NOTICE WARNING CONCERNING COPYRIGHT RESTRICTIONS:

The copyright law of the United States (title 17, U.S. Code) governs the making of photocopies or other reproductions of copyrighted material. Any copying of this document without permission of its author may be prohibited by law.

Transforming Static Data Structures to Dynamic Structures¹

James B. Saxe
Jon Louis Bentley²
Department of Computer Science
Carnegie-Mellon University
Pittsburgh, Pennsylvania 15213

3 September 1979

Abstract

In this paper we will investigate *transformations* that serve as tools in the design of new data structures. Specifically, we study general methods for converting static structures (in which all elements are known before any searches are performed) to dynamic structures (in which insertions of new elements can be mixed with searches). We will exhibit three classes of such transformations, each based on a different counting scheme for representing the integers, and then use a combinatorial model to show the optimality of many of the transformations. Issues such as online data structures and deletion of elements are also examined. To demonstrate the applicability of these tools, we will study six new data structures that have been developed by applying the transformations.

¹This research was supported in part by the Office of Naval Research under Contract N00014-76-C-0370.

²Also with the Department of Mathematics.

Table of Contents

1. Introduction	1
2. Definitions and Notation	1
3. The Transformations	4
3.1. The Binary Transformation	4
3.2. Transformations with Fast Query Time	14
3.3. Transformations with Fast Insertion Time	22
3.4. Summary of the Transformations	26
4. Lower Bounds on Transformations	28
4.1. The Model of Computation	29
4.2. Computing F and G	31
4.3. Transforming History Diagrams to Trees	33
4.4. Tree Properties and their Relation to Performance	34
4.5. The Behavior of $L_k(n)$	37
4.6. Allowing the Number of Static Structures to Grow	43
4.7. Justification of the Restriction to Arboreal Transforms	47
4.8. Limitations on the Significance of the Lower Bounds	48
5. Online Transformations	51
6. Deletion	55
6.1. A Lower Bound	56
6.2. A Fast Special Case	58
7. Conclusions	61
Acknowledgements	62
References	62
I. A List of Decomposable Searching Problems	64

1. Introduction

The design of efficient data structures for searching problems is an important and difficult problem. In this paper we will investigate a class of *transformations* that aid in the design of such data structures, and illustrate the use of those transformations by describing a number of new structures that have been designed by applying the transformations.

Specifically, we will examine transformations that convert static structures (which are built once-for-all before any queries are asked) into dynamic structures (in which queries can be mixed with insertions, and perhaps deletions). These transformations are applicable to a class of problems we call the *decomposable searching problems*. The static-to-dynamic transformations discussed in this paper are only a few of the known transformations on decomposable searching problems; a complete paper describing all the transformations is currently being prepared (Bentley and Saxe [1979]). The static-to-dynamic transformations should, however, serve to illustrate many of the features of other transformations.

In Section 2 we will examine definitions and notations necessary for discussing the transformations. The transformations are discussed in Section 3, and a proof of their optimality is given in Section 4. Online data structures and deletion are the subjects of Sections 5 and 6, and conclusions are offered in Section 7.

2. Definitions and Notation

In this section we will review a number of basic concepts that have to do with searching problems and give a number of definitions that will be used throughout the paper. The casual reader may therefore skim most of this section; the only part he should read in detail is the definition of the decomposable searching problems.

We will use the term *searching problem* in a fairly restricted sense throughout this paper. Specifically, we refer to maintaining a set F of objects so that *queries* asking the relation of a new object x to set F can be answered quickly. The best

known example of a query is what we call a *Member Query*: "Is x a member of F ?". If F were a set of reals, we might be interested in the *Nearest Neighbor* query of "what is the distance from x to the point in F closest to it?". The general query is that a question containing a variable of type $T1$ is asked of a set of elements of type $T2$, with an answer that is of type $T3$. In a Member query, $T1$ and $T2$ are the same, and $T3$ is boolean. In a Nearest Neighbor query, both $T1$ and $T2$ are real, and $T3$ is a nonnegative real. In the general case, the query Q can be viewed as a function mapping a $T1$ and a set of $T2$'s to a $T3$, or

$$Q: T1 \times 2^{T2} \rightarrow T3.$$

Throughout this paper we will identify a *searching problem* by its query; a *solution* to a searching problem is a data structure that allows the query to be answered quickly.

In this paper we will study data structures for a class of searching problems called the *decomposable searching problems*. A searching problem with query operation Q is decomposable if there exists an efficiently computable binary operator $\square$ satisfying the condition

$$Q(x, A \cup B) = \square[Q(x, A), Q(x, B)].$$

(Note that this definition implies that $\square$ is both associative and commutative.) For example, the member searching problem is decomposable because

$$\text{Member}(x, A \cup B) = \vee[\text{Member}(x, A), \text{Member}(x, B)],$$

and (distance to) nearest neighbor searching is decomposable because

$$\text{NN}(x, A \cup B) = \min[\text{NN}(x, A), \text{NN}(x, B)].$$

We will investigate a number of decomposable searching problems throughout this paper; a list of many of them can be found in Appendix I. All of the transformations that we will see later in this paper are applicable for precisely the decomposable searching problems. They exploit decomposability by partitioning a set into subsets, and answer a query by computing answers on the subsets and then using the $\square$ operator to combine those subanswers to yield a solution to the entire problem. Note that the $\square$ operator is essential in this strategy.

There are two subclasses of the decomposable searching problems that will be of special interest later in the paper. The first subclass consists of those problems whose $\square$ operator has a "zero" (or "sticky") element; that is, there exists some element z such that for any element x ,

$$\square(z, x) = z.$$

For example, false is a zero for $\wedge$, and true is a zero for $\vee$. A second class that will be of interest consists of the problems for which the $\square$ operator has an inverse (for example, if $\square$ is addition, its inverse is subtraction). We will examine both of these subclasses of the general decomposable searching problems in detail later in the paper.

We will make a distinction between two types of data structures for solving searching problems. A *static* structure is built once and then searched many times; insertions and deletions of elements are not allowed. To describe the performance of the static structure A we give three functions of N , the number of elements in the set represented by A :

$P_A(N)$ = the *preprocessing time* required to build A ,

$Q_A(N)$ = the *query time* required to perform a search in A , and

$S_A(N)$ = the *storage* required to represent A .

(Unless explicitly noted otherwise, throughout this paper we will deal only with worst-case cost functions.) A second type of data structure is the *dynamic* structure. This structure is initially empty, and the three operations available on it are for inserting a new element, for deleting a current element, and for performing a search. We analyze the performance of the dynamic structure B by giving the functions

$I_B(N)$ = the *insertion time* for B ,

$D_B(N)$ = the *deletion time* for B ,

$Q_B(N)$ = the *query time* required to perform a search in B , and

$S_B(N)$ = the *storage* required to represent B .

Later in this paper we will want to "mix apples and oranges" and compare the

performance of the static structure A with that of the dynamic structure B. To facilitate such comparisons we define the "insertion" time for the static structure A as

$$I_A(N) = P_A(N) / N,$$

which is the cost of building an N-element structure amortized over the N elements it represents. Likewise we define the cost of "preprocessing" the dynamic structure B to be

$$P_B(N) = \sum_{1 \leq i \leq N} I_B(i).$$

3. The Transformations

In this section we will investigate transformations that convert a static data structure for a decomposable searching problem into a dynamic data structure. We will restrict ourselves to the special case of dynamic structures that support only the operations of *inserting* a new element and *searching* to answer a query; we will return to the issue of deletion in Section 6.

3.1. The Binary Transformation

In this subsection we will examine a static-to-dynamic transformation that is based on the binary representation of the integers. We will study the transformation by first examining its application to the particular problem of nearest neighbor searching in the plane, and then discussing its more general properties.

In nearest neighbor searching we must organize a set of N points in the plane so that subsequent queries can tell the distance from the query point x to its nearest neighbor in the set. Therefore, objects of type T1 and T2 are points in $\mathbb{R}^2$, and those of type T3 are positive reals. (For ease of discussion we consider only the problem of finding the distance to the nearest neighbor and not the point realizing that distance.) Note that nearest neighbor searching is decomposable because it satisfies

$$NN(x, A \cup B) = \min[NN(x, A), NN(x, B)].$$

Lipton and Tarjan [1977] have described an elegant static data structure for nearest neighbor searching (which we will call LT) with performance

$$\begin{aligned} P_{LT}(N) &= O(N \lg N), \\ Q_{LT}(N) &= O(\lg N), \text{ and} \\ S_{LT}(N) &= O(N). \end{aligned}$$

Many applications, however, call for dynamic nearest neighbor searching, and the Lipton-Tarjan structure does not appear to be suitable for a modification that would facilitate insertions. We will now investigate a new structure (called DNN for dynamic nearest neighbor) that uses the Lipton-Tarjan static structure only as a subroutine, rather than trying to modify the structure. The DNN structure that we will describe is the best known structure for performing dynamic nearest neighbor searching in the plane.

The DNN structure will consist of a set of LT's: that is, the elements (points) currently stored in the DNN will be partitioned into subsets that are themselves represented by LT's. When there is one element in the DNN, there is an LT containing that single element. When the second element is inserted, that LT is discarded and a new LT of size two is created. At the arrival of the third element, a new LT of size one is created. This process continues so that when there are N elements represented by the DNN, there are LT's corresponding to all of the one bits in the binary representation of N . For example, when there are 79 elements in the DNN, there are LT's of size 64, 8, 4, 2 and 1. When the 80-th element is inserted, the four smallest structures are discarded and a new structure of size 16 is built. At any time in this process the distance to the nearest neighbor of a query point x can be found by locating its nearest neighbors in each of the LT's (using the $O(\lg N)$ algorithm) and taking the minimum of the distances; it is here that we make essential use of decomposability.

This scheme is illustrated pictorially in Figure 3.1 by a diagram commonly used to represent binary counting. The vertical axis in that figure denotes the number of elements currently in the dynamic structure. Each rectangle (square) represents a

particular static LT structure; for example, note the four by four square that comes into existence at time four and is then replaced at time eight. The LT structures in existence at time T can be found by drawing a horizontal line that intersects the vertical axis at T ; for example, at time seven there are three structures in existence -- of sizes four, two and one. We will find later that this type of diagram (which we call a "history diagram") is a handy way of representing transformations.

Figure 3.1. The binary transform.

It is easy to analyze the performance of the DNN structure given that we know the performance of the LT structure. Since the LT requires linear storage and the DNN just partitions its elements into LT's, the DNN will also require linear storage. A DNN of N elements will keep at most $\lg(N+1)$ LT's (each of size not greater than N), so the query time of a DNN is bounded above by $\lg(N+1)$ times the cost of querying an LT. The cost of inserting an element into a DNN is more difficult to analyze; note that while inserting the 1023-rd element is essentially free, the 1024-th element is very expensive, since a new structure of size 1024 must be built. We will therefore count the cost of inserting the first N elements into an initially empty structure, which is exactly $P_{\text{DNN}}(N)$. We will perform this analysis only for the case that $N = 2^j - 1$, and discuss later the value of the function for other N . If we have inserted $2^j - 1$ elements, then we have built one LT structure of size 2^{j-1} , two LT structures of size 2^{j-2} , and 2^{k-1} structures of size 2^{j-k} . (This is a trivial property of binary counting.) The total cost of inserting these elements is therefore

$$P_{DNN}(2^j-1) = 1 \cdot P_{LT}(2^{j-1}) + 2 \cdot P_{LT}(2^{j-2}) + \dots + 2^{j-1} \cdot P_{LT}(1).$$

For N a power of two we can rewrite this as

$$P_{DNN}(N-1) = 1 \cdot P_{LT}(N/2) + 2 \cdot P_{LT}(N/4) + \dots + (N/2) \cdot P_{LT}(1).$$

We know that $P_{LT}(N) = O(N \lg N)$, which implies that $P_{LT}(N) \leq cN \lg N$, for some positive constant c . Substituting this into the above equation yields

$$\begin{aligned} P_{DNN}(N-1) &\leq c[1 \cdot (N/2 \lg N/2) + 2 \cdot (N/4 \lg N/4) + \dots + (N/2) \cdot (1 \lg 1)] \\ &= (cN/2) \cdot [\lg N/2 + \lg N/4 + \dots + \lg 1] \\ &\leq (c/2)N \lg^2 N \\ &= O(N \lg^2 N). \end{aligned}$$

This completes our analysis of the DNN structure, establishing the following.

New Data Structure 1: (Dynamic Nearest Neighbor)

The DNN structure for dynamic nearest neighbor searching in the plane has performances

$$\begin{aligned} P_{DNN}(N) &\leq P_{LT}(N) \cdot \lg(N+1) = O(N \lg^2 N), \\ Q_{DNN}(N) &\leq Q_{LT}(N) \cdot \lg(N+1) = O(\lg^2 N), \text{ and} \\ S_{DNN}(N) &\leq S_{LT}(N) = O(N). \end{aligned}$$

Note that the cost of doing N pairs of Insert, Query operations in the DNN structure is proportional to $N \lg^2 N$; all other known dynamic nearest neighbor structures require $\Omega(N^2)$ time for the task.

