AN ALGORITHM FOR SEGMENT-DRAGGING AND ITS IMPLEMENTATION

Bernard Chazelle

CS-TR-077-87

January 1987

AN ALGORITHM FOR SEGMENT-DRAGGING AND ITS IMPLEMENTATION

by

Bernard Chazelle

Department of Computer Science

Princeton University

Princeton, NJ 08544

Abstract:

Given a collection of points in the plane, pick an arbitrary horizontal segment and move it vertically until it hits one of the points (if at all). This form of *segment-dragging* is a common operation in computer graphics and motion-planning. It can also serve as a building block for multidimensional data structures. This note describes a new approach to segment-dragging which yields a simple and efficient solution. The data structure requires $O(n)$ storage and $O(n \log n)$ preprocessing time, and each query can be answered in $O(\log n)$ time, where n is the number of points in the collection. The method is best understood as the end result of a sequence of transformations applied to a simple but inefficient starting solution.

This work was started while the author was a visiting professor at Ecole Normale Supérieure, Paris, France.

1. Introduction

This note is devoted to the following *segment-dragging* problem. Let $P = \{p_i : (x_i, y_i) | 0 \leq i < n\}$ be a collection of points in the plane, represented in a Cartesian system of coordinates (Ox, Oy). For any horizontal segment AB, with $A = (a, c)$, $B = (b, c)$ and $a < b$, we define $\texttt{hit}(AB) = i$, where p_i is the point of maximum ordinate, such that

$$a \leq x_i \leq b \qquad \text{and} \qquad y_i \leq c. \tag{1}$$

See Fig.1. If no p_i satisfies (1) then we assume that $\texttt{hit}(AB) = n$. If, on the contrary, more than one point have the largest ordinate among those verifying (1), we then pick the candidate with the smallest abscissa as $\texttt{hit}(AB)$. With these definitions at hand the problem can be formulated quite simply: implement the function `hit` *effectively*, using *little* preprocessing.

Segment-dragging is a common operation in computer graphics, especially for windowing tasks (Edelsbrunner et al [EOS]). It can also be useful for a window manager. As observed by Mitchell [M], segment-dragging is important in certain problems of motion-planning. Also, it is often a tool for more complex operations. For example, it can be used to test whether a query rectangular box is empty, given a collection of points in any fixed dimension.

In practice, dynamic solutions —where points can be added and removed— are often preferable to static ones. We restrict our attention to the static case, however, because the solution described in this paper falls in a certain family of data structures which can be efficiently dynamized. Unfortunately, the dynamic version is complicated and can hardly be called practical. On the contrary, our static solution is easy to implement and seems very efficient in practice. The size of the data structure is $O(n)$ and the time for its construction is $O(n \log n)$; moreover any query can be answered in $O(\log n)$ time. This is a new complexity result.

Previous work on segment-dragging includes a solution in $O(n \log n)$ space and $O(\log n)$ query time (Edelsbrunner et al [EOS]). As can be easily observed, segment-dragging is a special case of range searching for maximum. In the latter problem points are assigned weights, and a query

requests the point of maximum weight falling in a given rectangle. Thus, to perform segment-dragging, one can use the various algorithms for range searching described in (Chazelle [C]). In particular, this gives three solutions: one in $O(n)$ space and $O(\log^{1+\epsilon} n)$ query time, for any $\epsilon > 0$; another in $O(n \log\log n)$ space and $O(\log n \log\log n)$ query time; and a third one in $O(n \log^{\epsilon} n)$ space and $O(\log n)$ query time, again for any $\epsilon > 0$. A simpler, but less efficient, solution can be found in (Gabow et al [GBT]). Note that the algorithm which we present here is not a substitute for these solutions since it concerns a more restrictive problem. It uses a compaction scheme introduced in (Chazelle [C]), as well as a number of bit-twiddling tricks, in order to produce a data structure which is both simple and efficient.

We believe that our solution is best understood as the end result of a series of transformations applied to a simple but inefficient data structure. For the sake of clarity our discussion will follow this transformational approach. In the end we will obtain a complete implementation of the algorithm in standard *PASCAL*. Using techniques from (Chazelle [C]) the algorithm can be ported to a pointer machine, that is, all address calculations can be eliminated. Admittedly, much of its simplicity is lost in the process.

2. Preliminaries

In order to confine our discussion to the essential parts of the algorithm we choose to make a number of simplifying assumptions:

(i) A query segment AB may not contribute the smallest or largest coordinate in $P \cup \{A, B\}$.

(ii) A query segment may not share coordinates with the points of P and the latter may not share coordinates among themselves.

(iii) The number of points is of the form $n = 2^w$ $(w > 0)$.

(iv) A computer word can hold any integer between 0 and $n \log n$ (we choose this value for convenience, and others, such as n, would do just as well).

It is elementary to satisfy all these conditions. For (i), the query segment can be clipped, if needed. To ensure (ii) we can go into rank space using presorting. As to (iii), if necessary, we can always pad P with extra points at "infinity" to reach the nearest power of two. Finally, to satisfy (iv), we can use several words at once to emulate a longer computer word (if that is still not sufficient then time has come to buy yourself a new computer).

To illustrate our discussion we will use a running example throughout this paper. The points $p_0, \ldots, p_{n-1}$ will always be assumed to appear in increasing x-order. We have

$$P = \{(2,35), (5,17), (12,1), (22,13), (28,23), (34,3), (52,43), (63,15)\}.$$

3. Stage I: Getting Started

Our starting point is Bentley's *range tree* [B] applied to the set P. This gives us a complete binary tree T on n leaves: (i) the kth leaf from the left is associated with p_{k-1}; (ii) each node v is associated with the list $L(v)$ made of the points appearing at the leaves (at or) below v, sorted in increasing y-order. Figure 2 illustrates this notion. Note that we have conveniently replaced the name of each point by the value of their ordinate. There is no harm in this substitution since by assumption no two ordinates can be the same. Nodes of T have been labelled $\overline{a}, \ldots, \overline{o}$ in symmetric order. In addition, the leaves are also shown with the abscissae of the corresponding points.

