

Heuristics to Improve the BMW Method and Its Variants

Lídia Lizziane Serejo de Carvalho¹, Edleno Silva de Moura¹,
Caio Moura Daoud¹, Altigran Soares da Silva¹

Universidade Federal do Amazonas, Brazil
lidializz@gmail.com, {edleno, caiodaoud, alti}@icom.ufam.edu.br

Abstract. In this paper, we propose and evaluate heuristics to improve the performance of BMW and its variants. The proposed changes maintain the property of preserving the order of the top ranking results, while reduce query processing times and the amount of memory required for processing queries.

Categories and Subject Descriptors: H.3 [Information Storage and Retrieval]: Miscellaneous

Keywords: Information Retrieval, Query Processing, Inverted Indexes, Search Engines

1. INTRODUCTION

Query processors in search engines are implemented aiming to minimize costs with respect to the response time and memory usage. Their main goal is to take a query specified by the user and provide a ranking with the documents considered relevant to such query. To do so, they use Information Retrieval models that produce ranking functions, like the vector model [Salton et al. 1975] and the BM25 model [Robertson and Walker 1994], to compute a weight for each document given a query. In general, the systems return a ranking with the top- k documents with the higher weights, where k is a predefined number.

Ding and Suel (2011) proposed the Block-Max WAND (BMW) algorithm, an algorithm for processing queries that proved to be faster than the previously proposed methods for computing top- k ranking results. Variants of BMW method have been proposed and studied since then, including the BMW-CS [Rossi et al. 2013], which reduces the query processing time when compared to BMW, but has the disadvantages of increasing in the amount of memory required for processing queries and the possibility of losing results, thus changing the list of documents that appear at the top of the response for each query.

Algorithms such as the BMW use document discarding techniques (pruning techniques) that are directly related to the current heuristics that aim at reducing the processing costs. BMW maintains a minimum heap structure that stores the top- k documents during query processing. The smallest weight among all the documents present in the heap is used as a discarding threshold. Thus, while processing a query, only the scores of documents that have a chance to overcome the discarding threshold are computed.

The main objective of this work is to study and propose improvements to the BMW algorithm and its variant BMW-CS in order to reduce the query processing times and the amount of memory necessary to process queries with these algorithms. The most remarkable results are as follows: (i) reduction in the query processing times with an initial threshold pruning heuristic. (ii) reduction in memory usage due to a significant reduction in the size of an auxiliary structure used by the algorithms

Copyright©2014 Permission to copy without fee all or part of the material printed in JIDM is granted provided that the copies are not made or distributed for commercial advantage, and that notice is given that copying is by permission of the Sociedade Brasileira de Computação.

when processing queries, the skiplists. In both cases, the changes made have the advantage of not changing the initial results provided by the methods (BMW or BMW-CS). As BMW is an algorithm that preserves the top results for each query, our changes on the BMW also hold such property.

The remainder of this paper is structured as follows. The Section 2 presents the summary of the main related research. In Sections 3 e 4, the proposed strategies for the solution of the presented problem are explained. The Section 5 describes the experiments and the obtained results. Finally, in Section 6, we discuss the conclusions and directions for future research.

2. RELATED WORK

Several researchers have proposed strategies to improve the performance of query processing in search engines. The query processing methods found in the literature can be divided into two main categories: those that use TAAT (Term-At-A-Time) strategy and those that use DAAT (Document-At-A-time) strategy [Baeza-Yates et al. 2011]. In methods that use TAAT strategy, the inverted lists are sorted by the weight of the documents. Thus, the query is processed one term at a time. The biggest disadvantage is that this strategy requires large amounts of memory to store the scores achieved by each document for each term in the query. In the methods using DAAT strategy, the inverted lists are sorted by the *id* of the document, allowing the lists to be traversed in parallel, and the full score of a document can be computed at each iteration. Another feature of this strategy is the use of *skiplists* associated with each inverted list. These structures store pointers to entries in the inverted list, dividing it into blocks to facilitate access to documents, which typically have between 128 and 256 entries [Broder et al. 2003].

One of the most important works presented in recent years proposes a DAAT query processing method known as WAND (Weak AND or Weighted AND) [Broder et al. 2003]. In this method the highest score of each inverted list is stored during the document collection indexing. Knowing the maximum score that a document can achieve is important to avoid the cost of computing scores of documents that have no chance to change the top of ranking results. At query processing time, documents that are being evaluated, called pivots, are analyzed in two stages. First, the Max Score is obtained, maximum score of each list where the pivot has a chance to occur. When the Max Score is greater than the current pruning threshold, the inverted list of each term will be accessed so that the pivot document has its real score computed. When the Max Score does not exceed the pruning threshold, the pivot document is discarded and a new pivot is selected.

The BMW algorithm [Ding and Suel 2011] was created based on the WAND [Broder et al. 2003], but proposes an optimized solution that mainly covers a modification of skiplists. In BMW, for each input of skiplists, the highest score of the block is stored. The pruning strategy is similar to that adopted in the WAND method, with the advantage of a maximum score closer to the actual score of each document. With this advantage, the BMW method discards even more documents that do not have sufficient weight to be inserted on top of answers ranking. Pruning in the BMW occurs in two stages, (i) when the *Max Score* does not overcome the discarding threshold, as in the WAND method, and (ii) when the *Block Max Score* does not overcome the discarding threshold. The Block Max Score is computed based on the maximum score of the blocks in which the document can occur and is only checked when the Max Score exceeds the discarding threshold.

Discarding threshold is dynamically updated whenever a document has score higher than the current threshold. This document updates the heap of answers and the pruning threshold is updated with the weight of the smallest document in the heap.

