prefetching on SOSD using the Intel I7. The figures show the comparison between routines with and without explicit prefetching. In particular, we report **S-BS** in (A), **U-BS** in (B) and **U-EL** in (C). On the x-axis, we report the number of elements in the table and, on the y-axis, the mean query time in seconds. The vertical lines indicate the size of each cache memory level.

- **No learned indexing.** This scenario indicated the constraint of constant additional space with respect to the table to be searched into. The indication is to use **U-EL** as a *stand alone* routine, with prefetching on the Intel architectures and without it on the M1 architecture. For this latter, the table must fit in the cache memory, to avoid swapping.
- **learned indexing.** With the exclusion of **U-EL** that cannot be used by the current models, the experiments conducted so far provide the indication that prefetching is not really needed for routines that complete the search of a learned index. As far as the Apple M1 architecture is concerned, the fact that prefetching is of no advantage is evident, since we use only dataset fitting in the cache memory. As for the Intel architectures, the scenario is more complex. Indeed, **S-BS** never benefits from explicit prefetching, while **U-BS** achieves a substantial improvement with its use for datasets larger than the cache memory, that is datasets with a number of elements greater than about $4.19e6$ (see Figure 4B again). That is, the predicted interval of a model must be of length of at least $4.19e6$, in order to consider prefetching in the final search routine of that model. Now, it is worth recalling that, for models that are effective, large reduction factors are expected that, in turn, correspond to small tables to be searched into. Unless the original table is really large or the model particularly bad, the case of predicted interval lengths where prefetching can be useful with **U-BS** seems unlikely, as the experiment reported in Table 2 seems to indicate. As evident from that table, only in the case of the worst **RMI** on the **osm-L4** dataset, there may be some marginal benefit in using prefetching with **U-BS**. Therefore, in what follows, explicit prefetching is not used.

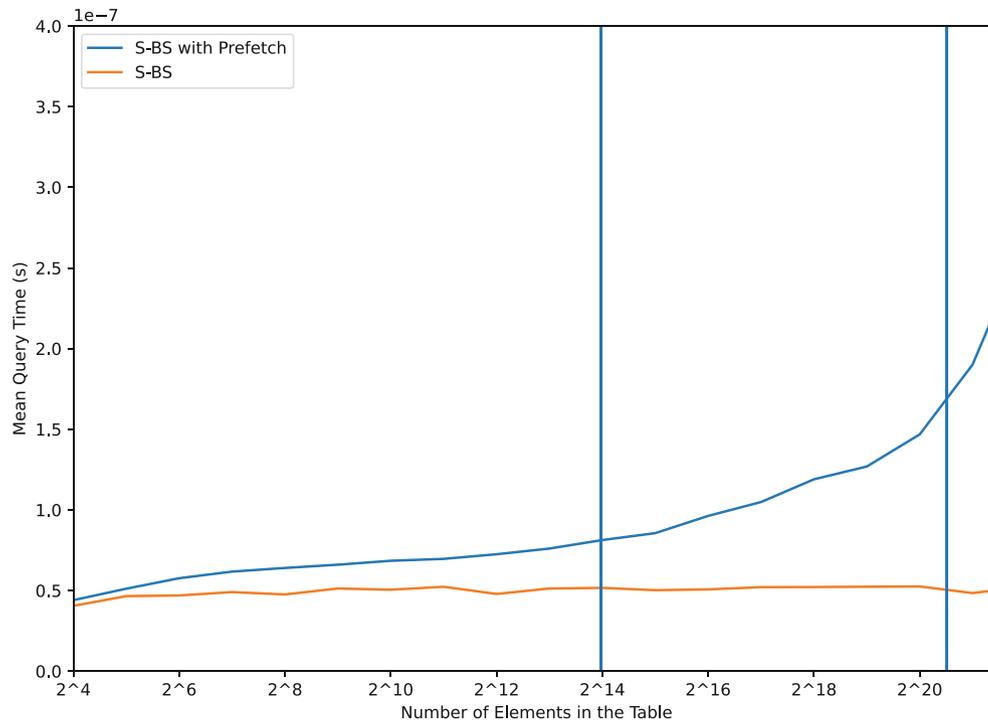


FIGURE 5 Comparison between S-BS with and without explicit prefetching on SOSD using the Apple M1. The figure shows the comparison between S-BS with and without explicit prefetching. The results for U-BS and U-EL are the same. On the x-axis, we report the number of elements in the table and, on the y-axis, the mean query time in seconds. The vertical lines indicate the size of each cache memory level.

TABLE 2 Worst-case average predicted search intervals length for L4 datasets

	RMI	PGM	RS
amzn32	$3.44e5 \pm 1.18e6$	$4.10e3 \pm 2.37e3$	$2.71e2 \pm 1.58e2$
amzn64	$6.19e4 \pm 5.88e5$	$2.05e3 \pm 2.37e3$	$1.35e2 \pm 1.57e2$
face	$8.59e4 \pm 4.19e5$	$4.10e3 \pm 2.37e3$	$5.31e2 \pm 3.08e2$
osm	$5.26e6 \pm 2.73e6$	$2.05e3 \pm 2.37e3$	$3.65e2 \pm 4.23e2$
wiki	$4.25e5 \pm 2.61e6$	$2.05e3 \pm 2.37e3$	$2.50e2 \pm 2.90e2$

Note: For each of the largest of the benchmark datasets and each model class, we have considered the ten model instances used in SOSD. For each, we have computed the average predicted interval length and its standard deviation over a query dataset obtained as described in Section 3.4. Then, we have taken the maximum average, which we report in the table for each dataset and model class, together with its standard deviation.

5 | EXPERIMENTS: SEARCHING USING LEARNED INDEXES, WITH OR WITHOUT SOSD

In order to better describe the experimental work presented in this Section, it is useful to recall that SOSD has been designed to provide an environment in which to evaluate the relative merits of existing and possibly future, learned index Models. Given a dataset, how models are trained is briefly described in Section 3.4. Although originally designed for benchmarking, SOSD can also be used to identify among the models it has available for a given dataset, the best performing one. The list of those Models is reported in Table D1 of the Appendix. Here we focus on average query time, although space may be of importance also.^{3,4} For a given dataset, this amounts to up to thirty different Models to choose from. The routine used for the final search stage is, by default, the **lower_bound** routine. Here we consider three search routines for the final stage, that is, U-BS, S-BS, and S-KS. As mentioned earlier, the **lower_bound** routine has been excluded since it is redundant with respect to the other ones. In summary, for a given dataset, one has up to ninety possible

Model configurations to choose from. Once such a Model configuration has been identified, and in view of the results reported in Section 4.1, it is not clear whether the actual deployment of the selected model in an application domain, for example, Databases¹¹ or Search Engines,¹² it should be executed within the **SOSD** platform or as a *stand-alone* software program. As a matter of fact, such a point has not been addressed in the Literature. The main goal of this part of our experimental work is to investigate the mentioned aspects, which also provide useful indications on which search routine to use.

5.1 | Selecting a best model and associated search routine via SOSD

For each dataset described in Section 3.4, we consider up to ninety possible model configurations. Then, for each model class and search routine so obtained, the Model with the best mean query time is chosen, based on its execution within **SOSD**. The results are reported in Figure 6A,C only for the **osm** dataset, while the remaining ones are reported in Figures D1–D4 of the Appendix. From those figures, we can extract the following findings regarding the search routines, for which we also provide a justification.

- **S-KS is the best.** It is self-evident from the results reported in the mentioned figures. Quite remarkably, they are consistent across all datasets, memory levels and architectures we have considered. This is a novel finding in this area, since only the **lower_bound** or the **S-BS** routine have been used for the terminal stage of searching in a learned index, as well documented in experimental studies prior to this one (see References 1 and 5). It is worth noting that the results reported here are coherent with the ones of the previous Section, in which we have evaluated **S-KS** as a generic binary search routine rather than as a terminal to a learned index. In order to explain such a coherence, we need to point out that the models resulted to be the best with **S-KS** provide a quite small predicted interval, on average: for instance on the **osm_L4** dataset and for the **RMI** model, the average predicted interval length is $4.41e + 2$ with a standard deviation of $2.73e + 3$. The full data regarding this point is not shown and is available upon request. For small tables, as far as the Apple M1 architecture is concerned and when the search routines are performed within **SOSD** without a prediction phase, Figures 3B and C4B of the Appendix report a wide margin in average query time between **S-KS** and the other routines, which apparently turns out to be preserved also in the learned indexing framework executed within **SOSD**. As for the Intel architectures, to better highlight the performance of such routines on small tables, we need to “zoom-in” in Figures 3A and C4A of the Appendix. The corresponding Figure C5A,B in the Appendix. It is evident that also in this case **S-KS** is always better than the other routines on small tables. It is worth noting that, compared to the results on the Apple M1, the margin in average query time between **S-KS** and the other routines is smaller, which is reflected in a smaller margin also in the learned indexing framework, executed within **SOSD**. A more refined analysis regarding such a coherence behaviour within **SOSD** would require individually measuring the prediction and search time for each query: as already mentioned in Section 3.4, this would yield unreliable results.
- **U-BS and S-BS.** In contrast to the case of **S-KS**, the results for **S-BS** and **U-BS** indicate that there is no clear winner between those two procedures. Indeed, which procedures to pick seems to depend, for both architectures, on the memory level in which the input table fits, the input table itself and the model we are using. Those latter results do not contradict the ones described in Section 4.1. In fact, as already mentioned in the discussion regarding **S-KS**, the best Models have a quite small average predicted interval length and, for small tables, the performance of **S-BS** and **U-BS** within **SOSD** is quite close.

