

IBM Research Report

Query Indexing with Containment-Encoded Intervals for Efficient Stream Processing

Kun-Lung Wu, Shyh-Kwei Chen, Philip S. Yu
IBM Research Division
Thomas J. Watson Research Center
P.O. Box 704
Yorktown Heights, NY 10598



Research Division
Almaden - Austin - Beijing - Haifa - India - T. J. Watson - Tokyo - Zurich

Query Indexing with Containment-Encoded Intervals for Efficient Stream Processing

Kun-Lung Wu, Shyh-Kwei Chen and Philip S. Yu
IBM T.J. Watson Research Center
Hawthorne, NY 10532
{klwu, skchen, psyu}@us.ibm.com

Abstract

Many continual range queries can be issued against data streams. To efficiently evaluate continual queries against a stream, a main memory-based query index with a small storage cost and a fast search time is needed, especially if the stream is rapid. In this paper, we study a CEI-based query index that meets both criteria for efficient processing of continual interval queries. This new query index is an indirect indexing approach. It centers around a set of predefined virtual *containment-encoded intervals*, or CEIs. The CEIs are used to first decompose query intervals and then perform efficient search operations. The CEIs are defined and labeled such that containment relationships among them are encoded in their IDs. The containment encoding makes decomposition and search operations efficient; from the encoding of the smallest CEI containing a data point, the encodings of other containing CEIs can be easily derived. Closed-form formulae for the bounds of the average index storage cost are derived. Simulations are conducted to evaluate the effectiveness of the CEI-based query index and to compare it with alternative approaches. The results show that the CEI-based query index significantly outperforms existing approaches in terms of both storage cost and search time.

1 Introduction

Recently, many data stream applications have been recognized [1, 10]. For example, financial applications, network monitoring, security, telecommunications data management, sensor networks and other applications can be modeled as data streams. In a stream model, individual data items can be relational tuples with well-defined attributes, e.g., network measurements, call records, meta data records, Web page visits, sensor readings, and so on. These data items flow in the network in the form of streams continually and perhaps rapidly.

In order to monitor a data stream and take proper actions, if needed, a large number of range queries or filtering predicates can be created and evaluated continually against each data item in the incoming stream. For example, in a financial stream application, various continual range queries can be created to monitor the prices of stocks, bonds, or interest rates. In a sensor network stream application, continual range queries can be used to monitor the temperatures, flows of traffic or other readings.

One approach to quickly evaluating these continual range queries is to use a query index. Each data value in the incoming stream is used to search the query index and identify those relevant queries that contain the data value. Though conceptually simple, it is quite challenging to design such an interval query index in a streaming environment, especially when the stream is rapid. The interval query index is preferably main memory-based and it must have two important properties: low storage cost and excellent search performance. Low storage cost is important so that the entire query index can be loaded into main memory. Excellent search performance is critical so that the query index can handle a rapid stream.

The goal of interval query indexing is to help identify, from a set of query intervals, those that contain an incoming data value. This was referred to as a *stabbing query problem* [11], i.e., finding all the intervals that are stabbed by the data value. Fig. 1 shows an example of a stabbing query problem, where various intervals are drawn as horizontal line segments. The intervals that intersect with the vertical line drawn at a data value are those that contain that data value. For instance, query intervals q_1, q_3, q_7 and q_9 contain data value a_i while q_2, q_4, q_6, q_8 and q_9 contain data value a_j .

There have been existing approaches to interval indexing to help solve the stabbing query problem, such as segment trees, interval trees [2], R-trees [5], interval binary search trees (IBS-trees) [7] and interval skip lists (IS-lists) [8]. However, they are generally not suitable for streams. Segment trees and interval trees generally work well in a static environment, but are not adequate when it is necessary to dynamically add or delete intervals. Originally designed to handle spatial data, such as rectangles, R-trees can be used to index intervals. However, as indicated in [8],




Figure 1: Example of a stabbing query problem.

when there is heavy overlap among intervals, search performance degenerates quickly. Moreover, R-trees are primarily disk-based indexes.

IBS-trees and IS-lists were both designed for main memory-based interval indexing [7, 8]. They were the first dynamic approaches that can handle a large number of overlapping intervals. As with other dynamic search trees, IBS-trees and IS-lists require $O(\log(n))$ search time and $O(n \log(n))$ storage cost, where n is the total number of intervals. Moreover, as pointed out in [8], in order to achieve the $O(\log(n))$ search time, a complex “adjustment” of the index structure is needed after an insertion or deletion. The adjustment is needed to re-balance the index structure. More importantly, the adjustment makes it difficult to reliably implement the algorithms in practice. Previous studies [8] indicated that IS-lists perform better than IBS-trees and are easier to implement, even though dynamic adjustments of the interval skip lists are still needed.

There are other spatial indexing methods that can be used to handle one dimensional intervals [9, 4, 11]. However, most of them are designed specifically for multidimensional spatial objects. In addition, they tend to be secondary storage-based indexes. In contrast, a main memory-based index is usually preferred for stream processing, especially if the number of intervals indexed is large and the stream is rapid.

All the aforementioned query indexes are *direct indexing* approaches. The query index is built directly with the query boundaries and search is conducted by comparing a data value with the query boundaries maintained in the index. Even with search time of $O(\log(n))$, such as the IS-lists approach, the search time may not be fast enough to handle a rapid stream, especially if n is large.

In contrast, we propose an *indirect indexing* approach. It is based on the concept of *virtual*

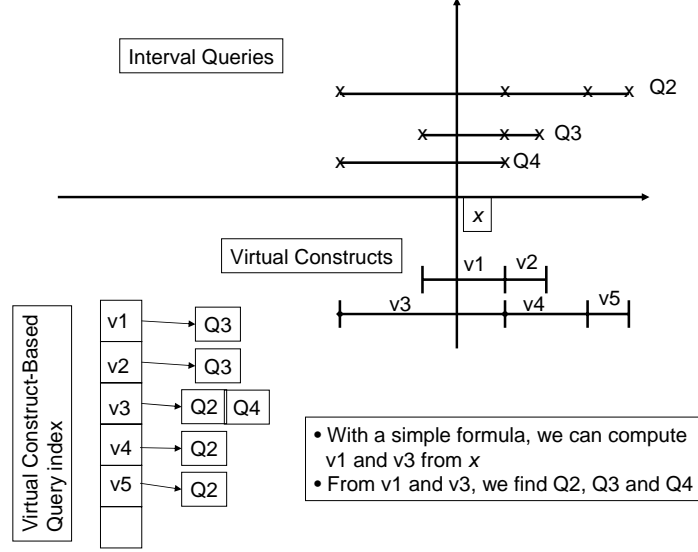


Figure 2: Example of a VC-based query index.

constructs (VCs). A set of virtual constructs is predefined and properly labeled. Each query is decomposed into one or more VCs and the query ID is inserted into the ID lists associated with the decomposed VCs. Search is conducted indirectly via the VCs. No comparison of a data value with any of the query boundaries is needed during a search. Because of the query decomposition, the search result is contained in the ID lists associated with the covering VCs of a data value. We define VCs in such a way that there is only a small and fixed number of VCs that can cover any data value. As a result, search time is independent of n , in terms of locating the covering VCs.

Fig. 2 shows an example of a VC-based query index. We show 5 virtual construct intervals, v_1, \dots, v_5 , which are used to decompose query intervals, Q_2, Q_3 and Q_4 . After decomposition, query IDs are inserted into the associated ID lists. To perform a search on data value x , we first find the covering VCs for x . In this case, they are v_1 and v_3 . The search result is contained in the ID lists associated with v_1 and v_3 . From the VC-based query index, we find Q_2, Q_3 and Q_4 .

Though conceptually simple, there are challenges in a VC-based query index. These include the definition and proper labeling of the virtual constructs, the decomposition of query intervals and the computation of the covering VCs that contain any data value during a search. These issues must be carefully addressed in order to have a query index that has both low storage cost and fast search time.

In this paper, we present a CEI-based query index that addresses all these challenging issues. It is memory-based and has the properties of low storage cost and fast search operation, partic-

ularly suitable for stream processing. Unlike the IS-lists approach, no re-balancing is required. Hence, it is much easier to implement than IS-lists. The new query index is centered around a set of predefined virtual *containment-encoded intervals*, or CEIs. We partition the range of interest of a numerical attribute into multiple segments. For each segment, a set of virtual CEIs are predefined. The CEIs are defined and labeled such that the containment relationships among them are encoded in their IDs. The containment encoding makes decomposition and search operations efficient because integer additions and logical shifts can be used to carry out most of the operations. Simulation studies show that the CEI-based query index outperforms the IS-lists approach in terms of both storage cost and search time.

The paper is organized as follows. Section 2 describes alternative virtual construct intervals and their associated challenges. Section 3 describes the details of the CEI-based query index, including containment-encoded intervals, query decomposition and search algorithm. Section 4 analyzes the average index storage costs for various VC-based query indexes. Section 5 presents performance studies based on simulations. Finally, Section 6 summarizes the paper.

