Published in final edited form as:

Appl Math Comput. 2020 May 1; 372: . doi:10.1016/j.amc.2019.125010.

The n -dimensional k -vector and its application to orthogonal range searching

David Arnas^{a,*}, Carl Leake^b, Daniele Mortari^b

^aUniversidad de Zaragoza, Valentin Cardenera 4, Huesca 22003, Spain

^bAerospace Engineering, Texas A&M University, College Station, TX 77843-3141, USA

Abstract

This work focuses on the definition and study of the n -dimensional k -vector, an algorithm devised to perform orthogonal range searching in static databases with multiple dimensions. The methodology first finds the order in which to search the dimensions, and then, performs the search using a modified projection method. In order to determine the dimension order, the algorithm uses the k -vector, a range searching technique for one dimension that identifies the number of elements contained in the searching range. Then, using this information, the algorithm predicts and selects the best approach to deal with each dimension. The algorithm has a worst case complexity of $\mathcal{O}(nd(k/n)^{2/d})$, where k is the number of elements retrieved, n is the number of elements in the database, and d is the number of dimensions of the database. This work includes a detailed description of the methodology as well as a study of the algorithm performance.

Keywords

Computer science; Computational geometry; Orthogonal range searching; Database structures; Multidimensional spaces

1. Introduction

Orthogonal range searching is a very common problem that arises in a wide variety of applications including engineering, physics, mathematics, computational science, economics, and statistics. The problem consists of finding all the elements from a database that are contained in a given orthogonal range. In its simplest form, orthogonal range searching requires determining whether or not each element of the database is included in a given searching range. This can be done in general by checking all the elements from the database individually using a brute force approach. However, following this procedure, the time required by the algorithm quickly increases with the size of the database, making the algorithm slower as the size of the database increases. Thus, and in order to improve this limiting behavior, a wide variety of methodologies have been proposed to deal with the problem. Examples include the projection method [1], the grid files method [2], the b -tree [3], the quad tree [4], the k -d tree [5], the R -tree [6], and the Kdb -tree [7]; however, there are

*Corresponding author. darnas@unizar.es (D. Arnas).

many others [8–11]. The multitude of techniques shows that range searching is an extensively studied subject in computer science [12,13].

In this work, we introduce the n -dimensional k -vector, a very efficient methodology for problems that require extensive orthogonal range searches in static databases with multiple dimensions. The n -dimensional k -vector is based on the idea of generating an auxiliary database during preprocessing that contains information about the element distribution in each dimension. Afterwards, during the searching process, the algorithm uses this information to provide the number of elements in the searching range for all the dimensions individually, as well as the positions of these elements in the database. That way, and for each search, the algorithm finds the sequence of dimensions that is most likely to minimize the number of range comparisons, starting with the dimension with the smallest number of elements retrieved. Then, in each subsequent dimension, the algorithm continues with a brute force approach or an intersection approach depending on the situation.

The n -dimensional k -vector is the evolution, for databases in multiple dimensions, of the so called k -vector, a range searching technique for one dimensional databases. The k -vector range searching algorithm can be thought of as a hash table like algorithm based on a bijective hash function. It was originally developed for the identification of stars observed by wide field-of-view star trackers using computationally limited processors on board spacecraft [14]. Since then, the k -vector has been repeatedly and successfully validated in space [15,16] becoming part of the state-of-the-art algorithm, Pyramid [17]. In addition, the one-dimensional k -vector has been successfully applied to solve other problems, including: inverting non-linear functions [18] (even Diophantine), generating random numbers with any prescribed distribution [19] (analytical or tabulated), solving for the intersection of nonlinear functions, identifying isosurfaces for level-set analysis, and finding gene sequences in long DNA chains [20].

The k -vector is based on the idea of describing the nonlinearities of a sorted database using a vector of integers of a chosen size called the k -vector [21] (a hash-like function). This is done by comparing the database distribution with a mapping function, for instance, a line. The key feature of the k -vector is that the searching time complexity of the algorithm is independent of the database size; it only depends on the nonlinearity of the sorted database with respect to the mapping function [22]. As with the binary search technique, the one-dimensional k -vector only works with sorted databases. For this reason, a direct extension of this technique to n -dimensions is not possible due to the lack of clear ordering in multidimensional spaces. Therefore, a different approach is required, which is the focus of this work, the n -dimensional k -vector.

This manuscript is organized as follows. First, an overview of the algorithm is introduced, which aims to provide a general idea of the methodology and its process. Then, a detailed exposition of the methodology is presented, including the structure of the database and the auxiliary databases required (the index array, the k -vector array, and the k -vector line array). Afterwards, the performance study of the algorithm is assessed in terms of complexity and speed, and compared with brute force and k -d tree, which are other common orthogonal range searching algorithms. Next some modifications to the one-dimensional k -vector

algorithm are presented that show up as a result of the development of the n -dimensional k -vector. Finally, some possible modifications to the algorithm are studied that reduce the memory required while maintaining the algorithm complexity.

2. General overview of the algorithm

The n -dimensional k -vector is a numerical algorithm specifically devised to perform orthogonal range searches in multidimensional static databases. The algorithm first finds the most convenient dimension in which to perform a projection of the database by estimating the number of expected retrieved elements in each dimension. This is done using three auxiliary databases that first estimate the number of elements in the searching range in each dimension, and then, identify which points are contained in the projection. Since up to this point the elements identified have only been checked in one dimension, a brute force approach follows to assess the remaining dimensions. Therefore, only a fraction of the database is checked during the searching process.

As such, the n -dimensional k -vector can be regarded as a modified projection method in one dimension, where an auxiliary database is used rather than a binary search. However, an additional idea is introduced in the algorithm. Instead of searching in the whole database at the same time, the n -dimensional k -vector performs multiple searches in subsets of the database, which are defined in such a way that they only contain elements in a known range from a given dimension selected during the preprocessing. This allows the n -dimensional k -vector to assess two dimensions at the same time: the first one due to the structure of the database, and the second one due to the projection method.

The algorithm requires a one-time preprocessing effort that prepares the necessary auxiliary databases. During this pre-processing effort, the algorithm first sorts the whole database in a chosen dimension, for instance, the last dimension. Then, this sorted database is separated into sub-databases, where each sub-database contains all the elements that are within a given range in the selected dimension in such a way that each sub-database contains approximately the same number of elements. Afterwards, each sub-database is sorted again in a different dimension, for example, the first dimension, to generate the final data structure of the algorithm. A flowchart of this process can be seen in Fig. 1.

