

*You're older than you've ever been and now you're even older
 And now you're even older
 And now you're even older
 You're older than you've ever been and now you're even older
 And now you're older still*

— They Might be Giants, "Older", *Mink Car* (1999)

1 Static-to-Dynamic Transformations

A **search problem** is abstractly specified by a function of the form $Q: \mathcal{X} \times 2^{\mathcal{D}} \rightarrow \mathcal{A}$, where \mathcal{D} is a (typically infinite) set of *data objects*, \mathcal{X} is a (typically infinite) set of *query objects*, and \mathcal{A} is a set of valid *answers*. A **data structure** for a search problem is a method for storing an arbitrary finite data set $D \subseteq \mathcal{D}$, so that given an arbitrary query object $x \in \mathcal{X}$, we can compute $Q(x, D)$ quickly. A *static* data structure only answers queries; a *dynamic* data structure also allows us to modify the data set by inserting or deleting individual items.

A search problem is **decomposable** if, for any pair of disjoint data sets D and D' , the answer to a query over $D \cup D'$ can be computed in constant time from the answers to queries over the individual sets; that is,

$$Q(x, D \cup D') = Q(x, D) \diamond Q(x, D')$$

for some commutative and associative binary function $\diamond: \mathcal{A} \times \mathcal{A} \rightarrow \mathcal{A}$ that can be computed in $O(1)$ time. I'll use \perp to denote the answer to any query $Q(x, \emptyset)$ over the empty set, so that $a \diamond \perp = \perp \diamond a = a$ for all $a \in \mathcal{A}$. Simple examples of decomposable search problems include the following.

- **Rectangle counting:** Data objects are points in the plane; query objects are rectangles; a query asks for the number of points in a given rectangle. Here, \mathcal{A} is the set of natural numbers, $\diamond = +$, and $\perp = 0$.
- **Nearest neighbor:** Data objects are points in some metric space; query objects are points in the same metric space; a query asks for the distance from a given query point to the nearest point. Here, \mathcal{A} is the set of positive real numbers, $\diamond = \min$, and $\perp = \infty$.
- **Triangle emptiness:** Data objects are points in the plane; query objects are triangles; a query asks whether any data point lies in a given query triangle. Here, \mathcal{A} is the set of booleans, $\diamond = \vee$, and $\perp = \text{FALSE}$.
- **Interval stabbing:** Data objects are intervals on the real line; query objects are points on the real line; a query asks for the subset of data intervals that contain a given query point. Here, \mathcal{A} is the set of all finite sets of real intervals, $\diamond = \cup$, and $\perp = \emptyset$.

1.1 Insertions Only (Bentley and Saxe* [3])

First, let's describe a general transformation that adds the ability to insert new data objects into a static data structure, originally due to Jon Bentley and his PhD student James Saxe* [3]. Suppose we have a static data structure that can store any set of n data objects in space $S(n)$, after $P(n)$ preprocessing time, and answer a query in time $Q(n)$. We will construct a new data structure with size $S'(n) = O(S(n))$, preprocessing time $P'(n) = O(P(n))$, query time $Q'(n) = O(\log n) \cdot Q(n)$, and *amortized* insertion time $I'(n) = O(\log n) \cdot P(n)/n$. In the next section, we will see how to achieve this insertion time even in the worst case.

Our data structure consists of $\ell = \lceil \lg n \rceil$ levels $L_0, L_1, \dots, L_{\ell-1}$. Each level L_i is either empty or a static data structure storing exactly 2^i items. Observe that for any value of n , there is a unique set of levels that must be non-empty. To answer a query, we perform a query in each non-empty level and combine the results. (This is where we require the assumption that the queries are decomposable.)

```

NEWQUERY(x):
  ans ← ⊥
  for i ← 0 to ℓ - 1
    if  $L_i \neq \emptyset$ 
      ans ← ans ◊ QUERY(x,  $L_i$ )
  return ans

```

The total query time is clearly at most $\sum_{i=0}^{\ell-1} Q(2^i) < \ell Q(n) = O(\log n) \cdot Q(n)$, as claimed. Moreover, if $Q(n) > n^\varepsilon$ for any $\varepsilon > 0$, the query time is actually $O(Q(n))$.