5.2 | Convenience of the execution of the best model within SOSD or stand-alone

For each experiment described above, we select the best performing model. As anticipated, we execute each of those models in a *stand-alone* setting, with the routines **S-BS**, **U-BS**, and **S-KS** as terminal for the search stage. The results are reported in Figure 7. It is evident that the convenience of using **SOSD** with the selected model rather than in a *stand-alone* setting is architecture-dependent. In particular, the experiments indicate that it is advisable to use **SOSD**, with **S-KS** as a terminal for the search stage, on the Apple M1 architecture. As for the Intel I7 architecture, we find

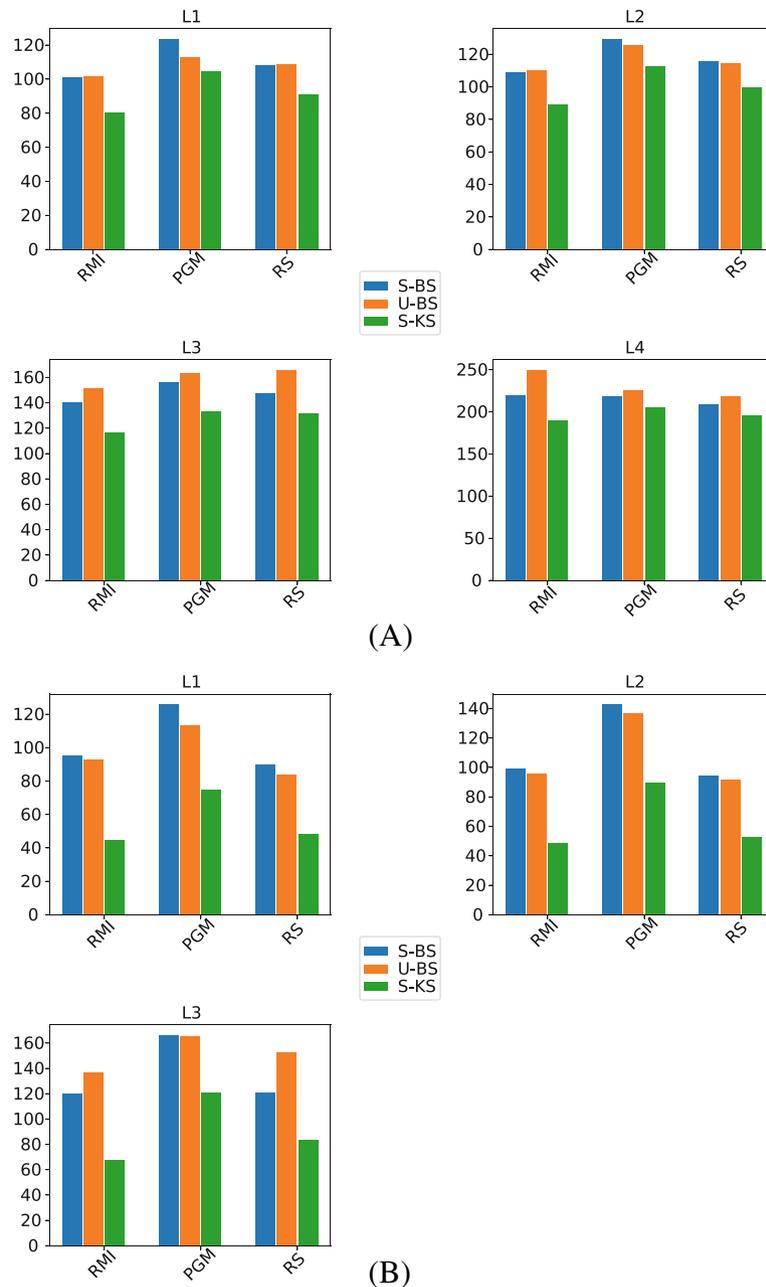


FIGURE 6 SOSD mean query times of best learned indexes on the osm dataset on the Intel I7 and the Apple M1. (A,B) report results using **SOSD** on the Intel I7 and Apple M1, respectively. For each model class, we report the mean query time in nanoseconds of the best learned indexes adopting in their last stage the routines described in Section 3.2.1. In particular, the blue bar is **S-BS**, the orange bar is **U-BS** and the green bar is **S-KS**.

that, for the largest datasets (level L4), it is advisable to use **SOSD**, again with **S-KS** as a terminal for the search stage. On the other memory levels, the performance is dataset-dependent. In particular, **SOSD** is better than the *stand-alone* settings for **face** and **wiki** datasets and worse for the remaining ones. As far as **S-BS** and **U-BS** is concerned, in *stand-alone* setting **U-BS** seems to be the method of choice. This is in agreement with the result reported in Section 4.1 (*Stand-alone* setting).

All of the above provides factual indications regarding the use of **SOSD**. As for their justification, a profiler analysis would be required for the Apple M1 architecture. Unfortunately, as already stated, this is not possible at this time with free software. As for the Intel architecture, the fact that **SOSD** is to be used on large datasets, may be attributed to the

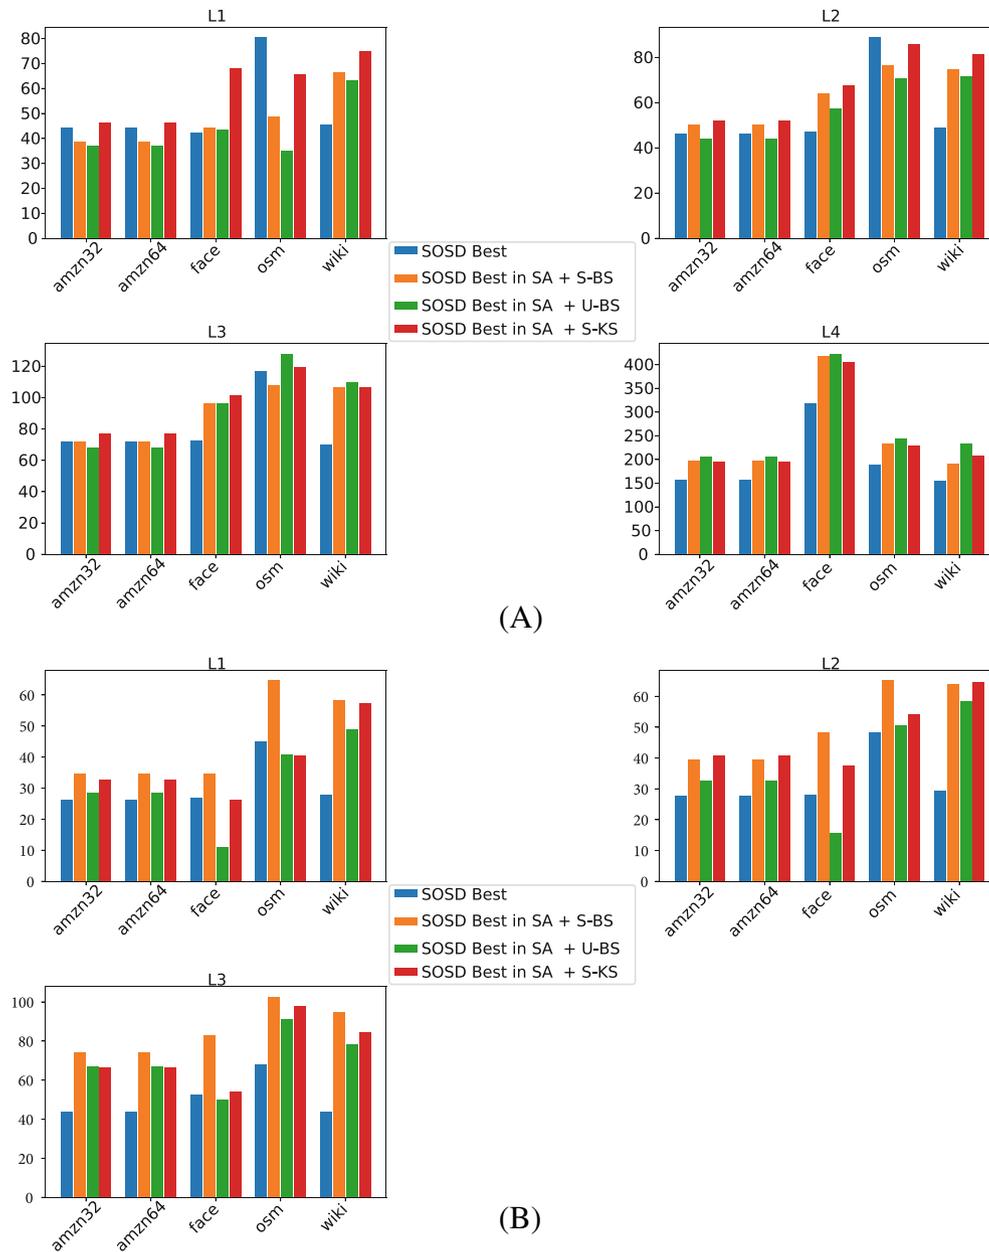


FIGURE 7 Comparison between the best model indicated by SOSD and its *Stand-alone* counterpart on the Intel I7 and the Apple M1. (A,B) report results for the Intel I7 and the Apple M1, respectively. Each group bar is relative to a dataset, while on the y-axis we report the mean query times in nanoseconds. In each group, the blue bar indicates the best model selected by **SOSD** (also named **SOSD Best**), while the next bars are the same model when used in *stand-alone* configuration (indicated as **SOSD best in SA**) together with **S-BS** (orange bar), **U-BS** (green bar) and **S-KS** (red bar) as terminal stage routines.

design and use of **SOSD**: the vast majority of the experiments conducted with **SOSD** are on large datasets. Moreover, as pointed out in Section 4.1 (profiler analysis), it seems to favour branchy code.

6 | CONCLUSIONS AND FUTURE DIRECTIONS

The main question we have addressed in this research is to provide indications on how the findings of Khuong and Morin¹² and Schulz et al.,¹³ regarding the choice of which binary search routines or variants are to be used on modern computer

architectures, can also be extended to the novel field of learned indexing, considering those routines for the final search stage. So far, for that stage, only the **lower_bound** routine has been considered. A summary of our results that can be useful both to designers and users of learned indexes is the following.