After the data structure is built, three auxiliary databases are generated: 1) the index array, 2) the k -vector array, and 3) the k -vector line array. The index array contains information about the ordering of points in the different dimensions of the database. This index is used to map the database from a sorted order in a given dimension to the ordering of the data structure. The k -vector array is used to quickly identify the database indexes that lay within a defined searching range in a given dimension. This is done by storing the indexes of the first elements, in a sorted database, whose components are above a prescribed set of known values in each dimension. These prescribed values are generated based on a linear mapping function, because a line can easily be inverted to obtain a one-to-one mapping. The elements required to invert this mapping function, the slope and the intercept, are stored in the k -vector line array. Thus, the k -vector line array is used to relate the searching range with the positions in the k -vector array. A summary of this preprocessing effort is presented in Fig. 2.

Once the preprocessing is finished, the n -dimensional k -vector is able to perform the orthogonal search. This is done by following these steps, which are also presented schematically in Fig. 3:

1. **Select a sub-database.** The algorithm selects one of the sub-databases in which the database is distributed.
2. **Apply mapping functions.** For each dimension of the problem, the algorithm applies the mapping functions to the user-specified searching ranges in order to obtain the approximate ranges in the k -vector array.
3. **Estimate the number of elements per dimension.** The algorithm accesses the k -vector array and retrieves the range of indexes that are within the approximate ranges. Based on the design of the k -vector array and one-dimensional k -vector searching process, the approximated range is guaranteed to contain all of the elements in the user-specified range; however, it may also contain extraneous elements due to the range approximation performed. That way, the algorithm can now easily compute the approximate number of elements contained inside the range projection in each dimension.
4. **Sort the dimensions.** All the dimensions are sorted based on the approximate number of points calculated during the previous step. This information is used to determine the dimension whose searching range contains the smallest number of elements. In addition, this information can be used later to define the order in which to evaluate the different dimensions of the problem during the brute force approach.
5. **Remove extraneous elements.** A dimension is selected based on the previous information. Then, a local search is performed at the extreme elements of the approximated range to remove any extraneous elements.
6. **Retrieve the elements.** Using the information provided by the k -vector, the subset of elements from the sub-database inside the range in the selected dimension is retrieved using the index array.
7. **Perform brute force.** A brute force approach is performed on the elements retrieved by the k -vector that checks the dimensions that have not been evaluated yet. Using the information previously computed, the order in which the dimensions are checked is defined by the approximate number of points in each dimension.
8. **Return to step 1.** Another sub-database is selected and the process is repeated until the whole database has been searched.

In the following sections, we describe the data structure of the algorithm, the auxiliary databases, and the searching process in more detail. After presenting these sections, discussing the algorithm complexity, and comparing it with other common orthogonal range searching algorithms, we study the special case of one dimensional databases and discuss possible modifications of the algorithm that focus on reducing the memory required for the algorithm.

3. Data structure

The data structure selected for the n -dimensional k -vector has two objectives. First, it positions the elements that belong to the same regions in the searching space into close or adjacent memory. Second, it provides a way to easily identify subsets of elements that are not contained in the searching range. This is done by performing two sorting processes in two different dimensions of the database. In that regard, and without loss of generality, we will assume that the first sorting process is performed in the last dimension of the database, while the second is performed in the first dimension of the database. This selection is only made to simplify the notation used.

The first step to generate the data structure is to sort all the elements of the database with respect to the last component in ascending order. Let d be the number of dimensions of the problem. Then, the database is sorted in such a way that:

$$x_{i-1}(d-1) \leq x_i(d-1) \leq x_{i+1}(d-1) \forall i \in \{1, \dots, n-2\}, \quad (1)$$

where $x_i(d-1)$ is the component in the last dimension of the point x_i , and n is the number of elements in the database (and thus, $i=0$ and $i=(n-1)$ are the first and last elements of the database).

Once this sorting process is performed, this database is separated into sub-databases of nearly equal numbers of elements. Let n_{db} be the number of sub-databases that are defined for a given database. Then, and since it is not always possible to distribute elements evenly in all sub-databases, we select the first sub-database to be the only one that has a different number of elements. That way, the number of points in a general sub-database is $n_p = \lfloor n/n_{db} \rfloor$, where $\lfloor x \rfloor$ is the largest integer less than the value of x (i.e. the floor function), while $n'_p = n - n_p(n_{db} - 1)$ is the number of elements in the first sub-database. By doing so, we are assuring that all the elements inside a given sub-database have their component d ranging between the last components of the first and last elements of that sub-database, because the database was originally sorted in the last dimension. Mathematically, this is written as,

$$x_p(d-1) \leq x_i(d-1) \leq x_q(d-1) \forall i \in \{p, \dots, q\}, \quad (2)$$

where p and q are the first and last elements of a given sub-database respectively. At this point, it is important to note that the choice of the number of sub-databases, n_{db} , has a large impact on the performance of the algorithm. This is due to the fact that n_{db} is directly related to the number of reference values generated for the last dimension (which control the precision of the range approximation during the searching process), but also because it determines the number of different searches that will be performed later (one for each sub-database). This means that n_{db} can be used to adapt the algorithm to the particularities of the problem considered.

The objective now is to provide a sorted structure in a different dimension. As defined before, each sub-database contains all the points inside a known range in the last dimension. Therefore, if another sorting process is performed in a different dimension, the elements from the same hyperprism will be located in adjacent positions in memory. In order to do

that, this second sorting is performed in each independent sub-database, but this time with respect to the first dimension. This means that the elements within a range in the first dimension are adjacent in memory for each sub-database. This property is used by the algorithm to improve its search performance when possible.

With that, the database structure of the algorithm is set. In the following section, the auxiliary databases are presented, which improve the search performance of the algorithm. However, these auxiliary databases are in fact optional, as it is possible to perform the range searching process with just the data structure as presented in Section 10.

4. Auxiliary databases

As part of the preprocessing effort of the n -dimensional k -vector, the algorithm generates a set of auxiliary databases whose main objective is to provide information on the database distribution. Later, during the searching process, the algorithm will use this information to quickly locate the elements from the database that lay inside a given searching range in one of the database dimensions.

The n -dimensional k -vector makes use of three auxiliary databases: the index array, the k -vector array, and the k -vector line array. These auxiliary databases are computed independently for each sub-database defined in the data structure. This means that during the searching process, each sub-database can be searched independently without further information. This property is interesting, as it may be useful for parallelizing the n -dimensional k -vector algorithm. In the following subsections, these auxiliary databases are studied in more detail.