2 Virtual construct-based query indexes

2.1 System model

There are two kinds of intervals discussed in this paper: (a) query or predicate intervals and (b) virtual construct intervals. Query intervals are real. They are defined by users over a numerical attribute to monitor a data stream. On the other hand, virtual construct intervals are created to decompose query intervals and perform fast search operations.

We focus on simple queries, where each query is specified with a single interval predicate on a numerical attribute. However, the result of this paper is applicable to complex queries, where each query is specified with a conjunction of multiple interval predicates on various attributes. For example, a CEI-based query index can be maintained for each attribute in a two-phase algorithm involving complex queries, such as the ones presented in [3, 13]. In the first phase, search operations against the CEI-based indexes are performed at individual attributes. The matched results are then merged in the second phase.

We assume that all query intervals are specified on the same attribute. The attribute can be of integer or non-integer type. Query intervals have two integer endpoints with an inclusive left endpoint and an exclusive right endpoint.¹ Namely, they can be specified as $[a, b)$, where

¹If the endpoints of a query interval are not integers, we can expand them to the nearest integers. The expanded interval is then decomposed and indexed. The VC-based indexing is still effective in identifying a set of candidate queries. Nevertheless, a final checking is needed to determine if the candidate queries indeed contain the data value.

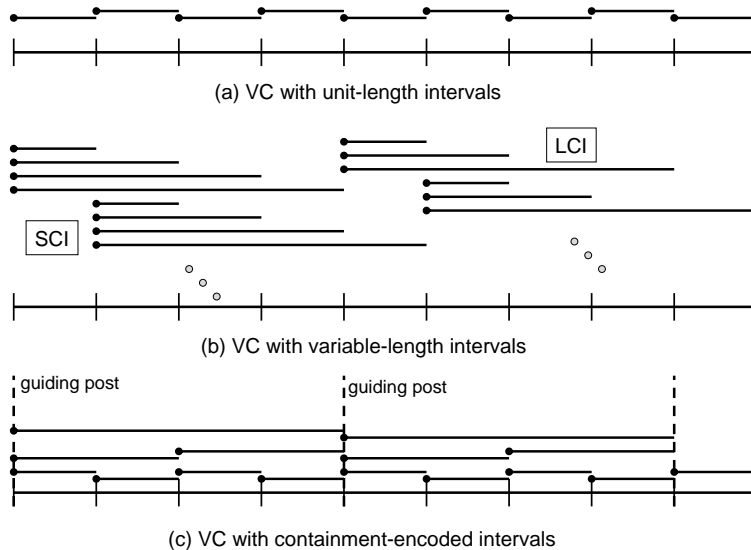


Figure 3: Example of three alternative virtual construct intervals.

a and b are integers and $a < b$. However, data values used for search can be any real numbers. Other design issues will be discussed on Section 3.4 for cases where both endpoints are inclusive or there is only one endpoint.

2.2 Alternative VCs and their challenges

In a VC-based query index, a query ID q is replicated by d times, where d is the number of VCs used in the decomposition of q . For each VC, a pointer is needed to maintain the associated ID list. Hence, the storage cost of a VC-based query index depends on two parameters: the total number of VCs defined, N , and the total number of query IDs inserted due to decomposition, D . Assume C is the number of VCs that can possibly cover a data value. The search time of a VC-based query index depends on C . Different VC-based indexes have different N 's, D 's and C 's, resulting in different storage costs and search times.

Fig. 3 shows three kinds of alternative virtual constructs and Table 1 provides a summarized comparison of the three. The first one is the simplest approach. It uses *unit-length intervals*, or UI (see Fig. 3(a)). The total number of VCs defined N equals r , where r is the scope or range of the attribute. With UIs, labeling and decomposition are obvious. We simply label the unit-length intervals sequentially. The decomposition is by partitioning a query interval into one or more unit-length UIs. Search time is extremely fast because there is only a single VC that can cover any given data value x . However, the storage cost can be high. This is because we need to use \bar{w} UIs to decompose a query interval on average, where \bar{w} is the mean length of a

Table 1: Comparison of three alternative virtual construct intervals.

	UI	LCI	CEI
VCs defined	r	$r(1 + \log(L))$	$2r - r/L$
covering VCs for any point	1	$2L - 1$	$\log(L) + 1$
labelling scheme	obvious	simple	simple, but non-obvious
decomposition algorithm	obvious	simple	simple, but non-obvious
storage cost	high	moderate	low
search time	fast	moderate/slow	fast

query interval. Hence, the cost for storing the ID lists becomes large. If \bar{w} is large, the UI-based scheme cannot be scaled to a large n because the storage cost can be too large to fit into main memory.

The second kind of VCs aims to reduce the storage cost by using multiple variable-length VCs. A larger-sized VC can be used for decomposition, reducing the number of VCs in the decomposition. Since a query ID is inserted into all the decomposed VCs, the cost for the ID lists is reduced. In its simplest form, we define for each integer attribute value a set of L virtual intervals with length of $1, 2, \dots, L$, where L is the maximal length of a VC. These L VCs all start at the same position but end at different positions. We call it a simple construct interval (SCI) approach (see Fig. 3(b)). The total number of VCs defined N equals rL . More VCs increase the storage cost for maintaining the pointers to the ID lists, offsetting the benefit of low storage cost achieved due to fewer decomposed VCs.

A more efficient variable-length VC approach would be a logarithmic construct interval (LCI) approach (see Fig. 3(b)). In LCI, the total number of VCs defined N equals $(1 + \log(L))r$ because we use $1 + \log(L)$ variable-length VCs for each integer attribute value. For the rest of the paper, we use LCI as the representative for the variable-length interval approach. Both labeling and decomposition are simple under the LCI-based approach. One can sequentially label the $(1 + \log(L))$ VCs with the same left endpoints. The decomposition tries to use as few LCIs as possible by always choosing the largest available LCI for decomposition beginning from the left endpoint of the query interval. It can be proved that the number of LCIs that can cover any data value x is $2L - 1$ [12]. The storage cost is moderate for LCI. However, the search time becomes moderate or slow, especially if a large L is chosen in order to reduce the storage cost.

In this paper, we describe in detail a third alternative approach using containment-encoded intervals (see Fig. 3(c)). It has the advantages of both low storage cost and fast search time. These advantages make CEI-based query index particularly suitable for a streaming environment. In CEI, the range of attribute values are partitioned into segments of length L , where L is an




Figure 4: Example of containment-encoded intervals and their local ID labeling.

integer that is a power of 2. Within each segment, we define $2L - 1$ containment-encoded intervals. As a result, the total number of VCs defined is only $2r - r/L$, and the number of covering VCs for any data value x is $1 + \log(L)$. However, the labeling scheme and decomposition algorithm are non-obvious. They must be labeled carefully so that both decomposition and search operations can be carried out efficiently. In the rest of the paper, we will present our efficient solutions which are also easy to implement.

3 CEI-based query indexing

3.1 Containment-encoded intervals

Fig. 4 shows an example of containment-encoded intervals and their local ID labeling. Assume the range of interest of a numerical attribute A is $[0, r)$. First, we partition r into r/L segments of length L , denoted as S_i , where $i = 0, 1, \dots, (r/L - 1)$, $L = 2^k$, and k is an integer. Here, we assume r is a multiple of L . If not, we can simply expand it and make it so without impacting the system performance. Segment S_i contains all the attribute values in $[iL, (i + 1)L)$. Namely, $S_i = \{x | x \in A \wedge iL \leq x < (i + 1)L\}$. Segment boundaries can be treated as guiding posts. Then, we define $2L - 1$ CEIs for each segment as follows: (a) Define 1 CEI of length L , containing the entire segment; (b) Define 2 CEIs of length $L/2$ by partitioning the segment; (c) Define 4 CEIs of length $L/4$ by partitioning the segment; (d) Continue the process until L CEIs of unit length ($L/L = 1$) are defined. For example, there are 1 CEI of length 8, 2 CEIs of length 4, 4 CEIs of




Figure 5: Example of a perfect binary tree and its labeling.

length 2 and 8 CEIs of length 1 in Fig. 4.

Property 1 *The total number of CEIs in a segment is $\sum_{i=0}^{k} 2^i = 2L - 1$, where $L = 2^k$.*

These $2L - 1$ CEIs are defined to have containment relationships among them in a special way. Every unit-length CEI is contained by a CEI of size 2, which is in turn contained by a CEI of size 4, which is in turn contained by a CEI of size 8, ... and so on.

Now, we describe the labeling of CEIs with containment encoding. The ID of a CEI has two parts: the segment ID and the local ID. For each segment, we assign $1, 2, \dots, 2L - 1$ to each of the $2L - 1$ CEIs as their local IDs. The local ID assignment follows the labeling of a perfect binary tree (see Fig. 5.) Local ID 1 is assigned to the CEI with length of L . Local ID 2 is assigned to the left CEI of length $L/2$ while local ID 3 is assigned to the right CEI of length $L/2$. Fig. 4 shows the assignment of local IDs to CEIs within a segment.

The global unique ID for a CEI in segment S_i , where $i = 0, 1, \dots, (r/L) - 1$, is simply computed as $l + 2iL$, where l is the local ID. Note that we assign $2L$ local IDs for every segment, even though there are only $2L - 1$ CEIs. Local ID 0 is not used.