- When no additional space with respect to the input table can be afforded, **U-EL** is the best choice both in **SOSD** and *stand-alone* settings. This result confirms the findings by Khuong and Morin¹² and extends them since we consider also k-ary Search.
- When learned indexing is to be used, for each model class considered in this research, **SOSD** returns the best models with **S-KS** as the terminal search stage. This fact holds for both the hardware architectures we have considered.
- When the choice between **SOSD** and *stand-alone* settings needs to be made to allow deployment in an Application Domain, such as Databases¹¹ or Search Engines,¹² several factors need to be considered.
 - On the Apple M1 architecture, **SOSD** with **S-KS** as the final search stage is to be preferred.
 - On the Intel I7 architecture, the choice depends on datasets and memory levels. Indeed, for datasets larger than cache memory, **SOSD** with again as final search stage **S-KS** is to be preferred, while, for the ones fitting the cache memory, such a choice is input data-dependent. In any case, for learned indexes to be executed in a *stand-alone* setting, **U-BS** as the final search stage seems to be the more convenient choice.

Among the many open problems that the new area of learned indexing poses, we mention two that are relevant to the research we have conducted. Although the Eytzinger Layout binary search is superior to the other routines we have considered, it cannot be used by the Index Models known so far. So, it would be very interesting to devise new models that can use such an excellent layout in the final search stage. The second problem is in relation to the extension of our study to Dynamic learned indexes, that is, the **PGM** and **Alex**.¹⁵ Although an extension to the **PGM** may be simple (such a Model uses a binary search routine as the final search stage, even in the Dynamic setting), it is not clear how to intervene in **Alex**, given the high level of engineering that has been deployed for the realization of that Model.

ACKNOWLEDGMENT

The authors are deeply indebted to the Associate Editor and the two Referees, for the very stimulating comments received during the entire revision process. This research is funded in part by MIUR Project of National Relevance 2017WR7SHH “Multicriteria Data Structures and Algorithms: from compressed to learned indexes, and beyond”. Additional support to Raffaele Giancarlo has been granted by Project INdAM–GNCS “Modellizzazione ed analisi di big knowledge graphs per la risoluzione di problemi in ambito medico e web”. Open Access Funding provided by Università degli Studi di Palermo within the CRUI-CARE Agreement.

DATA AVAILABILITY STATEMENT

The data that support the findings of this study are openly available in Datasets at <https://anonymous.4open.science/r/DatasetsAndSupplementaryMaterial-E534/README.md>.

AUTHOR CONTRIBUTIONS

All Authors contributed to the design of this research. DA wrote the software and performed the experiments. RG wrote the paper with the collaboration of the other authors.

ORCID

Giosué Lo Bosco  <https://orcid.org/0000-0002-1602-0693>

REFERENCES

1. Kraska T, Beutel A, Chi E, Dean J, Polyzotis N. The case for learned index structures. Proceedings of the 2018 International Conference on Management of Data, ACM, 2018:489–504.

2. Ao N, Zhang F, Wu D, et al. Efficient parallel lists intersection and index compression algorithms using graphics processing units. *Proc VLDB Endow*. 2011;4(8):470-481. doi:10.14778/2002974.2002975
3. Amato D, Lo Bosco G, Giancarlo R. Learned sorted table search and static indexes in small model space. *CoRR*. 2021;abs2107.09480.
4. Amato D, Lo Bosco G, Giancarlo R. Learned sorted table search and static indexes in small model space. In: *AIxIA 2021 - Advances In Artificial Intelligence*, Springer International Publishing 2022: 462-477.
5. Ferragina P, Vinciguerra G. The PGM-index: a fully-dynamic compressed learned index with provable worst-case bounds. *Pvldb*. 2020;13(8):1162-1175. doi:10.14778/3389133.3389135
6. Ferragina P, Vinciguerra G. Learned data structures. *Recent Trends in Learning from Data*. Springer International Publishing; 2020: 5-41.
7. Ferragina P, Lillo F, Vinciguerra G. On the performance of learned data structures. *Theoretical Computer Science*. 2021;871:107-120. doi:10.1016/j.tcs.2021.04.015
8. Maltry M, Dittrich J. A critical analysis of recursive model indexes. *Proceedings of the VLDB Endowment*. 2022;15(5):1079-1091. doi:10.14778/3510397.3510405
9. Kipf A, Marcus R, van Renen A, et al. RadixSpline: A single-pass learned index. Proceedings of the Third International Workshop on Exploiting Artificial Intelligence Techniques for Data Management, aiDM'20; ACM; 2020: 1-5.
10. Comer D. Ubiquitous B-tree. *ACM Computing Surveys (CSUR)*. 1979;11(2):121-137. doi:10.1145/356770.356776
11. Rao J, Ross KA. Cache conscious indexing for decision-support in main memory. Proceedings of the 25th International Conference on Very Large Data Bases, ACM; 1999: 78-89.
12. Khuong P-V, Morin P. Array layouts for comparison-based searching. *ACM Journal of Experimental Algorithmics*. 2017;22:1-39. doi:10.1145/3053370
13. Schulz L-C, Broneske D, Saake G. An eight-dimensional systematic evaluation of optimized search algorithms on modern processors. *Proceedings of the VLDB Endowment*. 2018;11(11):1550-1562. doi:10.14778/3236187.3236205
14. Kipf A, Marcus R, van Renen A, et al. SOSD: a benchmark for learned indexes. *NeurIPS Workshop on Machine Learning for Systems*. 2019. http://mlforsystems.org/neurips2019/accepted_papers.html
15. Ding J, Minhas UF, Yu J, et al. ALEX: An updatable adaptive learned index. Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data, SIGMOD'20. ACM. 2020: 969-984.
16. <https://github.com/globosco/A-modified-SOSD-platform-for-benchmarking-Branchy-vs-Branch-free-search-algorithms.git>.
17. <https://github.com/globosco/Stand-Alone-Platform-for-benchmarking-Branchy-vs-Branch-free-search-algorithms>.
18. <https://anonymous.4open.science/r/DatasetsAndSupplementaryMaterial-E534/README.md>.
19. Marcus R, Kipf A, van Renen A, et al. Benchmarking learned indexes. *Proc VLDB Endow*. 2020;14(1):1-13. doi:10.14778/3421424.3421425
20. Freedman D. *Statistical Models: Theory and Practice*. Cambridge University Press; 2005.
21. Knuth DE. *The Art of Computer Programming*. Sorting and Searching. Vol 3. Addison-Wesley Publishing Company; 1973.
22. Cormen TH, Leiserson CE, Rivest RL, Stein C. *Introduction to Algorithms*. 3rd ed. The MIT Press; 2009.
23. Schlegel B, Gemulla R, Lehner W. K-ary search on modern processors. Proceedings of the Fifth International Workshop on Data Management on New Hardware, DaMoN'09, New York, NY, USA. ACM. 2009; 52-60.
24. Kipf A *Personal Communication*; 2021.

How to cite this article: Amato D, Lo Bosco G, Giancarlo R. Standard versus uniform binary search and their variants in learned static indexing: The case of the searching on sorted data benchmarking software platform. *Softw Pract Exper*. 2023;53(2):318-346. doi: 10.1002/spe.3150

APPENDIX A. ALGORITHMS, CODE, AND SOFTWARE PLATFORMS

A.1 Binary search and its variants: Additional algorithms

Details about the routines mentioned in Section 3.2.1 are given here. In particular:

- **Lower_bound Function.** The routine is given in Algorithm 3.
- **k-ary Search.** The routines are given in Algorithms 4-5.
- **Eytzinger Layout.** The sorted table is now seen as stored in a virtual complete balanced binary search Tree. Such a tree is laid out in breadth-first search order in an array. An example is provided in Figure A1. Also, in this case, we adopt a Branch-free version with prefetching of the binary search procedure corresponding to this layout. It is reported in Algorithm 6.

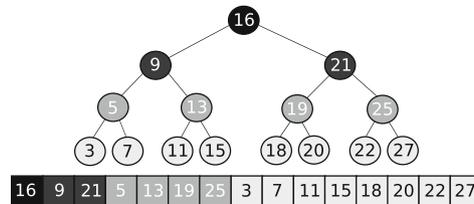


FIGURE A1 An example of Eytzinger layout of a table with 15 elements. See also Reference 12

Algorithm 3. Lower_bound template

```

1: ForwardIterator lower_bound (ForwardIterator first, ForwardIterator last, const T& val){
2:   ForwardIterator it;
3:   iterator_traits<ForwardIterator>::difference_type count, step;
4:   count = distance(first,last);
5:   while (count>0){
6:     it = first; step=count/2; advance (it,step);
7:     if (*it<val){
8:       first=++it;
9:       count-=step+1;
10:    }
11:    else count=step;
12:  }
13:  return first;
14: }
```

A.2 Index model classes in SOSD: Additional details

In this section details about the learned indexes mentioned in Section 3.2.2 are given. In particular:

- **RMI.**¹ It is a multi-stage model. When searching for a given key and starting with the first stage, a prediction at each stage identifies the model of the next stage to use for the next prediction. This process continues until a final stage model is reached. This latter is used to predict the interval to search into. An example is helpful. Figure A2A describes a two-stage model. The first stage model (linear) is used to pick one of the stage 2 models (cubic). Finally, the selected model provides the final interval to search into.
- **RS.**⁹ It is a two-stage model. It uses user-defined approximation parameter ϵ . With reference to Figure A2C, a spline curve approximating the CDF of the data is built. Then, the radix table is used to identify spline points to use to refine the search interval. Indeed, for the key in the figure, the most significant three bits are used to identify two spline points. Then, a binary search is performed on the table of the spline points delimited by the points identified earlier. Such a binary search identifies two spline points such that a simple interpolation guarantees a prediction with error ϵ .
- **PGM.**^{5,7,t} It is also a multi-stage model, built bottom-up and queried top down. It also uses a user-defined approximation parameter ϵ , that controls the prediction error at each stage. With reference to Figure A2B, the table is subdivided into three pieces. A prediction in each piece can be done via a linear model guaranteeing an error of ϵ . A new table is formed by selecting the minimum values in each of the three pieces. This new table is possibly again partitioned into pieces, in which a linear model can make a prediction within the given error. The process is iterated until only one linear model suffices, as in the case in the figure. A query is processed via a series of predictions, starting at the root of the tree.