4.1. Index array

The index array is an auxiliary database that contains information on how the points are sorted with respect to the different dimensions of the problem. To generate it, a sorting process in ascending order must be done for each dimension. Let $j \in \{0, \dots, d-1\}$ be a given dimension of the problem, and let y_j be the array containing the components of all the points in that dimension. Then, we define s_j to be the components of y_j sorted in ascending order. This means that:

$$\begin{aligned} s_j(i) &\leq s_j(i+1), \forall i \in \{0, \dots, n-2\}; \\ s_j(i) &= y_j(I_j(i)), \forall i \in \{0, \dots, n-1\}; \end{aligned} \quad (3)$$

where I_j is the index array. The index array relates the positions of the elements from the original database y_j with respect to the sorted one y_j in the dimension j . The sorting and computation of I_j has a complexity of $\mathcal{O}(n \log n)$ and represents the most demanding process during the preprocessing of the n -dimensional k -vector. Moreover, it is important to note that the index array for any dimension is defined with respect to the same data structure, and thus, a common reference is shared for all the dimensions of the problem.

The sorting process is repeated for all the dimensions of the problem, storing the values of I_j to generate the complete index array I . Once the preprocessing is finished, the values s_j can

be erased to free up memory, as they are not used during the searching process. Note that since the number of points and the number of dimensions is the same as the original database, but the first dimension is already sorted in each sub-database, the size of the index array is $n(d-1)$ integer numbers.

4.2. K-vector array

The k -vector array is an auxiliary database that contains information on how the element components in each dimension are distributed. This auxiliary database requires defining a mapping function that will serve as a reference to which the database distribution will be compared. This mapping function performs a one-to-one application (a bijection) between the index values (the location of elements in the database) and a set of prescribed values in the searching space defined by the mapping function. In other words, the mapping function serves as a first approximation between the sorted values of the database elements and their relative location, that is, their index. In addition, the mapping function is selected in such a way that it is easy to invert to improve search performance.

In order to generate this auxiliary database, we first need to decide the size that this database will have. This is an important parameter not only because it will determine the size of this auxiliary database and the number of operations required to compute it, but also because it affects the search performance of the algorithm. In particular, having a large size implies that the retrieval of elements will be faster since the algorithm will have more resolution in the grid defined in each dimension. Conversely, having a lower size implies that the probability of obtaining extraneous elements from the range increases. This means that the size of the k -vector array can be used as a design parameter to adapt this methodology to fit the needs of the problem to be solved.

Let n_k be the size of the k -vector array for a given dimension $j \in \{0, \dots, d-1\}$, and let $z_j(i)$ with $i \in \{0, \dots, n_k-1\}$ be the mapping function defined for that dimension. Then, the k -vector array for that dimension, k_j , can be described as:

$$k_j(i) = r \mid s_j(r) < z_j(i) \leq s_j(r+1), \quad (4)$$

where $r \in \{0, \dots, n_p-2\}$ and n_p is the number of elements of the sub-database. This means that $k_j(i)$ stores the number of elements in the database smaller than the corresponding value of the mapping function at that same i -th index location, $z_j(i)$, or in other words, it contains the index of the first element that fulfills the relation presented in Eq. (4). This computation has to be done for all the dimensions of the problem, storing the information contained in k_j in the k -vector array, k . As it can be seen, the generation of the k -vector array only requires reading the database and comparing each element component with the auxiliary function. This leads to a complexity of $\mathcal{O}(nd)$ to generate the whole k -vector array, and a memory requirement of $n_k d$ integer numbers for each sub-database.

Fig. 4 graphically shows how the k -vector array, index, and sorted database are connected. In Fig. 4, the sorted database points are shown as asterisks, the discrete mapping function values are shown as short, horizontal lines, and the line that connects the discrete mapping function values is the mapping function. In addition, an example search query from $[a, b]$ is

shown via two long horizontal lines. The x -axis labels on top of the figure show the index value for each of the sorted data points, and the x -axis labels on the bottom of the figure show the k -vector array values for each of the discrete mapping function points. The y -axis gives the value of each of the sorted data points, and the value of each of the mapping function points.

4.3. K-vector line array

The k -vector line array is an auxiliary database that directly relates a given searching range and its associated boundary elements in the k -vector array. This means that both auxiliary databases (the k -vector array and the k -vector line array) always work in cooperation.

In order to generate the k -vector line array, the mapping function of the k -vector must first be defined. In general, the mapping function is selected as a line function due to its easy and fast inversion. In this case, the k -vector provides general information about the non-linearity of the database. The definition of this mapping line is performed by taking the value of the mapping function as the variable of the function, that is:

$$i = m_j z_j(i) + q_j, \forall i \in \{0, \dots, n_k - 1\}, \quad (5)$$

This is done in order to remove the requirement of inverting the mapping function during the searching process, thereby, improving the speed of the methodology. The mapping function has the following expression:

$$z_j(i) = \frac{i}{m_j} - \frac{q_j}{m_j}, \forall i \in \{0, \dots, n_k - 1\}, \quad (6)$$

which represents a line connecting the extreme elements of the sorted database for the dimension j . However, and in order to take into account possible rounding and machine errors, the range of this line is slightly extended to the point-pairs $(1, s_j(l) - \delta\epsilon)$ and $(n_k, s_j(u) + \delta\epsilon)$, where

$$\delta\epsilon = (n_k - 1)\epsilon, \quad (7)$$

ϵ is the relative machine precision, and $s_j(l)$ and $s_j(u)$ are the minimum and maximum values in the sub-database for each dimension. Therefore, m_j and q_j can be defined as:

$$\begin{aligned} m_j &= \frac{n_k - 1}{s_j(u) - s_j(l) + 2\delta\epsilon}, \\ q_j &= -\frac{n_k - 1}{s_j(u) - s_j(l) + 2\delta\epsilon}(s_j(l) - \delta\epsilon), \end{aligned} \quad (8)$$

which are the parameters that are stored in the k -vector line array. This means that, using a line as a mapping function, $2dn_{db}$ real numbers are required to be stored in this auxiliary database for each sub-database defined. Moreover, the algorithm complexity required to generate the complete k -vector line array is $\mathcal{O}(n_{db}d)$.

5. Searching process

The first step in the searching process of the n -dimensional k -vector is to determine the boundary indexes of the k -vector line array for the different ranges in each dimension. This is done using the mapping functions. Let $j \in \{0, \dots, d-1\}$ be a given dimension in which the range $[a_j, b_j]$ is defined. Then, the approximate boundary indexes in the k -vector line array are:

$$\begin{aligned} A_j &= \lceil m_j a_j + q_j \rceil, \\ B_j &= \lceil m_j b_j + q_j \rceil, \end{aligned} \quad (9)$$

where A_j and B_j the lower and upper boundaries in the k -vector array for the dimension j , and $\lceil x \rceil$ is smallest integer larger than x (i.e. the ceil function). A visual example of this process is shown in Fig. 4.