For illustrative purposes, assume a scenario as shown in Figure 1. The pruning threshold indicates that it is necessary a document with a score higher than 3.2 to upgrade the current answers ranking. The figure describes a state of the processing were documents 20, 40 and 90, associated with each query term “yellow green blue” will be evaluated, in that order. The inverted list of the first document

to be processed has a maximum score of 2.5, i.e., the document 20, which only occurs in the list of term “yellow”, does not have sufficient weight to overcome the pruning threshold. However, the document 40, which can occur in the inverted list of the term “yellow”, may have a score up to 4.0 ($2.5 + 1.5 > 3.2$), higher than the current pruning threshold. Since such a possibility exists, the algorithm checks the Block Max Score of the pivot (document 40). If the pivot occurs in the inverted list of the term “yellow” and it has the highest score among all documents of its block, on both lists, its maximum score will be 3.3 ($2.3 + 1.0 > 3.2$), therefore higher than the pruning threshold. After those verifications, the document 40 has its full score evaluated. In this example, the complete score of the document is 3.3 ($2.3 + 1.0$), which exceeds once more the threshold pruning, so the document 40 is added to heap of answers, updating the pruning threshold. At this stage, the pointers of the inverted lists are moved so that they point to the next document higher than the evaluated pivot, and then the algorithm proceeds repeating the process.

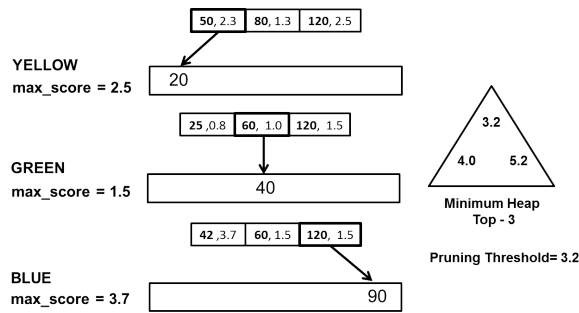


Fig. 1. Inverted lists for the query terms “yellow green blue”.

Some authors explore the partitioning approach of the inverted list in more than one tier to speed up query processing. One of the most recent, and also more successful partitioning strategy, is adopted by the BMW-CS algorithm [Rossi et al. 2013], which divides each inverted list into two tiers. The first tier has smaller size and is created using inputs that have the highest scores among the ones of each inverted lists, while the second tier has a larger index and contains the remaining entries.

The query processing is divided into two stages. The first, called the Candidate Selection, processes the first tier of the lists of the query terms. At this stage, the process is similar to the method used in the BMW, but also considers the maximum score of the document on the second tier. The value for the second tier is important, because there can be documents that are in the first layer, but also occur in the second layer to other query terms. The documents selected in the first stage are called candidates to be included in the answer. In the second stage, the second tier is traversed only to find occurrences of candidate documents in the second tier and compute the real score of such documents.

The BMW-CS method has the disadvantage of not preserving all the top-k answer results of a query, as documents that doesn't occur in the first tier of at least one of the query terms are not evaluated. In these cases, even if the documents have sufficient scores to be inserted into the answer, they will be discarded by the BMW-CS algorithm.

3. STUDY OF INITIALS THRESHOLDS FOR PRUNING

The pruning threshold is crucial for document discarding strategies to work correctly in the presented methods. At first, while the heap containing the top results is not full, the pruning threshold is zero. This means that, in the beginning, the pruning strategies are not applied by the BMW and its variants, since any considered document has a score higher than zero. After processing the first k documents, the methods obtains the first pruning threshold different from zero, starting the document discarding process that reduces the cost of query processing in the BMW algorithm and its variants.

As more documents are processed, the pruning threshold increases until achieve its maximum value, which is the minimum value among top results of the final ranking.

The presence of a low initial threshold makes that fewer documents are discarded at the beginning of the query processing in the BMW algorithm and its variants. Our first attempt to improve the algorithms was to study the performance of such algorithms when using as initial pruning threshold an estimated value that is both safe, guaranteed to be smaller than the maximum possible pruning threshold obtained at the end of processing, and large enough to impact the processing costs, having the closest possible value to the final pruning threshold.

The strategy for adoption of an initial threshold was applied in two scenarios: When we are interested in computing the top-10 results and the top-1000 for a given query. The two scenarios are considered typical in search systems and are usually studied in articles that deal with efficient query processing [Rossi et al. 2013]. To estimate the initial pruning threshold, we keep, during the indexing, the k -th highest score of each word of the collection in its corresponding entry in the inverted index structure. Given a query, it is estimated an initial minimum score value by taking the highest k -th score within the query terms. This choice ensures that the initial threshold will not discard answers that belong to the top k results, because it is known that there are at least k entries with a score at least equal to the chosen initial threshold. On the other hand, the adoption of this threshold allows to discard a priori entries that would have been taken into consideration when adopted an initial threshold equals zero.

In addition to propose strategies and try to determine an initial threshold, we also verified how much we can improve the performance with this type of strategy. To do this, the minimum score among k responses obtained from the query processing was stored. Then, the same queries were processed again, but starting the pruning threshold with its maximum value achieved by the end of query processing. We call this value as optimal threshold for being indicative of the highest gain we could get with the use of the initial pruning threshold strategy. The optimal threshold is only a reference value, since to obtain it we need to first to process the query, achieve the top k scores and then select the smaller among then as the optimal threshold.

4. BLOCKS OF SIZE VARIABLES

In the experiments presented by [Ding and Suel 2011] and [Rossi et al. 2013], fixed blocks of 128 documents were used in the creation of the skiplists. This means that for each 128 documents in an inverted list, an entry is created in the skiplist to accelerate the documents access and store discarding information, such as the highest score of each block (Block Max Score). As seen in the Figures 3(a), 3(b) and in the Table II, the smaller is the block size, the shorter is the processing time. This cost is directly related to the document search on the blocks to compute the real score.