A.3 Computer architecture and compilers: The production of branch-free code: Additional code

In this section, as indicated in Section 3.3, the assembly codes of the compiled routines are provided. In particular:

- The assembly code of standard binary search on the Intel I7 is provided in Code 2.

Algorithm 4. Implementation of standard k-ary search. The code is as in Reference 13

```

1: int StandardKarySearch(int *arr, int x, int start, int end, int k){
2:   int left = start, right = end;
3:   while (left < right)
4:   {
5:     int segLeft = left;
6:     int segRight = left + (right - left) / k;
7:     for (int i = 2; i <= k; ++i)
8:     {
9:       if (x <= arr[segRight]) break;
10:      segLeft = segRight + 1;
11:      segRight = left + (i * (right - left)) / k;
12:    }
13:    left = segLeft;
14:    right = segRight;
15:  }
16:  return left;
17: }
```

Algorithm 5. Implementation of uniform k-ary search. The code is as in Reference 13

```

1: int UniformKarySearch(int *arr, int x, int start, int end, int k){
2:   int left = start, right = end;
3:   while (left < right){
4:     int segLeft = left;
5:     int segRight = left + (1 * (right - left)) / k;
6:     for (int i = 2; i <= k; ++i)
7:     {
8:       int nextSeparatorIndex = left + (i * (right - left)) / k;
9:       segLeft = x > arr[segRight] ? segRight + 1 : segLeft;
10:      segRight = x > arr[segRight] ? nextSeparatorIndex : segRight;
11:    }
12:    left = segLeft;
13:    right = segRight;
14:  }
15:  return left;
16: }
```

Algorithm 6. Implementation of uniform binary search with eytzing layout and prefetching. The code is as in Reference 12. The version without prefetching is obtained by deleting line 5

```

1: int UniformEytzingerLayout(int *A, int x, int left, int right){
2:   int i = 0;
3:   int n = right;
4:   while (i < n){
5:     __builtin_prefetch(A+(multiplier*i + offset));
6:     i = (x <= A[i]) ? (2*i + 1) : (2*i + 2);
7:   }
8:   int j = (i+1) >> __builtin_ffs(~(i+1));
9:   return (j == 0) ? n : j-1;
10: }
```

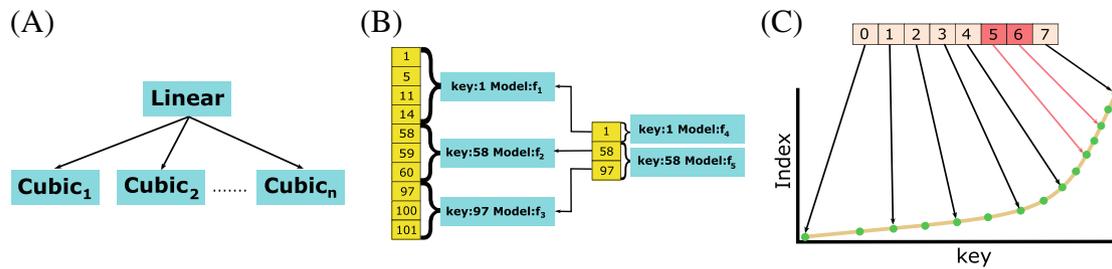


FIGURE A2 Examples of (A) RMI (B) PGM (C) RS. See also Reference 19

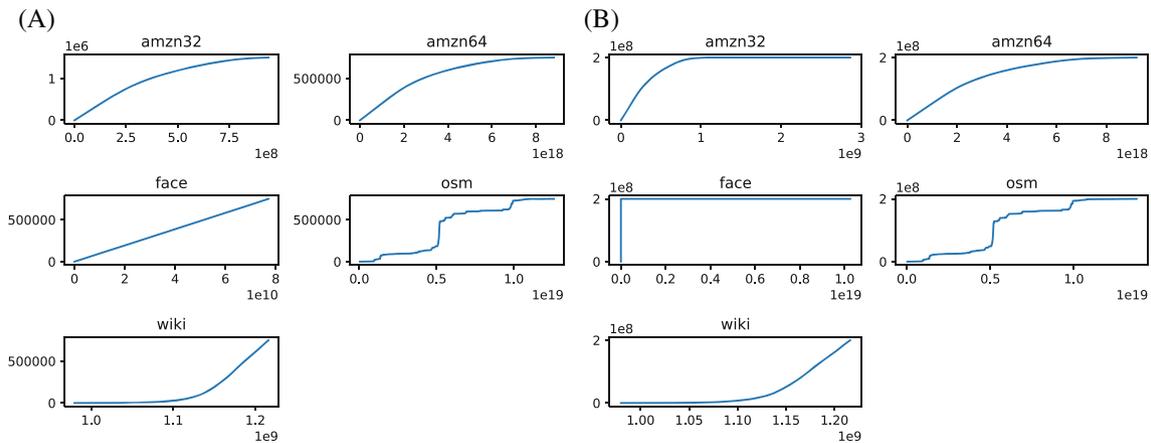


FIGURE B1 Datasets empirical CDF. For each dataset, we report the value of the elements on the x-axis and their position on the y-axis. The curve so obtained is analogous to the Empirical *CDF* of the Universe from which the data come. In particular, (A) is referred to the L3 memory level, while (B) to L4.

- The assembly codes of uniform binary search on the Intel I9 and the Apple M1 are provided in Code 3 and Code 4, respectively.
- The assembly code of lower_bound function on the Apple M1 is provided in Code 5.
- The assembly codes of the other branchy methods are here omitted and available to the reader upon request.

APPENDIX B. DATASETS AND INDEX MODEL TRAINING: ADDITIONAL CONSIDERATIONS

In this section, we report in Figure B1, the *CDF* of the datasets on L3 and L4 memory levels, described in Section 3.4. It can be noted that the curves of **face_L3** and **face_L4** show a significant difference due to the presence of some outliers in the L4 memory level which makes the representation not very accurate.

APPENDIX C. EXPERIMENTS: SEARCHING IN CONSTANT ADDITIONAL SPACE, WITH OR WITHOUT SOSD - ADDITIONAL RESULTS

C.1 Replication of The experiments by Khuong and Morin

We provide here the details of the experiments mentioned in Section 4.1. In particular:

- The comparison between the **lower_bound** and **S-BS** implementations is reported in Figure C1 for the case of the Intel I7 Architecture, and in Figure C2 for the Apple M1.
- The comparison of each of the considered routines in their SOSD and *stand-alone* implementations is reported in Figure C3
- The comparison among the remaining routines on the **osm** dataset is reported in Figure C4.

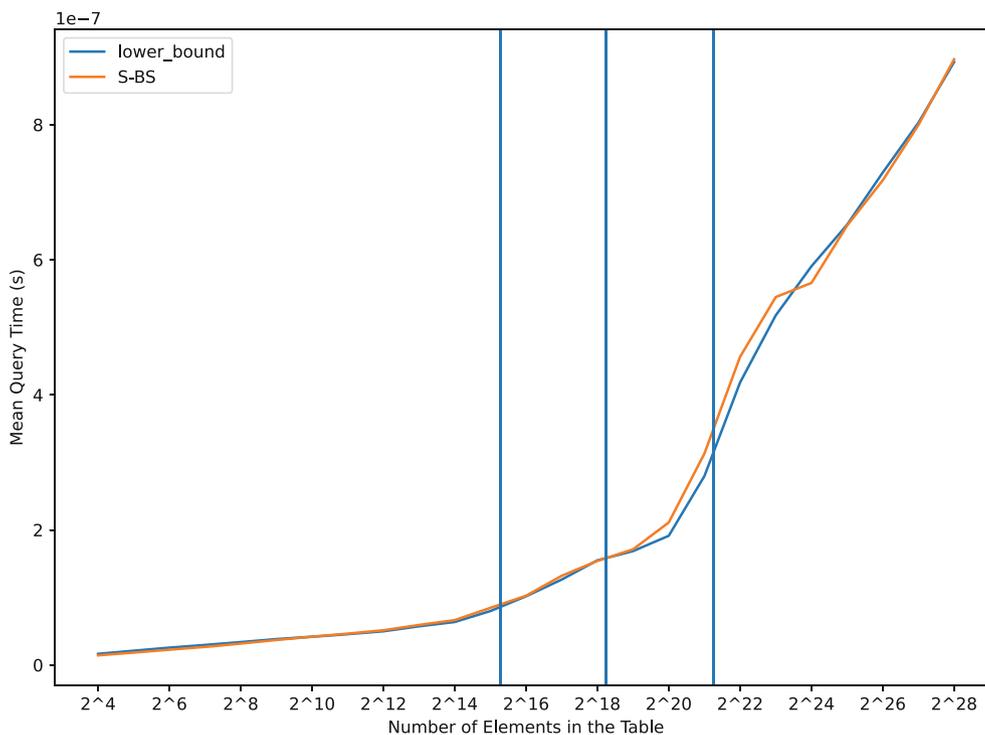


FIGURE C1 Mean query times comparison of the lower_bound function and S-BS on synthetic datasets on the Intel I7. The Figure reports only the comparison of the routines executed within the **SOSD** platform on Intel I7 since the results with Standard C++ implementations are the same. On the x-axis, we report the number of elements in the table and, on the y-axis, the mean query time in seconds. The vertical lines indicate the size of each cache memory level.