Instead of performing the original search defined by $[a_j, b_j]$, the algorithm is actually searching for the best approximation of the range using the prescribed values generated by the mapping function. Note that these approximated ranges are selected in such a way that the defined range $[a_j, b_j]$ is always contained in the approximation. Later, a small search will be performed on the extreme elements of the range, but for the moment, we are only interested in having an approximated value of the number of elements contained in the range of each dimension. In that regard, the number of points within this approximated range in the dimension j , $p(j)$, can be obtained using the following expression:

$$p(j) = k_j(B_j) - k_j(A_j). \quad (10)$$

Then, if the operations defined in Eqs. (9) and (10) are repeated for all the dimensions of the database, it is possible to sort the dimensions based on the number of points expected to be retrieved. This is done by applying a sorting algorithm to array p (of size d), to obtain a sorted array g in ascending order.

It is important to note that since the number of points in the searching range of each dimension is evaluated, there will be entire sub-databases where no retrieval will be performed. Particularly, these sub-databases correspond to the ones whose elements are outside the searching range of the last dimension, and thus, they can be removed from the searching range without any problem. This also means that the algorithm does not need to continue the searching process in these sub-databases. As a result, this property, derived from the data structure, allows performing an active search in two dimensions instead of just the one (which is provided by the k -vector methodology).

At this point, we already know which dimension has the smallest number of elements in its searching range. In general, this should be the dimension in which the algorithm should continue the searching process by using the information provided by the k -vector array. However, this is not always true. The reason for this is that checking the searching range is faster when the elements studied are adjacent in memory. Therefore, there could be cases where continuing the searching process in the first dimension of the problem is faster despite having to evaluate a larger number of elements compared with other dimensions. For

instance, with the test computers used by the authors, the ratio in speed between reading the database in an ordered process and accessing the database in a random order is on average:

$$r = \frac{3}{2}(\log_{10}(n_p) - 3), \quad (11)$$

where $n_p = 10^3$ is the number of points to be studied (or in the case of the n -dimensional k -vector, the number of elements in each sub-database). For this reason, the n -dimensional k -vector decides which is the most convenient dimension to continue the searching process based on the relation between $g(0)$ and $p(0)$ (the first components of the sorted and unsorted arrays for approximate number of elements per dimension). That way, two options are available. If $p(0) > g(0)$, the n -dimensional k -vector continues the searching process in the dimension with the smallest number of elements. Conversely, the algorithm continues the searching process in the first dimension of the data structure. Regardless of the dimension that is selected in this step, the process followed afterwards is completely equivalent.

Once the most convenient dimension is selected, the algorithm first checks the real range of the dimension selected, and then, it continues with the rest of dimensions of the problem. In the following lines we study each case independently due to their differences.

Let j be the dimension selected by the algorithm. Then, from the application of the n -dimensional k -vector, we know that all the elements in the searching range are contained in the subset of points defined by these indexes:

$$\begin{aligned} i &| (A_j \leq i < B_j) \text{ if } j = 0, \\ I_j(i) &| (A_j \leq i < B_j) \text{ if } j \neq 0. \end{aligned} \quad (12)$$

Therefore, we first have to remove the elements that are located outside the range $[a_j, b_j]$. In order to do that, and instead of checking all the elements individually, a linear search is performed on the extreme elements of the range by increasing the value of A_j and decreasing the value of B_j until the extreme elements are within the searching range. This means that for linear databases, on average, there will be n_p/n_k elements that must be removed. If the database is highly non-linear, or the size of the k -vector array is not large enough, this linear search can be substituted by a binary search. This binary search is performed between the elements $k_j(A_j)$ and $k_j(A_j + 1)$ for the extreme elements at the lower end of the searching range, and between elements $k_j(B_j - 1)$ and $k_j(B_j)$ for the extreme elements at the upper end of the searching range. Note also that the selection of the methodology to follow can be done during the searching process since we know the number of elements in the extremes (given by $[k_j(A_j + 1) - k_j(A_j)]$, and $[k_j(B_j) - k_j(B_j - 1)]$ respectively), and the expected number of elements for a linear distribution (n_p/n_k). Therefore, if the number of elements in the extremes is several times larger than the expected, the algorithm should change to a binary search in order to remove the possible extraneous elements and improve its performance.

The remaining dimensions are evaluated using a brute force approach, where the components in each dimension are checked following the ordering of dimensions defined by p and g . This is done to slightly improve the search performance of the algorithm since it reduces the average number of range comparisons in the elements retrieved.

6. Simple example of application

This section provides an example of the data structure, auxiliary databases, and searching process for a simple database consisting of ten elements and three dimensions. The original database considered is shown in Table 1, where the variables x , y , and z are used to denote the first, second, and third dimension respectively. The objective is to identify the elements within the range $[2, 8] \times [5, 6] \times [1, 3]$, that is, $2 \leq x \leq 8$, $5 < y \leq 6$, and $1 \leq z \leq 3$.

6.1. Preprocessing - data structure

First, the original database is sorted in the last dimension, as is shown in Table 2. Then, the original database sorted in the last dimension from Table 2 is divided into sub-databases (two sub-databases in this case) and re-sorted in the first dimension. These sub-databases are shown in Table 3. Table 3 also includes an element index that can be used to identify each element.

6.2. Preprocessing - auxiliary databases

After creating the sub-databases shown in Table 3, the auxiliary databases can be generated. The first auxiliary database is the index array. This data structure is used to sort each dimension of each sub-database in ascending order. Note that although the index array need not be generated for the first dimension because the sub-databases are sorted in the first dimension, it is included anyway for completeness. The index array is shown in Table 4.

Then, the k -vector line array can be generated according to Eq. (8), and the k -vector array can be generated according to Eq. (4). The k -vector line array, rounded to two decimal places, is shown in Table 5, and the k -vector array is shown in Table 6.

6.3. Searching process

The first step is to find the approximate boundaries of the searching range. By using Eq. (9) in the first sub-database:

$$\begin{aligned} A_x &= \text{floor}(1.00 \cdot 2 - 2.00) = 0 \rightarrow k_x(A_x) = 0; \\ B_x &= \text{ceil}(1.00 \cdot 8 - 2.00) = 6 \rightarrow k_x(B_x) = 5; \\ A_y &= \text{floor}(0.44 \cdot 5 + 0.00) = 2 \rightarrow k_y(A_y) = 2; \\ B_y &= \text{ceil}(0.44 \cdot 6 + 0.00) = 3 \rightarrow k_y(B_y) = 3; \\ A_z &= \text{floor}(1.00 \cdot 1 + 0.00) = 1 \rightarrow k_z(A_z) = 1; \\ B_z &= \text{ceil}(1.00 \cdot 3 + 0.00) = 3 \rightarrow k_z(B_z) = 4; \end{aligned}$$

which using Eq. (10) means that in dimension x there are 5 elements; in dimension y , 1; and in dimension z , 3 elements. Therefore, y is the dimension selected to perform the projection. In this projection there is only one possible candidate, the point corresponding to the position 2 in the sorted database with respect to dimension y , that is, the point with index $I_y(2) = 3$ and coordinates (5, 6, 2). Now, this point must be checked in the other dimensions, starting with dimension z (the dimension with the lowest number of elements excluding dimension y), and continuing with dimension x . That way, we conclude that the point with index 3 lies within the searching range.