Whenever the Max Score and the Block Max Score of the a pivot overcome the pruning threshold, the corresponding blocks will be accessed seeking the pivot document so that your real score is computed. Blocks are totally discarded when the Block Max Score of the pivot does not exceed the pruning threshold. Thus, it can be noticed that blocks with low Block Max Score have more probability to be discarded. In this case, we infer that adjacent blocks which have low value of Block Max Score could be integrated so that more documents could be discarded during an iteration of the query processor.

To determine whether the Block Max Score is low or not, the 1000-th highest score of each term in the collection is used as reference, making an association with the typical amount of inputs usually considered in the search results. Therefore, blocks with Block Max Score smaller than the reference value are considered blocks of low impact, and therefore it can merge with another block.

The Figure 2 illustrates an inverted list for a term A which is divided into n blocks, each one composed of 128 documents. In the example, the 1000-th highest score of the list has already been

Table I. Performance of BMW algorithm for fixed size blocks of 64, 128 and 256 documents. The number of blocks is represented by *.

		Pivots	Time(ms)
Blocks of 64 23.595.719*	Top-10	369.781.344	20
	Top-1000	1.002.528.576	51
Blocks of 128 11.798.404*	Top-10	426.335.552	22
	Top-1000	1.116.841.472	55
Blocks of 256 5.899.760*	Top-10	495.183.424	27
	Top-1000	1.223.783.040	59

computed and has value equal to 12. In this scenario, the blocks A_2 and A_3 are combined because they have maximum score smaller than the reference value 12. The blocks A_1 , A_4 and A_n remain as they were. The new block $A_2 + A_3$ will have 256 documents and Block Max Score equal to 8.




Fig. 2. Example of the inverted list with variable blocks.

Our hypothesis is that the heuristic of blocks union described above can be used to improve the pruning process in search engines such as BMW and BMW-CS. In the Experiments section, we show that our hypothesis is proven to work in practice and the heuristic proposed, when combined with the benefits of our proposal to estimate an initial pruning threshold, results in significant improvement in the performance indicators of the two algorithms.

5. EXPERIMENTAL EVALUATION

For the experiments, we used and modified the search system implemented by [Rossi et al. 2013]. The adopted system processes queries using either the algorithm BMW or the BMW-CS variant. We adopted the TREC GOV2 reference collection for the experiments. This consists of 25.205.179 million web pages crawled from the .gov domain in early 2004. The TREC GOV2 has 426 GB of text, composed of HTML pages and text extracted from pages in PDF and Postscript (PS) format. The full index has about 7 gigabytes of inverted lists and a vocabulary of about 4 million distinct terms. The TREC GOV2 collection was adopted due to the fact that it became a standard for studies that propose efficiency improvements in search engines, having been used in the experiments of the main studies found in the literature.

A set of 10.000 queries extract was randomly selected from the TREC 2006 efficiency queries. All stop-words of these queries have been removed. During the query processing, the entire index is loaded to memory in order to avoid any possible bias in the query processing time. All of these setup options were chosen for being similar to those adopted in previous studies [Strohman and Croft 2007; Ding and Suel 2011], which makes it easier to compare the studied methods. We ran the experiments at an Intel(R) Xeon(R), with X5680 Processor, 3.33GHz and 64GB of memory. All the experiments were performed in scenarios where the top-10 or top-1000 answers were returned. The recovery of the top-1000 answers was included to simulate an environment where the set top-1000 is used as input to a more sophisticated ranking method, as for example using a machine learning technique. The scenario of the top-10 answers was included to simulate a more usual scenario, where the user is only interested in a small list of results for his or her query. The algorithms were evaluated in terms of processing time and memory used.

The strategies were implemented in the BMW algorithm proposed by [Ding and Suel 2011] and the BMW-CS algorithm proposed by [Rossi et al. 2013]. These were developed in C++. Just as the original algorithms, we continued to use the probabilistic model Okapi BM25 as similarity function. For the BMW-CS algorithm, we varied the first tier size from 2 to 50 percent of the full index. We firstly ran the experiment with 2 percent, then 4, 6 and so on, always incrementing by 2 on each iteration of the experiment until 50 percent.

5.1 Results With Pruning Initial Threshold

The Figures 3(a) and 3(b) show gains by using the proposed strategies when compared to BMW algorithm without initial pruning threshold (Original BMW) for the top-10 and top-1000 queries, respectively. The figures also present the results of varying the size of the blocks generated. Experiments were executed with fixed size of 64, 128 and 256 documents. Fixed blocks of 64 documents result in smaller query processing time, but as indicated in the Table II, they form more blocks in the inverted index, requiring more memory to store the skiplists. In the opposite way, fixed blocks of 256 documents form fewer blocks, however result in a higher processing time. The larger the amount of formed blocks (number of entries in the skiplist structure), the higher is the cost of memory used in the processing of queries, as more space is necessary for storing these data.

The block size typically used in studies reported in the literature is of 128 entries. With this configuration, it is observed that the choice of an optimal initial threshold results in a gain of about 18% in the query processing time for the top-10 answers and 16% for the top-1000 answers, which would be the limit for any initial threshold estimation strategy. As can also be observed in the Figure 3, our heuristic of estimating an initial threshold also impacted the times in top-1000 answers, with a reduction of approximately 5.5% when compared to the original algorithm. The major gains occurred with blocks sized of 64 documents, where we had reductions of approximately 5.2% and 7.8% for top-10 and top-1000 answers, respectively.




Fig. 3. Performance of BMW algorithm on different approaches to initial pruning thresholds: Original BMW, BMW with tenth/thousandth highest score and BMW with the minimum score of the top-k answers to fixed blocks of 64, 128 and 256 documents.