With local IDs, all the CEIs in a segment can be organized as a perfect binary tree. The CEI with local ID 1 and length L is the root node, which contains two child CEIs of length $L/2$ whose local IDs are 2 and 3, respectively. The leaf CEIs have unit length and local IDs $L, L + 1, \dots$ to $2L - 1$. Figs. 4 and 5 show that, for $L = 8$, the leaf CEIs have local IDs from 8, 9, \dots , to 15. The labeling of CEIs based on a perfect binary tree has many nice properties. These properties make possible efficient decomposition and search for the CEI-based query index. Due to containment

encoding, almost all of the decomposition and search operations can be efficiently carried out.

Property 2 *There are $k + 1$ levels in the perfect binary tree formed with all the CEIs defined within a segment. Level i has 2^i CEIs, each with length $L/2^i$, for $i = 0, 1, \dots, k$. The local IDs of CEIs at level i starts from 2^i to $2^{i+1} - 1$.*

Property 3 *Any non-root CEI is fully contained by its parent.*

Property 4 *For any non-root CEI with a local ID l , the local ID of its parent can be computed by $\lfloor l/2 \rfloor$, or a logical right shift by 1 bit of the binary representation of l .*

Property 5 *For any non-leaf CEI with a local ID l , the local IDs of its children are $2l$ and $2l + 1$, respectively. Note that $2l$ can be computed by a logical left shift by 1 bit of the binary representation of l .*

Property 6 *For any data value $x \in A$, there are exactly $k + 1$ CEIs that contain it because there are $k + 1$ levels of non-overlapping CEIs in a segment.*

Property 7 *For any data value $x \in A$, the global ID of the unit-length CEI that contains x is $l + 2^k L$, where $i = \lfloor [x]/L \rfloor$, and $l = [x] - iL + L$.*

From Property 2, the unit-length CEIs in any segment have local IDs from $2^k = L$ to $2L - 1$. Hence, the local ID of the CEI containing x is $[x] - iL + L$, where $i = \lfloor [x]/L \rfloor$. Note that i can also be computed as $\lfloor [x]/L \rfloor$. However, it would involve a complex floating point division because x is a non-integer. In contrast, $\lfloor [x]/L \rfloor$ can be computed by first an integer conversion from a floating point number, $[x]$, and then a logical right shift by k bits because $L = 2^k$.

3.2 Query insertion and deletion

To insert a query interval, it is first decomposed into one or more CEIs, then its ID is inserted into the ID lists associated with the decomposed CEIs. The CEI-based query index maintains a set of query ID lists, one for each CEI. Similar to the inverted lists typically used in information retrieval, the query ID list associated with a CEI maintains all the query intervals that contain that CEI.

The decomposition algorithm is conceptually simple. It involves the concepts of guiding posts and remnant intervals. We use segment boundaries as guiding posts. For each query $q : [a, b)$, we define two boundary guiding posts: P_L and P_R . The left boundary guiding post $P_L = L \lceil a/L \rceil$ is the leftmost guiding post that is greater than or equal to a . The right boundary guiding post $P_R = L \lfloor b/L \rfloor$ is the rightmost guiding post that is smaller than or equal to b . These two

boundary guiding posts are used to compute two remnant intervals, if any. The left remnant interval, R_L , is defined as $[a, P_L)$. If P_L equals a , then R_L is empty. The right remnant interval, R_R , is defined as $[P_R, b)$. If P_R equals b , then R_R is empty.

Fig. 6 shows the pseudo code for the decomposition algorithm. First, we compute the boundary guiding posts, P_L and P_R , and the remnant intervals, R_L and R_R . Then, we check if P_L is greater than P_R . If yes, it means that q completely lies between 2 guiding posts without intersecting any. In this case, we call function **BestPartition** to decompose q . On the other hand, if $P_L < P_R$, then we insert the query ID q into the largest CEI of each segment that is completely contained by q between the two boundary guiding posts. This process repeats for $(P_R - P_L)/L$ times. After that, we use **BestPartition** to further decompose the two remnant intervals, if there is any.

There are more than one way to decompose an interval residing within a segment, or between two consecutive guiding posts. For example, we can always decompose it using consecutive unit-length CEIs. As a query ID is inserted into all the ID lists associated with the decomposed CEIs, using more CEIs means more storage cost for the index. The goal of **BestPartition** algorithm is thus to minimize the number of CEIs used in a decomposition. In other words, we want to use maximal-sized CEIs, if possible, for decomposition. **BestPartition** is based on the following property, which is derived from CEI labeling and can be easily observed from Fig 4.

Property 8 *Two CEIs at level i with consecutive local IDs can be replaced by a CEI at level $i - 1$ of double length only if the smaller ID of the two CEIs at level i is even.*

BestPartition can be viewed as a *merging* process, starting from the leftmost unit-length CEI at level k towards the root CEI at level 0. For any interval of length m , where $0 < m < L$, which lies between two consecutive guiding posts, it can be decomposed into m unit-length CEIs at level k . These unit-length CEIs have consecutive local IDs. The algorithm initializes a remnant interval R to be $[s, e)$ which is the interval to be decomposed. It starts with the local ID of the leftmost unit-length CEI. If the ID is an even number, then we try to replace two children CEIs with their parent CEI. Namely, we double the size of the CEI used in the decomposition. This is accomplished by dividing the ID by 2, or right shifting by 1 bit. Of course, if the right endpoint of the interval e is less than the right endpoint of the parent CEI, then we cannot use the parent CEI. In that case or if a local ID is an odd number, then we output a maximal CEI, remove the maximal CEI from R , and continue to decompose R from the local ID of the leftmost unit-length CEI for R . This process ends when R is empty.

Fig 7(a) shows an example of decomposing $[8, 14)$ using **BestPartition** algorithm. The query interval has 6 unit-length CEIs with consecutive local IDs from 8 to 13. The algorithm initializes

```

Decomposition ( $q, a, b$ ) {
   $P_L = L\lceil a/L \rceil$ ;  $P_R = L\lfloor b/L \rfloor$ ;
   $R_L = [a, P_L)$ ; // left remnant interval
   $R_R = [P_R, b)$ ; // right remnant interval
  if ( $P_L > P_R$ ) { //  $q$  lies between 2 consecutive posts
    BestPartition( $q, a, b, P_R$ );
  }
  else {
    if ( $P_L < P_R$ ) {
      for ( $p = P_L; p < P_R; p = p + L$ )
        insert( $2p + 1, q$ );
    }
    if ( $R_L \neq \phi$ ) BestPartition( $q, a, P_L, P_L - L$ );
    if ( $R_R \neq \phi$ ) BestPartition( $q, P_R, b, P_R$ );
  }
}

```

```

BestPartition ( $q, s, e, P_L$ ) {
   $R = [s, e)$ ; // initial remnant interval
   $l = s - P_L + L$ ;
  // local ID of leftmost unit-length CEI
  while (1) {
    if ( $(l \text{ is even}) \wedge (\text{CEI } \lfloor l/2 \rfloor \text{ does not exceed } e)$ )
       $l = l/2$ ; // local ID of parent CEI
    else {
      insert( $2P_L + l, q$ ); // a maximal-sized CEI
      remove CEI  $l$  from  $R$ ;
      if ( $R == \phi$ ) exit;
    }
    else
       $l = \text{local ID of leftmost unit-length CEI for } R$ ;
  }
}

```

Figure 6: Pseudo code for decomposition algorithm.

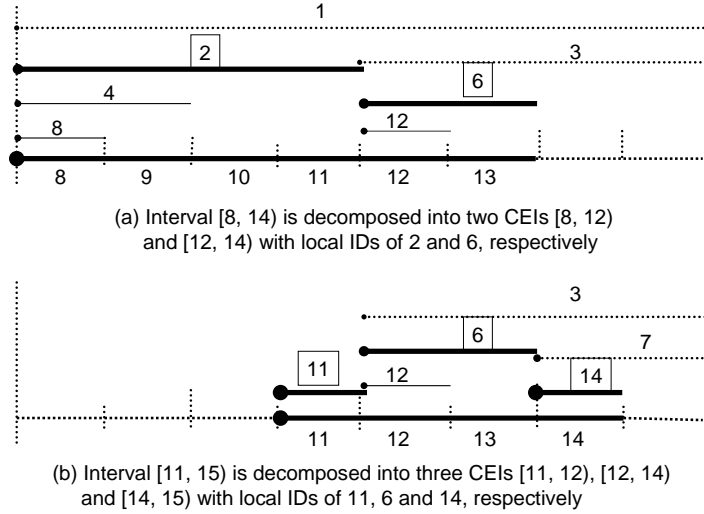


Figure 7: Examples of decomposition with **BestPartition**.