The insertion algorithm exactly mirrors the algorithm for incrementing a binary counter, where the presence or absence of each L_i plays the role of the i th least significant bit. We find the smallest empty level k ; build a new data structure L_k containing the new item and all the items stored in L_0, L_1, \dots, L_{k-1} ; and finally discard all the levels smaller than L_k . See Figure 1 for an example.

```

INSERT(x):
  Find minimum  $k$  such that  $L_k = \emptyset$ 
   $L_k \leftarrow \{x\} \cup \bigcup_{i < k} L_i$    $\llcorner$ takes  $P(2^k)$  time $\lrcorner$ 
  for i ← 0 to  $k - 1$ 
    destroy  $L_i$ 

```

During the lifetime of the data structure, each item will take part in the construction of $\lg n$ different data structures. Thus, if we charge

$$I'(n) = \sum_{i=0}^{\lg n} \frac{P(2^i)}{2^i} = O(\log n) \frac{P(n)}{n}.$$

for each insertion, the total charge will pay for the cost of building all the static data structures. If $P(n) > n^{1+\varepsilon}$ for any $\varepsilon > 0$, the amortized insertion time is actually $O(P(n)/n)$.

1.2 Lazy Rebuilding (Overmars* and van Leeuwen [5, 4])

We can modify this general transformation to achieve the same space, preprocessing, and query time bounds, but now with *worst-case* insertion time $I'(n) = O(\log n) \cdot P(n)/n$. Obviously we cannot get fast updates in the worst case if we are ever required to build a large data structure all at once. The key idea is to stretch the construction time out over several insertions.

As in the amortized structure, we maintain $\ell = \lceil \lg n \rceil$ levels, but now each level i consists of four static data structures, called *Oldest* _{i} , *Older* _{i} , *Old* _{i} , and *New* _{i} . Each of the ‘old’ data structures is either empty or contains exactly 2^i items; moreover, if *Oldest* _{i} is empty then so is *Older* _{i} , and if *Older* _{i} is empty then so is *Old* _{i} . The fourth data structure *New* _{i} is either empty or a partially built structure that will eventually contain 2^i items. Every item is stored in exactly one ‘old’ data structure (at exactly one level) and *at most* one ‘new’ data structure.

The query algorithm is almost unchanged.

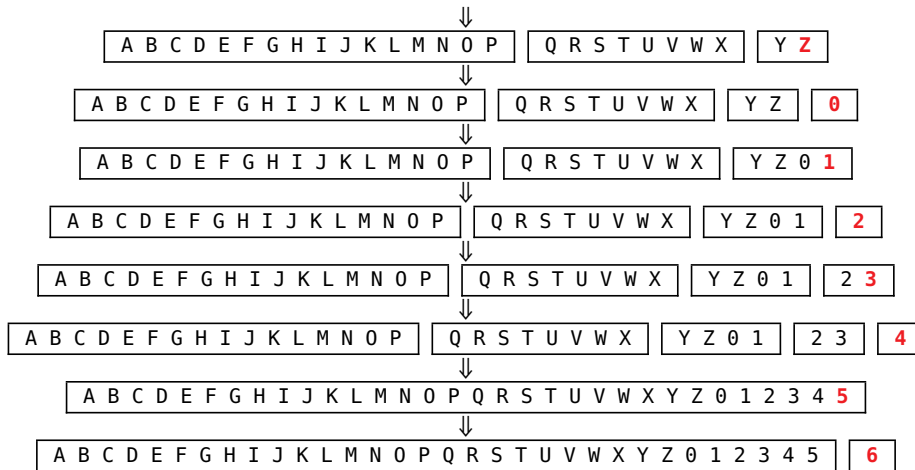


Figure 1. The 27th through 33rd insertions into a Bentley/Saxe data structure

```

NEWQUERY(x):
  ans ← ⊥
  for i ← 0 to ℓ - 1
    if Oldesti ≠ ∅
      ans ← ans ◊ QUERY(x, Oldesti)
    if Olderi ≠ ∅
      ans ← ans ◊ QUERY(x, Olderi)
    if Oldi ≠ ∅
      ans ← ans ◊ QUERY(x, Oldi)
  return ans
    
```

As before, the new query time is $O(\log n) \cdot Q(n)$, or $O(Q(n))$ if $Q(n) > n^\epsilon$.