Table II. Total number of entries of the skiplists for fixed blocks of 64, 128 and 256 documents.

Fixed 64	Fixed 128	Fixed 256
23.595.719	11.798.404	5.899.760

When looking at the experiments with BMW-CS algorithm, for fixed blocks of 64, 128 and 256 documents, and in the two tiers, the higher impacts of the method with an initial threshold strategy are obtained for top-1000 answers. Within this scenario, we calculated the average time between the variations of 18% and 50% of the first tier for all implementation options, as shown in Table III. The

table shows that the introduction of the thousandth as the initial threshold strategy reduces the query processing time by up to approximately 17%.

Table III. Average of the processing time in milliseconds calculated from the variation of 18% of the first layer, for top-1000 answers.

	Fixed 64	Fixed 128	Fixed 256
Original BMW-CS	54,65	57,94	59,94
BMW-CS with threshold	45,12	47,82	49,88
Reduction (%)	17,44	17,46	16,78

The Figures 4(a) and 4(b), respectively, show the behavior of the implementation of the initial threshold strategy and total blocks formed for the three scenarios when varying the size of first tier. It can be noticed that there are not much oscillation in the total of blocks as we increase the percentage of the first tier. Notice also that the use of fixed blocks of 64 entries implies in significant memory cost and, given its similarity in relation to the processing time with fixed blocks of 128 and 256 (Figure 4(a)), we considered as not good results, thus not a good option to be implemented. Such scenario had the approximate average of 23.6 thousand blocks, while blocks of 128 and 256 documents form about 12 thousand and 6 thousand blocks, respectively. Therefore, as is shown in Figure 4(a), since the processing time between the scenarios are quite similar, we can infer that the best implementation is one that has the lowest number of raised blocks, i. e., the scenario of fixed blocks with 256 entries.

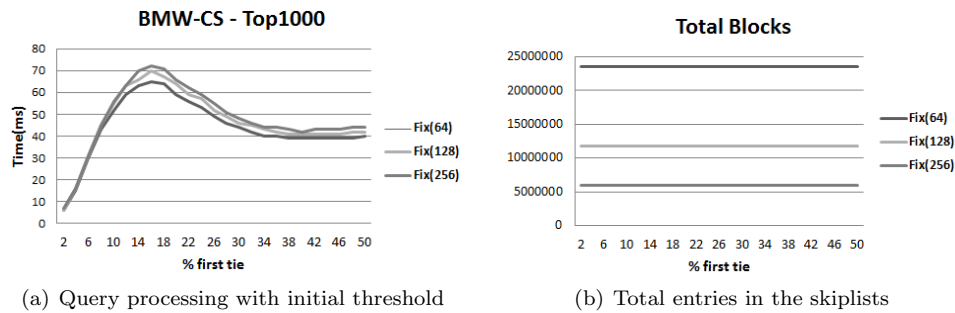


Fig. 4. Performance and total fixed blocks of the BMW-CS algorithm.

As already discussed in this section, the gain of processing time in relation to the introduction of the threshold for the BMW-CS algorithm was best seen for top-1000 answers. To better understand this behavior, the Figure 5 shows the experiment conducted with the BMW-CS algorithm with fixed blocks of 128 entries, with the first tier formed by 40% of the index. Exceptionally, this experiment involves other scenarios besides the top-10 and top-1000 answers returned in the query processing. For this one, the processing time for top-10, top-100, top-1000, top-5000 and top-10000 documents were recorded. By analyzing the graphic we can infer that the higher the amount of returned results (top) for each query, the larger is the impact of the insertion of the initial threshold. The initial threshold in the query processors, as already discussed, is only updated after the answers heap become completely filled. The initial threshold strategy produces better results for a higher top because when the top is small, for example heap that holds 10 answers, we will have few discards until the heap is full. However, when it is desired to obtain the top-5000 answers, for example, the heap will take much more time to be filled. The initial threshold, having a score of relevant value (in this case, the five thousandth), will bring more discards, accelerating the evaluation process until the heap have 5000 documents and the threshold is updated. Finally, even after the heap becomes filled, the bigger the number of answers computed, the bigger is the number of entries processed until the pruning threshold reaches the value of the k-th largest score among the query terms.




Fig. 5. Performance of the BMW-CS algorithm for fixed blocks with the first tier formed by 40% of the index.

5.2 Varying the Size of the Blocks

We started by evaluating the impact of varying the size of the blocks in the BMW. As a first strategy, we set a minimum and maximum size for the blocks. We varied blocks of at least 16, 32, 64, 128, 256 and 512 documents, composing the maximum of 1024 documents. The Figure 6 shows a comparison of this strategy with the Original BMW algorithm. We observed that the query processing time increases for both top-10 and top-1000 scenarios. When computing top-1000 results, and applying the variation with 512 blocks up to 1024 documents, the time increases from 52 milliseconds to a little over 60 milliseconds. On the other hand, we also see from the Figure that, despite the increase in query processing time, by using the strategy we managed to decrease to approximately 80% of the total number of blocks formed (Figure 6(c)) when compared with the Original BMW algorithm.