The binary transformation that we have just described for nearest neighbor searching is applicable to any decomposable searching problem: given a static data structure for a particular problem, a dynamic structure is achieved by keeping a set of static structures, each representing a set whose cardinality is a power of two. Insertion is accomplished by the same technique of binary counting. A query can be answered by querying all the static structures in existence at the time of the query, and combining the answers by repeated application of the $\square$ operator.

A computer program implementing the binary transform is sketched in Figure 3.2. It assumes the existence of a static structure S with operations $Query_S$, $Build_S$ and $Unbuild_S$ ($Unbuild_S$ returns the elements currently stored in the structure as a linked

list)³. The code implements a dynamic structure D providing routines Init_D (which initializes the structure to be empty), Insert_D , and Query_D . It implements the binary strategy by maintaining a one-way infinite array P with the invariant that $P[i]$ is either empty or contains a static structure of size 2^i . The variable High is an integer that is one greater than the last nonempty structure; $P[\text{High}]$ is always empty. Init_D initializes the structure to have this invariant. Query_D answers a query by iterating through the structures and combining the answers by the $\square$ operator. Insert_D can be understood most easily by considering incrementing a binary integer by one: to do so, we scan from right to left, changing ones to zeros until we come to the first zero (which we then make a one). An Alphard program very similar to the code in Figure 3.2 has been given by Bentley and Shaw [1979]; they also provide both a precise specification of the transform and a proof that the program accomplishes it.

The analysis of the general transformation is quite similar to the analysis of the DNN structure.⁴ Since at most $\lg(N+1)$ static structures exist for an N -element dynamic structure, if we assume the static query cost is monotone nondecreasing we have

$$Q_D(N) \leq Q_S(N) \cdot \lg(N+1).$$

To analyze the storage and processing costs we need the following definition: a function F is said to *grow at least linearly* if for every two positive integers, M and N , where $M < N$,

$$F(M)/M \leq F(N)/N.$$

A consequence of this definition is that if F is a function that grows at least linearly and A and B are positive integers, then

$$F(A+B) = A[F(A+B)/(A+B)] + B[F(A+B)/(A+B)] \geq F(A) + F(B).$$

³Throughout this paper we will retrieve a set of T2's from a structure by unbuilding the structure. In some applications it might be more efficient to store the set along with the structure.

⁴In the analysis of the transformed structure we will count only the costs incurred by operations on the original structure. Examination of the code in Figure 3.2 shows that the overhead costs for both Insert and Query are a small constant times $\lg N$.

```

proc InitD ←
    P[0] ←  $\phi$ ; High ← 0

proc InsertD(x) ←
    S ← {x}
    i ← 0
    while P[i]  $\neq \phi$  do
        S ← S  $\cup$  UnbuildS(P[i]) ; P[i] ←  $\phi$ 
        i ← i+1
    P[i] ← BuildS(S)
    if i=High then
        High ← High+1; P[High] ←  $\phi$ 

func QueryD(x) ←
    A ← QueryS(x,P[0])
    for i ← 1 to High-1 do
        A ←  $\square$ (A, QueryS(x, P[i]))
    return A

```

Figure 3.2. Sketch of code for the binary transform.

Since the dynamic structure partitions its elements among static structures without replication, if the storage cost S_S of the static structure grows at least linearly we have the relation

$$S_D(N) \leq S_S(N).$$

To analyze the processing cost we will first consider the case that N is a power of two; the reasoning used in our analysis of DNN shows that

$$P_D(N-1) = P_S(N/2) + 2P_S(N/4) + \dots + (N/2)P_S(1).$$

When P_S grows at least linearly, we know that $P_S(2i) \geq 2P_S(i)$ and we can use this fact inductively to show that

$$\begin{aligned} P_D(N-1) &\leq P_S(N/2) + P_S(N/2) + \dots + P_S(N/2) \\ &= P_S(N/2) \cdot \lg N. \end{aligned}$$

We will now use a less accurate (but more general) analytic technique to establish the value of $P_D(N)$ for N not one less than a power of two. Note that after N elements have been inserted, any particular element has been in at most $\lg(N+1)$ distinct static structures. We will now show that for any transform, if every element has been built into at most k structures, then the static and dynamic processing costs are related by

$$P_D(N) \leq P_S(N) \cdot k.$$

(This immediately yields the corollary that

$$P_D(N) \leq P_S(N) \cdot \lg(N+1)$$

for the binary transform, for any positive N .) Consider the cost that any particular element, E , contributes to $P_D(N)$. Each time E is built into a new static structure of size M , we can assign it a share of that cost of $P_S(M)/M$. Because P_S grows at least linearly and M is less than or equal to N , we know that

$$P_S(M)/M \leq P_S(N)/N,$$

and we can therefore assign E this latter cost as an upper bound. Multiplying the number of distinct elements (N) by the number of times each is built into a static structure (less than k) by this cost yields the desired result.

To enable us to speak more precisely about transforms on data structures for decomposable searching problems, we need the following definition.

Definition 3.1: (Admissible transform)

A transformation on decomposable searching problems is said to be an admissible $(F(N), G(N))$ transform if it converts the static structure A to the dynamic structure B *assuming only the property of decomposability*, and the following relations hold between the cost functions:⁵

⁵To simplify the analysis, we will count only the costs of calls to operations on the static structure, and not the costs of bookkeeping operations nor the cost of combining the results of queries into different static structures. Careful examination of our algorithms will show that these extra costs add only a small constant factor (which does not depend on F or G) to the compute times. In most cases, this constant will approach unity as N increases. Similarly, the only storage we charge to the dynamic structure is that used for storing instances of the static structure. Again, this is generally the dominant cost.

$$\begin{aligned} Q_B(N) &\leq Q_A(N) \cdot F(N), \\ P_B(N) &\leq P_A(N) \cdot G(N), \text{ and} \\ S_B(N) &\leq S_A(N). \end{aligned}$$

We assume here that Q_A is monotone nondecreasing and that both P_A and S_A grow at least linearly.⁶

We can now state precisely the fact that the binary transform efficiently converts a static data structure to a dynamic structure as Theorem 3.1.

Theorem 3.1: (The binary transform)

The binary transform is an admissible $(\lg(N+1), \lg(N+1))$ transform.

Proof:

Given in the preceding text. QED.

To illustrate some "tricks" available in using the binary transform, let us consider its application to the member query problem using the data structure of a sorted array. Precisely, consider the static data structure for member searching that stores the elements in increasing order in an array (built by sorting the set), and answers a query by performing a binary search. The analysis of this structure (which we call SA , for sorted array) shows

$$\begin{aligned} P_{SA} &= O(N \lg N), \\ S_{SA} &= O(N), \text{ and} \\ Q_{SA} &= O(\lg N). \end{aligned}$$

Consider the dynamic member searching structure achieved by applying the binary transformation to SA : we always maintain a set of sorted arrays, each of size a power of two. A particularly efficient representation of this structure (which we will

⁶For cases where P_A , Q_A , and S_A do not satisfy these criteria, we may choose functions P'_A , Q'_A , and S'_A that (a) satisfy the criteria and (b) dominate P_A , Q_A , and S_A , respectively. The relations given above will then hold between the dynamic cost functions and P'_A , Q'_A , S'_A .

call BL, for binomial list⁷) is to store these sorted arrays sequentially in one large array, with the largest sorted segment (which we call a *run*) leftmost in the array. Two snapshots of a BL are shown in Figure 3.3; the vertical bars in the figure separate the runs in the array. By the analysis of SA and the effect of the binary transform, we can easily describe the complexity of the BL structure as follows

$$\begin{aligned} P_{BL} &= O(N \lg^2 N), \\ S_{BL} &= O(N), \text{ and} \\ Q_{BL} &= O(\lg^2 N). \end{aligned}$$

Note that very little storage is used by a BL: it requires only N array words for the elements, plus $\lg N$ bits to describe the the cardinality of the represented set.

12	19	23	27	38	41	43	47	27	43	29			
----	----	----	----	----	----	----	----	----	----	----	--	--	--

a.) An 11-element binomial list.

12	19	23	27	38	41	43	47	27	29	36	43			
----	----	----	----	----	----	----	----	----	----	----	----	--	--	--

b.) After inserting 36.

Figure 3.3. Snapshots of a binomial list.

There is a glaring deficiency in the obvious implementation of this structure: the obvious insertion routine inserts the 1024-th element by ignoring all the structure currently in the array and re-sorting from scratch. A far superior strategy for any insertion is to consider the inserted element as a one-element run, and merge that with the rightmost one-element run giving a two-element run. We then merge that with its neighbor, giving a four-element run, and so forth. The amount of work in building a new run in this scheme is linear in the size of the run, and the cost of inserting N elements is therefore $O(N \lg N)$. We have thus avoided paying the logarithmic penalty factor inherent in the binary transform by observing that runs

⁷ This structure was invented for this application by the use of the binary transform, and was then studied in detail by Bentley, Detig, Guibas and Saxe [1979]. The name is taken from its similarity to the binomial queue data structure of Vuillemin [1978].

can be efficiently merged.⁸

We can sometimes avoid paying the transform penalty of a logarithmic slowdown in query time. Specifically, we will consider the average cost of performing a *successful* member search in a BL (that is, a search that finds the element it was looking for). If we assume that each element in the array is equally likely to be searched for, then the probability of finding the desired element in the first run is at least one-half. Therefore, half the time we need never search the other runs. Likewise, at least one-half of the remaining times we find the desired element in the next structure, so the probability of searching the third run is less than one-fourth. Summing the cost of searching each run times the probability of performing the search, we find that a successful member search is expected to be at most twice as expensive in the BL as in the SA.

The arguments that we have just sketched have been given in detail by Bentley, Detig, Guibas and Saxe [1979], who describe the following data structure.

New Data Structure 2: (Binomial Lists)

The binomial list (BL) structure for dynamic member searching has performances

$$\begin{aligned}P_{BL}(N) &= O(N \lg N), \\Q_{BL}(N) &= O(\lg^2 N), \text{ and} \\S_{BL}(N) &= O(N).\end{aligned}$$

The linear storage used by this structure consists of exactly N array words and $O(\lg N)$ additional bits, which is minimal.

Bentley, Detig, Guibas and Saxe [1979] have investigated this structure in detail and have shown that it is optimal in a certain model of *minimum-storage* dynamic member searching. The BL structure provides an interesting point of comparison with the minimum-storage structure described by Munro and Suwanda [1979]; this

⁸Only constant extra space is required to merge consecutive runs in an array -- see Knuth [1973, Exercise 5.2.4.18].

structure performs substantially better than theirs by working in a different model of computation.

There is yet another circumstance in which the logarithmic cost penalties of applying the binary transform do not have to be paid: when the original cost functions are fast growing. Consider, for example, a static data structure with N^2 preprocessing time. Our previous analysis shows that for N a power of two, we will have

$$\begin{aligned} P_D(N-1) &= P_S(N/2) + 2P_S(N/4) + \dots + (N/2)P(1) \\ &= (N/2)^2 + 2(N/4)^2 + \dots + (N/2)1^2 \\ &= (N^2/2) \cdot [1/2 + 1/4 + \dots + 1/N] \\ &= O(N^2). \end{aligned}$$

Similar analyses show that the logarithmic penalty in processing cost is not incurred when the binary transform is applied to any static structure with preprocessing cost of $\Omega(N^{1+\epsilon})$, for any positive ϵ . Likewise, it can be shown that the logarithmic penalty in query time will not have to be paid for any static structure with query time of at least $\Omega(N^\epsilon)$.

This concludes our study of the binary transform. In the next two subsections we will see that this transform is but one of many possible ways of converting a static structure to a dynamic structure, at the cost of penalty factors in the preprocessing and query costs. As we study the other transforms and their performance, it is important to keep in mind that the penalty factors need not always be paid. In this subsection we have seen three ways of avoiding them: by merging structures instead of rebuilding them from scratch, by counting the average search time instead of the worst-case time (this is appropriate whenever the $\square$ operator has a zero element), and by performing separate analyses for fast-growing functions.

3.2. Transformations with Fast Query Time

The binary transform of the last subsection provides us with an example of an admissible $(\lg(N+1), \lg(N+1))$ transform, and we might wonder if we can do better. In this subsection we will investigate a class of transforms that have faster query

times than the binary transform at the cost of slower insertion time. Specifically, we will see that an admissible $(k, (k!N)^{1/k})$ transform exists for any positive integer k . We will study this transform by first investigating the case $k=2$, and then move on to the general case.

We will call the transform for the case $k=2$ the *triangular transform*, because it is based on the triangular numbers (that is, numbers of the form $\binom{n}{2}$). The transform is illustrated in Figure 3.4. Note that when 5 elements are in the dynamic structure, there are static structures of size 3 and 2; when the 6-th element is inserted, those structures are destroyed and a new structure of size 6 is created. At any point in the history of the dynamic structure, there will be at most two static structures in existence. The insertion algorithm creates a new "large" static structure at every triangular number; otherwise it inserts an element by unbuilding the smaller structure and building it into a new structure with one additional element. A query can be answered by searching the two static structures and combining the answers by the $\square$ operator.

Figure 3.4. The triangular transform.