The insertion algorithm passes through the levels from largest to smallest. At each level i , if both $Oldest_{i-1}$ and $Older_{i-1}$ happen to be non-empty, we execute $P(2^i)/2^i$ steps of the algorithm to construct New_i from $Oldest_{i-1} \cup Older_{i-1}$. Once New_i is completely built, we move it to the oldest available slot on level i , delete $Oldest_{i-1}$ and $Older_{i-1}$, and rename Old_{i-1} to $Oldest_{i-1}$. Finally, we create a singleton structure at level 0 that contains the new item.

```

AGE(i):
  if Oldesti = ∅
    Oldesti ← Newi
  else if Olderi = ∅
    Olderi ← Newi
  else
    Oldi ← Newi
  Newi ← ∅

LAZYINSERTION(x):
  for i ← ℓ - 1 down to 1
    if Oldesti-1 ≠ ∅ and Olderi-1 ≠ ∅
      spend  $P(2^i)/2^i$  time executing  $New_i \leftarrow Oldest_{i-1} \cup Older_{i-1}$ 
      if  $New_i$  is complete
        destroy  $Oldest_{i-1}$  and  $Older_{i-1}$ 
         $Oldest_{i-1} \leftarrow Old_{i-1}$ 
         $Old_{i-1} \leftarrow \emptyset$ 
    AGE(i)
  New0 ← {x}
  AGE(0)
    
```

Each insertion clearly takes $\sum_{i=0}^{\ell-1} O(P(2^i)/2^i) = O(\log n) \cdot P(n)/n$ time, or $O(P(n)/n)$ time if $P(n) > n^{1+\epsilon}$ for any $\epsilon > 0$. The only thing left to check is that the algorithm actually works! Specifically, how do we know that Old_i is empty whenever we call $AGE(i)$? The key insight is that the modified insertion algorithm mirrors the standard algorithm to increment a non-standard binary counter, where every bit is either 2 or 3, except the most significant bit, which might

be 1. It's not hard to prove by induction that this representation is unique; the correctness of the insertion algorithm follows immediately. Specifically, $\text{AGE}(i)$ is called on the n th insertion—or in other words, the i th 'bit' is incremented—if and only if $n = k \cdot 2^i - 2$ for some integer $k \geq 3$.

Figure 2 shows the modified insertion algorithm in action.

Exercise 1. *Suppose we increase the time in the third line of `LAZYINSERTION` from $P(2^i)/2^i$ to $P(2^i)/2^{i-1}$. Prove that the modified insertion algorithm is still correct, and that if we start with an empty data structure, every component Old_i is **always** empty.*

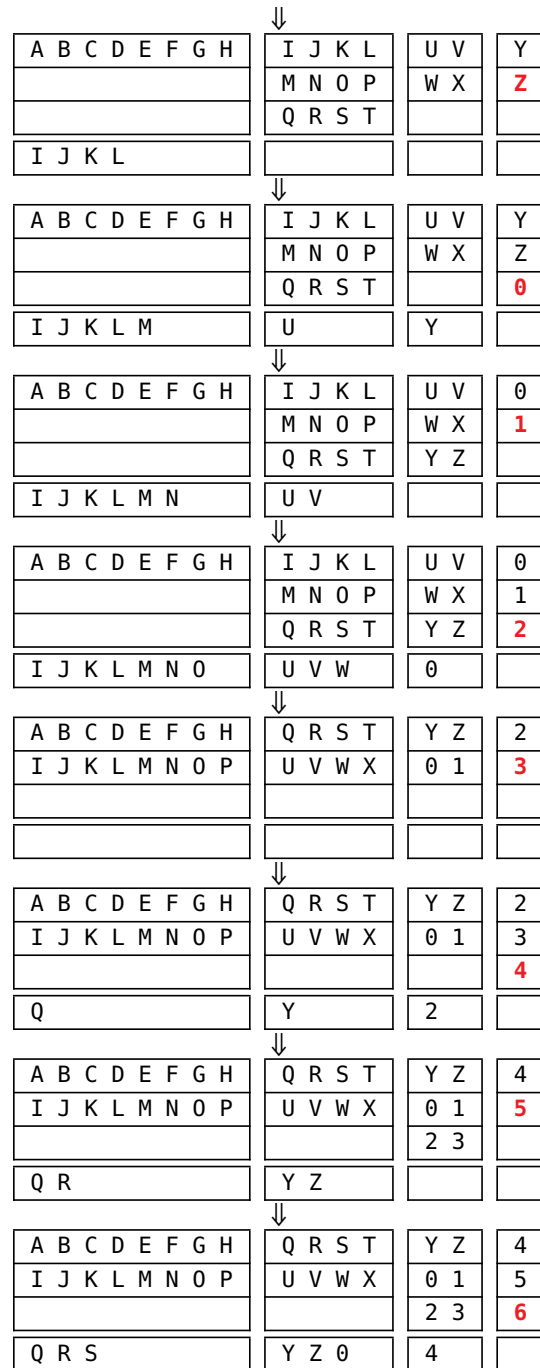


Figure 2. The 27th through 33rd insertions into a Overmars/van Leeuwen data structure