Additionally, the algorithm must also check the second sub-database. By using Eq. (9) again in this sub-database, we find that there are no elements in the projection of the searching range in dimension z (since $k_z(A_z) = k_z(B_z) = 5$). Thus, this sub-database does not contain any elements within the searching range. Therefore, the algorithm has finished the searching process.

7. Algorithm complexity

In order to estimate the algorithm complexity, we will assume a worst case scenario for the algorithm. Since the algorithm evaluates the number of elements in the range of each dimension, the worst case scenario for the n -dimensional k -vector happens when the projection of the searching range in every dimension contains the same number of elements. In that regard, it is important to remember that the algorithm performs a projection of the database in two selected dimensions, the last one through the database structure, and another selected by the k -vector. This means that, in general, the number of points that the algorithm must study is larger than the number of points in the searching range. Therefore, we need to evaluate this number of elements since it constitutes the largest computational effort made by the algorithm.

Let n be the number of elements of the database, d the number of dimensions of the database, and k the number of elements contained in the searching range. Then, and under the worst case scenario considered, the range in each dimension contains $n(k/n)^{(1/d)}$ elements, which are the ones that should be assessed if only a k -vector is applied. Note that this value represents a fraction of the number of elements considered by the k -vector based on the number of dimensions and the number of elements that have to be retrieved. However, the n -dimensional k -vector also arranges the database into n_{db} sub-databases based on the distribution in the last dimension, which allows the algorithm to remove the elements that are outside the searching range in the last dimension. This means that before applying the k -vector the size of the database is reduced to $n(k/n)^{(1/d)}$ elements, and thus, the k -vector only deals with this number of points instead of the whole database. Therefore, the number of elements that are evaluated by the algorithm is $n(k/n)^{(2/d)}$. Hence, due to the fact that the number of dimensions is d , the algorithm complexity is $\mathcal{O}(nd(k/n)^{(2/d)})$.

The previous study only considers the number of elements that are evaluated and not the number of operations required to obtain those elements. Therefore, we also study these computations in more detail and show that they represent a much smaller number of operations during the searching process. First, the algorithm requires transforming the searching range into the boundaries of the k -vector line array. This is done in $\mathcal{O}(d)$ operations. Once this is completed, the dimensions are sorted based on the expected number of points retrieved. Using an algorithm such as quicksort or mergesort implies a complexity of $\mathcal{O}(d \log d)$. Finally, the algorithm must check if extraneous elements are contained in the data retrieved by the k -vector for the chosen dimension. If a linear search is performed in the extreme elements, this will require an average of $n/(n_{db}n_k)$ operations. Note that if the size of the k -vector array is equal to the size of the database, $n = n_{db}n_k$, there is only one extraneous element on average. Conversely, if a binary search is performed, the algorithm will take an

average of $\log n/(n_{db}n_k)$ operations. As it can be seen, all these processes are orders of magnitude smaller than the algorithm complexity derived before, and thus, they are considered negligible in the calculation of the algorithm complexity.

8. Performance of the n -dimensional k -vector

The search times presented in this section were gathered in C using the Query Performance Counter function on a Windows 10 operating system using an Intel(R) Core(TM) i7-7700 CPU running at 3.60 GHz with 16 GB of RAM. Moreover, in order to be conservative in the evaluation of the search performance of the n -dimensional k -vector, the worst case scenario for this method has been considered; that is, there are approximately the same number of elements in the range of each dimension. This implies that the n -dimensional k -vector cannot benefit from evaluating the dimensions in a particular order. In addition, the number of elements used in the k -vector array for these tests was set to be a tenth the size of the number of elements in the database ($n_k = n/10$).

Figs. 5–8 compare how the number of dimensions affects search time differences between n -dimensional k -vector and the brute force and k -d tree algorithms. These figures were generated on a database of one million elements with a number of dimensions that ranged from one dimension to 20 dimensions. The horizontal axis of Figs. 5–8 shows the number of dimensions of the database and the vertical axis shows the n -dimensional k -vector speedup. The n -dimensional k -vector speedup is quantitatively defined as the search time of another algorithm (i.e. brute force or k -d tree) divided by the search time of the n -dimensional k -vector algorithm. If this number is greater than one, the n -dimensional k -vector is faster than the algorithm being compared against; if it is less than one, then the n -dimensional k -vector is slower. In particular, Fig. 5 was created using a range search that retrieved approximately 0.01% of the database, and Fig. 6 retrieved approximately 0.1% of the database. Fig. 7 retrieved approximately 1% of the database, and Fig. 8 retrieved approximately 10% of the database.

Figs. 5–8 show that the n -dimensional k -vector is faster than the other two algorithms in all instances tested except for the seven dimensional case in Fig. 5. In this instance, the n -dimensional k -vector speedup was 0.999, meaning that the k -d tree search time was almost identical to that of the n -dimensional k -vector. In general, the n -dimensional k -vector speedup for brute force increases as the number of dimensions decreases. The n -dimensional k -vector speedup for k -d tree initially decreases as the dimensionality of the database increases, but then the n -dimensional k -vector speedup starts increasing again starting somewhere between dimensions three and seven. Then, the n -dimensional k -vector speedup asymptotically approaches a final value around ten. In general, the n -dimensional k -vector speedup for brute force increases as the percentage of elements retrieved from the database decreases. The n -dimensional k -vector speedup for k -d tree has the opposite trend; however, the effect that percentage of elements retrieved from the database has on the n -dimensional k -vector speedup for k -d tree is not nearly as great as the effect it has on the n -dimensional k -vector speedup for brute force.

Fig. 9 compares the search time of the three algorithms, n -dimensional k -vector, brute force, and k -d tree, as a function of the number of elements in the database; the database has six dimensional elements, and the search queries were designed to return approximately five percent of the database. Each data point and associated error bar in Fig. 9 was created by performing the same search query 100 times for each algorithm. The mean of the 100 search queries is shown as a point, and the error bars associated with each point extend to the slowest and fastest search time for each query¹

Fig. 9 shows that the n -dimensional k -vector has the fastest average search time of the three algorithms for every database size; the brute force method is the second fastest, and k -d tree is the slowest. The difference between the average n -dimensional k -vector search time and the brute force time is typically less than an order of magnitude, but the difference between n -dimensional k -vector and k -d tree is larger than an order of magnitude in some instances. The error bars follow similar trends to the average values.