Given the horizontal segment AB, with $A = (a, c)$ and $B = (b, c)$, we are now ready to *drag it down.* To begin with, we decompose AB into $O(\log n)$ so-called *canonical* pieces. To do so, we determine the leaf α (resp. β) whose corresponding abscissa immediately precedes a (resp. follows b). Let v be the nearest common ancestor of α and β. Next, we collect all the nodes $v_1, v_2, \ldots, v_p$ that are either right children of nodes on the path from the left son of v to α or left children of nodes on the path from the right son of v to β. Finally, for each $i = 1, \ldots, p$, we compute the entry of $L(v_i)$ whose corresponding ordinate immediately precedes the value c. Of all the points gathered in this process, we keep the one of largest ordinate; its index is precisely $\texttt{hit}(AB)$. If no such point

is to be found, we have $\texttt{hit}(AB) = n$. Turning to our running example, let's try out the method for $a = 3$, $b = 37$, $c = 20$. We find $\alpha = \overline{a}$, $\beta = \overline{m}$, and $v = \overline{h}$. This gives $p = 3$ and

$$\{v_i | 1 \leq i \leq p\} = \{\overline{c}, \overline{f}, \overline{j}\},$$

from which we derive three candidate points: p_1, p_3, and p_5. The winner is p_1, therefore $\texttt{hit}(AB) = 1$.

Although we have not yet said anything about implementation, it is clear that by using binary search whenever appropriate we have here all the elements of a solution in $O(n \log n)$ space and $O(\log^2 n)$ query time. Our next task will be to bring down the storage to $O(n)$.

4. Stage II: Adding Fault-Tolerance

When faced with the task of trimming down a data structure, it is tempting to remove a fraction of the entries, and to do so in a uniform way in order to limit the damage. This will fail miserably here, however, because a single missing entry in T might entail a painful recovery operation. Looking around locally to retrieve the missing data simply won't do. This is all the more frustrating as the data structure is so evidently redundant. The remedy is to make it *fault-tolerant.* This term is not used here in the traditional sense. We say that a data structure is fault-tolerant if most of its elementary components can be reconstructed from their neighbors. For example, given two arrays of integers $A[1..n]$ and $B[1..n]$, where $B[1] = A[1]$ and $B[i] = B[i-1] + A[i]$ for $i = 2, \ldots, n$, the array B is clearly fault-tolerant; note, however, that a random permutation of its entries will likely not be so.

For our purposes here, we use a compaction scheme used previously in (Chazelle [C]). The basic idea is to replace the various occurrences of a given piece of data in T by single-bit *traces.* For example, making the convention once and for all that 0 means *left* and 1 means *right*, we replace 17 by 0 in $\overline{d}$ because 17 also appears in the left child of the node (observe that it cannot appear in both children); similarly 17 in $\overline{b}$ becomes 1 because it is found in $\overline{c}$. The transformed data structure in our running example appears in Fig.3 (ignore the numbers on the sides for the time being). Obviously, no trace need be stored at the leaves.

The ordinate of the query AB being 20, its *location* in $L(\overline{h})$ is 4: this means that $L(\overline{h})$ has precisely (4+1) entries less than or equal to 20. For example, the locations of -5, 35, 90 are respectively -1, 6, and 7. One key observation is that to find the location of 20 in $\overline{d}$ we simply have to count the number of 0s among the first 5 traces in the list at $\overline{h}$ and subtract 1 (the answer is 2); similarly we locate 20 in $\overline{l}$ by counting the number of 1s in the same sublist and subtracting 1, which gives 1.

Each operation can be carried out in constant time if we supplement traces with cumulative sums of 1s, as shown in Fig.3. Note that cumulative sums of 0s are not needed because they can be readily obtained by subtracting the sum of 1s from the current location in the list.

We will now try to convince the reader that the new data structure supports segment-dragging just as well as the range tree discussed in the previous section. Of course, we still need two separate lists giving the points of P in x-order and y-order. With the latter we locate the ordinate of the query ($c = 20$) in the root $\overline{h}$ (location 4). Next, we jump to $\overline{d}$ using cumulative sums (location 2) and to $\overline{f}$ (location 1). From $\overline{d}$ we also jump to $\overline{b}$ (location 0) and then on to $\overline{c}$ (location 0). Finally from $\overline{h}$ we jump to $\overline{l}$ (location 1) and to $\overline{j}$ (location 0). To summarize, we now have three candidate points

$$\overline{c} : \text{location } 0$$

$$\overline{f} : \text{location } 1$$

$$\overline{j} : \text{location } 0,$$

which we must run off against each other to determine the winner. Observe that so far the running time is within $O(\log n)$.

The problem which we are now facing is to *identify* a trace, that is, reconnect it to its corresponding point (note that traces have yet to be used in our discussion). Referring to a trace as a pair (node, location), let's show how to identify, say, $(\overline{h}, 5)$. The sixth trace of $\overline{h}$ being 1 we must pursue our search with the right child of $\overline{h}$, which is $\overline{l}$. We transfer locations as usual by determining the number of 1s among the first 6 traces of $\overline{h}$ and subtracting 1, which yields $(\overline{l}, 2)$. The new trace

being now 0, we proceed to the left and to do so we subtract from the current location the number of 1s among the first 3 traces at $\overline{l}$, which gives the pair $(\overline{j}, 1)$. Finally, a similar operation leads to $(\overline{i}, 0)$. The identification is complete after $O(\log n)$ steps.

Returning to our running example, we identify $(\overline{c}, 0)$, $(\overline{f}, 1)$, and $(\overline{j}, 0)$ to find the names of the candidates, p_1, p_3, and p_5. The first of them, p_1, is immediately determined as the winner. The new data structure occupies $O(n \log n)$ space and supports segment-dragging in $O(\log^2 n)$ time per query. In the next section we will see that this transformation was not as futile as it may seem right now. Despite appearances the new data structure is, for our purposes, much better than the previous one.

5. Stage III: Removing Redundancy

What we have gained by transforming T as we did is *fault-tolerance*. Indeed, a missing cumulative sum can be easily recovered by local examination. A sum can be recovered from another one, d away, by looking at d traces. This suggests trimming down the structure by keeping only one cumulative sum out of $w = \log n$ (why this choice of w will soon be obvious). In our running example, $w = 3$, so at node $\overline{h}$ we only keep locations 0, 3, and 6. Since a computer word has room for w bits (actually more, by our convention that it can assume values between 0 and $n \log n$), we can take the first w traces of $\overline{h}$ and store them as one word. This yields $2 = 010_2$, with the rest of the sequence giving $5 = 101_2$ and $1 = 01_2$. Figure 4 illustrates the reduced data structure. Before, we had nw pairs in the structure; now we only have n of them. Actually, a little more because each node may generate an extra pair if its *size* (a term from now on referring to the number of traces at the node) is not a multiple of w. At any rate, we now have a data structure of size $O(n)$ (with a rather small constant of proportionality).