1.3 Deletions via (Lazy) Global Rebuilding: The Invertible Case

Under certain conditions, we can modify the logarithmic method to support deletions as well as insertions, by periodically rebuilding the entire data structure.

Perhaps the simplest case is when the binary operation \diamond used to combine queries has an inverse $\bar{\diamond}$; for example, if $\diamond = +$ then $\bar{\diamond} = -$. In this case, we maintain two insertion-only data structures, a *main* structure M and a *ghost* structure G , with the invariant that every item in G also appears in M . To insert an item, we insert it into M . To delete an item, we *insert* it into G . Finally, to answer a query, we compute $Q(x, M) \bar{\diamond} Q(x, G)$.

The only problem with this approach is that the two component structures M and G might become much larger than the ideal structure storing $M \setminus G$, in which case our query and insertion times become inflated. To avoid this problem, we rebuild our entire data structure from scratch—building a new main structure containing $M \setminus G$ and a new empty ghost structure—whenever the size of G exceeds half the size of M . Rebuilding requires $O(P(n))$ time, where n is the number of items in the new structure. After a global rebuild, there must be at least $n/2$ deletions before the next global rebuild. Thus, the total amortized time for each deletion is $O(P(n)/n)$ plus the cost of insertion, which is $O(P(n) \log n/n)$.

There is one minor technical point to consider here. Our earlier amortized analysis of insertions relied on the fact that large local rebuilds are always far apart. Global rebuilding destroys that assumption. In particular, suppose M has $2^k - 1$ elements and G has $2^{k-1} - 1$ elements, and we perform four operations: insert, delete, delete, insert.

- The first insertion causes us to rebuild M completely.
- The first deletion causes us to rebuild G completely.
- The second deletion triggers a global rebuild. The new M contains $2^{k-1} - 1$ items.
- Finally, the second insertion causes us to rebuild M completely.

Another way to state the problem is that a global rebuild can put us into a state where we don't have enough insertion credits to pay for a local rebuild. To solve this problem, we simply scale the amortized cost of deletions by a constant factor. When a global rebuild is triggered, a fraction of this pays for the global rebuild itself, and the rest of the credit pays for the first local rebuild at each level of the new main structure, since $\sum_{i=0}^{\lg n} P(2^i) = O(P(n))$.

We can achieve the same deletion time in the worst case by performing the global rebuild lazily. Now we maintain three structures: a static main structure M , an *insertion* structure I , and a ghost structure G . Most of the time, we insert new items into I , delete items by inserting them into G , and evaluate queries by computing $Q(x, M) \diamond Q(x, I) \bar{\diamond} Q(x, G)$.

However, when $|G| > (|M| + |I|)/2$, we freeze M and G and start building three new structures M' , I' , and G' . Initially, all three new structures are empty. Newly inserted items go into the new insertion structure I' ; newly deleted items go into the new ghost structure G' . To answer a query, we compute $Q(x, M) \diamond Q(x, I) \diamond Q(x, I') \bar{\diamond} Q(x, G) \bar{\diamond} Q(x, G')$. After every deletion (that is, after every insertion into the new ghost structure G'), we spend $8P(n)/n$ time building the new main structure M' from the set $(I \cup M) \setminus G$. After $n/8$ deletions, the new static structure is complete; we destroy the old structures I, M, G , and revert back to our normal state of affairs. The exact constant 8 is unimportant, it only needs to be large enough that the new main structure M' is complete before the start of the next global rebuild.

With lazy global rebuilding, the *worst-case* time for a deletion is $O(P(n) \log n/n)$, exactly the same as insertion. Again, if $P(n) = \Omega(n^{1+\epsilon})$, the deletion time is actually $O(P(n)/n)$.