Figs. 10 and 11 compare the n -dimensional k -vector speedup to the percentage of the database retrieved from a database of one million elements. Fig. 10 was generated using a two-dimensional database, and Fig. 11 was generated using a ten-dimensional database. Each of the data points used to calculate the n -dimensional k -vector speedup in these two figures was generated by measuring the search time of 1000 random searches.

Figs. 10 and 11 show that n -dimensional k -vector is faster than the other two algorithms in every case except for when 10⁻⁴% of the database is retrieved in Fig. 10. In this instance, the n -dimensional k -vector speedup is 1.0, meaning that the n -dimensional k -vector and k -d tree had the same search time. The n -dimensional k -vector speedup for brute-force decreases in both figures as the percentage of the database that is retrieved increases. In addition, the n -dimensional k -vector speedup for brute force in ten dimensions is less than in the two dimensional case. In the two dimensional case (see Fig. 10), the n -dimensional k -vector speedup for k -d tree initially increases, and then remains approximately constant as the percentage of elements retrieved increases. In the ten-dimensional database (see Fig. 11), the n -dimensional k -vector speedup for k -d tree increases as the percentage of elements retrieved increases, except for the final data point, 100% of the database, where the n -dimensional k -vector speedup decreases.

8.1. Preprocessing performance

In addition to the search performance of the n -dimensional k -vector, the time the algorithm requires to generate the database structure and the auxiliary databases has also been studied. To that end, a direct comparison has been performed between the preprocessing times of the n -dimensional k -vector and k -d tree. Fig. 12 shows the results of this study for a database of one million elements and a number of dimensions ranging from 1 to 20. As it can be seen, the preprocessing of the n -dimensional k -vector takes less time than the k -d tree if the number of dimensions is low (5 dimensions or less). However, as the number of dimensions of the database increases, the n -dimensional k -vector requires more time than the k -d tree to

¹The slowest and fastest search times were chosen for the error bars rather than the 95% confidence interval of the t -distribution, as the 95% confidence interval of the t -distribution is too small to be visible in the figure.

perform the preprocessing. This is due to the increased number of database sorting processes that the n -dimensional k -vector must perform as the number of dimensions increases.

9. The one dimensional k -vector

Although in general the process described for the case of the n -dimensional k -vector can be followed without any modification when dealing with one dimensional databases, due to its special particularities and properties, it is interesting to study in more detail the case of the one dimensional k -vector. The k -vector is an already known methodology that has been used in several applications [17–20]. However, in this section we include some important modifications compared to previous works that come as a result of the development of the n -dimensional k -vector. Examples of these modifications include the database structure, the definition of the mapping function, and an alternative way to remove the extraneous elements. Additionally, a study of the algorithm complexity is also included. We study these modifications in more detail in the following subsections.

9.1. Preprocessing

Dealing with one dimensional databases introduces two key differences during the preprocessing effort of the algorithm. First, the algorithm does not require generating sub-databases, because searching in multiple dimensions is no longer required. Second, the index array does not provide any useful information since the data structure is already sorted in the only dimension of the problem. This means that the preprocessing effort can be greatly simplified for the case of one dimensional databases.

Therefore, the database structure is based on sorting the whole database in ascending order, and does not require any additional processes. Then, the mapping function is defined over this sorted structure, and k -vector and k -vector lines arrays are generated for the whole database (no sub-database is defined). This implies that the k -vector array is composed of n_k integer numbers, while the k -vector line array becomes just two real numbers, the slope of the mapping function m , and its intercept q .

9.2. Searching process

The searching process in the one dimensional k -vector is straightforward. The searching range is first approximated as a boundary range in the k -vector line array using Eq. (9), and then a search is performed on the extreme elements of the range to check for possible extraneous elements. As in the case of the n -dimensional k -vector, this can be done by a linear search, or if the database is highly non-linear, with a binary search. This allows the algorithm to retrieve all the elements contained in the searching range without further operations.

9.3. Algorithm complexity

In this subsection we deal with the evaluation of the algorithm complexity of the one dimensional k -vector. In that regard, two different scenarios are considered, the average case, and the worst case scenario for the algorithm.

In an average case, there are n/n_k elements in a given range $[z(i), z(i+1)] \forall i \in \{1, \dots, n_k - 1\}$ defined by the mapping function. This means that the expected number of extraneous elements during the searching process is n/n_k . Then, if a linear search is performed, the algorithm will require an average of n/n_k operations to finish the searching process, while a binary search applied at the k -vector array boundaries has a complexity of $\mathcal{O}(\log(n/n_k))$. Note that when following this methodology, it is possible to maintain the speed of the algorithm regardless of the number of elements contained in the database as long as the relation n/n_k is maintained.

In the worst case scenario for the methodology, all the elements of the database except one are contained in a particular range $[z(i), z(i+1)]$ of the searching space. This implies that for a linear search would require, on average, the evaluation of $(n-1)/2$ elements from the database, which greatly deteriorates the search performance of the algorithm. For this reason, a linear search in these cases is not recommended. Instead, and since this is a very non-linear database, a binary search approach must be performed as commented in the previous subsection. That way, a complexity of $\mathcal{O}(\log n)$ is obtained where the process is equivalent to doing a pure binary search.

10. Reducing the memory required

There are some applications where it is useful to reduce the amount of memory required for the algorithm. Examples include databases with a very large number of elements and dimensions, or the use of systems with limitations in memory. For this reason, in this section we discuss the possibilities of memory reduction in the n -dimensional k -vector methodology.

In particular, two approaches are considered in the following subsections. The first one is based on removing the necessity of the index array by fixing the dimension in which the projection is performed during the searching process. The second approach deals with the possibility of removing the k -vector and the k -vector line arrays by substituting them for a binary search technique. Note that performing these modifications produces a reduction in the search performance of the algorithm, but not in the algorithm complexity.

10.1. Removing the index array

The index array is the auxiliary database that generally requires more memory, because it is the same size as the database. Therefore, it is the first candidate to focus on in order to reduce the amount of memory required. If the index array is removed, it is no longer possible to select the most convenient dimension for the algorithm to perform a projection in, as it is no longer possible to define an effective ordering of the dimensions. Therefore, in this case the n -dimensional k -vector focuses on just two dimensions, which correspond to the ones used in the definition of the database structure. That way, the k -vector and k -vector line arrays are only defined in the first dimension of the problem, and thus, they have size $n_{db}n_k$ and $2n_{db}$ respectively.