R to be $[8, 14)$. It then starts from the leftmost unit-length CEI for R , which has local ID 8, and successively tries to use its parent CEI by dividing the local ID by 2. The if-condition that checks l is true when the local IDs are 8 and 4, but it is not true when the local ID is 2 since its parent CEI has a right endpoint exceeding 14. Therefore, it finds a maximal-sized CEI with local ID 2. This CEI is removed from R and the new remnant becomes $[12, 14)$. Hence, the algorithm continues with the leftmost unit-length CEI for R , which now has local ID 12. This time the process stops when the local ID is 6 because its parent CEI (local ID 3) has a right endpoint exceeding 14. Hence, the second maximal-sized CEI is the one with local ID 6. After that, the algorithm stops because there is no more remnant interval, i.e., $R = \phi$.

Fig 7(b) shows the decomposition of $[11, 15)$. The interval has 4 unit-length CEIs with local IDs from 11 to 14. Initially, R is $[11, 15)$. The algorithm starts from the leftmost unit-length CEI whose local ID is 11. Since 11 is an odd number, it finds one maximal-sized CEI with local ID 11 and this CEI is removed from R . Hence, R becomes $[12, 15)$. It continues with the leftmost unit-length CEI with local ID 12. It finds another maximal-sized CEI with local ID 6. Finally, it finds the third maximal-sized CEI with local ID 14.

Note that the **BestPartition** algorithm incrementally computes the local IDs for ancestor CEIs and the leftmost unit-length CEI for the remnant interval R , if any. The computation involves only additions and logical shifts by 1 bit. Also note that decomposition means partitioning the query intervals into a set of CEIs. A decomposition is minimal when the number of CEIs in the decomposition is minimal among all the possible decompositions. We don't want




Figure 8: Example of two decompositions.

the decomposed CEIs to be overlapping because it can result in duplicated query IDs in a search result. The storage cost would be higher as well because a query ID would be inserted into more ID lists.

Lemma 1 *Assume there is a decomposition, D_i , that includes a CEI c_i . If there is another decomposition, D_j , that includes a CEI c_j and CEI c_j is an ancestor of c_i , then D_i cannot be a minimal decomposition.*

Proof: Assume decomposition D_i is minimal. Fig. 8 shows an example of two decompositions D_i and D_j . D_i contains CEI c_i and D_j contains CEI c_j and c_j is an ancestor of c_i . From the definition of CEIs, there is no CEI of length less than $|c_j|$ that can cross the vertical lines at either endpoints of CEI c_j due to containment property. As a result, there are more than one CEI besides c_i that is needed to cover the region of c_j . Hence, D_i cannot be minimal, which is a contradiction. \square

Lemma 2 *A minimal decomposition is unique.*

Proof: Suppose there are two decompositions, D_1 and D_2 , that are different and are both minimal. If we sort both decompositions based on the left endpoints of the decomposed CEIs. There must exist a point where the first mismatch between D_1 and D_2 occurs if we scan the decompositions from left to right. The mismatch is due to different lengths since both have the same left endpoint. Hence, the small-sized CEI is a descendant of the the large-sized CEI. From Lemma 1, the one with the small-sized CEI cannot be minimal. Therefore, there is only one unique minimal decomposition. \square

Theorem 1 BestPartition *decomposes any query interval of length less than L into a minimal number of CEIs.*

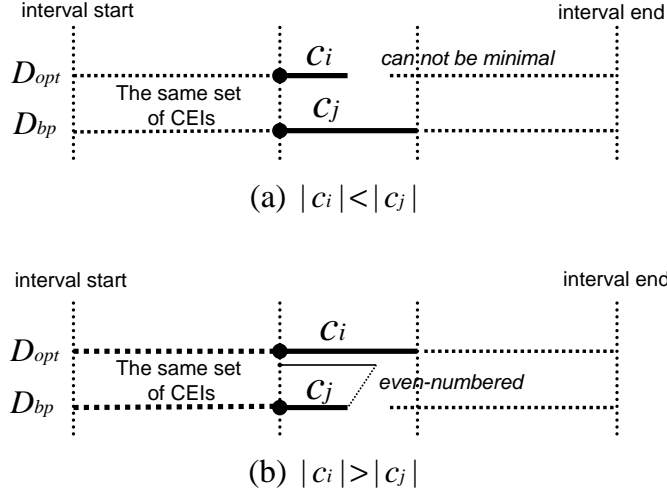


Figure 9: Scenarios of minimal decomposition relative to the decomposition by **BestPartition**.

Proof: Suppose D_{opt} is the minimal decomposition. Let D_{bp} be the decomposition generated by the **BestPartition** algorithm. Consider both decompositions as sorted based on the left endpoints of the decomposed CEIs. If CEIs c_i and c_j , where $c_i \in D_{opt}$ and $c_j \in D_{bp}$, are the first CEIs that show a mismatch between the two decompositions. There are two cases to consider (see Figs. 9(a) and (b)). For the first case, $|c_i| < |c_j|$. Since c_j is an ancestor of c_i , based on Lemma 1, decomposition D_{opt} cannot be minimal. A contradiction occurs. For the second case, $|c_i| > |c_j|$. In this case, c_i must be an ancestor of c_j . Since both CEIs have the same left endpoint, c_j must have an even-numbered local ID, based on Property 8. This is true for any CEI that is an ancestor of CEI c_j and a descendant of CEI c_i as shown in Fig. 9(b). If CEI c_i has an even local ID, then its parent CEI must exceed the right endpoint of the query interval. Otherwise, D_{opt} would not be minimal. Hence, CEI c_i must either have an odd local ID or have a parent CEI whose length exceeds the right endpoint of the query interval. Under either condition, the **BestPartition** algorithm does not stop at CEI c_j . It would continue skipping the ancestor of CEI c_j until it reaches CEI c_i . Therefore CEI c_i , instead of c_j , should be in decomposition D_{bp} . Hence, CEIs c_i and c_j cannot be the first point where mismatch occurs. A contradiction occurs. As a result, **BestPartition** indeed finds the minimal decomposition. \square

Fig. 10 shows an example of a CEI-based query index. It shows the decomposition of four query intervals: $Q1, Q2, Q3$ and $Q4$ within a specific segment containing CEIs of $c1, \dots, c7$. $Q1$ completely covers the segment, and its ID is inserted into $c1$. $Q2$ lies within the segment and is decomposed into $c5$ and $c6$, the largest CEIs that can be used. $Q3$ also resides within the

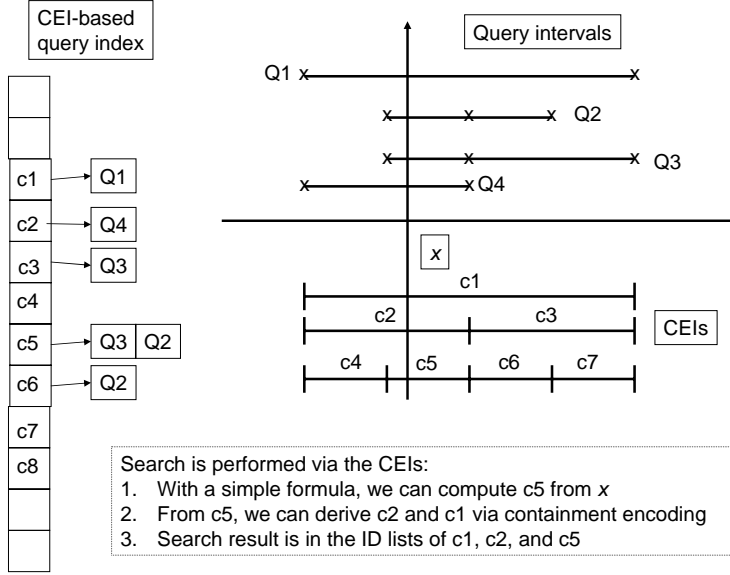


Figure 10: Example of CEI-based query indexing.

segment, but its right endpoint coincides with a guiding post. As a result, we can use $c3$, instead of $c6$ and $c7$ for decomposition. Similarly, $c2$ is used to decompose $Q4$. As shown in Fig. 10, query IDs are inserted into the ID lists associated with the decomposed CEIs.

To delete a query interval, it is decomposed to one or more CEIs. Then the query ID is deleted from the ID lists associated with the decomposed CEIs.

3.3 Searching the query index

Fig. 11 shows the pseudo code for the search algorithm. It first computes the ID of the segment that contains a data value x . From Property 7, this ID is computed via $\lfloor \lfloor x \rfloor / L \rfloor$. Then, it computes the local ID of the leaf CEI that contains x , which is $\lfloor x \rfloor - iL + L$. Based on Property 6, there are exactly $k + 1$ CEIs that contain any data value x . Hence, the search results are stored in the $k + 1$ ID lists associated with these CEIs. Because of the containment encoding (see Property 4), the local IDs of these $k + 1$ CEIs can be computed by successively dividing that of the leaf CEI by 2.