Of course, the compaction makes query-answering a little more difficult. To simulate random access to a given trace, we divide its location by w to find the word which contains it. Then we truncate the word accordingly. Recall that the kth least significant bit of a number M is given by the expression $\lfloor M/2^{k-1} \rfloor - 2\lfloor M/2^k \rfloor$. The floor of an integer division is in the usual repertoire of

a *RAM*, but exponentiation is not (the same is true in *PASCAL*). This is not a problem, however: because $k \leq w$, we have $2^k \leq n$, therefore we can comfortably store the first w powers of two in an array of size w. This takes care of traces. To retrieve a cumulative sum, given its location, we first approximate it by reading the nearest entry stored. Then we add the missing 1s by looking directly into the word M that contains the current trace. After appropriate truncation, we finally read the desired value in an array **one**$[0..n-1]$, where **one**$[i] = \#1$s in the binary representation of i. All these operations take constant time. To illustrate the discussion, let's determine the fifth trace of $\overline{h}$ and let's also jump from $(\overline{h}, 4)$ to $\overline{l}$. The desired trace appears in the kth (trace) word of $\overline{h}$ as the lth least significant bit, where $k = 1 + \lfloor \frac{4}{3} \rfloor$ and $l = 3k - 4$: its value is 0. To jump to $\overline{l}$ we compute the number of 1s among the first 5 traces of $\overline{h}$. As indicated above we take the kth (trace) word, 101_2, and bring the current trace (the lth least significant bit) to the lowest order position. To do so, we divide 101_2 by 2^{l-1}, which gives 10_2. Then we look up **one**$[10_2] = 1$, to which we add the kth cumulative sum explicitly stored (i.e.,1). Subtracting 1 from the result, we finally obtain the desired location in $\overline{l}$, that is, 1. Note that caution must be used if the current trace lies in the last (trace) word and the size of the node is not a multiple of w. As it turns out we will circumvent this difficulty altogether in the actual implementation.

To summarize, we have described a data structure of linear size which supports segment-dragging in $O(\log^2 n)$ query time. To cut down the query time we must avoid the identification of every candidate point. The idea is to have the candidates race against each other as soon as they are discovered rather than wait until they are all available. The crux is that although two traces from distinct nodes are incomparable without identification, their ordinates can be immediately compared if the traces belong to the same node.

6. Stage IV: Avoiding Repeated Identifications

For convenience, we pursue our discussion with reference to the data structure before the trimming phase. Note that for our purposes the two structures are functionally equivalent. To illustrate what follows we return to our running example. Recall that the nodes of interest (where the traces of

the candidate points are to be found) are the right (resp. left) children of the nodes of the path from $\overline{d}$ (resp. $\overline{l}$) to $\overline{a}$ (resp. $\overline{m}$). The two paths play symmetrical roles, so let's restrict ourselves to the one from $\overline{d}$ to $\overline{a}$. According to the method described earlier, we locate the candidates by their traces in $\overline{f}$ and $\overline{c}$, giving locations 1 and 0, respectively. Not much can be decided on the basis of these traces alone because they belong to distinct nodes. If they were in the same list, however, things would be different. For example, the candidates have traces at $\overline{h}$ (locations 2 and 4, respectively) from which, on the basis of location alone, the second candidate appears as the winner. This suggests a two-pass approach — one, top-down, to collect traces of candidates and another, bottom-up, to run them off against each other. A more elegant method consists of keeping a current winner while going down and then avoiding the second pass altogether.

At node $\overline{d}$ the query ordinate (20) falls at location 2 (i.e., there are 3 points at $\overline{d}$ with a lower ordinate). Rather than jumping to node $\overline{f}$ we can simulate this process locally while staying at node $\overline{d}$. Indeed, it suffices to locate the first trace equal to 1 at or below location 2. We find location 1, which becomes our current candidate. At this point we have two locations to keep track of: the query location (value 2) and the candidate location (value 1). If these locations coincided then we could stop (why?). Since they do not, we can only infer that the winner has a trace between these two locations. In our example, the only way for the candidate to lose out later on would be for location 2 to hold the future winner (which, in fact, *is* the case). Next we jump to node $\overline{b}$ from both query and candidate locations. This takes us to locations 0 and -1, respectively. Since the trace at location 0 is precisely 1, the current candidate is updated to that location. For two reasons at least we can now stop: the winner has been found and also we are reaching the bottom of the tree. We can now identify the candidate from its trace. (Dealing with the path from $\overline{h}$ to $\overline{m}$ in the same manner and keeping the higher point gives the final answer.)

Note that at each update of the candidate location it is important to record the new value separately, otherwise it will be lost at the next jump. This shows that in general the candidate location does *not* hold the trace of the current candidate. For example, at node $\overline{b}$ the candidate location is -1 before the update, which corresponds to no point of P whatsoever, whereas the

candidate is p_3 (discovered at node $\overline{d}$).

As a matter of fact the candidate location will often hold the trace of a point with lower ordinate than the current candidate. For this reason it is imperative to update the candidate *only* when the candidate location effectively moves up. If there is a tie we apply the rule of seniority: the older candidate wins. Of course, updates of this kind must take place only at the proper *turns*: left turns from $\overline{h}$ to $\overline{a}$ and right turns from $\overline{h}$ to $\overline{m}$.

To complete our discussion we must implement a primitive which locates the next trace equal to 1 at or below a given location. One simple solution consists of providing each trace with the number of consecutive 0s right below. Naturally, by symmetry, we must also keep the number of consecutive 1s. To mix these two structures into one, we simply keep the length of the run ending at each trace. For example, at node $\overline{d}$ we store the list 1,2,1,2, and at node $\overline{h}$ we keep 1,1,1,1,1,1,1,1. All these lists require $O(n \log n)$ storage, but fortunately they are highly *fault-tolerant*. Thus we can sparse them out by discarding all but a fraction $1/w$. This gives 1,2 in the case of $\overline{d}$ and 1,1,1 in the case of $\overline{h}$. To recover the missing information we just need to keep an array `trail`$[0..n-1]$, where `trail`$[i]$ indicates the length of the rightmost run in the binary representation of i.

