

Splay Trees, Davenport-Schinzel Sequences, and the Deque Conjecture

Seth Pettie
The University of Michigan

Abstract

We introduce a new technique to bound the asymptotic performance of splay trees. The basic idea is to transcribe, in an indirect fashion, the rotations performed by the splay tree as a Davenport-Schinzel sequence, none of whose subsequences are isomorphic to a fixed *forbidden subsequence*. We direct this technique towards Tarjan’s *deque conjecture* and prove that n deque operations take only $O(n\alpha^*(n))$ time, where $\alpha^*(n)$ is the minimum number of applications of the inverse-Ackermann function mapping n to a constant. We are optimistic that this approach could be directed towards other open conjectures on splay trees such as the *traversal* and *split* conjectures.

1 Introduction

Sleator and Tarjan proposed the *splay tree* [29] as a self-adjusting alternative to traditional search trees like red-black trees and AVL-trees. Rather than enforce some type of balance invariant, the splay tree simply adjusts its structure in response to the access pattern by rotating accessed elements towards the root in a prescribed way; see Figure 1. By letting the access pattern influence its own shape, the splay tree can inadvertently learn to perform optimally on a variety of access patterns. For example, the *static optimality* theorem states that splay trees are no worse than any fixed search tree. The *working set*, and *dynamic finger* theorems show that the access time is logarithmic in the distance to the accessed element, where *distance* is either temporal (working set [29]) or with respect to key-space (dynamic finger [9, 8]). Sleator and Tarjan went a step further and conjectured that splay trees are, to within a constant factor, just as efficient as any dynamic binary search tree, even one that knows the whole access sequence in advance. Despite a fair amount of attention over the years, this *dynamic optimality* conjecture is still open. In fact, there is currently no non-trivial (i.e., sub-logarithmic) bound on the competitiveness of splay trees. The difficulty of this problem stems from the fact that splay trees were deliberately designed

not to “play along” with our standard notions of good structure. Whether *good structure* refers to a balance invariant [14] or a more subtle structural invariant [10, 33, 15], *any* type of structural property must allow some tree configurations and precludes others, or perhaps regulate the speed with which one tree configuration could be transformed into another [14]. Thus, one could argue that *if* there is an online dynamically optimal search tree, its behavior must resemble that of splay trees: haphazard and eager to respond to accesses with major structural adjustments.

The revived interest in the dynamic optimality question is largely due to the invention of *tango trees* [10]. By appealing to the interleave lower bound of Wilbur [34] Demaine, Harmon, Iacono, and Pătraşcu [10] proved that tango trees perform within a $O(\log \log n)$ factor of the optimal dynamic search tree. Tango trees make use of red-black trees but it is easy to see that just about any standard binary search tree could be used as a black box in its place. Wang et al. [33] (see also [15]) showed that if splay trees are employed instead of red-black trees it is possible to have $O(\log \log n)$ -competitiveness and retain some properties of splay trees, such as $O(\log n)$ amortized time per access and sequential access in linear time. Wang et al. also extended their data structure to handle insertions and deletions.

If one’s immediate goal is to prove that splay trees are simply $o(\log n)$ -competitive, it suffices to show that they run in $o(n \log n)$ time on any class of access sequences for which the optimal binary search tree runs in $O(n)$ time. There is currently no “theory” of access sequences whose inherent complexity is linear. It is, therefore, not too surprising that all the major open conjectures on splay trees (corollaries of dynamic optimality) concern sequences whose optimal complexity is linear. Whether one’s goal is modest or ambitious, i.e., proving sub-logarithmic competitiveness or the full dynamic optimality conjecture, the first step must be to understand how splay trees behave on very easy access sequences. We restate below three unresolved conjectures on the opti-




Figure 1: Splay trees’ restructuring heuristics. After accessing an element x the tree rotates it to the root position by repeatedly applying a *zig-zig*, *zig-zag*, or *zig* as appropriate. On the left is the zig-zig case: x and its parent y are both left children (or both right children); the edges (y, z) and (x, y) are rotated in that order. On the right is the zig-zag case; the edges (x, y) and (x, z) are rotated in that order. Not depicted is the zig case, when y is the tree root and the edge (x, y) is rotated.

mality of splay trees [32, 29, 26].

Deque Conjecture Tarjan [32] conjectured that all double-ended queue operations¹ (push, pop, and their symmetric counterparts inject and eject) take $O(1)$ amortized time if implemented with a splay tree. A push makes the root of the splay tree the right child of a new vertex and a pop splays the leftmost leaf to the root position and deletes it. Inject and eject are symmetric.

Traversal Conjecture Sleator and Tarjan [29] conjectured that for two binary search trees S and T (defined on the same node set) accessing the elements in T by their preorder number in S takes linear time.

Split Conjecture Lucas conjectured [26] that any sequence of *splittings* in a splay tree takes linear time. A split at x consists of splaying x to the root and deleting it, leaving two splay trees, each subject to more splittings.

Sundar [30] established a bound of $O(n\alpha(n))$ on the time taken for n deque operations, where α is the inverse-Ackermann function. Lucas [26] showed that when the initial splay tree is a path (each node a left child), n split operations take $O(n\alpha(n))$ time. Notice that the split conjecture subsumes a special case of the deque conjecture, where only pops and ejects are allowed. We are aware of no published work concerning the traversal conjecture as such, though Chaudhuri and Höft [7] did prove it for the very special case when $S = T$.