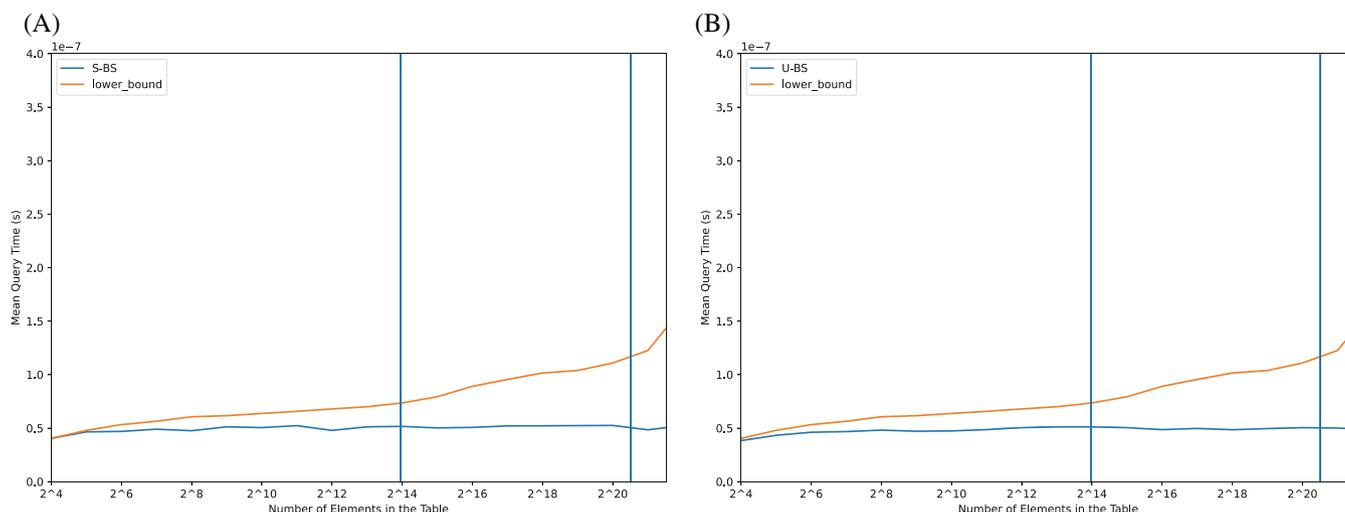


FIGURE C2 Mean query times comparison of lower_bound function with S-BS and U-BS on synthetic datasets on the Apple M1. Figures report only the comparison of the routines executed within the **SOSD** platform on Apple M1 architecture, since the results with a Standard C++ implementation are the same. We report the comparison between **lower_bound** function and **S-BS** in (A), while in (B) the comparison between **lower_bound** function and **U-BS**. On the x-axis, we report the number of elements in the table and, on the y-axis, the mean query time in seconds. The vertical lines indicate the size of each cache memory level.

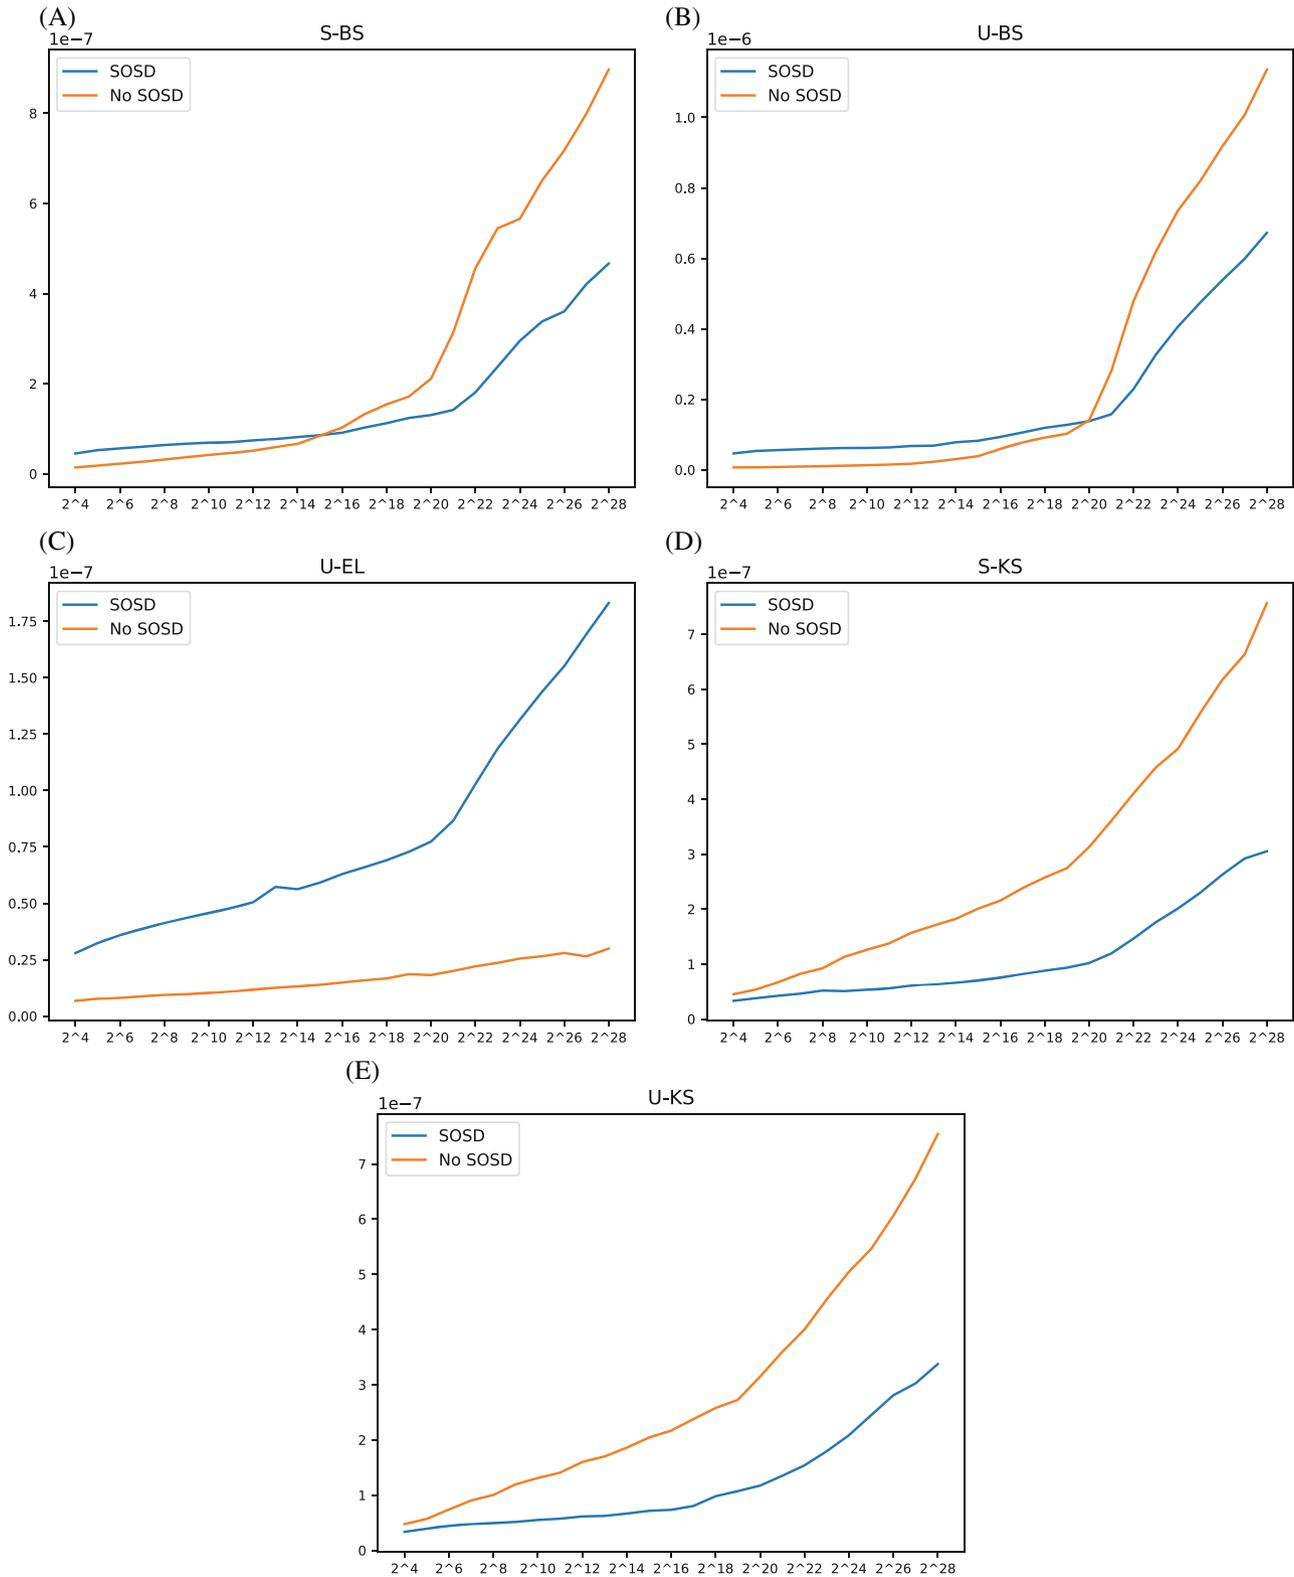


FIGURE C3 Mean query times comparison between SOSD and stand-alone implementations on the Intel I7. For each method, that is, S-BS (A), U-BS (B), U-EL (C), S-KS (D), U-KS (E), we report SOSD (blue curve) and *stand-alone* (orange curve) implementations for the case of Intel I7 architectures. The x-axis reports number of elements on the Table, while y-axis the mean query time in seconds.

Profiler analysis

In order to look more closely into the results in Section 4.1, we make use of the **profiler** on the routines to monitor the following parameters.

- **Percentage of Retiring operation.** That is the percentage of the pipeline slots used to perform useful work. It is divided into light operations, which are instructions that require at most one micro-operation, and heavy, which require two or more of such operations. A high percentage of this parameter implies a very effective use of the CPU.
- **Percentage of Front-end operation.** That is the percentage of pipeline slots where the processor is not able to fetch instructions to the Back-end part. This issue can be caused, for example, by stalls due to instruction-cache misses.
- **Percentage of Back-end operation.** That is the percentage of pipeline slots where the Back-end part is not able to deliver micro-operation for the retiring. This issue can be caused, for example, by stalls due to a lack of resources.
- **Percentage of Bad Speculation operation.** That is the percentage of pipeline slots where a wrong pipeline of instructions is loaded, due to a mispredicted branch.