The removal of the index array, however, leads to a new problem. Since we have removed the information about the last dimension of the database from the k -vector array, we can no

longer discard entire sub-databases during the searching process, which increases the number of elements that have to be evaluated. This difficulty can be overcome by storing the values of the smallest components in the last dimension of each sub-database in an array. Then, during the searching process, a binary search will be performed in this array to identify which sub-databases need to be searched. This modification does not change the overall algorithm complexity, because the binary search is of order $\log n_{db}$, which is much smaller than the other effects.

10.2. Removing the k -vector array and k -vector line array

It is also possible to reduce the size of the auxiliary database by removing the k -vector array and k -vector line array. This can be done regardless of whether the index array is removed or not. The idea of this modification is the following. Instead of using the k -vector to locate the boundary points of the searching range in each dimension, a binary search technique is applied, using the information of the sorted database in each dimension provided by the index array or the database structure. This means that for each dimension, two binary searches of complexity $\mathcal{O}(\log n)$ must be performed that correspond to the two boundary values of the range. Once these elements are located, the process continues as described previously in the searching process.

It is important to note that although the computed speed of this modification is worse than the original methodology, its complexity is maintained. In particular, this modification has a complexity of $\mathcal{O}\left(d \log n + nd(k/n)^{(2/d)}\right)$, where the second term is dominant. Therefore, the order of magnitude of the algorithm complexity is still $\mathcal{O}\left(nd(k/n)^{(2/d)}\right)$, the same as in the original methodology. In fact, having removed the k -vector from the process, the resultant algorithm is reminiscent of a modified projection method applied to two dimensions.

10.3. Search performance of the modification

In this subsection we compare the search performance of the n -dimensional k -vector with its modified version as described in this section. In particular, we consider the case with no index array, no k -vector array, and no k -vector line array. Fig. 13 shows the results of this comparison for a database of one million points as a function of the number of dimensions, where approximately 1% of the database was retrieved. As it can be seen, the original n -dimensional k -vector outperforms the modified version regardless of the number of dimensions of the database, where the speed differences come from the advantage that the n -dimensional k -vector has due to its auxiliary databases. The most noticeable difference is when dealing with databases in two dimensions, because the n -dimensional k -vector's auxiliary databases provide a better map of the search space. This difference becomes less relevant as the number of dimensions of the database increases. Note too that this behavior confirms the algorithm complexity study performed in this section. Moreover, when comparing Fig. 13 with Fig. 7, we can see that this modified version of the n -dimensional k -vector is still faster than a brute force or a k -d tree approach for the cases studied.

11. Conclusions

This work presents the n -dimensional k -vector algorithm, a methodology to perform orthogonal range searching in static databases with multiple dimensions. The n -dimensional k -vector is shown to be fast, easy to implement, and adaptable to the memory requirements of the problem to be solved. In addition, performance comparisons between the n -dimensional k -vector and other common algorithms from the literature, k -d tree and brute force approach, are shown. In the scenarios studied, the algorithm has proven to be faster than the aforementioned methodologies. Moreover, for all the tests performed during the creation of this article, the n -dimensional k -vector was always faster than a brute force approach regardless of the number of elements retrieved, the size of the database, or the dimensionality of the database. These tests included the retrieval of the whole database and problems with more than one hundred dimensions.

The basic idea behind the n -dimensional k -vector is to find the dimension where a projection can be performed such that the expected number of points to be evaluated in each dimension is minimized. To achieve that goal, an auxiliary database is generated that can obtain a good approximation of the number elements that meet the searching criteria in each dimension, which in turn can be used to determine which dimension to perform the projection in. Furthermore, the database structure of the n -dimensional k -vector, its auxiliary databases, and the whole searching process can be defined completely using vectorial notation, which simplifies the formulation of the method and eases the implementation of the algorithm in the majority of systems.

The n -dimensional k -vector requires a one time preprocessing effort that first rearranges the database following the database structure presented, and then generates the auxiliary databases. Therefore, although not impossible, the capabilities of this methodology for dynamic databases (databases that have elements added or removed over time) are reduced, because, in general, these structures must be updated when a database modification is performed. Nevertheless, for static databases, this is not a limitation, and the algorithm is able to work with the original preprocessing indefinitely.

In addition, the n -dimensional k -vector easily adapts to the memory available for any given problem. In that sense, several possibilities of modification for the algorithm have been presented. These modifications even included the case where no auxiliary database was defined, and thus, only the original sorted database was required. For a general application, the size of the k -vector and k -vector line arrays can adapt to the memory available for each particular application. That way, if more memory is available, the number of operations performed can be reduced. This effect is due to the fact that the initial range approximation from the k -vectors is closer to the actual searching range. Furthermore, if only an approximated range is required for a given application, then the n -dimensional k -vector becomes even more efficient for two reasons. First, the algorithm no longer requires checking the extreme elements of the points retrieved by the k -vectors. Second, it is possible to set the values of the approximated ranges through the mapping function. This provides perfect control over the actual searching ranges of the algorithm.

Although the obvious application of this algorithm is obtaining multidimensional elements inside any orthogonal searching range in massive databases, the algorithm can be used in other applications such as isosurface identification, function inversion, enhancing compressive sampling, and random sampling. Moreover, based on the results obtained, the algorithm is especially useful when a large proportion of the elements from the database are retrieved. Examples of these types of applications include isosurface identification and the computation of communications coverage in space missions. Nevertheless, in the problems where the number of elements retrieved is always very small, such as the nearest neighbor problem, the n -dimensional k -vector has also proven to be very fast, obtaining a search performance comparable with other methodologies such as a k -d tree. Such applications using the n -dimensional k -vector will be studied in future works.

Acknowledgments

David Arnas acknowledges the support of the Spanish Ministry of Economy and Competitiveness, Project No. ESP2017-87113-R (AEI/FEDER, UE); and the Aragon Government and European Social Fund (group E24_17R). Carl Leake acknowledges the support of a NASA Space Technology Research Fellowship [NSTRF 2019], Grant #: 80NSSC19K1152.