The search algorithm is simple and efficient. For each search, there is an integer conversion from a floating point number. The rest involves only integer additions and logical shifts. No complex floating point multiplications or divisions are needed. As a result, the search operation is extremely fast. The search time is independent of n , where n is the total number of query intervals maintained, in terms of finding the $k + 1$ ID lists. However, if reporting time is also

```

Search ( $x$ ) { //  $x \in A$  and  $x$  is a non-integer number
   $i = \lfloor x/L \rfloor$ ; // segment ID
   $l = \lfloor x \rfloor - iL + L$ ;
    // local ID of the leaf CEI that contains  $x$ 
  for ( $j = 0; j \leq k; j = j + 1$ ) {
     $c = 2iL + l$ ; // global ID of the CEI
    if ( $\text{IDList}[c] \neq \text{NULL}$ ) { output( $\text{IDList}[c]$ ); }
     $l = l/2$ ; // local ID of parent
  }
}

```

Figure 11: Pseudo code for the search algorithm.

included, the search time increases as n increases. This is because there are more IDs stored in each ID lists, on average, as n increases.

As an example, to search with a data value x in Fig. 10, we first compute the local ID of the unit-length CEI that contains it. In this case it is $c5$. Then, from $c5$, we can compute the local IDs of all its ancestors that contain $c5$. In this case, they are $c2$ and $c1$. As a result, the search result is contained in the 3 ID lists associated with $c1, c2$ and $c5$. We can verify from Fig. 10 that the result indeed contains $Q1, Q2, Q3$ and $Q4$.

3.4 Other system issues

So far, we have assumed $q : [a, b)$. When the right endpoint is also inclusive, i.e., $[a, b]$, we need to use additional r point CEIs. These points are maintained separately. They will be useful only if the attribute is of integer data type. During decomposition, the query ID is additionally inserted into the ID list associated with the point CEI representing b . At search, we also report the IDs stored in the ID list associated with the integer data value.

A query interval may be an open-ended interval, i.e., $A > 4$. In this case, a query ID can be inserted into r/L CEIs in the worst case. If L is small, the storage cost can be high. Fortunately, we can choose a large L to reduce the storage cost. In fact, we can set $L = r$. In that case, the search time may degrade a bit because we need to collect the search result from exactly $k + 1$ ID lists, where $k = \log(L) = \log(r)$.

In practice, the CEI-based query index is naturally suited for parallel processing. We can control both storage cost and search time by choosing a relatively large L and by properly separating r into partitions. One machine can then be used to process a partition.

4 Storage cost analysis

Here, we analyze the average storage costs of the three VC-based query indexes. The storage cost can be expressed as $4N + 8n\mu$ bytes, where N is the total number VCs defined, n is the total number of queries and μ is the average number of decomposed VCs per query interval. We assume that a query ID and a pointer require 4 bytes each. N is available from Table 1. However, μ is dependent on the starting position and the length of a query interval. For the analysis, the query length w is uniformly distributed between 1 and W and the starting endpoint is also uniformly distributed. The first term represents the storage cost for the headers while the second term for the ID lists. For the UI-based query index, the average of unit-length intervals decomposed per query, denoted as μ_{UI} , is $(1+W)/2$ because each query interval is decomposed into w unit-length intervals and $1 \leq w \leq W$. We will now derive μ_{CEI} and μ_{LCI} as follows.

4.1 CEI storage cost

Definition 1 Let B_w^k be the k -bit binary representation of w , and $|B_w^k|$ be the number of 1's in B_w^k .

Lemma 3 If a query interval q , whose length $w < L$, has either its left or right endpoint aligned to a guiding post, it will be decomposed into $|B_w^k|$ CEIs by the **BestPartition** algorithm.

Proof: Since q is aligned to a guiding post, it must have an endpoint that is a multiple of L . Thus q is either in the form of $[iL, iL + w)$ or $[iL - w, iL)$, for $1 \leq w < L$ and i is an integer. It is easy to verify that the **BestPartition** algorithm finds CEIs of distinct lengths of 2's powers in decreasing order for the case of left-endpoint alignment and in increasing order for the case of right-endpoint alignment. Since every bit in B_w^k represents a distinct 2's power, there is a 1-to-1 mapping between the bit positions and lengths of decomposed CEIs. Therefore, q can be decomposed into $|B_w^k|$ CEIs. \square

Lemma 4 $\sum_{w=1}^{2^d-1} |B_w^d| = d * 2^{d-1}$.

The sum represents the total number of 1's in all of the possible distinct d -bit binary representations. Since each bit position has 2^{d-1} 1's among all possible d -bit binary representations, there are $d * 2^{d-1}$ 1's. When $d = k = \log(L)$, we have $\sum_{w=1}^{2^k-1} |B_w^k| = k2^{k-1} = \frac{L}{2} \log(L)$.

Now we are ready to derive μ_{CEI} , the average number of decomposed CEIs per query interval. We enumerate all of the possible query intervals based on their left endpoints and lengths. More formally,

$$\mu_{CEI} = \frac{\sum_{p \in [0, r)} \sum_{w \in (0, W]} t(p, w)}{rW},$$

where $t(p, w)$ is the number of CEIs used in decomposing the interval with a left endpoint at p and a length w . Note that query intervals and left endpoints are represented with the style of $[a, b)$, while interval lengths are specified with the style of $(c, d]$.

Due to similar patterns, we enumerate all query intervals whose left endpoints are within a specific common segment and whose lengths are between 1 and W . Without loss of generality, we pick two arbitrary successive guiding posts (say λL and $(\lambda + 1)L$, where λ is an integer) as the common segment. Namely, $p \in [\lambda L, (\lambda + 1)L)$. For simplicity, the common segment is chosen such that $(\lambda + 1)L + W < r$. There are L possible positions for the left endpoint of a query interval, and for each position, there are up to W query intervals to be decomposed. In other words,

$$\mu_{CEI} = \frac{\sum_{p \in [\lambda L, (\lambda + 1)L)} \sum_{w \in (0, W]} t(p, w)}{LW}.$$

Since all left endpoints are in the common segment, we can use the shorthand notation $T_{(0, W]}$ to represent the total number of decomposed CEIs. Hence, $T_{(0, W]} = LW \mu_{CEI}$. For this analysis, we only consider the case $W > L$ because, if $W < L$, a smaller L can always be chosen.

Let βL be the greatest multiple of L that is smaller than or equal to W , i.e., $\beta = \lfloor \frac{W}{L} \rfloor$. Since we can categorize intervals of lengths from 1 to W into different groups based on multiples of L , and $(0, W] = \bigcup_{0 \leq i < \beta} (iL, (i + 1)L] \cup (\beta L, W]$, we have

$$T_{(0, W]} = \sum_{i=0}^{\beta-1} T_{(iL, (i+1)L]} + T_{(\beta L, W]}.$$

Therefore,

$$\sum_{i=0}^{\beta-1} T_{(iL, (i+1)L]} \leq T_{(0, W]} \leq \sum_{i=0}^{\beta} T_{(iL, (i+1)L]} \quad (1)$$

We will derive a closed-form formula for

$$\sum_{i=0}^{\beta-1} T_{(iL, (i+1)L]}.$$

When $1 \leq i \leq \beta - 1$, all of the enumerated intervals have lengths greater than L , which must cross the right guiding post of the common segment. If there are either one or 2 remnant intervals due to the decomposition process, they must be aligned with some guiding posts. Therefore, Lemmas 3 and 4 can be used to prove the following Lemma:

Lemma 5 $T_{(L, 2L]} = \frac{L^2}{2} + \frac{3L}{2} + L^2 \log L$.

Proof: As shown in Fig. 12, each integer point in the common segment is the left endpoint of L query intervals of lengths $L + 1, L + 2, \dots, 2L$ respectively. These intervals must be decomposed

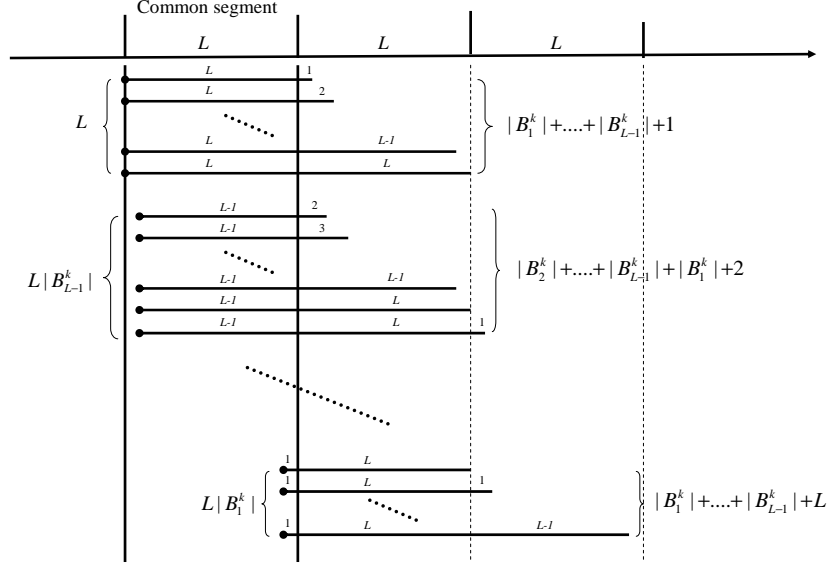


Figure 12: Enumeration of query intervals for computing $T_{(L,2L)}$.

into 2 or 3 sub-intervals by 1 or 2 guiding posts. These sub-intervals must have at least 1 endpoint aligned to some guiding posts. The number of decomposed CEIs inside the common segment is $L + L|B_{L-1}^k| + L|B_{L-2}^k| + \dots + L|B_1^k| = L + L \sum_{i=1}^{L-1} |B_i^k|$. The number of decomposed CEIs outside the common segment is $L \sum_{i=1}^{L-1} |B_i^k| + \sum_{i=1}^L i$, where the second sum is for counting the number of decomposed CEIs with length L . Therefore,

$$T_{(L,2L)} = 2L \sum_{i=1}^{L-1} |B_i^k| + L + \frac{L(L+1)}{2} = \frac{L^2}{2} + \frac{3L}{2} + L^2 \log L.$$

□

Lemma 6 $T_{(iL,(i+1)L)} = T_{(L,2L)} + (i-1)L^2, i = 1, 2, \dots, \beta - 1$.