Figure 3. A high-level view of the deletion structure for invertible queries, during a lazy global rebuild.

1.4 Deletions for Non-Invertible Queries

To support both insertions and deletions when the function \diamond has no inverse, we have to assume that the base structure already supports *weak deletions* in time $D(n)$. A weak deletion is *functionally* exactly the same as a regular deletion, but it doesn't have the same effect on the *cost* of future queries. Specifically, we require that the cost of a query after a weak deletion is no higher than the cost of a query before the weak deletion. Weak deletions are a fairly mild requirement; many data structures can be modified to support them with little effort. For example, to weakly delete an item x from a binary search tree being used for simple membership queries (Is x in the set?), we simply mark all occurrences of x in the data structure. Future membership queries for x would find it, but would also find the mark(s) and thus return FALSE.

If we are satisfied with amortized time bounds, adding insertions to a weak-deletion data structure is easy. As before, we maintain a sequence of levels, where each level is either empty or a base structure. For purposes of insertion, we pretend that any non-empty level L_i has size 2^i , even though the structure may actually be smaller. To delete an item, we first determine which level contains it, and then weakly delete it from that level. To make the first step possible, we also maintain an auxiliary dictionary (for example, a hash table) that stores a list of pointers to occurrences of each item in the main data structure. The insertion algorithm is essentially unchanged, except for the (small) additional cost of updating this dictionary. When the total number of undeleted items is less than half of the total 'size' of the non-empty levels, we rebuild everything from scratch. The amortized cost of an insertion is $O(P(n) \log n/n)$, and the amortized cost of a deletion is $O(P(n)/n + D(n))$.

Once again, we can achieve the same time bounds in the worst case by spreading out both local and global rebuilding. I'll first describe the high-level architecture of the data structure and discuss how weak deletions are transformed into regular deletions, and then spell out the lower-level details for the insertion algorithm.

1.4.1 Transforming Weak Deletions into Real Deletions

For the moment, assume that we already have a data structure that supports insertions in $I(n)$ time and *weak* deletions in $D(n)$ time. A good example of such a data structure is the *weight-balance B-tree* defined by Arge and Vitter [1].

Our global data structure has two major components; a main structure M and a *shadow copy* S . Queries are answered by querying the main structure M . Under normal circumstances, insertions and deletions are made directly in both structures. When more than half of the elements of S have been weakly deleted, we trigger a global rebuild. At that point, we freeze \bar{S} and begin building two new clean structures M' and S' . The reason for the shadow structure is that we cannot copy from S while it is undergoing other updates.

During a global rebuild, our data structure has four component structures M , S , M' , and S'




Figure 4. A high-level view of the deletion structure for non-invertible queries, during a lazy global rebuild

and an *update queue* U , illustrated above. Queries are evaluated by querying the main structure M as usual. Insertions and (weak) deletions are processed directly in M . However, rather than handling them directly in the shadow structure S (which is being copied) or the new structures M' and S' (which are not completely constructed), all updates are inserted into the update queue U .

M' and S' are incrementally constructed in two phases. In the first phase, we build new data structures containing the elements of S . In the second phase, we execute the stream of insertions and deletions that have been stored in the update queue U , in both M' and S' , in the order they were inserted into U . In each phase, we spend $O(I(n))$ steps on the construction for each insertion, and $O(P(n)/n + D(n))$ steps for each deletion, where the hidden constants are large enough to guarantee that each global rebuild is complete well before the next global rebuild is triggered. In particular, in the second rebuild phase, each time an update is inserted into U , we must process and remove at least two updates from U . When the update queue is empty, the new data structures M' and S' are complete, so we destroy the old structures M and S and revert to 'normal' operation.

1.4.2 Adding Insertions to a Weak-Deletion-Only Structure

Now suppose our given data structure does not support insertions or deletions, but does support weak deletions in $D(n)$ time. A good example of such a data structure is the *kd-tree*, originally developed by Bentley [2].