The triangular structure is very easy to analyze. Because at most two static structures exist at any time, the dynamic query cost is given by

$$Q_D(N) \leq 2Q_S(N).$$

If we assume that the static storage requirements grow at least linearly, we know that the dynamic structure does not use more storage. To analyze the insertion time, consider the case that a total of $\binom{M}{2}$ elements have been inserted. It is easy to prove by induction that no element has been built into more than M structures (the proof is based on the recurrence for the triangular numbers). In general, if N elements have been inserted, no single element has been built into more than $(2N)^{1/2}$ static structures. By the arguments in the previous subsection, this implies

$$P_D(N) \leq P_S(N) \cdot (2N)^{1/2}.$$

These arguments together establish the following theorem.

Theorem 3.2: (The triangular transform)

The triangular transform is an admissible $(2, (2N)^{1/2})$ transform.

Proof:

Given in the preceding text. QED.

Just as the binary transform is isomorphic to the binary representation of the integers, so is the triangular transform isomorphic to a representation of the integers based on triangular numbers. (This system is called the "binomial number system" by Knuth [1968, Exercise 1.2.6.56].) Specifically, an integer N is represented by a pair of integers i and j (with $i > j$) by the expression

$$N = \binom{i}{2} + \binom{j}{1}.$$

Note that both i and j are less than twice the square root of N ; this explains the processing cost of the transform. The general transform, which we will call the *k-binomial transform*, is based on a straightforward generalization of this scheme, in which an integer is (uniquely) represented as the sum of k binomial coefficients, whose lower parts are the integers 1 through k . This counting scheme is illustrated for the cases $k=2$ and $k=3$ in Figure 3.5. Row 15 of the table is interpreted as follows: in the 2-binomial representation, 15 is the sum of 15 and 0, or $\binom{6}{2}$ and $\binom{0}{1}$. In the 3-binomial representation, 15 is the sum of 10, 3 and 2, or $\binom{5}{3}$, $\binom{3}{2}$ and $\binom{2}{1}$.

Integer	(₂)	(₁)	Integer	(₃)	(₂)	(₁)
0 = 0+0	1	0	0 = 0+0+0	2	1	0
1 = 1+0	2	0	1 = 1+0+0	3	1	0
2 = 1+1	2	1	2 = 1+1+0	3	2	0
3 = 3+0	3	0	3 = 1+1+1	3	2	1
4 = 3+1	3	1	4 = 4+0+0	4	1	0
5 = 3+2	3	2	5 = 4+1+0	4	2	0
6 = 6+0	4	0	6 = 4+1+1	4	2	1
7 = 6+1	4	1	7 = 4+3+0	4	3	0
8 = 6+2	4	2	8 = 4+3+1	4	3	1
9 = 6+3	4	3	9 = 4+3+2	4	3	2
10 = 10+0	5	0	10 = 10+0+0	5	1	0
11 = 10+1	5	1	11 = 10+1+0	5	2	0
12 = 10+2	5	2	12 = 10+1+1	5	2	1
13 = 10+3	5	3	13 = 10+3+0	5	3	0
14 = 10+4	5	4	14 = 10+3+1	5	3	1
15 = 15+0	6	0	15 = 10+3+2	5	3	2
16 = 15+1	6	1	16 = 10+6+0	5	4	0
17 = 15+2	6	2	17 = 10+6+1	5	4	1
18 = 15+3	6	3	18 = 10+6+2	5	4	2
19 = 15+4	6	4	19 = 10+6+3	5	4	3
20 = 15+5	6	5	20 = 20+0+0	6	1	0
21 = 21+0	7	0	21 = 20+1+0	6	2	0
22 = 21+1	7	1	22 = 20+1+1	6	2	1

Figure 3.5. 2-binomial and 3-binomial counting.

With the example of Figure 3.5 as background, we can now describe k -binomial counting more precisely. We will use an array $D[1..k]$ to store the upper parts of the binomial coefficients. The invariant of this counting scheme has two parts: first, the represented integer is given by

$$N = \binom{D[k]}{k} + \binom{D[k-1]}{k-1} + \dots + \binom{D[1]}{1},$$

and secondly, each coefficient $D[i]$ satisfies the condition

$$D[i] > D[i-1]$$

for $2 \leq i \leq k$. We can initialize the array to represent zero by assigning each $D[i]$ to have the value $i-1$; we will also find it handy to assume that the value of $D[k+1]$ is "infinity". The code for incrementing an integer by one is then as follows.

```

D[1] ← D[1] + 1
i ← 1
while D[i] = D[i+1] do
  D[i+1] ← D[i+1] + 1
  D[i] ← i-1
  i ← i+1

```

It is easy to prove by induction that this code correctly implements the above counting scheme.

It is straightforward to modify the above counting scheme to yield an admissible transform. To do so we will retain the array D (with the same invariant as above), and add an array P[1..k] of static structures. The number of elements in P[i] is always $\binom{D[i]}{i}$. The code for this k-binomial transform is given in Figure 3.6, and Figure 3.7 illustrates the 3-binomial transform.

Figure 3.7 The 3-binomial transform.

The correctness of the code can be proven by induction, and its analysis establishes the following theorem.

Theorem 3.3: (The k-binomial transform)

The k-binomial transform is an admissible $(k, (k!N)^{1/k})$ transform.

```

proc InitD ←
  for i ← 1 to k do
    D[i] ← i-1; P[i] ← ϕ
  D[k+1] ← ∞

proc InsertD(x) ←
  D[1] ← D[1]+1; S ← UnbuildS(P[1]) ∪ {x}; P[1] ← ϕ
  i ← 1
  while D[i] = D[i+1] do
    D[i+1] ← D[i+1] + 1; S ← S ∪ UnbuildS(P[i])
    D[i] ← i-1; P[i] ← ϕ
    i ← i+1
  P[i] ← BuildS(S)

func QueryD(x) ←
  A ← QueryS(x, P[1])
  for i ← 2 to k do
    A ← □(A, QueryS(x, P[i]))
  return A

```

Figure 3.6. Code for the k-binomial transform.

Proof:

Since at most k structures exist at any one time, we have

$$Q_D(N) \leq Q_S(N) \cdot k.$$

Since the space requirement for the static structure grows at least linearly with the number of elements, the dynamic structure can be no more expensive.

To bound the processing time of the dynamic structure, we will investigate the maximum number of structures into which any element may be built during the first N insertions. Note that after N insertions, we have

$$\begin{aligned}
 N &\geq \binom{D[k]}{k} \\
 &\geq (D[k]-k+1)^k / k!.
 \end{aligned}$$

implying

$$D[k] \leq (k!N)^{1/k} + k - 1.$$

This, together with the invariant that

$$D[k] > D[k-1] > \dots > D[1] \geq 1$$

implies that each $D[i]$ satisfies

$$0 \leq D[i] - i \leq (k!N)^{1/k} - 1$$

for $1 \leq i \leq k$. Finally, we note that whenever a structure is discarded and its elements rebuilt into a new structure, the difference between the upper and lower parts of the binomial coefficient giving the size of the structure increases by one; that is, a structure of size

$$\binom{m}{i}$$

is always replaced by a structure of size

$$\binom{m+1}{i}$$

or of size

$$\binom{m+2}{i+1}.$$

This implies that no element is ever built into more than $(k!N)^{1/k}$ static structures, from which it follows that

$$P_D(N) \leq P_S(N) \cdot (k!N)^{1/k}.$$

QED.

Note that for all positive k , $k!^{1/k} < k$. For large k , Stirling's approximation gives⁹ $k!^{1/k} \sim k/e$.

To illustrate the application of the binomial transforms, we will consider the problem of range searching. In this problem, the stored set contains points in a d -dimensional space, and a query asks for all points with each dimension in a specified range. (Note that this problem is decomposable with the $\square$ operator

⁹We use the notation, " $A \sim B$ " as a shorthand for " $|A-B| = o(B)$ ".

interpreted as U.¹⁰) Bentley and Maurer [1978] give a structure for static range searching (SRS) with performances

$$\begin{aligned} Q_{SRS}(N) &= O(\lg N), \\ P_{SRS}(N) &= O(N^{1+\delta}), \text{ and} \\ S_{SRS}(N) &= O(N^{1+\delta}) \end{aligned}$$

for any fixed $\delta > 0$. By choosing, for example, $k = \lceil 2/\epsilon \rceil$ and $\delta = \epsilon/2$, we can apply the k -binomial transform to achieve the following structure.

New Data Structure 3: (Dynamic Range Searching)

A dynamic range searching (DRS) structure supporting insertions and queries for point sets in d -space with performance

$$\begin{aligned} Q_{DRS}(N) &= O(\lg N), \\ P_{DRS}(N) &= O(N^{1+\epsilon}), \text{ and} \\ S_{DRS}(N) &= O(N^{1+\epsilon}) \end{aligned}$$

can be achieved for any fixed $\epsilon > 0$ and positive integer d .

Such a structure is useful for range searching in a situation in which the number of queries is known to exceed greatly the number of insertions. Specifically, if the number of insertions in a set of N insertions and queries were known to be $O(N^p)$ for some $p < 1$, then this structure would allow the operations to be processed in $\theta(N \lg N)$ time. The best performance for this task prior to this structure was achieved by Lueker [1978]; his structure required $\theta(N \lg^d N)$ time.

It is important to observe that the penalties incurred by the k -binomial transform need not always be paid. Just as in the binomial transform, they can occasionally be avoided by merging static structures, by counting the expected query cost, or by performing separate analyses for fast-growing functions.

¹⁰ In order to implement (multiset) union as a constant-time operation, we ask that a query return a tree whose leaves are the points within the specified range. Two such trees can be combined in constant time by allocating a new root node containing pointers to the two trees.

3.3. Transformations with Fast Insertion Time

In the last subsection we investigated a set of transforms that only slightly increase the query time at the cost of greatly increasing the processing time. In this subsection we will study a class of structures *dual* to those, which only slightly increase the processing time and greatly increase the query time. Specifically, we will see that there exists an admissible $(k(k!N)^{1/k}, k)$ transform for any positive integer k . As before, we will first investigate the case that $k=2$, and then turn to the general case.

The *dual triangular transform* is illustrated pictorially in Figure 3.8(a). At time 9, there are 6 structures (of sizes 1, 2, 3, 1, 1, and 1); when the 10-th element is inserted it is combined with the last three structures to create a new static structure of size 4. In general, when the $\binom{M}{2}$ -th element is inserted, M elements are combined together to form a static structure of size M ; other elements are kept in singleton structures as they are inserted. Since each element is built into only two static structures (the large and the singleton), we know that

$$P_D(N) \leq 2P_S(N).$$

It is easy to show that at most $2(2N)^{1/2}$ static structures exist at any time, so we have

$$Q_D(N) \leq Q_S(N) \cdot 2(2N)^{1/2}.$$

These facts together imply the following theorem.

Theorem 3.4: (The dual triangular transform)

The dual triangular transform is an admissible $(2(2N)^{1/2}, 2)$ transform.

Proof:

Given in the preceding text. QED.

That this transform is dual to the triangular transform of Subsection 3.2 is intuitively clear from Figure 3.8(a). To make the duality more precise we will study

(a) The dual triangular transform.

(b) The dual 3-binomial transform.

Figure 3.8. Dual binomial transforms.

the dual triangular transform from the viewpoint of the triangular-number counting scheme of the last subsection. The history of the dynamic structure is shown in a tabular form in Figure 3.9. The eighth row shows that when 8 elements are in the dynamic structure, there are 5 static structures: three "large" structures (of size 1, 2, and 3) and two "small" structures (each of only one element). In general, if the number in the "large" column is $\binom{M}{2}$, then there are large structures of size 1, 2, 3, ..., $M-1$. The number in the "small" column gives the number of unit-sized static structures. Note that the entries in the number column are identical to the 2-binomial counting depicted in Figure 3.5.

Structures		Number	
Large	Small	Large	Small
()	()	0	0
(1)	()	1	0
(1)	(1)	1	1
(1,2)	()	3	0
(1,2)	(1)	3	1
(1,2)	(1,1)	3	2
(1,2,3)	()	6	0
(1,2,3)	(1)	6	1
(1,2,3)	(1,1)	6	2
(1,2,3)	(1,1,1)	6	3
(1,2,3,4)	()	10	0
(1,2,3,4)	(1)	10	1
(1,2,3,4)	(1,1)	10	2

Figure 3.9. History of the dual triangular transform.

This duality carries through to the k-binomial transform. For the case of the dual 3-binomial transform, each element will be built into at most three static structures (which we call small, medium and large). All small structures have exactly one element, medium structures have an integer number of elements, and large structures contain a triangular number of elements. At any point in the history of the transform, each set of existing small, medium and large structures contains structures of adjacent sizes. The following table shows the history of the dual 3-binomial transform from the insertion of the fourth through the tenth elements; a history diagram of the dual 3-binomial transform appears in Figure 3.8(b).

N	Structures			Populations		
	Large	Med	Small	Large	Med	Small
4	(1,3)	()	()	4	0	0
5	(1,3)	(1)	()	4	1	0
6	(1,3)	(1)	(1)	4	1	1
7	(1,3)	(1,2)	()	4	3	0
8	(1,3)	(1,2)	(1)	4	3	1
9	(1,3)	(1,2)	(1,1)	4	3	2
10	(1,3,6)	()	()	10	0	0

The extension of this strategy from the dual 3-binomial transform to the dual k -binomial transform is straightforward. The code of Figure 3.6 is modified so that instead of containing a static structure of $\binom{D[i]}{i}$ elements, $P[i]$ now contains a list of structures of sizes

$$\binom{D[i]-1}{i-1}, \binom{D[i]-2}{i-2}, \dots, \binom{i-1}{i-1}.$$

Note that the sum of the sizes of the structures is $\binom{D[i]}{i}$. This allows us to establish the following theorem.