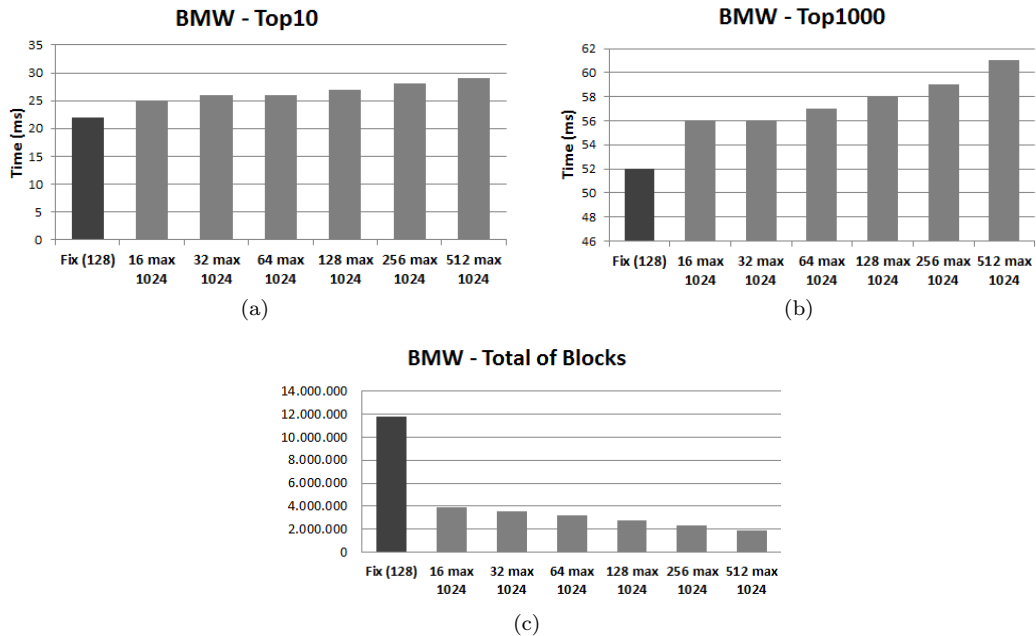


Fig. 6. Performance of the BMW algorithm with the use of the approaches of variation of the blocks size (maximum of 1024 documents).

Another strategy we have experimented was to reduce the maximum size of 1024 to 128 or 256

documents. The graph presented in the Figure 7 shows the results achieved with these new variations. There is a reduction in processing time for indexes formed by blocks with a maximum of 128 documents for the top-10 and top-1000 answers. Moreover, it can be noticed that we can maintain similar processing time to the Original BMW when using the maximum size of 256 documents per block. In Figure 7(c) it can be seen that with this implementation option it was possible to reduce by up to approximately 44% the total number of blocks formed. For example, generating blocks with minimum size of 32 and maximum of 256 documents resulted in the same query processing time obtained by using the Original BMW (52 milliseconds). However, the number of blocks generated by this choice was 35% lower, generating 7.676.292, against 11.798.404 blocks from the Original BMW. This is a significant reduction in the amount of memory to store information generated for each block.

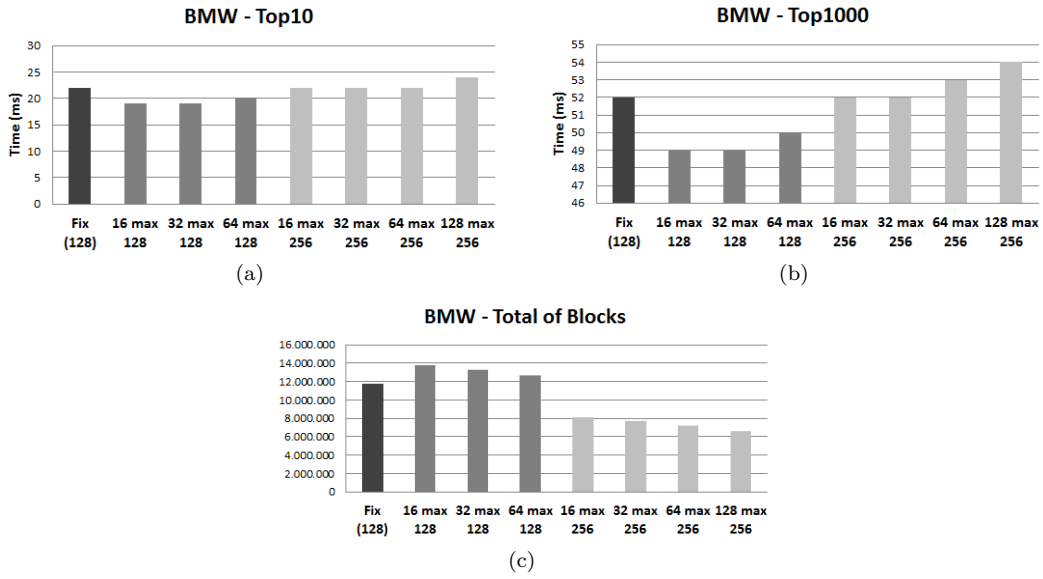


Fig. 7. Performance of the BMW algorithm with the use of the approaches of variation of the blocks size (maximum of 128 and 256 documents).

We analyzed some queries in order to identify which blocks were accessed during the stage of query processing, and thus create an index with a structure where adjacent blocks, not accessed, were aggregated into a single block. The objective was to handle queries to create a block structure where we would know *a priori* which of them would be discarded during the evaluations. At the end, we would have an indication of the time in which a query would be processed when applying the strategy of varying blocks size.

The Figure 8 describe a hypothetical scenario of the processing of a query “A B C” with fixed blocks of 128 documents. At the first moment (Figura 8(a)), the lowest document of the answers heap presents score equal to 14 and the selected pivot is the document 3097. The document 3097 is present in the inverted lists of the terms of A, B and C, and we have the value of the scores of each block to which it belongs. Since the sum of the maximum score of the block ($3.6 + 6.4 + 3.3$) of all occurrences of the pivot document is not higher than the minimum heap (current threshold), a new pivot will be selected. When the new pivot is selected (document 4000), the sum of the maximum score of block in all instances of this document in the index also does not overcome the current threshold. The algorithm then follows seeking possible candidates to be returned as response.