Although our discussion has been kept informal and special cases have been conveniently avoided, we nevertheless have all the ingredients at our disposal for building a working algorithm. The data structure is of size $O(n)$; as we shall see shortly, it can be constructed in time $O(n \log n)$. Answering a query takes $O(\log n)$ time. In the next section we describe the preprocessing in detail. Save for minor differences it will follow precisely the outline given above.

7. The Preprocessing

We use two constants: *maxsize* specifies the maximum size of the arrays (a *PASCAL* requirement) and *maxint* is the largest integer used. For simplicity we declare all arrays to be of the same type.

<<< INSERT 1 >>>

Aside from n and $w = \log n$, we also have the global variables

<<< INSERT 2 >>>

In our running example, we have $n = 8$ and $w = 3$. The arrays X and Y give the coordinates of the points of P sorted in increasing x-order. For convenience, we set a sentinel $Y[n]$ to $-maxint$ (in our example, $Y[n] = 0$ will do); for this reason we have the condition $maxsize \geq n$.

$$X = 2, 5, 12, 22, 28, 34, 52, 63$$

$$Y = 35, 17, 1, 13, 23, 3, 43, 15, 0$$

$$Ysort = 1, 3, 13, 15, 17, 23, 35, 43.$$

We represent T as an implicit data structure, that is, one without pointers. Traces, for example, are all stored in a single array of nw bits. The first n bits are the traces of the parents of the leaves from left to right (level 1). The next n bits are the traces at level 2, etc. The array is broken up into pieces of length w, each fitting into a computer word (note how fortunate it is that the number of traces should be a multiple of w!) To store cumulative sums we do not use a local reference system for each node. Instead, we maintain sums of 1s with respect to the beginning of a single array. The same holds true of the list of runs. (This justifies our assumption that a computer word should be able to represent any integer up to nw; as mentioned earlier, this assumption can be easily relaxed.) What motivates this decision is the fact that the boundaries of the lists of traces do not necessarily coincide with word boundaries in the arrays. We are now ready to define the remaining variables.

(a) `bit`: For $i = 0, \ldots, n-1$, the w least significant bits of `bit`$[i]$ are the traces from position iw to position $(i+1)w - 1$. In our running example, the traces being

$$100110101100010101010101,$$

we have

$$\texttt{bit}[0..n-1] = 4, 6, 5, 4, 2, 5, 2, 5.$$

(b) `one`: For $i = 0, \ldots, n-1$, `one`$[i]$ gives the number of 1s among the first $(i+1)w$ traces, that is, in the binary representation of `bit`$[0..i]$.

$$\texttt{one}[0..n-1] = 1, 3, 5, 6, 7, 9, 10, 12.$$

(c) `tally`: For $i = 0, \ldots, n-1$, `tally`$[i]$ is equal to the number of 1s in the binary representation of i.

$$\texttt{tally}[0..n-1] = 0, 1, 1, 2, 1, 2, 2, 3.$$

(d) `shift`: For $i = 0, \ldots, w$, `shift`$[i] = 2^i$.

$$\texttt{shift}[0..w] = 1, 2, 4, 8.$$

(e) `suffix`: For $i = 0, \ldots, n-1$, `suffix`$[i]$ indicates the length of the longest suffix in the binary representation of `bit`$[0..i]$ that constitutes a run (i.e., with all bits equal).

$$\texttt{suffix}[0..n-1] = 2, 1, 1, 2, 1, 1, 1, 1.$$

(f) `trail`: For $i = 0, \ldots, n-1$, `trail`$[i]$ is defined as `suffix`$[i]$, but now with respect to the binary representation of i over w bits. Unlike the entries of `suffix`, note that no value in `trail` can exceed w.

$$\texttt{trail}[0..n-1] = 3, 1, 1, 2, 2, 1, 1, 3.$$

To complete this section we give some code to build the arrays above. (No attempt has been made to optimize the code). The preprocessing consists of four procedure calls:

<<< INSERT 3 >>>

1. Initialize: Sets the values of n, w, X, and Y; initializes *Ysort* to Y and clears the array `bit` to 0.

2. BitDef: Computes the arrays *Ysort* and `bit`. The procedure is a variant of mergesort. It consists of two nested loops. The outer one takes us from one level in the merge tree to the next one above. The control variable, *level,* indicates how many leaves descend from any node at the current level. The value of *level* grows geometrically from 2 to $2^w = n$. The inner loop performs all the merges at the current level. In the running example, we have four merges at the first level, two at the second level, etc. Right before the merging starts, l is set to the leftmost entry to be merged and r is set to the rightmost ($r = l + level - 1$). Between merges, l increases by the value *level.* Turning now to the merging proper, the variable m indicates the rightmost location of the first list. The three variables l, r, and m are boundary delimiters; on the contrary, i and j are the running variables which control the merging (Fig.5). The array T is used as temporary storage (necessary because the sorting is not in place).

Before merging we copy the values of interest into T (setting T[$r+1$] to *maxint* to be used as a sentinel). The details of the merging are straightforward. Concurrently, we accumulate traces in the array `bit`. The current location in `bit` is called *index.* It is updated by keeping track of a counter *fill*, which oscillates between 0 and w. To add a new trace we first shift the bits of `bit`[*index*] to the left to make room for the newcomer (bit[index]:= 2×bit[index]). Then, if j is incremented, we set its lower bit to 1 (bit[index]:= bit[index]+1).

<<< INSERT 4 >>>

3. TallyDef: Computes the array `tally` by enumerating the bits of each i between 0 and $n-1$.

<<< INSERT 5 >>>

4. TrailDef: Computes the array `trail` in linear time. The idea is to find all the integers t such that `trail`$[t] = k - 1$, for $k = 2, \ldots, w+1$. The case $k = w+1$ is taken care of separately at the beginning. For $k \leq w$, we simply enumerate all possible prefixes of length $w - k$ (variable i), to which we tack on either 1 followed by $k-1$ zeros or 0 followed by $k-1$ ones. Note that the latter number is derived from the former by subtracting 1 from it.

<<< INSERT 6 >>>

5. OneSufDef: Computes the arrays **one** and **suffix**, using the arrays **tally** and **trail**, respectively.

<<< INSERT 7 >>>

6. ShiftDef: Computes the array **shift**.

<<< INSERT 8 >>>

8. Answering a Segment-Dragging Query

We proceed bottom-up, starting off with a few utilities which will be handy later on. The function *weight* computes cumulative sums. Given a positive integer p, it returns the number of 1s among the first p traces in **bit**. Let $b_0, \ldots, b_{wn-1}$ be the traces in **bit**$[0..n-1]$: we have $weight(p) = \sum_{0 \leq i < p} b_i$. In our running example, we have $weight(p) = 3, 4, 9$, for $p = 5, 7, 18$, respectively. The implementation is straightforward: locate the word of **bit** that contains b_p and truncate it accordingly to get the answer via **tally** and possibly **one**.