Proof: The $(i-1)$ in the second term is due to additional $i-1$ segments, and each extra segment contributes L^2 more CEIs of length L . □

Now, we need to derive $T_{(0,L]}$. However, $T_{(0,L]}$ is hard to derive because an interval of length less than or equal to L may be totally contained within the common segment without crossing any guiding post. Fortunately, we observed that intervals of lengths from 1 to L can be categorized further into intervals of lengths in the ranges of $(L/2, L]$, $(L/2^2, L/2]$, $(L/2^3, L/2^2]$, \dots , $(2, 4]$, $(1, 2]$, and $(0, 1]$. In other words,

$$(0, L] = \bigcup_{1 \leq j \leq k} \left(\frac{L}{2^j}, \frac{L}{2^{j-1}} \right] \cup (0, 1].$$

As there are L distinct CEIs of length 1 within the common segment, we have $T_{(0,1]} = L$ and

$$T_{(0,L]} = \sum_{j=1}^k T_{(\frac{L}{2^j}, \frac{L}{2^{j-1}}]} + L.$$

Now we need to compute $T_{(\frac{L}{2^j}, \frac{L}{2^{j-1}}]}$, for $1 \leq j \leq k$. We partition the common segment into 2^j sub-segments, and apply the same counting method to the leftmost or first sub-segment, i.e., $[\lambda L, \lambda L + \frac{L}{2^j})$. In other words, we look at intervals with left endpoints in the first sub-segment and lengths from $\frac{L}{2^j} + 1$ to $\frac{L}{2^{j-1}}$. More formally, we compute the number of decomposed CEIs as $\sum_{p \in [\lambda L, \lambda L + \frac{L}{2^j})} \sum_{w \in (\frac{L}{2^j}, \frac{L}{2^{j-1}}]} t(p, w)$. In fact, we treat each sub-segment as a scaled-down version of $T_{(L, 2L]}$. However, instead of working with the full k bits, we are now using only the $k - j$ least significant bits (LSB's). Since left endpoints distribute evenly on all sub-segments, we can multiply the number obtained using the first sub-segment by the number of sub-segments (2^j). However, there are over-counts and we need to offset them.

For example, as shown in Fig. 13, for $j = 1$ the common segment is partitioned into 2 sub-segments $[\lambda L, \lambda L + L/2)$ and $[\lambda L + L/2, \lambda L + L)$. We evaluate the $(\frac{L}{2})^2$ intervals that have left endpoints in the first sub-segment and with various lengths from $L/2 + 1$ to L . There are $(\frac{L}{2})^2$ similar intervals but with left endpoints in the second sub-segment. Both groups of intervals share the same pattern, and therefore we can just count the first group. There are 2 CEIs of length $L/2$ in the sub-segment, $[\lambda L, \lambda L + L/2)$ and $[\lambda L + L/2, \lambda L + L)$. Consider the vertical dotted line at $\lambda L + L/2$. There are no CEIs of lengths smaller than L that can cross this vertical line. The number of CEIs decomposed from the left-hand side is

$$L/2 + |B_{L/2-1}^{k-1}|L/2 + |B_{L/2-2}^{k-1}|L/2 + \dots + |B_1^{k-1}|L/2,$$

and from the right hand side is

$$L/2(|B_1^{k-1}| + \dots + |B_{L/2-1}^{k-1}|) + \sum_{i=1}^{L/2} i.$$

Since $\log(L/2) = k - 1$, only $k - 1$ LSB's are involved. Adding both numbers (left and right), multiplying the sum by 2 due to the repeating pattern of the second sub-segment, and offsetting the over-count by 1, we get

$$\begin{aligned} T_{(L/2, L]} &= 2\left[\frac{L}{2} + \sum_{i=1}^{L/2} i + L\left(\sum_{i=1}^{L/2-1} |B_i^{k-1}|\right)\right] - 1 \\ &= 2\left[\frac{(L/2)^2}{2} + \frac{3}{2}\left(\frac{L}{2}\right) + \left(\frac{L^2}{4}\right)\log\left(\frac{L}{2}\right)\right] - 1. \end{aligned}$$




Figure 13: Enumeration of query intervals for computing $T_{(L/2, L)}$.

Note that from Lemma 4, when $d = k - 1$, we have $\sum_{i=1}^{L/2-1} |B_i^{k-1}| = \frac{L}{4} \log(\frac{L}{2})$. The over-count happens for interval $[\lambda L, \lambda L + L]$, which requires 2 CEIs of length $L/2$ (from left and right) to decompose, while we can use just one CEI of length L .

In general, the over-count is 2^{j-1} . There are 2^j CEIs of length $L/2^{j-1}$, where half of them are over-counted by 1 each. This is because we use 2 CEIs of length $L/2^{j-2}$ to decompose each of the 2^j CEIs within and outside of the sub- or common segment. Therefore, we have

$$T_{(\frac{L}{2^j}, \frac{L}{2^{j-1}}]} = 2^j \left[\frac{(\frac{L}{2^j})^2}{2} + \frac{3}{2} \left(\frac{L}{2^j} \right) + \left(\frac{L}{2^j} \right)^2 \log \left(\frac{L}{2^j} \right) \right] - \frac{2^j}{2}.$$

We then have the following lemma:

Lemma 7 $T_{(\frac{L}{2^j}, \frac{L}{2^{j-1}}]} = \frac{L^2}{2^{j+1}} + \frac{3}{2}L + \frac{L^2}{2^j} \log L - \frac{jL^2}{2^j} - 2^{j-1}$.

From these lemmas, we can have the following theorem by summing up intervals of various lengths.

Theorem 2

$$\begin{aligned} \sum_{i=0}^{\beta-1} T_{(iL, (i+1)L]} &= T_{(0, L]} + T_{(L, \beta L]} \\ &= \left\{ L^2 \log L + \frac{3}{2}L \log L - \frac{3}{2}L^2 + \frac{3}{2}L + 1 \right\} \\ &\quad + \left\{ (\beta - 1) \left[\frac{3}{2}L + L^2 \log L - \frac{L^2}{2} \right] + L^2 \frac{\beta(\beta - 1)}{2} \right\}. \end{aligned}$$

Since β is in the order of $O(W/L)$, we have μ_{CEI} in the order of

$$\sum_{i=0}^{\beta-1} \frac{T_{(iL, (i+1)L]}}{WL} \approx \frac{L^2 \log L + \frac{W}{L} L^2 \log L + \frac{W^2}{L^2} L^2}{WL} \approx O(\log L + W/L). \quad (2)$$

4.2 LCI storage cost

Similar to the CEI case, we enumerate all of the possible query intervals based on the left-end points and their lengths. We also derive the average number of decomposed LCIs for a query interval, denoted as μ_{LCI} . More formally,

$$\mu_{LCI} = \frac{\sum_{p \in [0, r]} \sum_{w \in (0, W]} t'(p, w)}{rW},$$

where $t'(p, w)$ is the number of LCIs used in decomposing the interval with a left endpoint p and a length w .

Due to similar patterns, every point in $[0, r - W]$ has W intervals of lengths $1, 2, \dots,$ and W , respectively. Note that we exclude the boundary condition at the right hand side to simplify the calculation. Therefore, we are actually computing $\frac{\sum_{w \in (0, W]} t'(w)}{W}$, where $t'(w)$ is the number of LCIs used to decompose any interval with a length w . We can similarly define $T'_{(0, W]}$ as a shorthand representation for $\sum_{w \in (0, W]} t'(w)$. Instead of looking at the left endpoints within a segment as in the CEI scheme, we look at left endpoints matching any arbitrary point.

$$T'_{(0, W]} = \sum_{i=0}^{\beta-1} T'_{(iL, (i+1)L]} + T'_{(\beta L, W]}.$$

$$\sum_{i=0}^{\beta-1} T'_{(iL, (i+1)L]} \leq T'_{(0, W]} \leq \sum_{i=0}^{\beta} T'_{(iL, (i+1)L]} \quad (3)$$

We will derive a closed-form formula for $\sum_{i=0}^{\beta-1} T'_{(iL, (i+1)L]}$.

For LCI approach, $T'_{(0, L]}$ is easy to derive as shown in the following lemma.

Lemma 8 $T'_{(0, L]} = 1 + \frac{L}{2} \log L$.