To add support for insertions, we modify the lazy logarithmic method. As before, our main structure consists of $\lg n$ levels, but now each level consists of *eight* base structures $New_i, Old_i, Older_i, Oldest_i, SNew_i, SOLD_i, SOlder_i, SOLdest_i$, as well as an deletion queue D_i . We also maintain an auxiliary dictionary recording the level(s) containing each item in the overall structure. As the names suggest, each active structure $SFoo_i$ is a shadow copy of the corresponding active structure Foo_i . Queries are answered by examining the active old structures. New_i and its shadow copy $SNew_i$ are incrementally constructed from the shadows $SOlder_{i-1}$ and $SOLdest_{i-1}$ and from the deletion queue D_i . Deletions are performed directly in the active old structures and in the shadows that are *not* involved in rebuilds, and are inserted into deletion queues at levels that are being rebuilt. At each insertion, if level i is being rebuilt, we spend $O(P(2^i)/2^i)$ time on that local rebuilding. Similarly, for each deletion, if the appropriate level i is being rebuilt, we spend $O(D(2^i))$ time on that local rebuilding. The constants in these time bounds are chosen so that each local rebuild finishes well before the next one begins.

Here are the insertion and deletion algorithms in more detail:




Figure 5. One level in our lazy dynamic data structure.

```

AGE(i):
  if Oldesti = ∅
    Olderi ← Newi; SOlderi ← SNewi
  else if Olderi = ∅
    Olderi ← Newi; SOlderi ← SNewi
  else
    Oldi ← Newi; SOLDi ← SNewi
    Newi ← ∅; SNewi ← ∅

```

```

LAZYINSERT(x):
  for i ← ℓ - 1 down to 1
    if Oldesti-1 ≠ ∅ and Olderi-1 ≠ ∅
      spend O(P(2i)/2i) time building Newi and SNewi from SOldesti-1 ∪ SOlderi-1
      if Newi and SNewi are complete
        destroy Oldesti-1, SOldesti-1, Olderi-1, and SOlderi-1
        Oldesti-1 ← Oldi-1; Oldi-1 ← ∅
        SOldesti-1 ← SOLDi-1; SOLDi-1 ← ∅
    else if Di ≠ ∅
      spend O(P(2i)/2i) time processing deletions in Di from Newi and SNewi
    if Di = ∅
      AGE(i)
  New0 ← {x}; SNew0 ← {x}
  AGE(0)

```

```

WEAKDELETE(x):
  find level i containing x
  if x ∈ Oldesti
    WEAKDELETE(x, Oldesti)
    if Olderi ≠ ∅
      Add x to Di+1
      Spend O(D(2i+1)) time building Newi+1 and SNewi+1
    else
      WEAKDELETE(x, SOldesti)
  else if x ∈ Olderi
    WEAKDELETE(x, Olderi)
    Add x to Di+1
    Spend O(D(2i+1)) time building Newi+1 and SNewi+1
  else if x ∈ Oldi
    WEAKDELETE(x, Oldi); WEAKDELETE(x, SOLDi)

```

1.4.3 The Punch Line

Putting both of these constructions together, we obtain the following worst-case bounds. We are given a data structure that the original data structure requires space $S(n)$, can be built in time $P(n)$, answers decomposable search queries in time $Q(n)$, and supports weak deletions in time $D(n)$.

- The entire structure uses $O(S(n))$ space and can be built in $O(P(n))$ time.
- Queries can be answered in time $O(Q(n) \log n)$, or $O(Q(n))$ if $Q(n) > n^\epsilon$ for any $\epsilon > 0$.
- Each insertion takes time $O(P(n) \log n/n)$, or $O(P(n)/n)$ if $P(n) > n^{1+\epsilon}$ for any $\epsilon > 0$.
- Each deletion takes time $O(P(n)/n + D(n) \log n)$, or $O(P(n)/n + D(n))$ if $D(n) > n^\epsilon$ for any $\epsilon > 0$.

References

- [1] L. Arge and J. S. Vitter. Optimal external memory interval management. *SIAM J. Comput.* 32(6):1488–1508, 2003.
- [2] J. L. Bentley. Multidimensional binary search trees used for associative searching. *Communications of the ACM* 18:509–517, 1975.
- [3] J. L. Bentley and J. B. Saxe*. Decomposable searching problems I: Static-to-dynamic transformation. *J. Algorithms* 1(4):301–358, 1980.
- [4] M. H. Overmars*. *The Design of Dynamic Data Structures*. Lecture Notes Comput. Sci. 156. Springer-Verlag, 1983.
- [5] M. H. Overmars* and J. van Leeuwen. Worst-case optimal insertion and deletion methods for decomposable searching problems. *Inform. Process. Lett.* 12:168–173, 1981.

*Starred authors were PhD students at the time that the work was published.