The results are summarized in Tables C1 and C2.

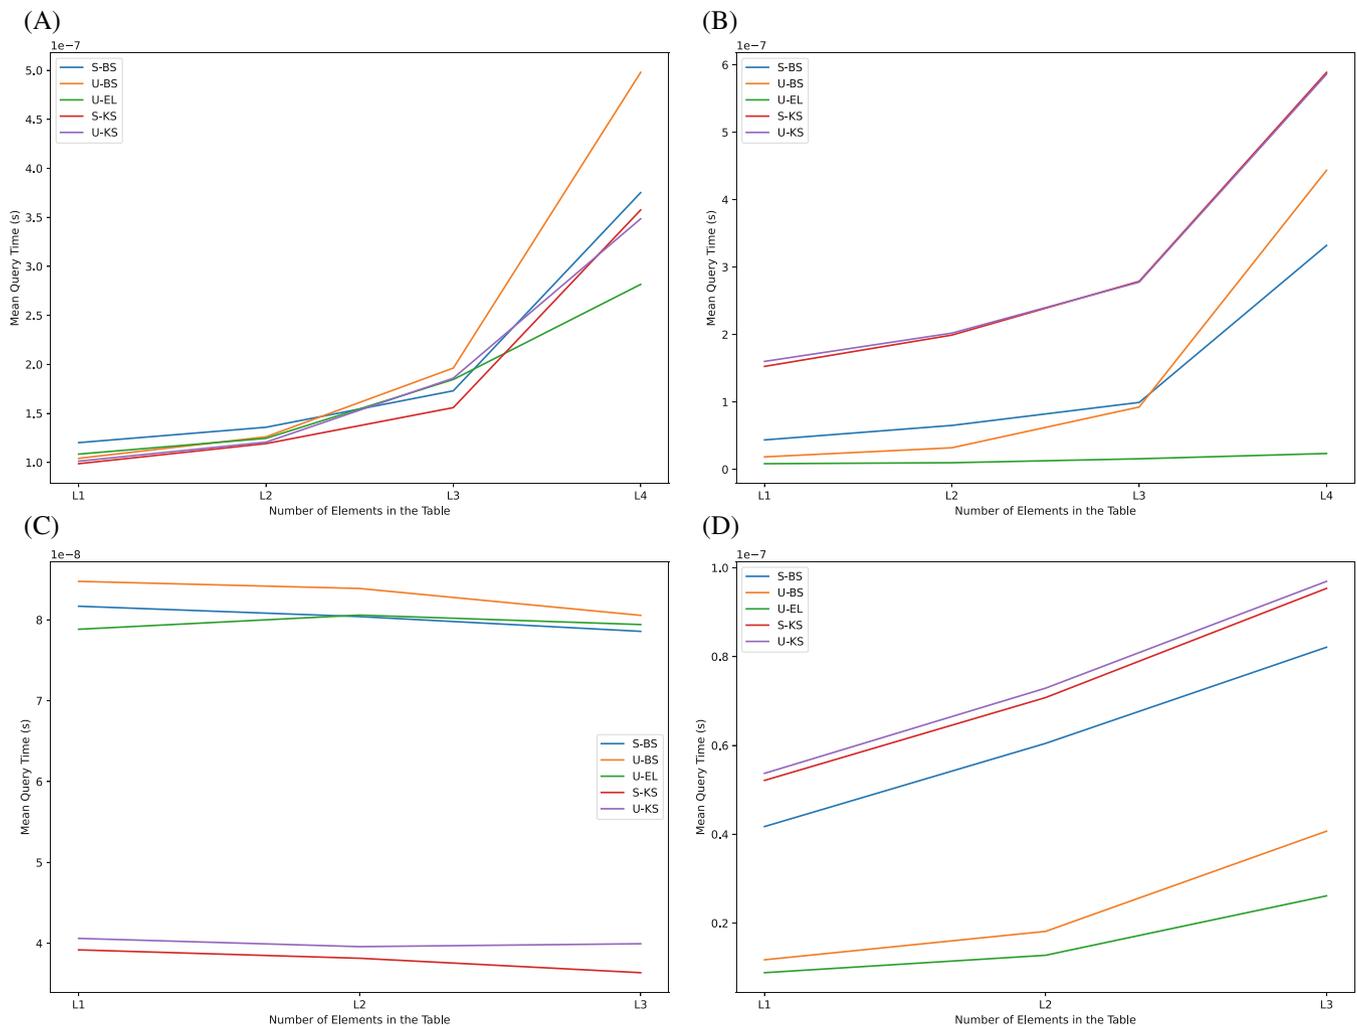


FIGURE C4 Mean query times on osm datasets on Intel I7 (A,B) and Apple M1 (C,D) architectures. (A,C) show only the comparison of the routines executed within the **SOSP** platform, while (B,D) with a Standard C++ implementation. On the x-axis, we report the hierarchical memory level in which datasets fit and, on the y-axis, the mean query time in seconds.

Code 2. Assembly code of standard binary search on Intel I7 (only main loop). No prefetching is used

```

000000000001229 <_ZL22Standard_Binary_SearchPmmii>:
.
.
124d: jg 12bc <_ZL22Standard_Binary_SearchPmmii+0x93>
124f: mov -0x24(%rbp),%edx
1252: mov -0x28(%rbp),%eax
1255: add %edx,%eax
1257: mov %eax,%edx
1259: shr $0x1f,%edx
125c: add %edx,%eax
125e: sar %eax
1260: cltq
1262: mov %rax,-0x8(%rbp)
1266: mov -0x8(%rbp),%rax
126a: lea 0x0(,%rax,8),%rdx
1272: mov -0x18(%rbp),%rax
1276: add %rdx,%rax
1279: mov (%rax),%rax
127c: cmp %rax,-0x20(%rbp)
1280: jae 128e <_ZL22Standard_Binary_SearchPmmii+0x65>
1282: mov -0x8(%rbp),%rax
1286: sub $0x1,%eax
1289: mov %eax,-0x28(%rbp)
128c: jmp 1247 <_ZL22Standard_Binary_SearchPmmii+0x1e>
128e: mov -0x8(%rbp),%rax
1292: lea 0x0(,%rax,8),%rdx
129a: mov -0x18(%rbp),%rax
129e: add %rdx,%rax
12a1: mov (%rax),%rax
12a4: cmp %rax,-0x20(%rbp)
12a8: jbe 12b6 <_ZL22Standard_Binary_SearchPmmii+0x8d>
12aa: mov -0x8(%rbp),%rax
12ae: add $0x1,%eax
12b1: mov %eax,-0x24(%rbp)
12b4: jmp 1247 <_ZL22Standard_Binary_SearchPmmii+0x1e>
12b6: mov -0x8(%rbp),%rax
12ba: jmp 12c1 <_ZL22Standard_Binary_SearchPmmii+0x98>
12bc: mov -0x28(%rbp),%eax
.
.

```

C.2 Pros and Cons of Prefetching

This Section presents the details of the profiler analysis indicated in Section 4.2. A deeper analysis with the profiler, as just done in the previous Section, is reported in Tables C3–C4. Such analysis allows us to see that the utilisation of pipeline slots is relatively identical with or without explicit prefetching on a dataset that is small enough to be stored in cache memory, for example, uniform with 2^{10} elements. On the other hand, it can be seen that if the dataset resides in central memory, for example, uniform with 2^{24} elements, the percentage of stalls in the back-end component of the CPU is higher without the use of explicit prefetching. In particular, as shown in Table C4, 70% of these are due to a data cache miss problem that explicit prefetching manages to avoid.

Code 3. Assembly code of uniform binary search on Intel I9 (only main loop). The predicated instruction is line 1475 (in bold)

```
000000000001450 <_Z21Uniform_Binary_SearchPmmmm>:
.
.
1463: jbe 1482 <_Z21Uniform_Binary_SearchPmmmm+0x32>
1465: nopl (%rax)
1468: mov %rcx,%rdx
146b: shr %rdx
146e: lea (%rax,%rdx,8),%r8
1472: cmp (%r8),%rsi
1475: cmovae %r8,%rax
1479: sub %rdx,%rcx
147c: cmp $0x1,%rcx
1480: ja 1468 <_Z21Uniform_Binary_SearchPmmmm+0x18>
.
.
```

Code 4. Assembly code of uniform binary search on Apple Arm M1 (only main loop). The predicated instruction is line 10000371c (in bold)

```
00000001000036f8 <__Z21Uniform_Binary_SearchPyyyy>:
.
.
100003708: b.lo 0x10000372c <__Z21Uniform_Binary_SearchPyyyy+0x34>
10000370c: lsr x10, x9, #1
100003710: add x11, x8, x10, lsl #3
100003714: ldr x12, [x11]
100003718: cmp x12, x1
10000371c: csel x8, x8, x11, hi
100003720: sub x9, x9, x10
100003724: cmp x9, #1
100003728: b.hi 0x10000370c <__Z21Uniform_Binary_SearchPyyyy+0x14>
.
.
```

TABLE C1 Profiling of S-BS, U-BS, U-EL, and S-KS on the Intel I7

Type	S-BS		U-BS		U-EL		S-KS	
	SOSD	NO SOSD	SOSD	NO SOSD	SOSD	NO SOSD	SOSD	NO SOSD
Retiring	34.2–95.3	18.7–95.7	38.1–100	36.8–100	17.0–100	23.8–100	33.5–88.1	25.6–32.4
Front-end	22.6	16.5	15.8	8.9	25.2	37.5	21.8	42.3
Back-end	28.5	21.9	41.1	51.3	29.7	25.5	28.2	15.1
Bad Spec.	14.6	43.0	5.0	3.0	28.2	13.2	16.5	17.1

Note: The results refer to a synthetic table of 2¹⁰ elements. We consider **S-KS** only, since its performance is analogous to **U-KS**. The first row reports the methods and the second row, for each method, indicates the execution environment, that is, with or without **SOSD**. The next rows represent hardware operation types. All the values are reported in percentage. In particular, the second value on the Retiring row is the percentage of light operations.