<<< INSERT 9 >>>

Similarly, given a positive integer p and an integer $b \in \{0, 1\}$, $run(b, p)$ returns the length of the maximal run of the form $b_{p-k}, \ldots, b_{p-1}, b_p$, where $b_i = b$ $(p - k \leq i \leq p)$. In our example, we have $run(1, 12) = 0$ and $run(0, 12) = 3$. Note that we always have $run(0, p) \times run(1, p) = 0$. The implementation is similar to *weight*, only slightly more complicated. The only subtle point is that **trail**$[j]$ may fail to give the desired answer when b and j are both 0. Indeed, in that case, **trail**$[j] = w$, so we lose the fact that j is a truncated word. Of course, this problem does not have a symmetric instantiation for 1 (why?)

<<< INSERT 10 >>>

Next, we attack the problem of jumping from the trace list of a node to the trace list of one of its children. Before going any further we must address an important question: how do we specify a node? It is tempting to say that this is not really an issue because we can directly derive the

boundaries of a trace list in `bit` given any location in it. This is true, and will actually be used for identification purposes. However, it is inadequate for jumping. Why? Hint: think of what happens if the location is the highest in the trace list of a node distinct from the root. To deal with this problem we characterize a node by specifying explicitly the location of its first trace. For example, the root is assigned the value $(w-1)n$, the leftmost leaf is given the value $-n$, and its parent is assigned the value 0. Note that the left and right children of node $k \geq 0$ are assigned the values $k-n$ and $k-n+c/2$, respectively, where c denotes the size of node k. We easily derive that $c = 2^{\lfloor k/n \rfloor + 1}$.

We are now ready to implement the function *trans* which, given a branching direction — *branch*= 0 (left) or 1 (right) —, a *node*, and a location p, jumps to the child (left or right, depending on *branch*. Recall that, by definition, a number y has location p at *node* if, within the trace list of *node*, b_p corresponds to the point of maximum ordinate less than or equal to y.

<<< INSERT 11 >>>

Given a location p in a trace list the function $id(p)$ returns the index of the point corresponding to the trace at location p. To execute this function, first we determine the value of the trace variable (b), and then we call the function *trans* as often as needed. As soon as the current location becomes negative, we add n and we are done.

<<< INSERT 12 >>>

Next on our list, we have a function *path* which, given an array $A[0..n-1]$ of integers sorted in increasing order and an integer q, returns the largest index i such that $q \geq A[i]$. If no such index exists, $path(A, q) = 0$. The algorithm uses a simple binary search. (The entry $A[n]$, although accessible, is irrelevant).

<<< INSERT 13 >>>

The function *path* will serve two purposes. On the one hand it will allow us to search for the query ordinate in *Ysort*. On the other hand, it will indicate the sequence of turns in T which must be taken in order to reach the nodes α and β and thus decompose the query segment.

Recall that v is the nearest common ancestor of α and β. The contest between candidates proceeds from v to α and from v to β, successively. A flag *dir* is set to 0 in the first case and 1 in the second case. What arguments should be passed to a function *cand* so that it can start off the contest? Aside from *dir* and the initial location p at the root, we also need the sequence of turns from the root to α (if $dir = 0$) or β (if $dir = 1$). This is encoded in the bits of an integer $path < n$. Finally we also need v; actually its size, denoted *break*, will do just as well. The function uses a few local variables: *best* holds the location of the current candidate. This is not always the same as the candidate location, denoted *try* (go back to Section 6 if this subtle difference no longer makes sense...) Initially, *try* is set to the value $(w-1)n-1$ (one slot below the first trace of the root) and *best* is set to -1. This negative value is handy to check if in the end a valid candidate has been found. A variable *size* keeps track of the current node size. It is used to determine when the race should start ($size < break$). When the branching is appropriate we test whether *best* should be updated ($try < tmp = p{-}run(dir, p)$). Note that *tmp* may not always correspond to a valid candidate. At worst, however, *try* is only one slot below the bottom of the current trace list. Therefore, if we have $tmp > try$, then *try* will automatically correspond to a well-defined candidate point. To conclude each loop we jump from p and *try* to the corresponding locations in the appropriate child. This child is derived from *path* and *size*.

<<< INSERT 14 >>>

Finally, the function *hit* can be implemented. Its parameters are respectively the left abscissa, the right abscissa, and the ordinate of the query segment. At the outset, the function *path* is invoked three times to find (i) the initial query location (p), (ii) the left path (*left*), and (iii) the right path (*right*). Next, a *while* loop is used to determine the size (*break*) of the nearest common ancestor v, and candidates are determined via two calls on *cand*. The answer readily follows.

<<< INSERT 15 >>>

9. Conclusions

We have completely described a new algorithm for dragging a segment amidst a set of points in the plane. The data structure requires linear storage and can be used to answer queries in logarithmic time. Experiments have shown that the algorithm performs very well in practice.

The purpose of this paper has been two-fold: to present a new complexity result, and to show that the underlying algorithm can be implemented with little effort. We also hope that the transformational approach taken in explaining the data structure has been helpful to the reader. We close with an interesting open problem: designing a dynamic version of the data structure which is both simple and efficient. At present, we are able to achieve both criteria, but not simultaneously.

REFERENCES

[B] Bentley, J.L. *Multidimensional divide and conquer*, Comm. ACM, 23, 4 (1980), 214–229.

[C] Chazelle, B. *A functional approach to data structures and its use in multidimensional searching*, Brown Univ. Tech. Rep. No. CS–85–16, Sept. 1985. Preliminary version in Proc. 26th Annu. IEEE Symp. on Foundat. of Comput. Sci., Oct. 1985, pp. 165–174.

[EOS] Edelsbrunner, H., Overmars, M.H., Seidel, R. *Some methods of computational geometry applied to computer graphics*, Computer Vision, Graphics, and Image Processing, 28 (1984), pp. 92–108.

[GBT] Gabow, H.N., Bentley, J.L., Tarjan, R.E. *Scaling and related techniques for geometry problems*, Proc. 16th Annu. ACM Symp. on Theory of Comput., pp. 135–143, 1984.