In this hypothetical scenario, the blocks which contains the document 3097 and 4000, for the term A, were not read, but the remaining blocks have been accessed. To handle this query, we join the

blocks that have not been accessed, because we intuitively think that the algorithm may skip more documents (blocks formed by more entries) reducing processing costs. The Figure 8(c) illustrates this new scenario where the query “A B C” presents a structure with variable blocks. At this point of the query processing, the pivot document 3097 will be evaluated. With the blocks structure modified, have that all occurrences of the document in all query terms result in a total value of maximum score of 15.1 ($5.4 + 6.4 + 3.3$). This value is higher than the current threshold. This indicates that the pivot document will have its complete score computed, i. e., the algorithm will scan a larger block looking the pivot document, which will affect its performance.

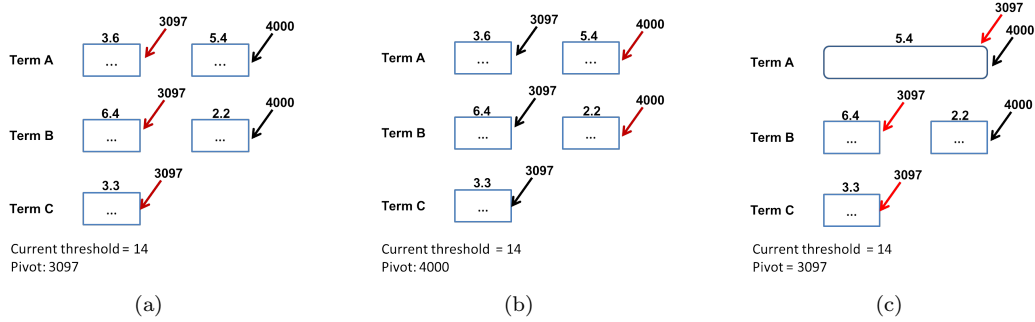


Fig. 8. Hypothetical Scenario A B C.

Changes in the size of the blocks can produce side effects on the performance of the algorithms, as they also change the maximum score values of the blocks and, consequently, affect also the document pruning strategies, either on the BMW and BMW-CS algorithms. Nonetheless, as seen in this section, we may achieve a reduction in the total number of blocks formed without having a negative impact on query processing times. But bringing together the use of variable size blocks and the initial threshold insertion strategy, we obtained gains in processing time and, at the same time, reduction in the size of the skiplists used to represent documents blocks when processing queries.

The Figure 9 shows the study conducted when using fixed blocks with less and more documents (64, 128 and 256), in the two tiers of the inverted list in the BMW-CS algorithm. All implementation options of fixed blocks demonstrated and explained in this study used the initial pruning threshold strategy. We joined three variations on each chart, choosing those that best define the algorithm's behavior when using fixed size blocks. It is important to remember that, like the BMW, the less are in each block, more entries will be in the skiplists. This indicates that there will be greater memory costs, given that more blocks will be created.

The Figures 9(a), 9(c) and 9(e) indicate that blocks formed with fewer documents reduce query processing time. However, when analyzing memory costs added up to a medium processing time, we realize that fewer entries in the skiplists, i. e., blocks formed with more document, between the variations presented, generate a better setting for the algorithm. We also note that the difference of the processing time between all exposed variations (Figure 9) differ by a few milliseconds. This shows that we can reduce memory costs without causing much impact on the algorithm processing time. In the worst case, the blocks formation with few documents, for example, 64 documents in both tiers, have a total of about 24 million entries in the skiplists, while for blocks with 256, in both tiers, have about 6 million entries. By implementing fixed blocks with 128 documents in the two tiers we can achieve median costs in relation to the total number of formed blocks (average of approximately 11 million) for each percentage of the first tier.

We also conducted experiments to study the application of variable size blocks in the BMW-CS algorithm. We varied the size of the blocks from 64 up to 1024 documents in both tiers and analyzed the algorithm performance for each percentage of variation of the first tier, to top-10 and top-1000