References

- [1]. Bentley JL, Friedman JH, Data structures for range searching, *ACM Comput. Surv. (CSUR)* 11 (4) (1979) 397–409, doi:10.1145/356789.356797.
- [2]. Nievergelt J, Hinterberger H, Sevcik KC, The grid file: an adaptable, symmetric multikey file structure, *ACM Trans. Database Syst. (TODS)* 9 (1) (1984) 38–71, doi:10.1145/348.318586.
- [3]. Bayer R, McCreight E, Organization and Maintenance of Large Ordered Indexes, in *Software Pioneers*, Springer, Berlin, Heidelberg, 2002, pp. 245–262, doi:10.1007/978-3-642-59412-0_15.
- [4]. Finkel RA, Bentley JL, Quad trees a data structure for retrieval on composite keys, *Acta Inform.* 4 (1) (1974) 1–9, doi:10.1007/BF00288933.
- [5]. Bentley JL, Multidimensional binary search trees used for associative searching, *Commun. ACM* 18 (9) (1975) 509–517, doi:10.1145/361002.361007.
- [6]. Guttman A, R-trees: a dynamic index structure for spatial searching, *ACM* 14 (2) (1975) 47–57, doi:10.1145/971697.602266.
- [7]. Robinson JT, The KDB-tree: a search structure for large multidimensional dynamic indexes, in: *Proceedings of the 1981 ACM SIGMOD International Conference on Management of Data, AMC*, 1981, pp. 10–18, doi:10.1145/582318.582321.
- [8]. Willard DE, New data structures for orthogonal range queries, *SIAM J. Comput.* 14 (1) (1985) 232–253, doi: 10.1137/0214019.
- [9]. Lueker GS, A data structure for orthogonal range queries, in: *19th Annual Symposium on Foundations of Computer Science (SFCS 1978)*, IEEE, 1978, pp. 28–34, doi:10.1109/SFCS.1978.1.
- [10]. Alstrup S, Brodal GS, Rauhe T, New data structures for orthogonal range searching, in: *Proceedings 41st Annual Symposium on Foundations of Computer Science, IEEE*, 2000, pp. 198–207, doi:10.1109/SFCS.2000.892088.
- [11]. Arya S, Mount DM, Approximate range searching, *Comput. Geometry* 17 (3–4) (2000) 135–152, doi:10.1016/S0925-7721(00)00022-5.
- [12]. Agarwal PK, Erickson J, Geometric range searching and its relatives, *Contemp. Math.* 223 (1999) 1–56.
- [13]. Chazelle B, Lower bounds for orthogonal range searching: I. The reporting case, *J. ACM (JACM)* 37 (2) (1990) 200–212, doi:10.1145/77600.77614.

- [14]. Mortari D, A fast on-board autonomous attitude determination system based on a new star-ID technique for a wide FOV star tracker, AAS/AIAA Space Flight Mechanics Meeting, Austin, TX, 1996.
- [15]. Crew RVG, Dorty J, Experience with the pyramid algorithm, Tech. rep, MIT Center for Space Research, Cambridge MA, 02139 USA, 2002 .
- [16]. Brady T, Tillier C, Brown R, Jimenez A, Kourepenis A, The inertial stellar compass: a new direction in spacecraft attitude determination, in: Proceedings 16th Annual USU Conference on Small Satellites, 2002 Paper SSC02-II-1
- [17]. Mortari D, Samaan M, Bruccoleri C, Junkins JL, The pyramid star identification technique, ION Navig. 51 (3) (2004) 171–183, doi:10.1002/j.2161-4296.2004.tb00349.x.
- [18]. Arnas D, Mortari D, Nonlinear function inversion using k -vector, Appl. Math. Comput. 320 (2018) 754–768, doi:10.1016/j.amc.2017.10.009.
- [19]. Arnas D, Leake C, Mortari D, Random sampling using k -vector, Comput. Sci. Eng 21 (1) (2019) 94–107, doi:10.1109/MCSE.2018.2882727.
- [20]. Mortari D, Rogers J, A k -vector approach to sampling, interpolation, and approximation, AAS J. Astronaut. Sci 60 (3) (2015) 686–706.
- [21]. Mortari D, Memory adaptive k -vector, in: AAS 14-207, 2014 AAS/AIAA Space Flight Mechanics Meeting Conference, Santa Fe, NM, Jan. 26–30, 2014.
- [22]. Mortari D, Neta B, K -vector range searching techniques, Adv. Astronaut. Sci 105 (Pt. I) (2000) 449–464.

Fig. 1.
Generation of the database structure.

Fig. 2.
Preprocessing process.

Fig. 3.
Searching process.

Fig. 4.
One dimensional k -vector example.

Fig. 5. N -dimensional k -vector speedup versus number of dimensions for a database with one million elements and approximately 0.01% of the database being retrieved.

Fig. 6. N -dimensional k -vector speedup versus number of dimensions for a database with one million elements and approximately 0.1% of the database being retrieved.

Fig. 7.
 N -dimensional k -vector speedup versus number of dimensions for a database with one million elements and approximately 1% of the database being retrieved.

Fig. 8. N -dimensional k -vector speedup versus number of dimensions for a database with one million elements and approximately 10% of the database being retrieved.

Fig. 9. Search time versus number of elements for a six dimensional database with five percent of the database being retrieved.

Fig. 10.
 N -dimensional k -vector speedup versus percentage of the database being retrieved for a two-dimensional database with one million elements.

Fig. 11.
 N -dimensional k -vector speedup versus percentage of the database being retrieved for a ten-dimensional database with one million elements.

Fig. 12. N -dimensional k -vector preprocessing speedup versus a k -d tree for a database with one million elements.

Fig. 13. N -dimensional k -vector speedup versus algorithm not using auxiliary databases for a database with one million elements and approximately 1% of the database being retrieved.

Table 1

Original database.

x	6	9	0	2	4	3	5	1	8	7
y	9	3	2	7	1	0	6	8	4	5
z	1	9	5	3	4	0	2	8	6	7

NASA Author Manuscript

NASA Author Manuscript

NASA Author Manuscript

Table 2

Original database sorted in last dimension.

x	3	6	5	2	4	0	8	7	1	9
y	0	9	6	7	1	2	4	5	8	3
z	0	1	2	3	4	5	6	7	8	9

NASA Author Manuscript

NASA Author Manuscript

NASA Author Manuscript

Table 3

Sub-databases.

index	0	1	2	3	4	index	5	6	7	8	9
x	2	3	4	5	6	x	0	1	7	8	9
y	7	0	1	6	9	y	2	8	5	4	3
z	3	0	4	2	1	z	5	8	7	6	9

NASA Author Manuscript

NASA Author Manuscript

NASA Author Manuscript

Table 4

Index array.

x	0	1	2	3	4	x	5	6	7	8	9
y	1	2	3	0	4	y	5	9	8	7	6
z	1	4	3	0	2	z	5	8	7	6	9

Table 5

K -vector line array.

	m	q		m	q
x	1.00	-2.00	x	0.44	0.00
y	0.44	0.00	y	0.67	-1.33
z	1.00	0.00	z	1.00	-5.00

NASA Author Manuscript

NASA Author Manuscript

NASA Author Manuscript

Table 6

K -vector array.

x	0	1	2	4	5	x	5	7	7	7	10
y	0	2	2	3	5	y	5	7	8	9	10
z	0	1	2	4	5	z	5	6	7	9	10

NASA Author Manuscript

NASA Author Manuscript

NASA Author Manuscript