[M] Mitchell, J.S.B. *Private Communication*, 1986.

```
type
   table = array[0..maxsize] of integer;

                    << INSERT 1 >>


var
   n, w: integer;
   X, Y, Ysort, bit, one, tally, shift, suffix, trail : table;

                    << INSERT 2 >>


 Initialize;
 BitDef;
 TallyDef;
 TrailDef;
 OneSufDef;
 ShiftDef;
                    << INSERT 3 >>
```

```
procedure BitDef;
  var
    fill, index, level, i, j, k, l, m, r : integer;
    T:table;
begin
  fill := 0; index := 0; level := 2;
  while level <= n do
    begin
      l := 0;
      while l < n do
        begin
          r := l + level - 1; m := (l + r) div 2;
          for k := l to r do T[k] := Ysort[k];
          T[r + 1] := maxint; i := l; j := m + 1; k := l - 1;
          while (i <= m) or (j <= r) do
            begin
              bit[index] := 2 * bit[index]; k := k + 1;
              if (T[j] < T[i]) or (i > m) then
                begin
                  Ysort[k] := T[j]; j := j + 1;
                  bit[index] := bit[index] + 1
                end
              else
                begin
                  Ysort[k] := T[i]; i := i + 1
                end;
              fill := fill + 1;
              if fill = w then
                begin
                  index := index + 1; fill := 0
                end
            end;
          l := l + level
        end;
      level := 2 * level
    end
end;
```

<< INSERT 4 >>

```
procedure TallyDef;
   var
      i, j, t : integer;
begin
   for i := 0 to n - 1 do
      begin
         j := i; tally[i] := 0;
         while j > 0 do
            begin
               t := j; j := j div 2; tally[i] := tally[i] + t - 2 * j
            end
      end
end;
```

<< INSERT 5 >>

```
procedure TrailDef;
   var
      i, j, k, l, t : integer;
begin
   trail[0] := w; trail[n - 1] := w; l := 2;
   for k := 2 to w do
      begin
         l := 2 * l; j := n div l - 1;
         for i := 0 to j do
            begin
               t := i * l + l div 2; trail[t] := k - 1; trail[t - 1] := k - 1
            end
      end
end;
```

<< INSERT 6 >>

```
procedure OneSufDef;
   var
      i : integer;
begin
   one[0] := tally[bit[0]]; suffix[0] := trail[bit[0]];
   for i := 1 to n - 1 do
      begin
         one[i] := one[i - 1] + tally[bit[i]]; suffix[i] := trail[bit[i]];
         if (suffix[i] = w) and not odd(bit[i] + bit[i - 1]) then
            suffix[i] := suffix[i - 1] + w
      end
end;
```

<< INSERT 7 >>

```
procedure ShiftDef;
   var
      i : integer;
begin
   shift[0] := 1;
   for i := 1 to w do   shift[i] := 2 * shift[i - 1]
end;
```

<< INSERT 8 >>

```
function weight (p : integer) : integer;
   var
      i, j, z : integer;
begin
   i := p div w; j := bit[i] div shift[(i + 1) * w - p]; z := tally[j];
   if i > 0 then   z := z + one[i - 1];
   weight := z
end;
```

<< INSERT 9 >>

```
function run (b, p : integer) : integer;
   var
      i, j, z : integer;
begin
   z := 0; i := p div w; j := bit[i] div shift[(i + 1) * w - p - 1];
   if j - 2 * (j div 2) = b then   z := trail[j];
   if (j = 0) and (b = 0) then   z := p - i * w + 1;
   if (z = p - i * w + 1) and (i > 0) then
      if bit[i - 1] - 2 * (bit[i - 1] div 2) = b then   z := suffix[i - 1] + z;
   run := z
end;
```

<< INSERT 10 >>

```
function trans (branch, node, p : integer) : integer;
   var
      size : integer;
begin
   size := 2 * shift[node div n];
   if branch = 1 then
      trans := node - n + weight(p + 1) - weight(node) - 1 + size div 2
   else
      trans := p - n + weight(node) - weight(p + 1)
end;
```

<< INSERT 11 >>

```
function id (p : integer) : integer;
   var
      b, j, size : integer;
begin
   repeat
      size := 2 * shift[p div n]; b := p div w;
      j := bit[b] div shift[(b + 1) * w - p - 1]; b := j - 2 * (j div 2);
      p := trans(b, (p div size) * size, p)
   until p < 0;
   id := p + n
end;
```

<< INSERT 12 >>

```
function path (A : table; q : integer) : integer;
   var
      k, l, r : integer;
begin
   l := 0; r := n - 1;
   while l < r do
      begin
         k := 1 + (l + r) div 2;
         if q < A[k] then  r := k - 1 else  l := k
      end;
   path := l
end;
```

<< INSERT 13 >>

```
function cand (dir, p, path, break : integer) : integer;
   var
      b, best, node, size, tmp, try : integer;
begin
   best := -1; node := (w - 1) * n; try := node - 1; size := n;
   repeat
      b := (2 * path) div size - 2 * (path div size);
      if (size < break) and (b = dir) then
         begin
            tmp := p - run(dir, p);
            if try < tmp then
               begin
                  try := tmp;
                  best := try
               end
         end;
      p := trans(b, node, p); try := trans(b, node, try);
      size := size div 2; node := node - n + b * size
   until (size = 1) or (try = p);
   if best >= 0 then  cand := id(best)  else  cand := n
end;
```

<< INSERT 14 >>

```
function hit (a, b, c : integer) : integer;
   var
      p, i, j, left, right, break : integer;
begin
   p := path(Ysort, c) + (w - 1) * n; break := n;
   left := path(X, a); right := path(X, b) + 1;
   while (2 * left) div break = (2 * right) div break do
      break := break div 2;
   i := cand(0, p, left, break);
   j := cand(1, p, right, break);
   if Y[i] < Y[j] then  hit := j  else hit := i
end;
```