Proof: Each point can be the left endpoint of intervals of lengths $1, 2, \dots, L$. Since the decomposing algorithm starts from the left-hand side and uses an LCI as large as possible, the number of LCIs decomposed for an interval of length w is thus $|B_w^k|$. Therefore, we have $T'_{(0, L]} = 1 + \sum_{w=1}^{L-1} |B_w^k| = 1 + \frac{L}{2} \log L$. \square

Lemma 9 $T'_{(iL, (i+1)L]} = 1 + \frac{L}{2} \log L + iL$.

Table 2: Comparison of query decomposition for VC-based query indexes.

Virtual construct	average number of decomposed virtual constructs per query interval
UI	$\mu_{UI} = \frac{1+W}{2}$
LCI	$\mu_{LCI} \geq (\beta + \frac{1}{2}\beta L \log(L) + \frac{1}{2}\beta(\beta - 1)L)/W$ $\mu_{LCI} \leq (\beta + 1 + \frac{1}{2}(\beta + 1)L \log(L) + \frac{1}{2}(\beta + 1)\beta L)/W$
CEI	$\mu_{CEI} \geq (\beta L^2 \log(L) + \frac{1}{2}(\beta^2 - 2\beta - 2)L^2 + \frac{3}{2}L \log(L) + \frac{3}{2}\beta L + 1)/(WL)$ $\mu_{CEI} \leq ((\beta + 1)L^2 \log(L) + \frac{1}{2}(\beta^2 - 3)L^2 + \frac{3}{2}L \log(L) + \frac{3}{2}(\beta + 1)L + 1)/(WL)$

The i in the third term is due to additional i segments, and each extra segment contributes L more LCIs of length L . Therefore, we have

$$T'_{(iL, (i+1)L]} = T'_{(0, L]} + iL = 1 + \frac{L}{2} \log L + iL.$$

Summing the numbers up, we have the following Theorem.

Theorem 3 $\sum_{i=0}^{\beta-1} T'_{(iL, (i+1)L]} = \beta + \frac{\beta L}{2} \log L + \frac{\beta(\beta-1)}{2} L.$

Since β is in the order of $O(W/L)$, we have μ_{LCI} in the order of

$$\sum_{i=0}^{\beta-1} \frac{T'_{(iL, (i+1)L]}}{W} \approx \frac{\frac{W}{L} L \log L + \frac{W^2}{L^2} L}{W} \approx O(\log L + W/L). \quad (4)$$

4.3 Comparisons of three query indexes

Table 2 summarizes the closed-form formulae for the average number of decomposed virtual constructs per query interval for the three VC-based query indexes. Both the lower and upper bounds for μ_{CEI} and μ_{LCI} are shown. To verify the accuracy of our analyses, we ran simulations and plotted the total number of decomposed virtual constructs. Figs. 14(a) and 14(b) show the upper and lower bounds and simulation results for the CEI-based and LCI-based indexes, respectively. Fig. 14(c) compares all three indexes. For this experiment, $W = 200, n = 100,000$ and $r = 1,048,576$. These figures show that (1) the analyses are accurate; (2) the number of total decomposed virtual constructs is significantly larger for the UI-based index than the other two indexes; and (3) the CEI-based index uses more virtual constructs than the LCI-based index in query decomposition.

5 Simulation studies

Simulations were conducted to evaluate and compare the CEI-based index with the LCI-based index, the UI-based index and the IS-lists approach [8]. We focused on storage cost and search




Figure 14: Total number of decomposed virtual constructs for (a) CEI-based; (b) LCI-based and (c) all three VC-based query indexes.

time because they are two of the most important metrics in determining the effectiveness of a query index for stream processing. A low storage cost is desirable so that the entire index can be loaded into main memory. A fast search time is important so that a rapid stream can be properly handled. The IS-lists approach represents the state of the art direct indexing approach. It has $O(\log(n))$ search time and $O(n \log(n))$ storage cost. In contrast, the search times, in terms of finding the ID lists from the covering VCs, of the UI-based, LCI-based and CEI-based indexes are independent of n . However, in our simulations, the search time included the report time. As a result, it generally increases as n increases. This is because on average the ID lists contain more query IDs. We implemented the three alternative VC-based indexes. For the IS-lists approach, we downloaded from the Web an implementation from the authors [6] and used it to run the same sets of input data. The simulations were conducted on an RS6000 model 43P machine running AIX 5.1.

Attribute values range from 0 to r , with 1,048,576 as the large r and 65,536 as the small r . The starting point of a query interval was randomly chosen between 1 and $r - 1$ with a length of w randomly chosen between 1 and W . Namely, $\bar{w} = W/2$. Various W 's were used, with a default 200. A total of n query intervals were generated and inserted. In the experiments, n was varied from 5,000 to 640,000. After insertion, we performed 50,000 random searches and computed the average search time and storage cost. The data value for search was a floating point number chosen randomly between 1 and r . The search time included the report time, which involves reporting the query IDs in the ID lists. If there are large number of IDs stored in the ID lists, the report time will be large as a result.

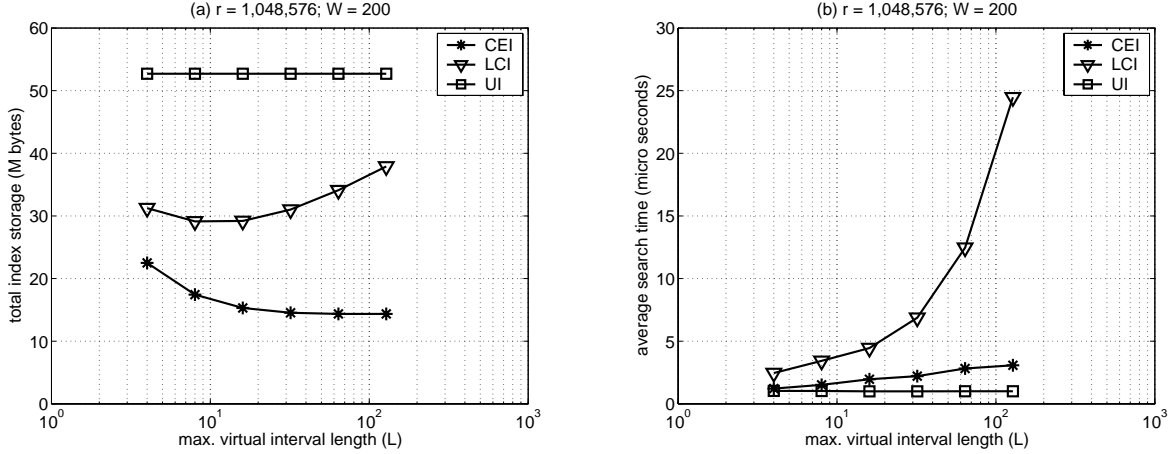


Figure 15: The impacts of L , when r is large, on (a) total index storage cost; (b) average search time.

5.1 Impact of maximum interval length (L)

We first examine the impact of the maximum interval length L on the three alternative VC-based indexes in terms of storage cost and average search time. For this set of experiments, we compare the UI-based, LCI-based and CEI-based indexes. Figs. 15(a) and (b) show the storage cost and average search time, respectively, of the three indexes when r is large. In contrast, Figs. 16(a) and (b) show storage cost and average search time, respectively, when r is small. For the entire set of experiments, a default W of 200 was used. The total number of query intervals inserted was 100,000.

As expected, the UI-based index has a high storage cost and fast search time for a large r and a small r . Moreover, the storage cost and search time are independent of L because it used unit-length virtual intervals. The CEI-based index has both low storage cost and fast search time for both r 's. In general, the average search time increases as L increases. This is particularly true for the LCI-based index because its search time increases much more quickly than the CEI-based index as L increases. This can be seen from the formula for C described in Table 1. Hence, it argues for a small L for the LCI-based index from the search time perspective. However, if L is too small, the storage cost tends to increase for both schemes. This is because more query IDs are stored in the ID lists because the total number of decompositions D is larger with a smaller L (see Table 1). Moreover, when r is large, the storage cost of the LCI-based index also increases as L increases (see Fig. 15(a)). However, under all cases, the storage cost of the CEI-based index decreases as L increases.