Theorem 3.5: (The dual k -binomial transform)

The dual k -binomial transform is an admissible $(k(k!N)^{1/k}, k)$ transform.

Proof:

Because each element is built into at most k static structures, it is clear that the processing cost increases by at most a factor of k . The analysis used in the proof of Theorem 3.3 shows that each of the k classes of structures contains at most $(k!N)^{1/k}$ distinct structures at any point. Therefore at most $k(k!N)^{1/k}$ static structures exist at any time, providing the upper bound on the query time penalty. QED.

To illustrate the application of this transformation we will again consider the problem of range searching in a d -dimensional point set. Bentley and Maurer [1978] describe a second structure for range searching (which we will call SRS') with properties

$$\begin{aligned} Q_{SRS'}(N) &= O(N^\delta), \\ P_{SRS'}(N) &= O(N \lg N), \text{ and} \\ S_{SRS'}(N) &= O(N), \end{aligned}$$

for any fixed $\delta > 0$. By choosing, for example, $k = \lceil 2/\epsilon \rceil$ and $\delta = \epsilon/2$, we can apply the dual k -binomial transform to achieve the following structure.

New Data Structure 4: (Dual Dynamic Range Searching)

A dynamic range searching (DRS') structure supporting insertions and queries for point sets in d -space with performance

$$\begin{aligned} Q_{DRS'}(N) &= O(N^\epsilon), \\ P_{DRS'}(N) &= O(N \lg N), \text{ and} \\ S_{DRS'}(N) &= O(N) \end{aligned}$$

can be achieved for any fixed $\epsilon > 0$ and positive integer d .

Note that this structure is appropriate when there are many more insertions than queries; it reduces the cost of the computation of certain sequences of N insert and query operations (analogous to those discussed at the end of Subsection 3.2) from the $O(N \lg^d N)$ time required by Lueker's [1978] method to $O(N \lg N)$.

3.4. Summary of the Transformations

In this section we have seen a number of different static-to-dynamic transformations on data structures for decomposable searching problems. We will now spend just a moment reviewing these transformations. The transformations themselves are summarized in Figure 3.10.

<u>Transformation</u>	<u>Query Factor</u>	<u>Processing Factor</u>
k -binomial	k	$(k!N)^{1/k}$
Binary	$\lg(N+1)$	$\lg(N+1)$
Dual k -binomial	$k(k!N)^{1/k}$	k

Figure 3.10. Summary of transformations.

There are many other transformations besides those that we have already investigated. A simple way of achieving a new transformation is by isomorphism to a particular number system (counting scheme). This is illustrated in Figure 3.11 for the radix-3 number system (ternary counting). Part (a) of that figure shows the ternary transform: each static structure is of size either a power of three or twice a power of three, and corresponds to either a one or a two in the ternary representation of the number of elements in the dynamic structure. This transform is an admissible $(\lceil \log_3 N \rceil, 2\lceil \log_3 N \rceil)$ transform.¹¹ Its dual is shown in part (b) of the figure; every structure in the dual is of size a power of three, and there are 0, 1 or 2 structures for any power of three, corresponding to the appropriate digit in the ternary expression of the integer size of the structure. This is an admissible $(2\lceil \log_3 N \rceil, \lceil \log_3 N \rceil)$ transform. This scheme can be extended to radix-k counting to yield a primary $(\lceil \log_k N \rceil, (k-1)\lceil \log_k N \rceil)$ transform and a dual $((k-1)\lceil \log_k N \rceil, \lceil \log_k N \rceil)$ transform. An interesting open problem is to examine other counting schemes (such as Fibonacci counting) for their properties as transforms.

Figure 3.11. Radix-3 transformations.

It is now easy to state formally the relationship of the primary and dual transforms derived from a particular counting scheme. In the primary transform, there is a single structure corresponding to each digit, whereas in the dual

¹¹ This and the following claims about radix-k transforms assume $N > 1$.

transform each digit corresponds to a set of structures that are the "carries" from its right neighbor (the units digit is a set of structures of size one).

The transformations of this section together provide a powerful set of tools for designing new data structures for both particular applications and as a component in larger algorithms. To design a dynamic structure in a given context, the algorithm designer first designs a static structure (which is usually much easier than designing a dynamic structure), and then applies one of the transformations to achieve an efficient dynamic structure. Which transformation he uses depends on the relative efficiency of the static preprocessing and query costs and on the expected frequency of insertions and queries.

As we mentioned before, the cost penalties of the transformations need not always be paid. One can often avoid them by merging static structures, by analyzing the average query time, or by performing separate analyses for fast-growing cost functions.

4. Lower Bounds on Transformations

Our main goal in this section is to prove the optimality, in a certain sense, of some of the transformations discussed in Section 3. Our path to this goal will have many steps, and the reasons for each step might not be clear in advance. To aid the reader, we now briefly sketch the contents of this section.

In Subsection 4.1 we define the model of computation which we will use throughout the rest of the section. We also advise the reader that the use of this model implies certain limitations on the applicability of the results we will obtain. In Subsections 4.2 through 4.4 we show a method for representing an initial sequence of insertions under some transform as a binary tree, and show how the efficiencies of transformations are related to properties of the corresponding trees. To achieve the correspondence between transforms and trees, we restrict our attention to a class of transforms which we call the *arboreal* transforms. In Subsection 4.5 we state and solve a recurrence relating the various tree properties defined in

Subsection 4.4, and interpret this result as it applies to the k-binomial transformations. We then extend the basic result to answer questions about other transformations (including the binary transformation) in Subsection 4.6. In Subsection 4.7 we discuss the justification of the restriction to arboreal strategies, and in Subsection 4.8 we return to explore the limitations (implied by our model) of the preceding results, showing a number of cases in which our "lower bounds" can be beaten by going outside the model.

4.1. The Model of Computation

The most important assumption of our model is that the transformations under consideration are not allowed to use any specific knowledge about the original problem or static structure except for the fact that the problem is decomposable. It therefore remains plausible for any particular decomposable searching problem, P , that there exists a dynamic data structure for P having performance better than that produced by applying any optimal static-to-dynamic transform to any static structure for P . For example, AVL trees (see Knuth [1973]) provide a dynamic data structure for member searching with

$$\begin{aligned} P_{AVL} &= O(N \lg N), \\ S_{AVL} &= O(N), \text{ and} \\ Q_{AVL} &= O(\lg N). \end{aligned}$$

The results of this section imply that no dynamic structure with this efficiency can be obtained (in the worst case) by applying a general transform to a static structure for member searching; the efficiency of AVL trees depends on particular properties of the member searching problem other than decomposability (in particular, the ability to maintain the structural invariant under rotation).

Our model of computation is that we have three operations, Build, Query, and $\square$, whose inner workings we may not examine. Build works with performance P_S to create static structures. Query works with performance Q_S to search the structures created by Build. The $\square$ operator is guaranteed to have the property

$$\square(\text{Query}(x, \text{Build}(A)), \text{Query}(x, \text{Build}(B))) = \text{Query}(\text{Build}(A \cup B))$$

The only way to answer a query is by applying Query one or more times to structures created by Build and then combining the results using $\square$. It is assumed that P_S grows at least linearly and that Q_S is monotone non-decreasing.

To measure the computation costs (P_D and Q_D) associated with a dynamic structure, we will charge only for the computation time of calls to Build and Query. It should be noted that these costs will generally be the dominant parts of the total costs of the dynamic algorithms. In any case, this approximation is certainly acceptable for the purpose of establishing lower bounds on the costs of dynamic algorithms.

Our goal in the search for efficient transformations is to minimize simultaneously the penalty functions

$$F(N) = \max_{1 \leq i \leq N} Q_D(i)/Q_S(i) \text{ and}$$

$$G(N) = P_D(N)/P_S(N).$$

The bulk of this section will be devoted to showing limits on just how far this process may be carried in the worst case. Our interpretation of the term "worst case" in this context is a bit tricky. We have already mentioned that we may assume no specific knowledge about the problem or the original static structure except for decomposability. It is also important to note that we do not allow ourselves to assume any specific knowledge about the *efficiency* of the underlying static structure, except that P is at least linear and Q is monotone non-decreasing. (Note, for example, that the improvements in F and G which occur for fast-growing P and Q are not examples of worst-case behavior, so there is no contradiction in the fact that our lower bounds deny the possibility of such improvements in the general case.)

The reader may find it helpful to think of the worst case as that in which P is linear and Q is constant, the intuition being that it is hardest for the dynamic structure's costs to approach the static structure's costs when the latter are as small as possible. Since we may not use any specific knowledge about the original

static problem or data structure, any solution to the dynamic problem must work by maintaining a collection of static structures. Whenever an element is inserted, a new structure must be created containing that element¹² and possibly some other elements. Also, some existing static structures may be thrown away. When a query is made to the dynamic structure, it is necessary to search some set of static structures which together contain all the elements inserted so far.

For the following analysis, we will place a few restrictions on the nature of the dynamic structures we will consider. We will return later to the problem of justifying these restrictions. Our first restriction is as follows:

Restriction 4.1: (Dynamic structures partition elements into static structures)

We assume that at any time there exists exactly one static structure containing each element which has been inserted so far. That is, the static structures partition the set of elements represented by the dynamic structure.

With the preceding assumptions in mind, we are now ready to move on to the first steps of our analysis.

4.2. Computing F and G

We now give some rules for determining the worst-case values of the penalty functions F and G associated with a particular strategy.

Definitions: (f and g)

Consider the history of a dynamic structure over the course of any number of insertions starting when the structure is empty. We define $f(N)$ as the maximum number of static structures existing after one of the first N insertions. We define $g(N)$ as the sum of the cardinalities of all sets of elements built into static structures created over the course of the first N insertions.

¹²While we may conceive of strategies in which new static structures are created by queries into the dynamic structure, we need not consider this possibility for this worst-case analysis, since P_S could grow much more rapidly than Q_S .

Note that, while the definitions of f and g actually depend on the specific transform used, the identity of the transform under consideration will always be clear from context. We may now bound F and G as follows:

Theorem 4.1: (f bounds F)

For any positive integer N , $F(N) \leq f(N)$.

Proof:

After any of the first N insertions (say the i -th), at most $f(N)$ static structures exist. To compute the cost of answering a query, we charge precisely for querying these structures. Since each of these structures has cardinality no larger than i , and since Q_S is monotone non-decreasing, the total cost is at most $f(N)Q_S(i)$. QED.

Theorem 4.2: (g/N bounds G)

For any positive integer, N , $G(N) \leq g(N)/N$.

Proof:

We note that any static structure built during the first N insertions will have cardinality no larger than N . Consider such a structure, S , having cardinality i . By the fact that P_S grows at least linearly, we may bound the cost of building S by the inequality

$$P_S(i) \leq iP_S(N)/N$$

Summing over all static structure, we get

$$P_D(N) \leq g(N)P_S(N)/N,$$

implying

$$G(N) = P_D(N)/P_S(N) \leq g(N)/N.$$

QED.

By the assumptions in Subsection 4.1, the preceding bounds are the tightest possible for the general case. We will therefore concern ourselves henceforth with

the problem of minimizing f and g rather than F and G .

4.3. Transforming History Diagrams to Trees

The transforms we discussed in Section 3 are all representable by history diagrams, such as those in Figures 3.1, 3.4, 3.7, 3.8, and 3.11. It is not the case, however, that all transforms are so representable; in order for a static structure to be represented as a (contiguous) rectangle in a history diagram, it is necessary that it be built from a set of elements which were inserted consecutively during the history of the structure. We now impose our second restriction on the class of dynamic structures to be considered:

Restriction 4.2: (Contiguity of static structures)

We will restrict our attention to transforms whose histories are representable by history diagrams.

Indeed, we will further restrict our attention to those history diagrams (such as the ones in Section 3) in which every rectangle reaches to the "diagonal" of the diagram. We may state this otherwise as

Restriction 4.3: (Eagerness of static structures)

We will restrict our attention to transforms in which each static structure is built as soon as all its elements have been inserted, and in which the elements of any discarded static structure are always built into a single new static structure (along with some additional elements).

Strategies which satisfy Restrictions 4.1, 4.2, and 4.3 will be called *arboreal* strategies for a reason that will soon become obvious.

Consider the history diagram for the first N insertions into a dynamic structure which is maintained by an arboreal strategy. Any such diagram induces a binary tree, as shown in Figure 4.1. We may draw this tree by tracing the left and upper edges of each rectangle in the diagram. The internal nodes of the tree will thus be at the upper left corners of the various rectangles; each internal node of the tree

corresponds to a (unique) static structure. We will now go on to study some relationships between the efficiencies of arboreal strategies and properties of their induced trees.

Figure 4.1. A history diagram and its induced tree.

4.4. Tree Properties and their Relation to Performance

We now introduce some basic vocabulary for discussing properties of binary trees.