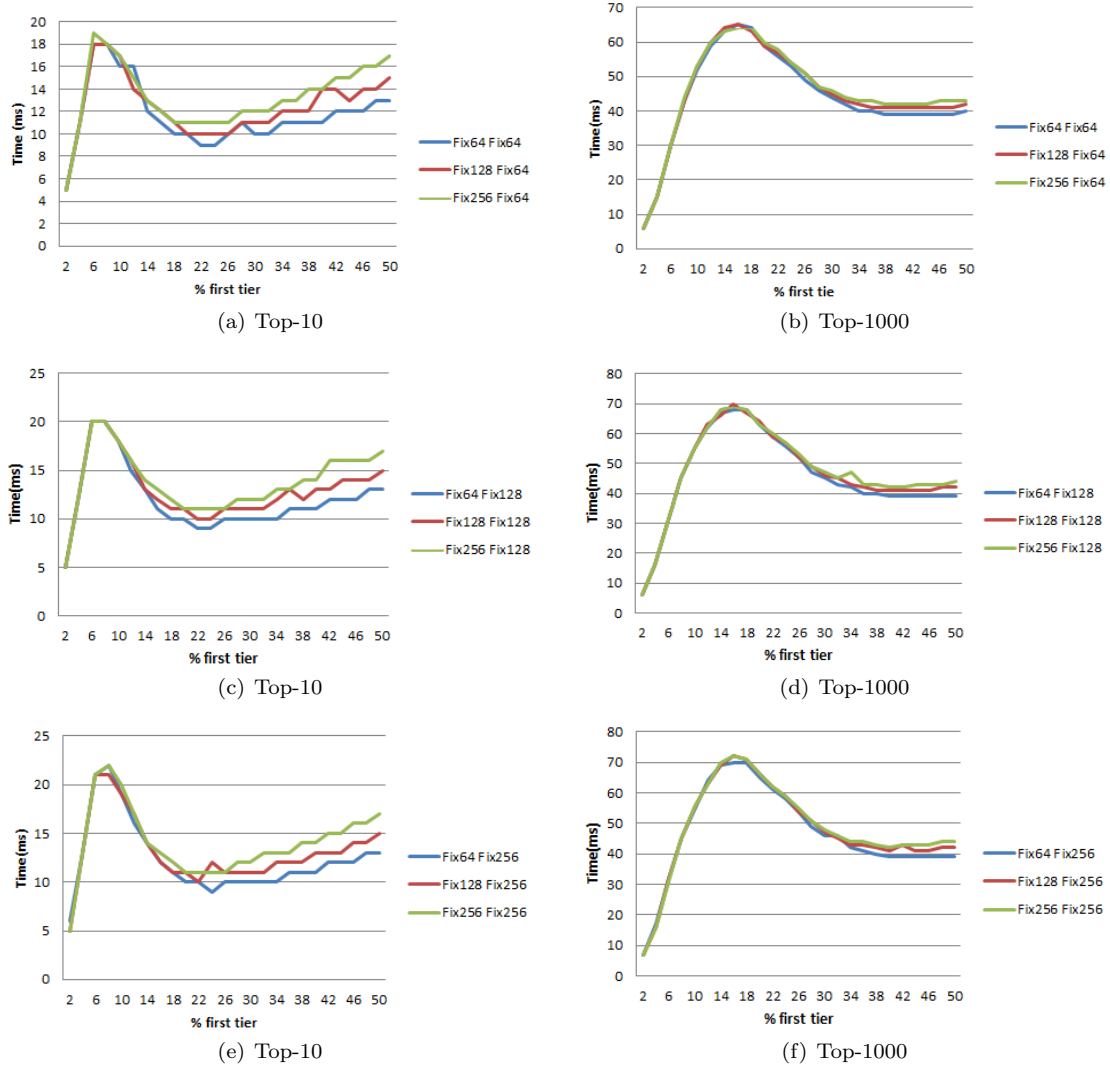


Fig. 9. Performance of the BMW-CS algorithm varying the number of fixed blocks in the two tiers of the index.

answers. All implementations options (with the first, second or both tiers variables) were compared to the BMW-CS algorithm using fixed blocks with 128 documents and the initial threshold strategy proposed by us. The experiments (Figure 10) show that we did not obtain gains in processing time for both the top-10 and top-1000 answers. The Table IV shows the average time between variations of 18% to 50% of the first layer for all implementations options in the top-1000 answers scenario. It was found then that the processing times obtained by the approaches are quite similar.

Taking into consideration that the processing time between the strategies of variables blocks presents no great variations, we can take into account, as in previous approaches, the possibility of reducing memory costs. The Figure 11 makes a comparison of the total of blocks along all implementation options, including that used in the Original BMW-CS algorithm. Also in the figure, it can be noticed that the BMW-CS algorithm with fixed blocks of 128 documents turns to be the worst case of the implementation, considering that this one results in the average of about 11.8 million of blocks formed.

The implementation option of varying the two tier with a minimum of 64 and maximum of 256 documents generates an approximate reduction of 43% on average when compared to the Original

Table IV. Variables Blocks - Average of the processing time in milliseconds calculated from the variation of 18% of the first layer.

	1 ^a variable tier	2 ^a variable tier	1 ^a and 2 ^a variable tiers
Min 64 Max 256	48,47	48,88	49,18
Min 64 Max 512	48,76	49,65	50,41
Min 64 Max 1024	49,00	50,12	51,35
Original BMW-CS: 57,94			
Original BMW-CS with threshold: 47,82			