The average search time of the CEI-based index increases much slower than the LCI-based

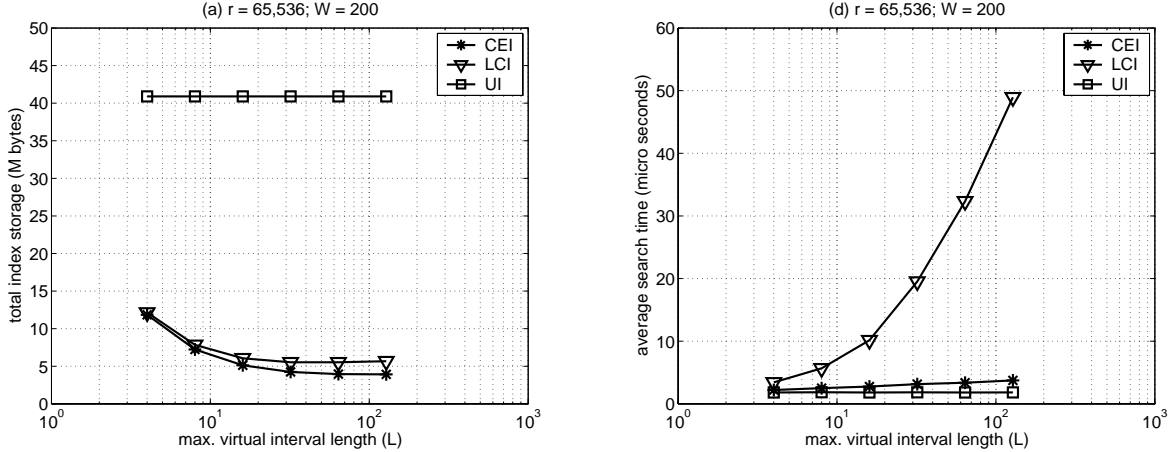


Figure 16: The impacts of L , when r is small, on (a) total index storage cost; (b) average search time

index as L increases. This is because the report time is the same for all different L 's since n remains the same for all the experiments. The difference is due to the computation of the IDs of the extra covering CEIs for a larger L . Because $C = 1 + \log(L)$, it increases much slower than that of the LCI-based index, which is $2L - 1$. Hence, for the CEI-based index, it argues for a larger L since the storage cost tends to decrease and the average search time increases only slightly as L increases. This result is important in practice. We can choose a larger L in order to minimize the storage cost without significantly degrading the average search time.

Under all cases, the CEI-based index performs better than the LCI-based index in terms of both storage cost and average search time. The advantage in average search time is particularly dramatic when L is large (see Fig 15(b) and Fig. 16(b)).

5.2 Comparisons of CEI-based index with IS-lists

Now we compare the CEI-based scheme with the UI-based, the LCI-based and the IS-lists approaches. Figs. 17(a) and (b) show the impact of the maximal query interval width, W , on the performance of the four schemes. Fig. 17(a) shows the storage cost while Fig. 17(b) shows the average search time. For this experiment, W was varied from 10 to 320, $L = 16$, $r = 65,536$ and $n = 50,000$ were used.

Both figures show that the UI-based index has high storage cost but fast search time for the entire range of W 's. The CEI-based has both low storage cost and fast search time. All three VC-based approaches outperform the IS-lists approach in terms of search time. For the smallest $W = 10$, the performance difference is almost 20 times. Namely, the average search time of




Figure 17: Impact of W on (a) index storage cost; (b) average search time.

the CEI-based index is only 1/20 that of the IS-based approach. The IS-lists approach also has a higher storage cost than both the LCI-based and CEI-based indexes. But, it has a smaller storage cost than the UI-based index when W is large (see Fig. 17(a)).

Figs 18(a) and (b) show the impact of n on the storage cost and average search time, respectively, when r is large. Figs. 19(a) and (b) show the index storage cost and average search time, respectively, when r is small. For this set of experiments, W was 200 and n was varied from 5,000 to 640,000 and $L = 16$.

From Fig. 18(b) and Fig. 19(b), the average search time of the CEI-based index outperforms both the LCI-based index and the IS-lists approach. Moreover, it is very close to the UI-based index, which has the best search time performance. The IS-lists approach, which is the best direct indexing approach, has the highest search time. Note that there are missing data points for the UI-based and the IS-lists approaches. This is because for those data points, n was too large for both schemes to be executed properly on our computer system, which is an AIX model 43P machine. Memory was exhausted by the indexes.

The storage cost of IS-lists approach is better for the case of a large r when n is relatively small. But, as n increases, this advantage disappears. This is because the UI-based, LCI-based and CEI-based indexes need to maintain pointers to the ID lists for every VC. When r is large and n is small, the storage cost for the pointers dominates the total storage cost, even though most of them are NULL. In contrast, the IS-lists approach does not have to maintain those pointers. However, for a small r , the storage cost advantage of the IS-lists approach also disappears.

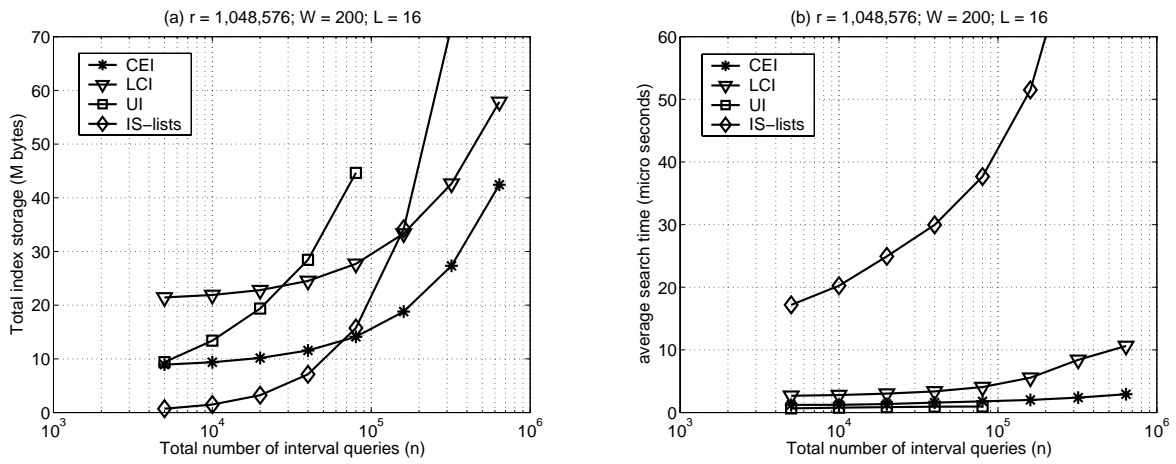


Figure 18: Impact of n , when r is large, on (a) index storage cost; (b) average search time.




Figure 19: Impact of n , when r is small, on (a) index storage cost; (b) average search time.

6 Summary

We have presented a CEI-based query indexing approach for efficient stream processing. It is main memory-based approach and has both low storage cost and fast search time, ideally suited for stream processing. The CEI-based query index is centered on a set of virtual containment-encoded intervals. The range of a numerical attribute is partitioned into segments. For each segment, a set of virtual containment-encoded intervals were carefully defined and properly labeled as a perfect binary tree. A query interval is first decomposed into one or more CEIs and the query ID is inserted into the ID lists associated with the decomposed CEIs. Search is carried out indirectly via these CEIs. There is no need to compare a data value with any of the query boundaries during a search. The containment encoding makes possible efficient decomposition and search operations. Almost all of the operations can be carried out with integer additions and logical shifts. Analyses and simulations were conducted to evaluate the performance of the CEI-based query index. The results show that, compared with other alternatives, the CEI-based query index does have low storage cost and fast search time and it outperforms other alternatives in both.

References

- [1] D. Carney, U. Cetintemel, M. Cherniack, C. Convey, S. Lee, G. Seidman, M. Stonebraker, N. Tatbul, and S. Zdonik. Monitoring streams – a new class of data management applications. In *Proc. of VLDB*, 2002.
- [2] H. Edelsbrunner. Dynamic data structures for othogonal intersection queries. Technical Report 59, Institute for Information Processing, Technical University of Graz, Graz, Austria, 1980.
- [3] F. Fabret, H. A. Jacobsen, F. Llirbat, J. Pereira, K. A. Ross, and D. Shasha. Filtering algorithms and implementation for very fast publish/subscribe systems. In *Proc. of ACM SIGMOD*, 2001.
- [4] V. Gaede and O. Günther. Multidimensional access methods. *ACM Computing Surveys*, 30(2):170–231, June 1998.
- [5] A. Guttman. R-trees: A dynamic index structure for spatial searching. In *Proc. of ACM SIGMOD*, 1984.
- [6] E. Hanson. ISlist.tar: A tar file containing C++ source code for IS-lists. <http://www-pub.cise.ufl.edu/~hanson/IS-lists/>.

- [7] E. Hanson, M. Chaaboun, C.-H. Kim, and Y.-W. Wang. A predicate matching algorithm for database rule systems. In *Proc. of ACM SIGMOD*, pages 271–280, 1990.
- [8] E. Hanson and T. Johnson. Selection predicate indexing for active databases using interval skip lists. *Information Systems*, 21(3):269–298, 1996.
- [9] C. P. Kolovson and M. Stonebraker. Segment indexes: Dynamic indexing techniques for multi-dimensional interval data. In *Proc. of ACM SIGMOD*, 1991.
- [10] S. R. Madden, M. A. Shah, J. M. Hellerstein, and V. Raman. Continuously adaptive continuous queries over streams. In *Proc. of ACM SIGMOD*, 2002.
- [11] H. Samet. *Design and Analysis of Spatial Data Structures*. Addison-Wesley, 1990.
- [12] K.-L. Wu, S.-K. Chen, P. S. Yu, and M. Mei. Efficient interval indexing for content-based subscription e-commerce and e-service. In *Proc. of IEEE Int. Conf. on e-Commerce Technology for Dynamic E-Business*, Sept. 2004.
- [13] K.-L. Wu and P. S. Yu. Efficient query monitoring using adaptive multiple key hashing. In *Proc. of ACM CIKM*, pages 477–484, 2002.