Definitions: (Tree properties)

Let T be a binary tree. Then $\text{leaves}(T)$ denotes the set of all leaves of T and $\text{nodes}(T)$ denotes the set of all internal nodes of T . The *weight* of T , denoted $|T|$, is defined as the cardinality of $\text{leaves}(T)$. For any internal node, a , of T the left and right sons of a are denoted $\text{lson}(a)$ and $\text{rson}(a)$, respectively. If a is a leaf of T , then the *right depth* of a , written $\text{rd}(a)$, is defined as the number of right branches along the path from the root of T to a . The *right height* of T , $\text{rh}(T)$, is the maximum right depth of any leaf of T . The *right path length* of T , $\text{R}(T)$, is defined as the sum of the right depths of all leaves of T . Left depth, left height, and left path length are defined analogously.

We will sometimes identify a (not necessarily internal) node, x , of a tree with the subtree rooted at x . For example, we may write $|x|$ to indicate the number of leaves which are descendants of x .

We now make the following observation:

Theorem 4.3: (Alternate characterization of left path length)

Let T be a tree. Then,

$$L(T) = \sum_{n \in \text{nodes}(T)} |\text{son}(n)|$$

Proof:

Consider any leaf, x , of T . We need only note that the left branches along the path from the root of T to x emanate precisely from those nodes of T whose left sons contain x . QED.

With this characterization of left path length in mind, we may now relate the trees induced by arboreal strategies to the penalty functions associated with those strategies.

Consider the tree in Figure 4.1(b). To each static structure created during the partial history represented by that tree, there corresponds a right (horizontal in the diagram) branch whose length (in the diagram) is proportional to the cardinality of that static structure. Moreover, for any internal node, n , of the tree, the length (in the diagram) of the right branch from n corresponds precisely to the number of leaves in the *left* son of n . By summing over all internal nodes of the tree, we establish the following result:

Theorem 4.4: (Relation of g to left path length)

Let N be a positive integer and let T be the tree induced from the history diagram representing the first N insertions into a dynamic structure maintained by some arboreal strategy. Then, $L(T) = g(N)$.

Proof:

Given in the preceding text. QED.

We may also characterize N and f in terms of tree properties:

Theorem 4.5: (Relation of N and f to tree properties)

Let N be a positive integer and let T be the tree induced from the history diagram representing the first N insertions into a dynamic structure maintained by some arboreal strategy. Then,

$$|T| = N + 1 \text{ and} \\ rh(T) = f(N).$$

Proof:

Inspection of Figure 4.1 will reveal that these results are obvious. QED.

The theorems proven so far in this section allow us to address the problem of "simultaneously minimizing" F and G by investigating a closely related problem about trees, namely that of "simultaneously minimizing" the right height and left path length of a tree with a fixed number of nodes. To discuss this more precisely, we make the following definition:

Definition: (Minimal left path length)

Let n and k be positive integers. We define

$$L_k(n) = \text{Min } \{L(T) \mid T \text{ is a tree such that } |T| = n \text{ and } rh(T) \leq k\}.$$

Since the only tree with zero right height is the tree of one node (which also has zero left path length), we also define

$$L_0(1) = 0.$$

By convention, we will regard $L_0(n)$ as "positive infinity" whenever $n > 1$. A tree with n leaves, right path length k , and left path length $L_k(n)$ will be called an *economical tree*.

In the next few pages, we will investigate the behavior of $L_k(n)$ as k and n vary, and then restate our findings in terms of lower bounds on worst-case penalty

functions.

4.5. The Behavior of $L_k(n)$

Consider a binary tree, T , with root node t . Let A and B be the subtrees rooted at $a=lson(t)$ and $b=rson(T)$, respectively. The weight, right height, and left path length of T may be recursively computed from properties of A and B by the relations

$$\begin{aligned} |T| &= |A| + |B|, \\ rh(T) &= \max(rh(A), rh(B)+1), \text{ and} \\ L(T) &= L(A) + |A| + L(B). \end{aligned}$$

From this, we obtain the following recurrence for $L_k(n)$:

Theorem 4.6: (Recurrence for $L_k(n)$)

Let n and k be any positive integers. Then,

$$L_k(n) = \begin{cases} 0 & n=1 \\ n-1+L_1(n-1) = \binom{n-1}{2} & k=1, n>1 \\ \min_{1 \leq i \leq n-1} [L_k(i) + i + L_{k-1}(n-i)] & k>1, n>1 \end{cases}$$

Proof:

The results for $k=1$ follow by considering the unique binary tree of any weight which has right height ≤ 1 . For the case $k>1$, we consider a tree, T (with root t) having weight $n>1$ and height k . Let t be the root of T . And let A and B be the subtrees rooted at $a=lson(t)$ and $b=rson(t)$, respectively. Then we must have:

$$\begin{aligned} 1 &\leq |A| \leq n, \\ |A| + |B| &= n, \\ rh(A) &\leq k, \text{ and} \\ rh(B) &\leq k-1. \end{aligned}$$

Moreover, if the left path length of T is to be minimal, the left path lengths of A and B must be minimal. That is, we must have

$$\begin{aligned} L(A) &= L_k(|A|) \text{ and} \\ L(B) &= L_{k-1}(|B|). \end{aligned}$$

These requirements are precisely captured by our recurrence. QED.

We now come to the principal theorem of this section, wherein the behavior of $L_k(n)$ is precisely characterized in terms of binomial coefficients.

Theorem 4.7: (Characterization of $L_k(n)$)

Let k and m be non-negative integers such that $k \leq m$, and let n be a positive integer satisfying

$$\binom{m}{k} \leq n \leq \binom{m+1}{k}$$

Then,

$$L_k(n) = k \binom{m}{k+1} + (m-k-1)N, \quad [1]$$

where

$$N = n - \binom{m}{k}.$$

Proof:

Our proof will proceed by induction on k and, for each fixed positive value of k , by induction on n .

Base Step: ($k = 0$)

In this case, we have

$$\binom{m}{0} = 1 = \binom{m+1}{0}.$$

This implies that $n = 1$, so the right hand side of [1] reduces to

$$\begin{aligned} 0 \binom{m}{0+1} + (m-0-1) \left(n - \binom{m}{0} \right) &= 0 + (m-1)(1-1). \\ &= 0 \\ &= L_0(1) \end{aligned}$$

Inductive Step: ($k > 0$)

We now must show that the theorem holds for any $k > 0$ assuming it holds for all smaller k . We proceed by induction on n . In doing this, we must take note of the interaction between m and n . Since k is positive, $\binom{m}{k}$ increases monotonically with m . Thus, the minimum possible value of n is $\binom{k}{k} = 1$, and for any positive value of n , there is at least one possible value for m (and occasionally there will be two).

Base Step: ($n = 1$)

In this case, we must have $m = k$, so the right hand side of [I] reduces to

$$\begin{aligned} k \binom{k}{k+1} + (k-k-1) \left[1 - \binom{k}{k} \right] &= k(0) + (-1)(1-1) \\ &= 0 \\ &= L_k(1). \end{aligned}$$

Inductive Step: ($n > 1$)

We first show that the right hand side of [I] gives an upper bound on $L_k(n)$. Note that

$$\binom{m-1}{k} + \binom{m-1}{k-1} = \binom{m}{k} \leq n \leq \binom{m+1}{k} = \binom{m}{k} + \binom{m}{k-1}.$$

We now pick a and b such that

$$\begin{aligned} \binom{m-1}{k} &\leq a \leq \binom{m}{k}, \\ \binom{m-1}{k-1} &\leq b \leq \binom{m}{k-1}, \text{ and} \\ a + b &= n. \end{aligned} \tag{II}$$

By Theorem 4.6, we have

$$\begin{aligned} L_k(n) &\leq L_k(a) + a + L_{k-1}(b) \\ &= k \binom{m-1}{k+1} + ((m-1)-k-1)(A) + \\ &\quad \binom{m-1}{k} + A + \\ &\quad (k-1) \binom{m-1}{k} + ((m-1)-(k-1)-1)(B) \\ &= k \left[\binom{m-1}{k+1} + \binom{m-1}{k} \right] + (m-k-1)(A+B) \\ &= k \binom{m}{k+1} + (m-k-1)N, \end{aligned}$$

where,

$$\begin{aligned} A &= a - \binom{m-1}{k}, \\ B &= b - \binom{m-1}{k-1}, \text{ and} \\ N &= n - \binom{m}{k}. \end{aligned}$$

This establishes that our expression is an upper bound on $L_k(n)$. To establish that this is also a lower bound, we must show that no other way of expressing n as the sum of two positive numbers, a and b , will give a smaller value for

$$L_k(a) + a + L_{k-1}(b) \tag{III}$$

To show this, we consider the effect on the value of expression [III] of increasing or decreasing a by steps of one.¹³ Suppose we start with a and b chosen to satisfy [II], and then start incrementing a and decrementing b by steps of 1. So long as a remains less than $\binom{m}{k}$ and b remains greater than $\binom{m-1}{k-1}$, the effect of each increment will be to increase $L_k(a) + a$ by $((m-1)-k-1) + 1 = m-k-1$ and to decrease $L_{k-1}(b)$ by $(m-1)-(k-1)-1 = m-k-1$, leaving the total value of [III] unchanged.¹⁴ However, as soon as either a or b exceeds the stated bound, one or more of the following things will happen:

1. The incremental growth of $L_k(a)$ will increase while the incremental shrinkage of $L_{k-1}(b)$ decreases or remains the same,
2. The incremental shrinkage of $L_{k-1}(b)$ will decrease while the incremental growth of $L_k(a)$ increases or remains the same, or
3. b will diminish to 0.

In any case, a smaller value for [III] will not be obtained. Similarly, if we start with a and b as in [II] and decrease the value of a while increasing b , then we will have zero or more steps at which [III] remains unchanged, zero or more steps where the increase in $L_{k-1}(b)$ exceeds the decrease in $L_k(a) + a$, and finally the step at which a diminishes to zero. Thus, the rules given in [II] give an optimal partitioning of n into a and b . This completes the induction step and the proof.

QED.

¹³In the following, we assume that $k > 1$. If $k = 1$ we must always take $b = 1$ (and $a = n-1$), since only then is $L_0(b)$ defined.

¹⁴The incremental changes given here are found by substitution into the second term of the right hand side of [I], under the induction hypothesis.

The use of the auxiliary variable, m , in expression [1] makes it a bit difficult to grasp intuitively what is being said about the effects of n and k on $L_k(n)$. To make the picture clearer, we will briefly study the asymptotic behavior of $L_k(n)$ as k remains fixed and n grows without bound. Consider first what happens as n ranges only over binomial coefficients of the form $\binom{m}{k}$. We note that

$$n = \binom{m}{k} \Rightarrow m-k+1 \leq (n/k!)^{1/k} \leq m.$$

So,

$$\begin{aligned} L_k(n) &= k \binom{m}{k+1} \\ &= kn(m-k)/(k+1) \\ &\sim [k/(k+1)]k!^{1/k} n^{1+1/k}. \end{aligned}$$

Since the growth of $L_k(n)$ is very well behaved,¹⁵ the preceding may be extended to cover all values of n .

Theorem 4.8: (Asymptotic behavior of $L_k(n)$)

Let k be any positive integer. Then,

$$L_k(n) \sim [k/(k+1)]k!(1/k)n^{1+1/k}.$$

Proof:

The result follows directly from the preceding text. QED.

By precisely characterizing $L_k(n)$, Theorem 4.7 gives us a bound on the efficiencies of arboreal static-to-dynamic transforms. Any such strategy which has $f(N) \leq k$ for all N must always have $g(N) \geq L_k(N+1)$. The asymptotic behavior of $L_k(n)$ given by Theorem 4.8, and our knowledge that Theorems 4.1 and 4.2 are the best possible within our model, tell us that whenever we have

$$F(N) < k$$

¹⁵ Given the values where n is of the form " m choose k ", we can find the exact values at all other n by linear interpolation.

for any positive integer k , we must also have

$$G(N) \geq L_k(N+1)/N \sim (k!N)^{1/k}.$$

This is the precisely behavior achieved by the k -binomial transforms, up to lower order terms. Note, however, that the exact lower bound is not always achievable. The reason for this is the consideration of *immutability of history*. If we know in advance that there will be exactly N insertions, then an optimal strategy can be devised by working backwards from an economical tree of weight $N+1$ and right height k . But if the total number of insertions to be made turns out to be larger, then a different strategy for the first N insertions may have been called for. Fortunately, the results of this restriction turn out not to be too severe, since the k -binomial strategies have efficiency very close to this theoretical limit. The following theorem shows that, for any k , the $G(N)$ achieved by the k -binomial transform is optimal (for $F(N) \leq k$) not only to within lower order terms but actually to within an additive constant of 1.

Theorem 4.9: (Optimality of k -binomial transforms)

For any positive integer, k , the k -binomial transform achieves

$$\begin{aligned} f(N) &\leq k \text{ and} \\ g(N) &\leq L_k(N+1) + N \end{aligned}$$

for all positive N .

Proof:

Examination of the optimal construction given in the proof of Theorem 4.7 shows that the k -binomial strategy achieves the optimal value of

$$f(N) = L_k(N+1)$$

when N is of the form

$$N = \binom{m}{k} - 1$$

for some $m \geq k$. For intermediate values of N , we need only note that, after the first N insertions under the k -binomial strategy, the sum of the cardinalities of all structures formed so far *except those in existence after the N -th insertion* (note that these latter must have a total cardinality of N) will never be greater

than $L_k(n)$. This fact may be established by induction on k , using the fact that values of $L_k(n)$ are given exactly by linear interpolation between points at which the k -binomial transform gives absolutely minimal values of $f(N)$. QED.