Code 5. Assembly code of C++ lower_bound function on Apple Arm M1 (only main loop). The predicated instruction are in lines 10000371c, 100003fa0 (in bold)

```

1: 0000000100003f68 <__Z6searchRNSt3__16vectorIyNS_9allocatorIyEEEEyPmmm>:
2: .
3: .
4: 100003f7c: b.eq 0x100003fa8 <__Z6searchRNSt3__16vectorIyNS_9allocatorIyEEEEyPmmm+0x40>
5: 100003f80: asr x9, x9, #3
6: 100003f84: lsr x10, x9, #1
7: 100003f88: add x11, x8, x10, lsl #3
8: 100003f8c: ldr x12, [x11], #8
9: 100003f90: mvn x13, x10
10: 100003f94: add x9, x9, x13
11: 100003f98: cmp x12, x1
12: 100003f9c: csel x9, x9, x10, lo
13: 100003fa0: csel x8, x11, x8, lo
14: 100003fa4: cbnz x9, 0x100003f84 <__Z6searchRNSt3__16vectorIyNS_9allocatorIyEEEEyPmmm+0x1c>
15: .
16: .

```

TABLE C2 Profiling of S-BS, U-BS, and S-KS on the Intel I7

Type	S-BS		U-BS		U-EL		S-KS	
	SOSD	NO SOSD	SOSD	NO SOSD	SOSD	NO SOSD	SOSD	NO SOSD
Retiring	19.1–75.4	5.5–100	16.4–76.2	6.8–79.4	7.2–62.5	30.8–94.8	9.1–71.4	15–46.2
Front-end	14.5	4.4	9.1	3	13.3	30.8	13.1	34.4
Back-end	51.3	61.4	71.9	89.2	51.9	21.0	45.9	32.5
Bad Spec.	15.2	28.7	2.7	1	27.0	17.0	31.9	18.1

Note: The results refer to a synthetic table of 2^{24} elements. The legend is as in Table C1.

TABLE C3 U-BS with or without explicit prefetching on Intel I7 pipeline usage percentage

Type	2^{10}		2^{24}	
	Prefetch	No prefetch	Prefetch	No prefetch
Retiring	38.4	39.2	24.4	19.7
Front-end	16.2	16.3	8.2	6.8
Back-end	40.4	39.4	64.4	71.2
Bad Spec.	5.1	5.1	2.7	2.3

Note: The first column reports CPU pipeline operation types. For each of them, we report the percentage of slots used.

APPENDIX D. EXPERIMENTS: SEARCHING USING LEARNED INDEXES, WITH OR WITHOUT SOSD - ADDITIONAL RESULTS

We report here the details of the experiments mentioned in Section 5. In particular, we provide the following.

- We report the result of Best learned indexes within **SOSD** other than **osm** in Figures D1–D4.
- The full list of best Models indicated by **SOSD** in Table D1.

TABLE C4 U-BS with or without explicit prefetching on Intel I7 back-end usage percentage

Type	Uniform 2 ¹⁰		Uniform 2 ²⁴	
	Prefetch	No prefetch	Prefetch	No prefetch
Memory	23.4	22.5	58.4	70.2
Core	76.6	77.5	41.6	29.8

Note: The first column reports the types of Back-end stalls. Memory indicates stalls occurred due to a data cache miss error, while Core indicates stalls due to an overloaded execution unit. For each of them, we report the percentage of slots used.

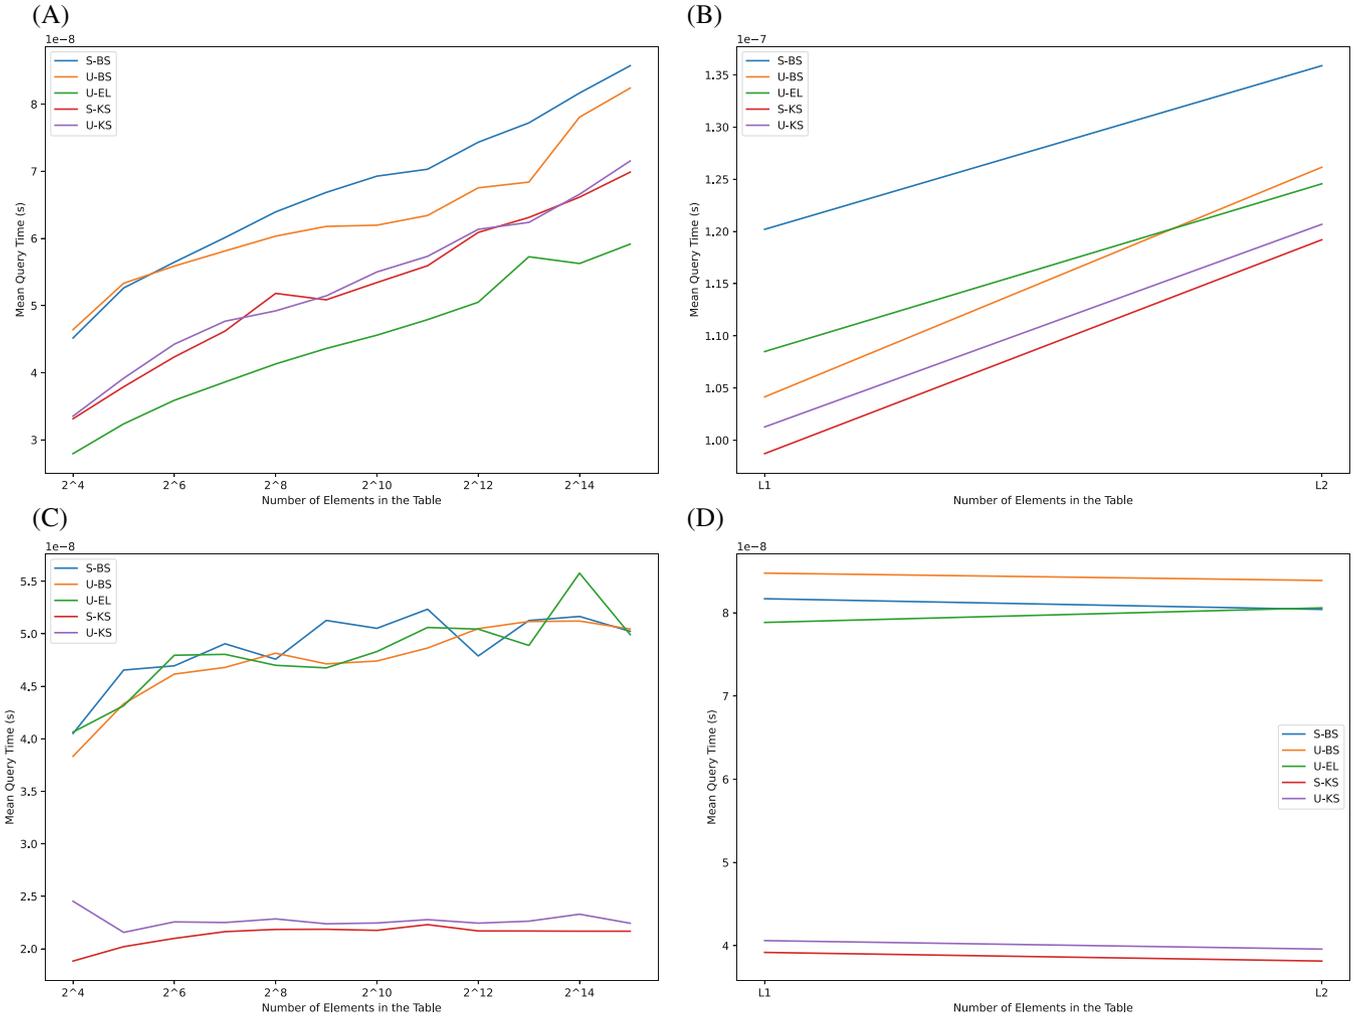


FIGURE C5 Zoom-in on mean query time with SOSD. (A,B) reports a zoom-in of Figure 3A and of the Figure C4A, respectively, while the (C,D) reports a zoom-in of Figure 3C and of the Figure C4C, respectively. The legend is as in the corresponding figures.

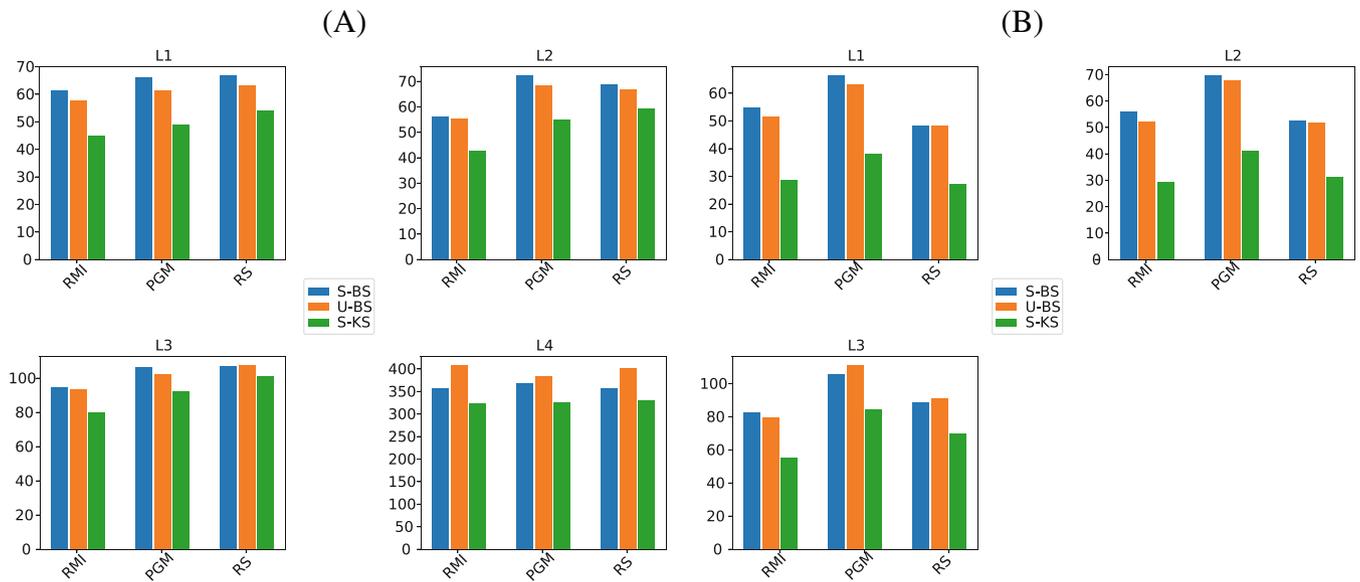


FIGURE D1 Mean query times of best learned indexes on amzn32 dataset on the Intel I7 and the Apple M1. (A,B) report results using **SOSD** on the Intel I7 and Apple Arm M1, respectively. For each model class, we report the mean query time in nanoseconds of the best learned indexes adopting in their last stage the routines described in Section 3.2.1. In particular, the blue bar is **S-BS**, the orange bar is **U-BS** and the green bar is **S-KS**.