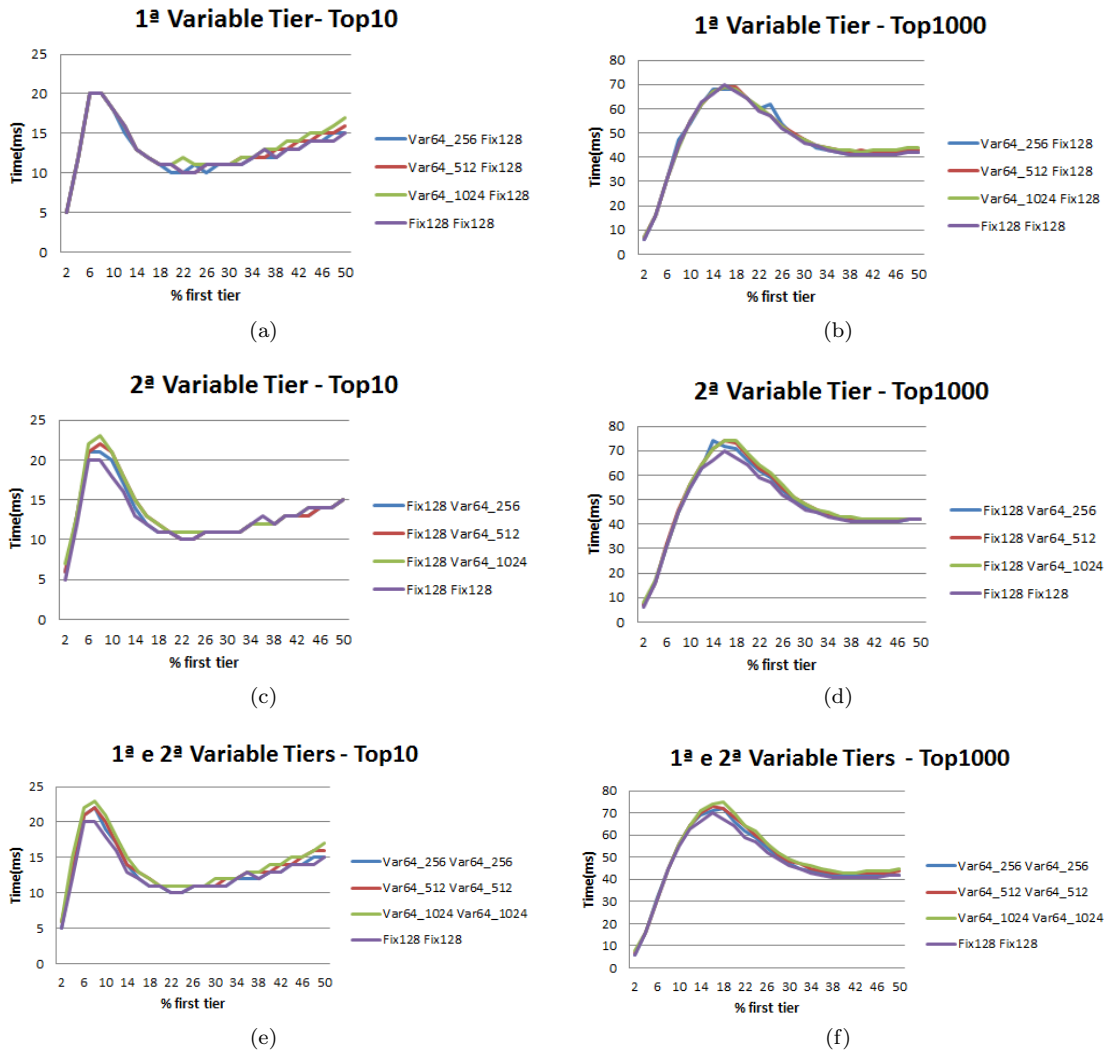


Fig. 10. Perfomance of the BMW-CS algorithm varying the number of blocks in the two tiers of the index.

BMW-CS. Such results are obtained with a little change in processing time. For example, in our experiments, the average processing time to consider all size thresholds studied for the first tier (Table IV) in the Original BMW-CS algorithm, that uses the initial threshold strategy, is of 47.82(ms), while we have an average time of 49.18 (ms) to vary the size of the blocks in two tiers with a minimum of 64 and maximum of 256 documents.




Fig. 11. Variable blocks - Total number of entries of skiplists for each percentage of the first tier.

6. CONCLUSIONS AND FUTURE WORK

The results obtained show that the estimation of an initial threshold when discarding documents can improve the performance of the BMW algorithm and its BMW-CS variant for both scenarios, the top-10 and top-1000 answers. The biggest gains on modifying the BMW for use on our initial pruning threshold strategy come when with fixed block of size 64 occur, with a reduction of about 5.2% and 7.8% in query processing times for top-10 and top-1000 answers, respectively. Even greater gains were obtained when the modified version of the BMW-CS algorithm were used, with a reduction of approximately 17% in the query processing times.

The strategy of changing the size of the blocks of the index structure did not result in the reduction of query processing times, but provided significant decrease in the number of skiplists entries, impacting then the amount of extra memory required by each method to process queries. For the BMW algorithm, the best configuration found in the experiments uses blocks with the maximum of 256 entries. In this case, the same query processing times of the original algorithm was maintained, but the total number of blocks created decreased by about 35%. In the BMW-CS algorithm, the option of varying the size of the blocks in the two tiers, with a minimum of 64 and maximum of 256 documents, our experiments have shown an average reduction of 43% when compared to the original BMW-CS, without significant changes in the query processing times.

For future work, it would be interesting to study the impact of the techniques here implemented on scenarios where the search engine index does not fit in the main memory. In such situations, gains of space in the skiplists may become more relevant and have more positive impact in the query processing times, since small skiplists have a greater chance of being completely placed into the main memory. It would also be interesting to study the impact of threshold detection heuristics when applied to other practical problems, such as document classification [Cristo et al. 2003].

REFERENCES

- BAEZA-YATES, R., RIBEIRO-NETO, B., ET AL. *Modern information retrieval*. ACM press New York, 2011.
- BRODER, A. Z., CARMEL, D., HERSCOVICI, M., SOFFER, A., AND ZIEN, J. Efficient query evaluation using a two-level retrieval process. In *Proceedings of the twelfth international conference on Information and knowledge management*. ACM, pp. 426–434, 2003.
- CRISTO, M., CALADO, P., DE MOURA, E. S., ZIVIANI, N., AND RIBEIRO-NETO, B. A. Link information as a similarity measure in web classification. In *String Processing and Information Retrieval, 10th International Symposium, SPIRE 2003, Manaus, Brazil, October 8-10, 2003, Proceedings*. pp. 43–55, 2003.
- DING, S. AND SUEL, T. Faster top-k document retrieval using block-max indexes. In *Proceedings of the 34th international ACM SIGIR conference on Research and development in Information Retrieval*. ACM, pp. 993–1002, 2011.
- ROBERTSON, S. E. AND WALKER, S. Some simple effective approximations to the 2-poisson model for probabilistic

- weighted retrieval. In *Proceedings of the 17th annual international ACM SIGIR conference on Research and development in information retrieval*. Springer-Verlag New York, Inc., pp. 232–241, 1994.
- ROSSI, C., DE MOURA, E. S., CARVALHO, A. L., AND DA SILVA, A. S. Fast document-at-a-time query processing using two-tier indexes. In *Proceedings of the 36th international ACM SIGIR conference on Research and development in information retrieval*. ACM, pp. 183–192, 2013.
- SALTON, G., WONG, A., AND YANG, C.-S. A vector space model for automatic indexing. *Communications of the ACM* 18 (11): 613–620, 1975.
- STROHMAN, T. AND CROFT, W. B. Efficient document retrieval in main memory. In *Proceedings of the 30th annual international ACM SIGIR conference on Research and development in information retrieval*. ACM, pp. 175–182, 2007.