4.6. Allowing the Number of Static Structures to Grow

So far in this section we have only considered minimizing $g(N)$ where $f(N)$ is bounded by a constant. In other words, we have considered only strategies which allow some fixed maximum number of static structures to exist at one time. In Section 3, however, we also investigated strategies (the binary and the dual k -binomial transforms) which allow the the number of static structures to grow without limit as the total number of elements in the dynamic structure increases. We will now, therefore, briefly investigate transforms which allow $f(n)$ to grow without bound.

To study the efficiency of transforms in which $f(N)$ is unbounded, we may consider the behavior of $L_k(n)$, where k is allowed to vary with n .¹⁶ We must be aware of two possible consequences of allowing k to grow:

- (1) For any particular k , n may never grow large enough for $L_k(n)$ to approach the asymptotic behavior given by Theorem 4.9.
- (2) Our previous caveat about the immutability of history may become more significant.

Since the asymptotic approach of $L_k(n) / [k/(k+1)]k!^{1/k}n^{1+1/k}$ to unity (as n grows and k remains constant) is from below, (1) may be ignored for the purpose of investigating upper bounds. Since the immutability of history can never make it easier to devise efficient transforms, this consideration may be ignored for the investigation of lower bounds. Because of these complicating factors, our results for transforms with unbounded f are less precise than those for bounded f . A few

¹⁶ In accordance with the notational conventions of this section, we have $k = f(n) = f(N+1)$, since the first N insertions always give a history diagram which induces a tree of weight $N+1$.

results are nonetheless worth noting. The first of these is the following.

Theorem 4.10a: (Optimality of the binary transform)

For any arboreal transform such that $f(N) = O(\lg N)$, $g(N) = \Omega(N \lg N)$.

Proof:

Since constraining the growth of f can only increase and never decrease the necessary growth of g , we need only consider the case where $f(N) = \theta(\lg N)$. We must show that $L_{f(N)}(N+1) = \Omega(N \lg N)$. We define the function M by

$$M(n,k) = \text{Max } \{m \mid \binom{m}{k} \leq n\}.$$

From the fact that $f(N) = \theta(\lg N)$, it follows that $M(N, f(N)) - f(N) = \theta(\lg N)$. This gives us

$$\begin{aligned} g(N) &\geq L_{f(N)}(N+1) \\ &\geq L_{f(N)}(N) \\ &\geq f(N) \binom{M(N, f(N))}{k+1} \\ &= [f(N)/(f(N)-1)] [M(N, f(N)) - f(N)] N \\ &= \theta(N \lg N) = \Omega(N \lg N). \end{aligned}$$

QED.

This result tells us that the binary transform is optimal in the sense that any transform that pays as small a penalty in search cost (within a constant factor) must pay at least as large a penalty in insertion (again within a constant factor); any arboreal transform which achieves $F(N) = O(\lg N)$ in the worst case must also pay $G(N) = \Omega(\lg N)$.¹⁷ The binary transform is also optimal in the sense that any transform which is actually cheaper (by more than a constant factor) for searches must be strictly more expensive (again by more than a constant factor) for insertions. We state this result more formally in the following theorem.

Theorem 4.10b: (Optimality of the binary transform)

¹⁷ This follows from the fact that Theorems 4.1 and 4.2 are the tightest results possible within our model.

For any arboreal transform such that $f(N) = o(\lg N)$, $g(N) = \omega(N \lg N)$.

Proof:

Let the function h be defined by

$$h(N) = (\lg N)/f(N).$$

From the hypothesis that $f(N) = o(\lg N)$, it follows that $h(N) = \omega(1)$. Moreover, since $M(N, f(N)) \leq \lg N$, we have $f(N) = o(M(N, f(N)))$, which means that the approximation in Theorem 4.8 remains valid.¹⁸ This gives us

$$\begin{aligned} g(N) &\geq L_{f(N)}(N+1) \\ &\geq L_{f(N)}(N) \\ &\sim [f(N)/(f(N)+1)] f(N)!^{1/f(N)} N^{1+1/f(N)} \\ &\sim [1](f(N)/e) N^{1/f(N)} N \\ &= [(\lg N)/(e h(N))] 2^{h(N)} N \\ &= \omega(N \lg N). \end{aligned}$$

QED.

This implies that any arboreal transform which achieves $F(N) = o(\lg N)$ in the worst case must also pay $G(N) = \omega(N \lg N)$.

In the preceding proof, we saw that the approximation given in Theorem 4.8 still serves to provide a lower bound on the growth of g even when f is allowed to grow without bound, provided that $f(N) = o(\lg N)$. The next natural question is whether this bound can always be achieved. It turns out that this is not always possible. If f grows in a very irregular manner, having sudden spurts of growth separated by intervals of almost no change, then the immutability of history will cause $g(N)$ to be much larger than $L_{f(N+1)}(N+1)$ for values of N immediately following the sudden increases. If f grows "smoothly" (the precise meaning of this term is implicit in the following theorem), however, this lower bound for $g(N)$ is very nearly obtainable. We state this result formally as follows.

¹⁸That is, consideration (1) may be disregarded.

Theorem 4.11 (Optimizing g for slowly growing f)

Let h be a differentiable function such that

$$\begin{aligned} h(x) &= \omega(1) \text{ and} \\ h'(x) &= o(1/x). \end{aligned}$$

Then, there exists a transform having

$$\begin{aligned} f(N) &\leq \lceil h(N) \rceil \text{ and} & [I] \\ g(N) &\sim (h(N)/e)N^{1+1/h(N)}. & [II] \end{aligned}$$

Moreover, given [I], [II] is optimal up to lower order terms.

Proof:

A structure having the performance described may be formed by a process of "cutting and pasting" from the history diagrams of the various k -binomial strategies. We omit the details for brevity and for the sake of keeping the reader awake. The optimality of [II], given [I], is implicit in the proof of Theorem 4.10b. QED.

Our results for transforms in which $f(N) = \omega(\lg N)$ are much less complete. In particular, we know that the performance of the dual k -binomial transforms falls substantially short of the bound given by the inequality

$$g(N) \geq L_{f(N)}(N+1).$$

We conjecture that this is an inevitable penalty of the immutability of history, and that the dual binomial transforms are in fact optimal in some strong sense, similar to that of Theorem 4.9 for the ordinary binomial transforms. The problem of finding optimal transforms in which $f(N)$ grows faster than $\lg N$ but slower than n^ϵ for any positive ϵ remains open.¹⁹

¹⁹We may equivalently view this as the problem of optimizing f when $g(N)$ grows asymptotically faster than N but slower than $N \lg N$.

4.7. Justification of the Restriction to Arboreal Transforms

In Subsections 4.1 and 4.3, we introduced three restrictions which together constrained our investigation to arboreal transforms. While we conjecture that arboreal strategies are optimal, in the sense that for any non-arboreal transform there exists an arboreal transform which is at least as good (given the "black box" model described in Subsection 4.1), we have not yet found a rigorous proof. In this subsection, we will summarize our reasons for considering each of the restrictions reasonable.

Restriction 4.1 forbids the existence of multiple structures containing the same element. Our intuition is that any strategy which permits such overlapping structures can be improved by omitting the shared elements from all but one of the overlapping structures. To justify this intuition would require careful examination of the consequences of this omission when that one structure is finally destroyed. We may also forbid overlapping structures on the grounds that transformations which allow them cannot be optimal for space in the worst case. An even more serious objection is that there are a number of problems which satisfy the definition of decomposability only when the unions involved are of disjoint sets.

Our intuitive justification for Restriction 4.2 (contiguity of static structures) is the belief that a partial history which does not satisfy this restriction can be turned into one that does, at no cost in $f(N)$ or $g(N)$, by a kind of "permutation of the names of the elements." To show this would justify the restriction at least for the cases where f is bounded or grows slowly and smoothly, so that the immutability of history is not a significant problem.

For Restriction 4.3, we can actually give a rigorous justification, at least over the class of transforms which already satisfy Restrictions 4.1 and 4.2. We express this in the following theorem:

Theorem 4.12: (Optimality of eager strategies)

Let N be a positive integer. For any partial history consisting of the first N

insertions and satisfying Restrictions 4.1 and 4.2, there exist a partial history which also satisfies Restriction 4.3 and which has $f(N)$ and $g(N)$ no greater than those for the original partial history.

Proof:

Any partial history which satisfies the first two restrictions may be represented by a history diagram. We may insure that the rectangle in the upper left corner of the diagram represents a structure which is formed as soon as all its elements become available, for any diagram which does not have this property can be transformed at no cost into one that does. The construction is as follows:

Let R be the upper left rectangle in the diagram. Consider the leftmost rectangle immediately below R . If it is wider than R , then we extend it upwards to the top of the diagram, obliterating R ; if it is narrower than R , then we extend R downwards by one step. This process is repeated until the property holds.

But now the rest of the diagram (excluding the upper-left rectangle) must consist of zero, one, or two staircase-shaped pieces to which the same process may be applied recursively, finally yielding a diagram satisfying Restriction 4.3. No step in this process increases either the total preprocessing cost or the maximum number of simultaneously existing structures, so Restriction 4.3 has been formally justified. QED.

4.8. Limitations on the Significance of the Lower Bounds

The lower bounds we have derived in this section are based on the model of computation given in Subsection 4.1. Before concluding the section, we will mention some of the limitations which this implies for the applicability of our results.

We have already mentioned that it is often possible to obtain superior dynamic data structures for individual decomposable problems (e.g., Member) by using specific properties of those problems. Another assumption on which our lower bounds depend is that Theorems 4.1 and 4.2 are the strongest possible results of their kind, because we assume no knowledge about the performance of the original static algorithm. As we saw at the end of Subsection 3.1 the penalty factors, $F(N)$

and $G(N)$, may be greatly reduced (from $\theta(\lg N)$ to $O(1)$ in the example of Subsection 3.1 if the cost functions of the static structure are already fast-growing. We now present some results concerning a slightly different way of lowering the penalty functions given fast-growing cost functions for the original static structure.

Suppose we are given a static structure for a decomposable searching problem having preprocessing cost $P_S(N)$ and query cost $Q_S(N)$. We will make only the usual assumption about Q_S --that it is monotone non-decreasing. We will, however, make the assumption that $P_S(N)$ not only grows at least linearly with N , but is actually $\theta(N^2)$. If we apply the 2-binomial (triangular) transform, will obtain a dynamic structure having cost functions, P_D and Q_D , which satisfy

$$Q_D(N) \leq 2Q_S(N) \text{ and} \\ P_D(N) = \theta(N^{5/2}).$$

The reader is advised to go through the exercise of verifying the latter assertion. The penalty factor in preprocessing is given by

$$G(N) = P_D(N)/P_S(N) = \theta(N^{1/2}),$$

which is at most a constant factor improvement over the worst-case result given in Theorem 3.2. We appear to get negligible compensation for the fact that the preprocessing cost is already much more than linear. If we look a little more carefully, however, we may notice an interesting phenomenon.

In the triangular strategy, we maintain two structures, a large one, having cardinality $O(N)$, and a small one, having cardinality $O(N^{1/2})$. If we break down $P_D(N)$ into the cost of forming all the large structures built during the first N insertions and the cost of forming all the small structures built during the first N insertions, we find that the large structures have a total cost of $\theta(N^{5/2})$, while the total cost of the small structures is only $\theta(N^2)$. If P_S had been linear, then the costs of the two families of structures would have been equal within a constant factor, each being $\theta(N^{3/2})$. The present disparity suggests that it might be better to merge the small structures into the large ones less frequently. And, indeed, if we adopt the strategy of rebuilding all the elements into a single structure only when

the size of the small structure would exceed $N^{2/3}$, we achieve a dynamic structure having

$$Q_D(N) \leq 2Q_D(N) \text{ and} \\ P_D(N) = \theta(N^{7/3}) = O(N^{1/3}P(N))$$

(as the reader may again wish to verify), the total preprocessing cost being split evenly (within a constant factor) between the two families of structures. The preceding results may be generalized to arbitrary polynomial preprocessing costs and arbitrary binomial transforms, as shown in the following theorem.

Theorem 4.13: (Shift-of-strategy speed-ups)

Let k be an arbitrary positive integer and let r be a real number²⁰ greater than

1. Suppose that we are given a static structure for a decomposable searching problem with cost functions satisfying the following criteria:

$$Q_S(N) \text{ is monotone non-decreasing,} \\ P_S(N) \text{ grows at least linearly, and} \\ P_S(N) = \omega(N^r).$$

Then, a dynamic data structure can be constructed such that

$$Q_D(N) \leq kQ_S(N) \text{ and} \\ P_D(N) = O(N^R P_S(N)),$$

where

$$R = (r-1)/(r^k-1).$$

Proof:

We maintain a set of structures satisfying the following invariants:

- (1) After any insertion there are at most k static structures.
- (2) Let j be a positive integer. After the N -th insertion, the cardinality, C_j , of the j -th largest structure (if there are at least

²⁰The nit-picking reader will delight in noting that it is not quite correct to allow r to be an arbitrary real number. In order for the desired transform to be implementable, r must be Turing computable. Even then, if r is very expensive to compute, the bookkeeping costs may kill us. Similar considerations apply to the function h in Theorem 4.11.

j structures in existence) satisfies

$$C_j \leq N^{(r^k - r^j)/(r^k - 1)}.$$

When an element is inserted, we see how many structures already exist. If there are fewer than k , we simply build the new element into a static structure of cardinality one. If k structures already exist, we rebuild the smallest structure to include the new element. We then repeatedly (zero or more times) merge the smallest two structures until (2) is satisfied. We leave it to the reader to verify that this strategy achieves the advertised performance. QED.