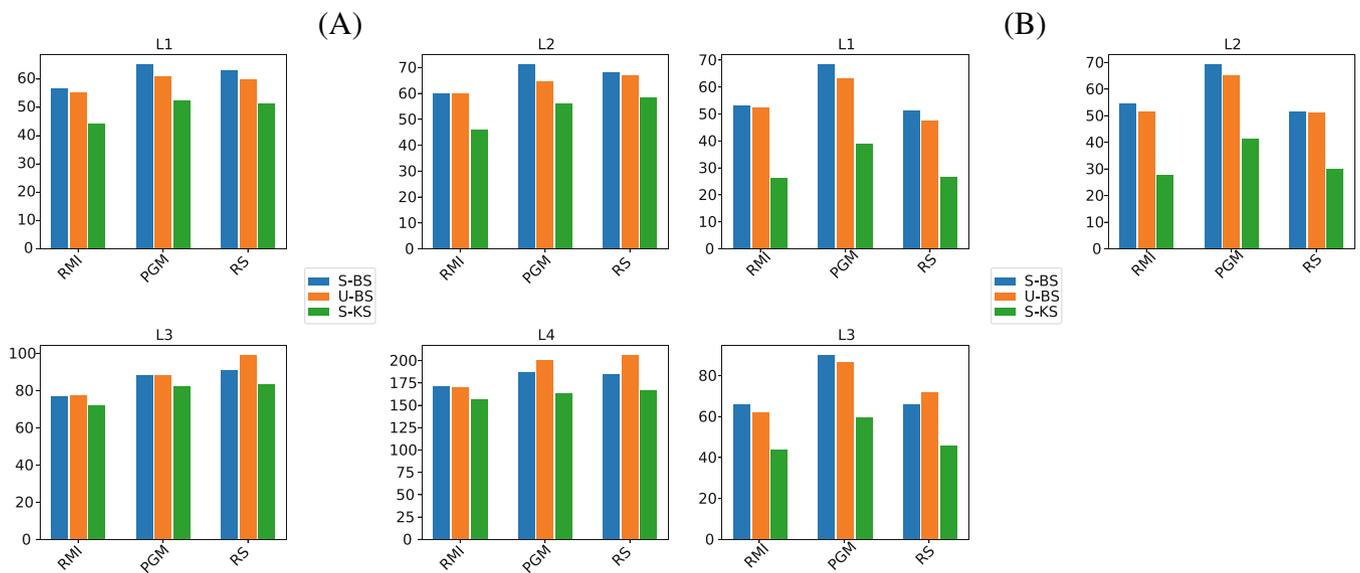


FIGURE D2 Mean query times of best learned indexes on amzn64 dataset on the Intel I7 and the Apple M1. The legend is as in Figure D1.

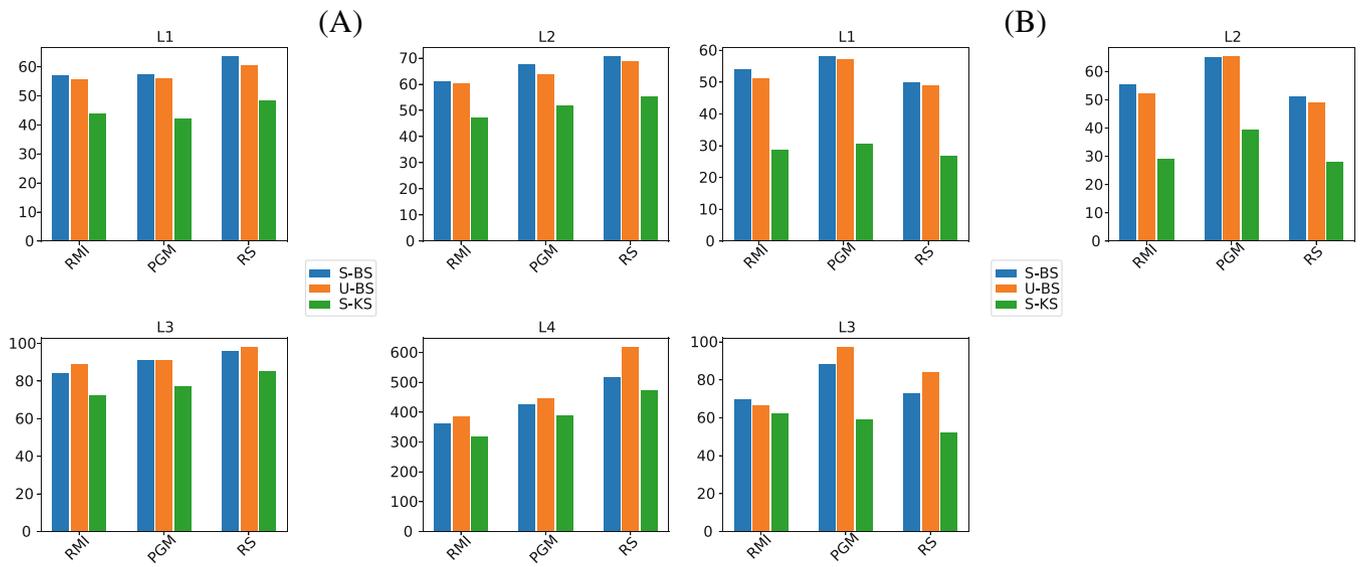


FIGURE D3 Mean query times of best learned indexes on face dataset on the Intel I7 and the Apple M1. The legend is as in Figure D1.

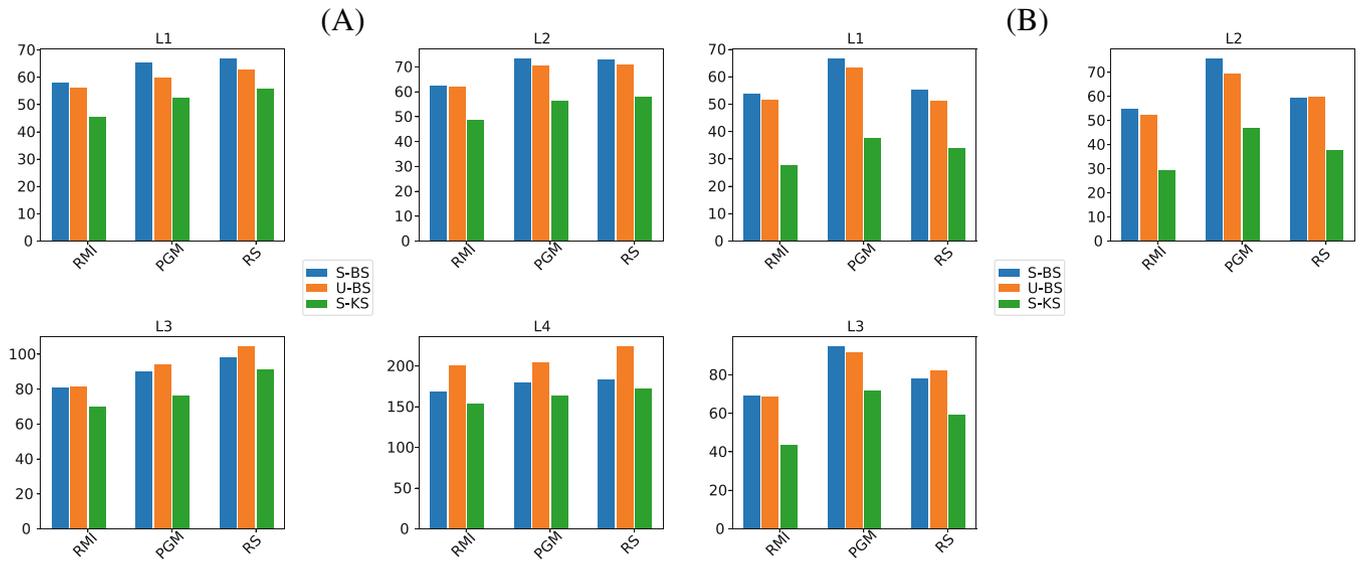


FIGURE D4 Mean query times of best learned indexes on wiki dataset on the Intel I7 and the Apple M1. The legend is as in Figure D1.

TABLE D1 SOSD best models

Intel I7				
	L1	L2	L3	L4
amzn32	RMI	RMI	RMI	RMI
amzn64	RMI	RMI	RMI	RMI
face	PGM	RMI	RMI	RMI
osm	RMI	RMI	RMI	RMI
wiki	RMI	RMI	RMI	RMI
Intel I9				
	L1	L2	L3	L4
amzn32	RMI	RMI	RMI	RMI
amzn64	RMI	RMI	RMI	RMI
face	RMI	RMI	RMI	RMI
osm	RMI	RMI	RMI	RMI
wiki	RMI	RMI	RMI	RMI
Apple M1				
	L1	L2	L3	L4
amzn32	RMI	RMI	RMI	-
amzn64	RMI	RMI	RMI	-
face	RS	RS	RS	-
osm	RMI	RMI	RMI	-
wiki	RMI	RMI	RMI	-

Note: For each dataset and each memory level, the table reports the class of the best performing model indicated by SOSD.