<< INSERT 15 >>

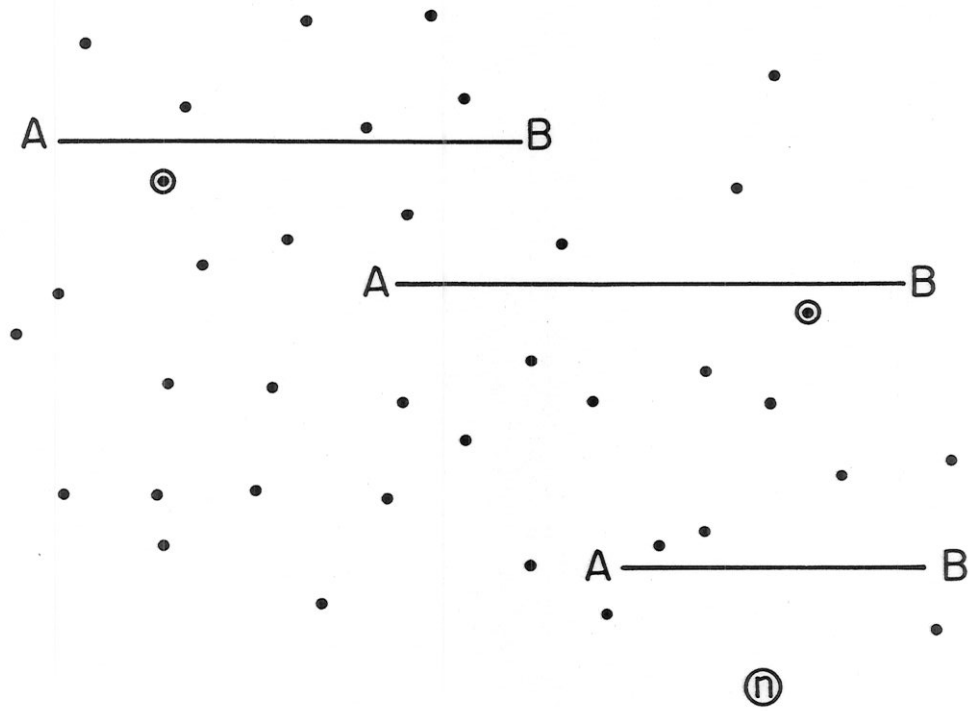

FIGURE 1

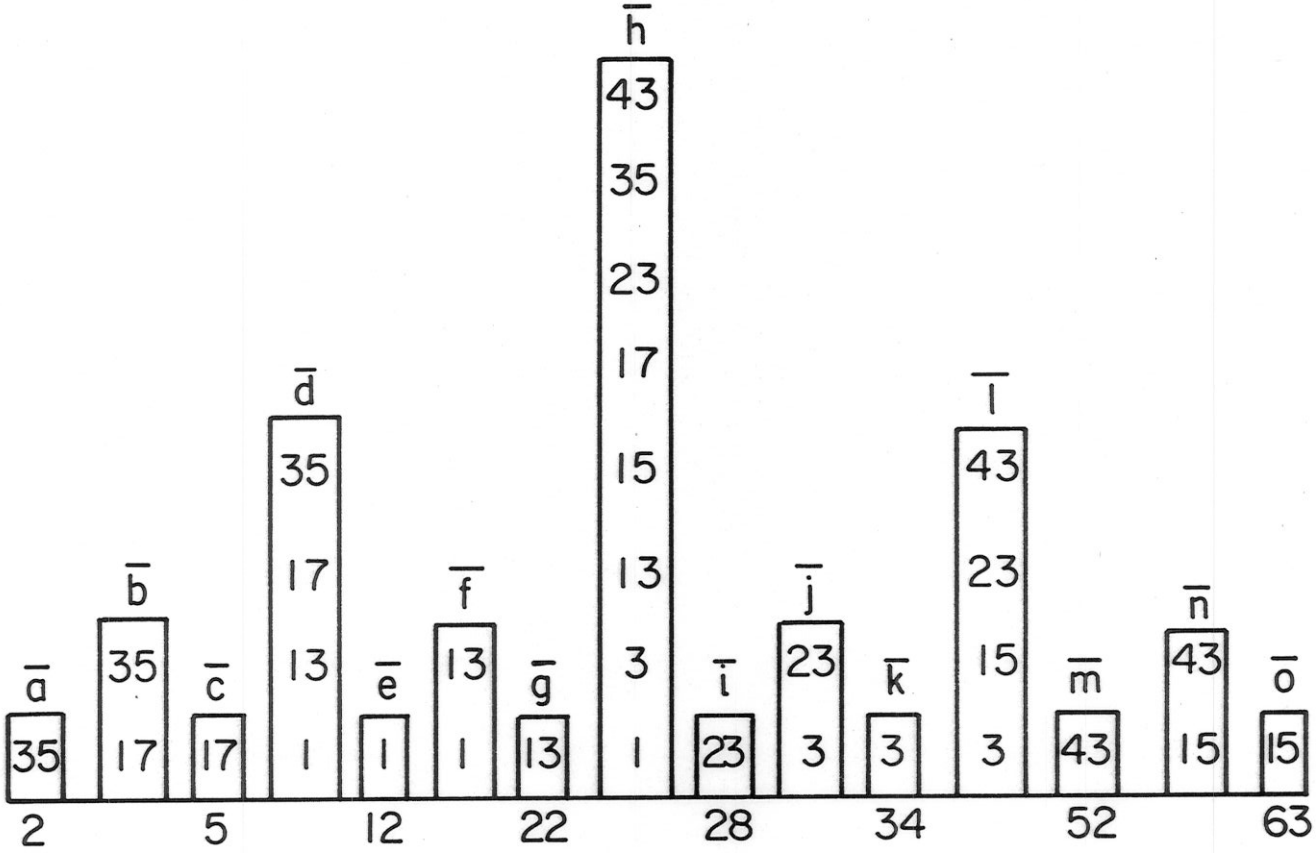
$\bar{a}$
35
$\bar{b}$
35
17
$\bar{c}$
17
$\bar{d}$
35
17
13
1
$\bar{e}$
1
$\bar{f}$
13
1
$\bar{g}$
13
$\bar{h}$
43
35
23
17
15
13
3
1
$\bar{i}$
23
$\bar{j}$
23
3
$\bar{k}$
3
$\bar{l}$
43
23
15
3
$\bar{m}$
43
$\bar{n}$
43
15
$\bar{o}$
15
2
5
12
22
28
34
52
63