In any strategy based on the construction in the previous proof, the total preprocessing will be divided evenly (up to constant factors) among k families of structures. We conjecture that this gives optimal P_D within a constant factor (which may depend on r and k). Needless to say, similar improvements are available, both in preprocessing time and in query time, for a number of other transformations, given sufficiently fast-growing cost functions. Only a small fraction of the possibilities have been explored.

5. Online Transformations

All of the transforms in Section 3 have the property that some insertions are very cheap while others are very expensive. For example, in the binary transform the 1023-rd insertion is much less costly than the 1024-th. While this situation is quite acceptable in certain applications (such as when the total cost of accessing a structure throughout an entire algorithm is counted), it is prohibitive in others (such as online data bases). In this section we will show how the transforms in Section 3 can be modified to amortize the cost of building static structures over the time of many insertions.

In Section 4, we worked on the principle that any static structure might as well be formed as soon as all its elements became available, since the cost of building it would eventually have to be paid anyway. While this is reasonable if we are

concerned only with the total cost of all insertions, it is inappropriate if we wish to make sure that no individual insertion is inordinately expensive. Figure 5.1 shows a strategy which is similar to the binary strategy of Subsection 3.1, except that each structure of cardinality C is completed at the end of the C -th insertion that all its elements are available, rather than at the end of the first such insertion. A structure, s , is said to be *pending* during the N -th insertion if the all elements of s become available at or before the beginning of the N -th insertion and s is completed during the N -th insertion or later. (The x 's in Figure 5.1 denote the structures that are pending during the eighth insertion). A structure of cardinality C will therefore be pending during exactly C insertions.

Figure 5.1. The online binary transform.

To limit the work done in any insertion step, we require that $1/C$ of the work required to build any structure of size C be performed during each of the C steps in which that structure is pending.²¹ We call the resulting transformation the *online*

²¹The exact means by which this is insured are left unspecified. We may modify the static algorithm to include appropriate breakpoints (generally an easier task than totally reworking the algorithm into a dynamic algorithm by *ad hoc* methods), or we could assume that we can determine the required computation time in advance (at negligible cost) and set a hardware interrupt. For our present purposes, we will assume that the ability to partition the compute time of a call to Insert is available by magic. It should also be noted that the partitioning of the work into equal parts will not be exact in practice; this will lead to slightly greater insertion times than those we are about to advertise.

binary transformation. Analysis of this transform's performance yields the following theorem.

Theorem 5.1: (The on-line binary transformation)

Suppose we are given a static structure, S , for a decomposable problem such that

- (1) $Q_S(N)$ is monotone non-decreasing,
- (2) A structure of cardinality N may be built by N calls, each of cost $I_S(N)$ (recall that $I_S(N)$ is defined as $P_S(N)/N$),
- (3) $I_S(N)$ is monotone non-decreasing,
- (4) The space used at any point during the formation of a static structure is at most $S_S(N)$, and
- (5) $S_S(N)$ grows at least linearly.

Then, there exists a dynamic structure, D , such that

$$\begin{aligned} Q_D(N) &\leq 2\lceil \lg(N+1) \rceil Q_S(N), \\ I_D(N) &\leq \lceil \lg N \rceil I_S(N), \text{ and} \\ S_D(N) &\leq 3S_S(N) \end{aligned}$$

(recall that $I_D(N)$ is the worst-case time to insert the N -th element in a dynamic structure).

Proof:

By assumption (2), application of the online binary transform is well-defined. We will now show that the resulting dynamic algorithm has the stated performance. We first note that all structures which are either active (completed but not yet discarded) after the N -th insertion or pending during the N -th insertion have cardinalities which are exact powers of two and which are $\leq N$. Moreover, there are never more than two active structures of any given cardinality. This and assumption (1) justify the claim about Q_D . Similarly, assumption (3) and the fact that there is never more than one pending structure of any cardinality together justify the claim about I_D . Finally, we note that the sum of the cardinalities of all structures active and pending after the N -th insertion is no more than $3N$ (N for the active structures and no more

than $2N$ for the pending structures). Together with assumptions (4) and (5), this fact justifies the claim about S_D . QED.

To illustrate the application of the online binary transformation, we will consider the problem of d -dimensional maxima searching. A vector is said to be maximal with respect to a set of vectors if no vector in the set is greater than the given vector in all coordinates. Preparata [1978] has given a data structure SMS for d -dimensional maxima searching with performances

$$\begin{aligned} P_{SMS}(N) &= O(N \lg^{d-2} N), \\ S_{SMS}(N) &= O(N \lg^{d-2} N), \text{ and} \\ Q_{SMS}(N) &= O(\lg^{d-2} N), \end{aligned}$$

for any $d \geq 3$. Applying the online binary transform to this structure yields the following.

New Data Structure 5: (Dynamic Maxima Searching)

For any fixed $d \geq 3$ there exists a dynamic data structure DMS for d -dimensional maxima searching with performance

$$\begin{aligned} I_{DMS}(N) &= O(\lg^{d-1} N), \\ O_{DMS}(N) &= O(\lg^{d-1} N), \text{ and} \\ S_{DMS}(N) &= O(N \lg^{d-2} N). \end{aligned}$$

This structure has the same performance as Lueker's [1979], but is substantially easier to code and prove correct; his structure, however, also supports deletions. (The two structures were discovered independently.)

The other transforms we have studied may also be modified to give online versions, as shown by the examples in Figure 5.2. The *online triangular transform*, shown in Figure 5.2(a), gives the performance

$$\begin{aligned} I_D(N) &\leq (2N)^{1/2} I_S(N), \\ O_D(N) &\leq 3O_S(N), \text{ and} \\ S_D(N) &\leq 2S_S(N). \end{aligned}$$

Similarly, the *online dual triangular transform*, shown in Figure 5.2(b), achieves

Figure 5.2. Online triangular transforms.

$$I_D(N) \leq 2I_S(N),$$

$$Q_D(N) \leq 3(2N)^{1/2}Q_S(N), \text{ and}$$

$$S_D(N) \sim S_S(N).$$

Determination of good lower bounds for the penalty factors associated with online transformations remains an open problem.

6. Deletion

So far in this paper we have considered dynamic data structures that support only insertions and queries. In this section we will present two results dealing with data structures that support deletions and their realization by decomposable transforms. In Subsection 6.1 we present a negative result that says that, in general, it is impossible to achieve by a transform a data structure that efficiently supports deletions. In Subsection 6.2 we will examine a transformation that efficiently achieves deletion, but is applicable only to a subset of the decomposable searching problems.

6.1. A Lower Bound

In this subsection we will study a lower bound on the efficiency of performing deletion in a structure achieved by a decomposable transformation. As with all lower bound proofs, it is important that we accurately define our model of computation, which is very similar to that used in Section 4. We assume that there is a static structure S with operations Build and Query, which have performances P_S and Q_S , respectively. The function P_S grows at least linearly, and Q_S is positive and monotone nondecreasing. There is no way to answer a query other than by using the Query subroutine (on a structure built by Build) and the $\square$ operator. The only costs that we will count are those of P_S , Q_S , and a constant cost for computing $\square$.

To state the lower bound precisely, we need some definitions. For a dynamic structure with deletions (which we call DD) we will define the functions $I_{DD}^*(N)$, $D_{DD}^*(N)$, and $Q_{DD}^*(N)$ for the insertion, deletion and query costs, respectively. To strengthen our result, we let these costs denote not the worst-case times, but rather the average cost (over a distribution that we will make precise in the proof of the theorem). We are now ready to state and prove the primary theorem of this subsection.

Theorem 6.1: (Expense of deletion)

For any dynamic structure with deletions (which we call DD) obtained by a transformation applicable to all decomposable searching problems, there exists a sequence of insertions, deletions and queries for which

$$[Q_{DD}^*(N)] \cdot [I_{DD}^*(N) + D_{DD}^*(N) + Q_{DD}^*(N)] = \Omega(N).$$

Note that this implies that at least one of the insertion, deletion and query costs requires at least $\Omega(N^{1/2})$ time.

Proof:

We will prove this theorem by considering a "steady state" in which there is a structure of size N , and a sufficiently long string of repeated query, delete, and insert operations is performed. After M repetitions of these operations,

the structure will still be of size N , and a total of M queries will have been performed. Each query that is performed must examine some collection of static structures whose total size is at least N (so that each element of the set is represented in the query); assume that $C^*(N)$ such structures are examined on the average. We therefore know that at least half the queries examine no more than $2C^*(N)$ static structures each (if more were examined, then the average would be too high), and in these cases the largest structure examined must contain at least $N/(2C^*(N))$ elements.

Consider now an adversary who causes each deletion in the sequence to be deleted from the largest existing static structure -- because of our model of computation, this structure must now be discarded. For sufficiently long sequences of operations, static structures must be created as often as they are deleted. The costs of building the static structure must therefore be paid in insertion, deletion, and query costs, yielding

$$I_{DD}^*(N) + D_{DD}^*(N) + Q_{DD}^*(N) \geq (1/2) P_S(N/2C^*(N)).$$

(The right hand side is from the fact that at least one-half of the queries access a structure of size $N/2C^*(N)$, and the adversary always deletes that structure.) We also know that

$$Q_{DD}^*(N) = \Omega(C^*(N)),$$

because each structure queried costs at least some constant. Multiplying these two inequalities yields

$$\begin{aligned} [Q_{DD}^*(N)] \cdot [I_{DD}^*(N) + D_{DD}^*(N) + Q_{DD}^*(N)] \\ &= \Omega(C^*(N) \cdot P_S(N/2C^*(N))) \\ &= \Omega(P_S(N)) \\ &= \Omega(N). \end{aligned}$$

The last two inequalities both follow from the fact that P_S grows at least linearly. QED.

Maurer and Ottmann [1979] describe a static-to-dynamic transformation with deletion that comes close to achieving this lower bound by always keeping approximately $N^{1/2}$ static structures, each of size approximately $N^{1/2}$. Fortunately, however, additional information can often be used to achieve more rapid deletion outside the model for which this lower bound holds. (Any such transform,

however, is not applicable to all decomposable searching problems.)

6.2. A Fast Special Case

Theorem 6.1 shows that any search for an efficient deletion transformation *for all decomposable searching problems* must be in vain. In this section we will see a transformation that does in fact efficiently support deletions as well as insertions, but is not applicable to all decomposable searching problems. We will investigate this transform by first studying a particular example, and then turn to the general case.

The particular problem that we will study is that of *counting* the number of times a given element occurs in a multiset. A suitable static structure for this problem is the sorted array, which we discussed in Subsection 3.1; it has performances

$$\begin{aligned}P_{SA}(N) &= O(N \lg N), \\S_{SA}(N) &= O(N), \text{ and} \\Q_{SA}(N) &= O(\lg N).\end{aligned}$$

We saw in that subsection that this structure can be transformed to yield the binomial list data structure that efficiently supports both insertions and member queries. It is a trivial modification to have it support count queries as well; the $\square$ operator is now plus rather than or.

Binomial lists can be modified to support deletion by keeping two binomial lists at all times, which we will call the *real* and the *ghost* structures. Each time an element is inserted, it is inserted into the real structure. When an element is deleted, we insert it into the ghost structure. To count the number of times an element occurs in the set, we count the number of times it occurs in the real structure and subtract from that the number of times it occurs in the ghost structure. We maintain the further invariant that the ghost structure always holds less than half as many elements as the real structure; when deletion of an element violates this invariant we destroy the ghost structure, unbuild the set of elements in the real structure and subtract all deleted elements from it, and finally rebuild that set into a new real structure (giving an empty ghost structure).

We must now analyze the performance of binomial lists with deletions. The cost of inserting an element and of performing a count search remain the same; they are respectively $O(\lg N)$ and $O(\lg^2 N)$. The "immediate" cost of deleting an element is $O(\lg N)$ (for performing the insertion into the ghost structure); we must also count, however, the cost of rebuilding the structure. The cost of rebuilding an $M/2$ -element real structure is incurred only after $M/2$ elements have been deleted; since the total cost is $O(M \lg M)$, we can assign each element a share proportional to $\lg M$. Thus the cost of deletion in an N -element set can be amortized to $O(\lg N)$.

The strategy of using real and ghost structures can be generalized to give a dynamic structure supporting deletions for any decomposable searching problem whose $\square$ operator has an inverse. The most common case is when $\square$ is plus, for which $\square^{-1}$ is minus. If $\square$ is and or or, then one can often transform the problem to involve plus instead (for instance, we could transform member queries to count queries, whose $\square$ operator is invertible). If $\square$ is multiset union, then this scheme works only when the size of the answer set for the ghost structure is much smaller than the size of the total answer set (and this is often not the case). Finally, if $\square$ is min or max, this scheme is usually impossible to apply.