FIGURE 2

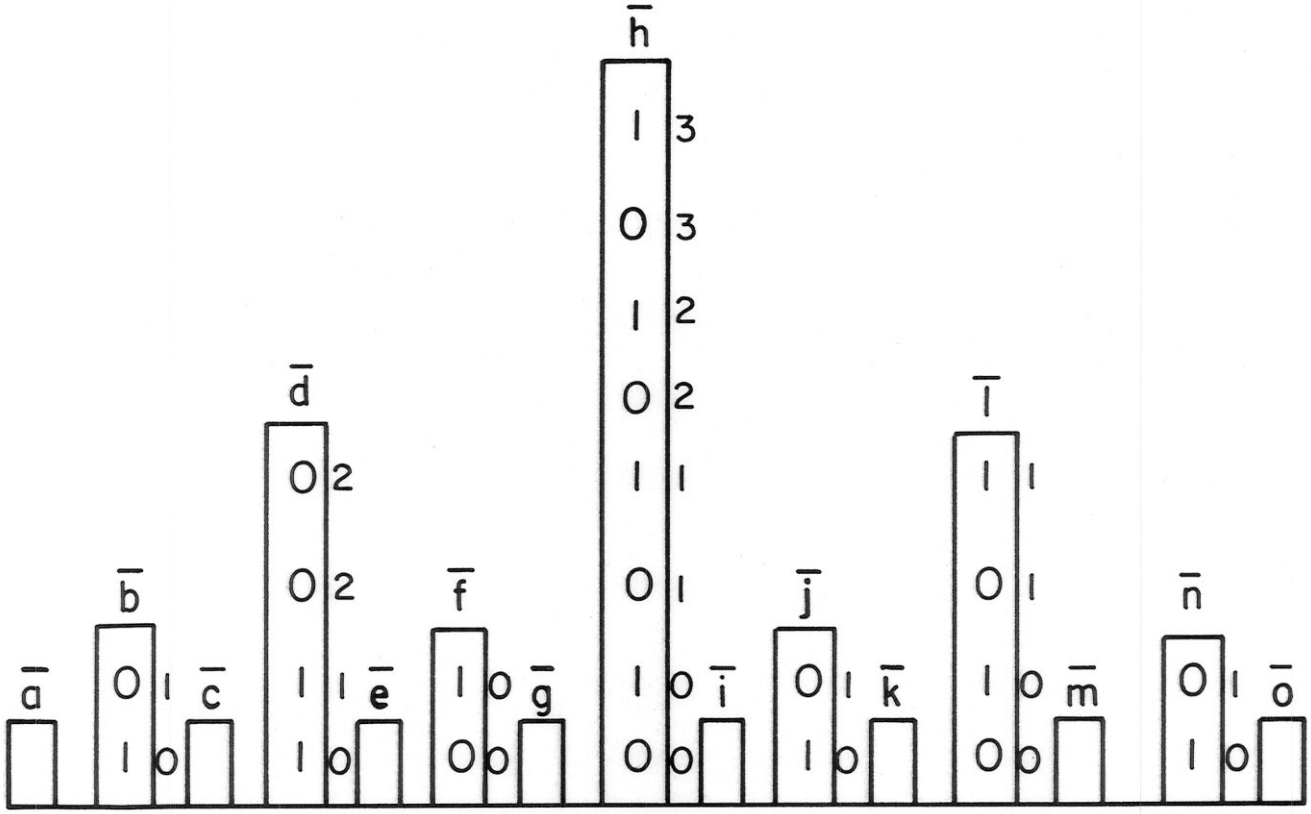
$\bar{a}$
$\bar{b}$
0 1
1 0
$\bar{c}$
$\bar{d}$
0 2
0 2
1 1
1 0
$\bar{e}$
$\bar{f}$
1 0
0 0
$\bar{g}$
$\bar{h}$
1 3
0 3
1 2
0 2
1 1
0 1
1 0
0 0
$\bar{i}$
$\bar{j}$
0 1
1 0
$\bar{k}$
$\bar{l}$
1 1
0 1
1 0
0 0
$\bar{m}$
$\bar{n}$
0 1
1 0
$\bar{o}$

FIGURE 3

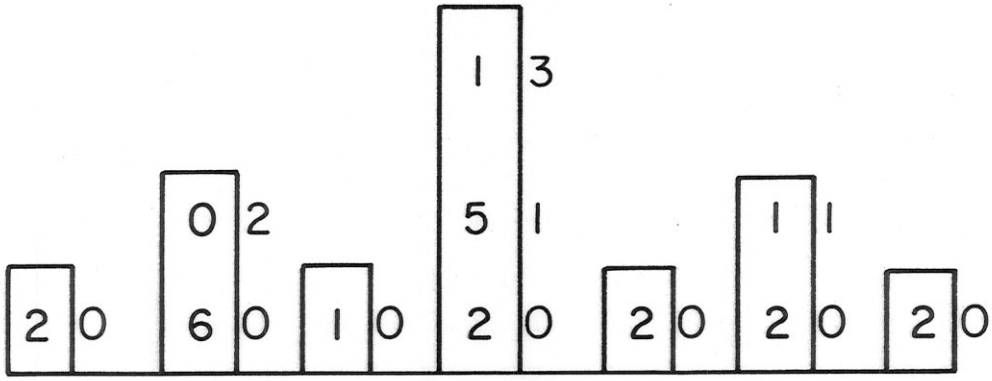

FIGURE 4

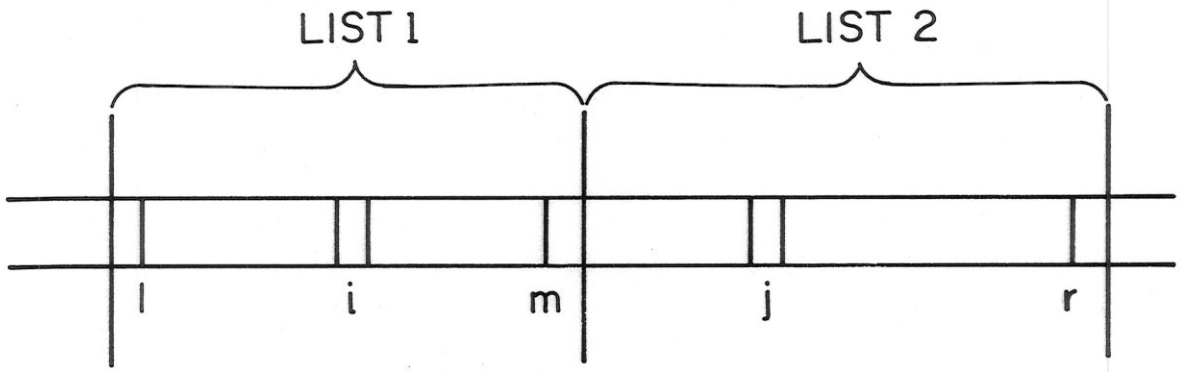

FIGURE 5