To describe the strategy more precisely we will need some notation to describe the efficiency of structures with deletions. If DD is a dynamic structure supporting deletions, we let $P_{DD}(M, N)$ denote the total insertion cost involved in a sequence of N insertions and M deletions in an initially empty structure. The function $Q_{DD}(M, N)$ denotes the cost of answering a query in a structure built by N insertions and M deletions. Finally, $D_{DD}(M, N)$ denotes the total time spent in processing deletions in a series of N insertions and M deletions, and $S_{DD}(M, N)$ denotes the maximum space required by the structure during the sequence. With this background we can describe the transformation supporting deletions precisely in the following theorem.

Theorem 6.2: (Transformations supporting deletions)

Assume that there exists an admissible $(F(N), G(N))$ transformation. Then, given any static structure S for a decomposable searching problem P such that

the inverse of the $\square$ operator for P is computable in constant time, it is possible to achieve a new structure DD with performances

$$\begin{aligned} S_{DD}(M,N) &\leq S_S(2(N-M)) + S_S(N-M), \\ P_{DD}(M,N) &\leq G(N) \cdot P_S(N), \\ Q_{DD}(M,N) &\leq F(2(N-M)) \cdot Q_S(2(N-M)) + F(N-M) \cdot Q_S(N-M), \text{ and} \\ D_{DD}(M,N) &\leq G(M) \cdot P_S(M) + P_S(2M). \end{aligned}$$

We assume here that Q_S is monotone nondecreasing and that both P_S and S_S grow at least linearly.

Proof:

The DD structure maintains two dynamic structures (each achieved by applying the admissible $(F(N), G(N))$ transform to S): the real structure and the ghost structure. Both structures are initially empty. To insert a new element into DD , insert it into the real structure. To answer a query, answer it on the real structure and subtract from that the answer on the ghost structure (using $\square^{-1}$). To delete an element, insert it into the ghost structure. If the ghost structure ever becomes half the size of the real structure, rebuild the real structure with only undeleted elements, and discard the current ghost structure.

The storage requirements of DD follow immediately from the superlinear growth of S_S . If a total of N insertions and M deletions have been performed, then at most $N-M$ elements are "really" stored in the structure. The ghost structure can therefore contain at most $N-M$ elements, and the real structure contains at most twice that number. The time spent on insertion is straightforward, and so is the query time. The time spent on deletion is at most that for inserting M elements into the ghost structure and then rebuilding the real structure; this latter action is never carried out on more than $2M$ elements. These facts together establish the theorem. QED.

There are two important things to note about the transformation of Theorem 6.2. The first is that it is not online in the sense of Section 5; as it stands, the expense of rebuilding the real structure and discarding the ghost structure must occasionally be paid in a single block of time. The second interesting thing to note is the fact that there is nothing magic about insisting that the ghost structure be one-half the

size of the real structure: we could just as well use any constant A in the range $(0,1)$. For small A , the query time decreases and the storage utilization is higher; for large A , the deletion time decreases.

As an application of this transformation, we will consider the problem of Empirical Cumulative Distribution Function (ECDF) searching in a set of N d -dimensional vectors. One vector is said to *dominate* another if it is greater than it in all components; an ECDF query asks for the number of vectors a given vector dominates. Bentley and Shamos [1977] describe a data structure for d -dimensional ECDF searching (for $d \geq 2$) with performances

$$\begin{aligned} P_{\text{ECDF}}(N) &= O(N \lg^{d-1} N), \\ S_{\text{ECDF}}(N) &= O(N \lg^{d-1} N), \text{ and} \\ Q_{\text{ECDF}}(N) &= O(\lg^d N). \end{aligned}$$

We can apply the binary transform of Section 3.1 and the transform of Theorem 6.2 to their structure to achieve the following.

New Data Structure 6: (Dynamic ECDF Searching)

It is possible to achieve a data structure for dynamic ECDF searching in which performing a sequence of N insertions and deletions requires $O(N \lg^d N)$ time. When containing N elements, the structure requires $O(N \lg^{d-1} N)$ space, and an ECDF query can be answered in $O(\lg^{d+1} N)$ time.

Lueker [1979] later used a different transformation on decomposable searching problems to achieve an (online) structure with performance identical to this, but with a logarithmic factor removed from the query time; his structure is more difficult to code and to prove correct, however.

7. Conclusions

We will now briefly review the contributions of this paper. The subject throughout has been general methods for converting static data structures to dynamic data structures. In Section 3 we saw three distinct classes of transformations, each based on a combinatorial representation of the integers. In Section 4 we saw that many of those transformations are optimal, in a very strong

sense. In Section 5 we considered structures in which each insertion must be handled very quickly; this is important in "online" applications. Our study of dynamic structures up to this point concentrated on structures that supported only insertions and queries; in Section 6 we investigated structures that also support deletions. We saw that although it is impossible to achieve efficient deletions in the general case, they can be achieved for an important subclass of the decomposable searching problems.

The contributions of this paper can be classified on three distinct levels. On the first level are the new data structures that we have seen. Each one is currently the best known structure for its task (with the exception of New Data Structure 6), and *each was discovered by conscious application of the transforms described in this paper*. On a second level are the transformations themselves; they are very interesting from a combinatorial viewpoint, and provide a useful addition to the algorithm designer's tool bag. On the third and final level is the new kind of result represented by the transformations: they are not just a single solution to a single problem, but rather a set of solutions to a broad class of problems.

Acknowledgements

The helpful comments of Kevin Brown, Michael Shamos and Andrew and Frances Yao are gratefully acknowledged.

References

- Bentley, J. L., D. Detig, L. Guibas and J. B. Saxe [1979]. An optimal data structure for minimal-storage dynamic member searching, in preparation.
- Bentley, J. L. and H. A. Maurer [1978]. "Efficient worst-case data structures for range searching," in Carnegie-Mellon Computer Science Report CMU-CS-78-136.
- Bentley, J. L. and J. B. Saxe [1979]. "Decomposable searching problems," in preparation. (Preliminary version by J. L. Bentley, *Information Processing Letters* 8, 5 (June 1979), pp. 244-251.)

- Bentley, J. L. and M. I. Shamos [1977]. "A problem in multivariate statistics: algorithm, data structure, and applications," *Proceedings of the Fifteenth Allerton Conference on Communication, Control and Computing*, pp. 193-201, September 1977.
- Bentley, J. L. and M. Shaw [1979]. "An Alghard specification of a correct and efficient transformation on data structures," *Specifications of Reliable Software Conference*, pp. 222-237, April 1979, IEEE.
- Dobkin, D. and R. J. Lipton [1976]. "Multidimensional searching problems," *SIAM J. Comp.* 5, 2 (June 1976), pp. 181-186.
- Knuth, D. E. [1968]. *The Art of Computer Programming, vol. 1: Fundamental Algorithms*, Addison-Wesley, Reading, Mass.
- Knuth, D. E. [1973]. *The Art of Computer Programming, vol. 3: Sorting and Searching*, Addison-Wesley, Reading, Mass.
- Lipton, R. J. and R. E. Tarjan [1977]. "Applications of a planar separator theorem," *Eighteenth Symposium on the Foundation of Computer Science*, pp. 162-170, October 1977, IEEE.
- Lueker, G. [1978]. "A data structure for orthogonal range queries," *Nineteenth Symposium on the Foundations of Computer Science*, pp. 28-34, October 1978, IEEE.
- Lueker, G. [1979]. "A transformation for adding range restriction capability to dynamic data structures for decomposable searching problems," UCI Technical Report #129, February 1979, 34 pp.
- Maurer, H. A. and T. Ottmann [1979]. "Dynamic solutions of decomposable searching problems," Report 33, Institut fur Informationsverarbeitung, Technische Universitat Graz, June 1979.
- Munro, J. I. and H. Suwanda [1979]. "Implicit data structures," *Eleventh Symposium on the Theory of Computing*, pp. 108-117, April 1979, ACM.
- Preparata, F. P. [1978]. A new approach to planar point location, University of Illinois Coordinated Science Laboratory Report R-829, September 1978, 24 pp.
- Rivest, R. L. [1976]. "Partial match retrieval algorithms," *SIAM J. Comp.* 5, 1 (March

1976), pp. 19-50.

Shamos, M. I. [1978]. Computational Geometry, unpublished Ph.D. Thesis, Yale University, New Haven, Connecticut.

Vuillemin, J. [1978]. "A data structure for manipulating priority queues," *Comm. of the ACM* 21, 4, April 1978, pp. 309-315.

Willard, D. [1978]. Predicate-oriented Database Search Algorithms, Harvard Aiken Computation Laboratory Report TR-20-78.

I. A List of Decomposable Searching Problems

Throughout the body of this paper we have examined a number of operations on decomposable searching problems. In this appendix we will list some (twenty-three) searching problems that have the property of decomposability. For each problem we will note its $\square$ operator in square brackets.

The most common kind of searching problems are those defined on totally-ordered sets. We already saw that Member searching (which asks "is x an element of F ?") is decomposable with $\square$ operator V . Other examples are Successor (what is the least element in F greater than x ?) [$\min$], Predecessor [$\max$], Rank (how many elements in F are less than x ?) [$+$], and Count (how many elements in multiset F have value x ?) [$+$]. Two queries on ordered sets that have no query element are the priority queue operations Min [$\min$] and Max [$\max$]. These problems, their applications, and data structures for their solutions are discussed in depth by Knuth [1973].

Many of the problems that arise in database applications are decomposable. In this context, the set of elements is usually a file of records, each of which contains certain keys. An Exact Match query calls for a list of all records that have all keys equal to specified values [U]. A Partial Match query asks for all records that match on some subset of the keys [U]. Range queries ask for all records that have each key in a specified range of values [U]. Intersection queries specify a subset of the key space and ask for a list of all records in that subset (thus asking for the

intersection of the query space and the record set) [U]. Finally, Best Match queries specify an "ideal" record and a distance function (often the Hamming distance), and ask for the record in the set closest to the ideal [min]. These queries and data structures for answering them are discussed by Rivest [1976].

We saw in the body of the paper two decomposable searching problems that arise in statistics. Both of the problems are defined in terms of vectors domination (one vector is said to dominate another if it is greater in all coordinates). A Maxima query asks whether the query vector is dominated by any in the set [V]. The Empirical Cumulative Distribution Function (ECDF) query asks how many vectors a given vector dominates [+].

Examples of decomposable searching problems abound in computational geometry. Many queries are asked of sets of points in the plane or Euclidean k-space, including Nearest Neighbor (which point in the set is nearest the query point?) [min], Furthest Neighbor [max], and Near Neighbor (list all points within distance d of the query point) [U] queries. Other queries deal with more complicated objects. For example, we might wish to know whether a given point is in the intersection of a set of half-planes (this problem arises in linear programming); Feasible Region queries are decomposable with the $\wedge$ operator. Other queries include Rectangle Intersection (what rectangles in the set does this rectangle intersect?) [U] and Circle Intersection [U]. These queries and many others have been discussed in detail by Shamos [1978]. Dobkin and Lipton [1976] investigate a number of decomposable searching problems in multidimensional space; these include such queries as "is this point on any of the lines" [V] and "is this point on any of the hyperplanes" [V]. Many of the other problems that we have already mentioned can be cast in geometric terms; these include ECDF, Maxima and Range searching.

Convex Hull searching is a very interesting problem from the viewpoint of decomposability. In its simplest form--"Is point x within the convex hull of point set F?"--it is simple to prove that it is not decomposable, since whenever F contains at least two points we can partition F and specify x so that x is not in the hull of

either part but either is or is not in the hull of the union. If we ask instead the query "what does the hull of the set look like from here?" (the answer being either an assertion that the query point is within the hull or a pair of angles giving the extremal points of the hull as "viewed" from the query point), the problem is now decomposable. The transforms described in this paper are therefore applicable to any data structure for Convex Hull searching, provided that that structure can be cheaply modified to answer the more complicated query. While this result is not of particular interest in itself (since it is easy to develop a fast *ad hoc* algorithm for dynamic Convex Hull searching), it indicates a possibly fruitful technique for extending the domain of applicability of the transforms, namely the identification of any searching problem P such that (1) P may be made decomposable by having the query provide some extra information and (2) known static algorithms for P can be altered to yield that extra information at low cost. The identification of other such "pseudo-decomposable" problems (and other decomposable problems in general) remains a open problem.

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER CMU-CS-79-141	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) TRANSFORMING STRATIC DATA STRUCTURES TO DYNAMIC STRUCTURES		5. TYPE OF REPORT & PERIOD COVERED Interim
		6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s) JAMES B. SAXE + JON L. BENTELY		8. CONTRACT OR GRANT NUMBER(s) N00014-76-C-0370
9. PERFORMING ORGANIZATION NAME AND ADDRESS Carnegie-Mellon University Computer Science Department Pittsburgh, PA 15213		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS
11. CONTROLLING OFFICE NAME AND ADDRESS Office of Naval Research Arlington, VA 22217		12. REPORT DATE September 1979
		13. NUMBER OF PAGES 69
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)		15. SECURITY CLASS. (of this report) UNCLASSIFIED
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE
16. DISTRIBUTION STATEMENT (of this Report) Approved for public release; distribution unlimited		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)		
18. SUPPLEMENTARY NOTES		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number)		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number)		

DD FORM 1 JAN 73 1473

EDITION OF 1 NOV 65 IS OBSOLETE
S/N 0102-014-6601

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)