Our Contributions. We introduce a new technique in the analysis of splay trees that is fundamentally different from all previous work on the subject [29, 32, 30, 9, 8, 14, 11]. The idea is to bound

the time taken to perform a sequence of accesses by *transcribing* the rotations performed by the splay tree as a Davenport-Schinzel sequence \mathcal{S} , i.e., one avoiding a specific forbidden subsequence. We apply this technique to the deque problem and show that n deque operations take $O(n\alpha^*(n))$ time, where $\alpha^*(n)$ is the number of applications of the inverse-Ackermann function mapping n down to a constant. This time bound is established by generating not one sequence \mathcal{S} but a hierarchy of sequences, each of which avoids subsequences isomorphic to *abababa*. Nearly tight bounds on the length of such sequences were given by Agarwal, Sharir, and Shor [2]. We believe that a generalized version of this technique should be useful in resolving other open conjectures on splay trees. For instance, a particularly clean way to prove the deque, split, or traversal conjectures would be to transcribe their rotations as a *generalized* Davenport-Schinzel sequence with length $O(n)$, i.e., one avoiding a fixed forbidden subsequence, not necessarily one of the form *abab...*. There is a large family of forbidden subsequences whose extremal function is known to be linear [22] as well as some rather simple forbidden subsequences, *abacabc* and *abacacb* for example, whose linear/non-linear status is still open.

Related Work. Iacono [17] defined a weaker notion of dynamic optimality called *key independent* optimality. One assumes that keys are assigned to elements randomly. The optimal cost of a sequence of operations is the expected optimal cost over all key assignments. Iacono showed that any data structure having the working set property is also optimal in the key independent sense. Blum et al. [5] defined another weaker notion of dynamic optimality called dynamic *search* optimality. In this cost model the (online) search tree can perform any number of rotations for free after each access, i.e., it only pays for

¹Also called a *deque*, pronounced “deck.”

actually doing searches. Georgakopoulos [14] showed that splay trees are competitive against a large class of dynamic *balanced* binary search trees, which can be “self adjusting” in the sense that they change shape in preparation for future searches. Iacono defined a *unified property* for search structures that subsumes the working set and dynamic finger properties. In his data structure [19] the access time for an element is logarithmic in its distance, where *distance* is a natural combination of temporal distance and key-space distance. Iacono’s data structure [19] is not a binary search tree and it is currently open whether *any* offline binary search tree has the unified property. In other words, it is not known to be a corollary of dynamic optimality.

Just after the invention of splay trees [29], Fredman et al. [13] invented the pairing heap as a kind of self-adjusting version of Fibonacci heaps. There is no obvious (and still interesting) analogue of dynamic optimality for priority queues, though Iacono [18] did show that pairing heaps possess an analogue of the working set property. See Fredman [12] and Pettie [27] for the best lower and upper bounds on the performance of pairing heaps.

List maintenance and caching algorithms (such as move-to-front or least-recently-used [28]) are sometimes described as being self adjusting heuristics. In these problems the (asymptotic) dynamic optimality questions are pretty well understood [28], though the leading constants have not yet been pinned down [3, 4].

Organization. In Section 2 we describe a known reduction [32, 30, 26] from the deque problem to a restrictive system of path compressions. In Section 3 we define some notation for Davenport-Schinzel sequences, path compressions, and slowly growing functions. Section 4 introduces a recurrence relation for a type of path compression system and shows how it can be analyzed easily using bounds on Davenport-Schinzel sequences. Section 5 gives the proof that n deque operations take $O(n\alpha^*(n))$ time.

2 Deque Operations and Path Compression Schemes

The relationship between deque operations on a splay tree and halving path compressions on an arbitrary tree was noted by Lucas [25, 26] and implicitly in [32, 30]. Let us briefly go through the steps of the reduction. At the beginning of a phase we divide the k -node splay tree into left and right halves of size $k/2$. The phase ends when one half has been deleted due to pops and ejects. We ignore the right half for now and look at the binary tree induced by the left half; call

it L and its root r .² The root of this tree may or may not correspond to the root of the whole splay tree. In any case, we imagine rotating the nodes on the right spine of L across the root until the root has only one (left) child; call this tree L' . Finally, we transform the binary tree L' into a general tree L'' as follows. The left child of a vertex in L' corresponds to its leftmost child in L'' and its right child in L' corresponds with its right sibling in L'' . See Figure 2. Notice that if v and its L -parent $p_L(v)$ lie on the left spine of L , and $p_L(v) \neq r$, rotating the edge $(v, p_L(v))$ corresponds to making v the leftmost child of its grandparent in L'' . If $p_L(v) = r$ rotating $(v, p_L(v))$ makes v the root of L but does not change the structure of L' or L'' . In an almost symmetric way, if $v, p_L(v)$, and $p_L^{(2)}(v) \neq r$ lie on the right spine of L then rotating $(v, p_L(v))$ in L corresponds to rotating $(p_L^{(2)}(v), p_L(v))$ in L' and making $p_L^{(2)}(v)$ the leftmost child of its grandparent (v) in L'' . Notice that $v = p_{L'}(p_L(v)) = p_{L'}^{(2)}(p_L^{(2)}(v))$. Observation 2.1 summarizes the relationship between deque operations and path compressions; see also [26, 30].

OBSERVATION 2.1. *A pop operation corresponds to a halving path compression in L'' that begins at the leftmost leaf and terminates at r , followed by a deletion of the leftmost leaf and a possible root relocation from r to its leftmost child. A push operation causes a new leaf v to be added as the leftmost child of r in L'' , followed by a root relocation from r to v . An eject operation corresponds to a halving path compression originating at r and terminating at some ancestor (not necessarily the root of L''), followed by a possible root relocation from r to $p_{L'}(r)$. An inject operation has no effect on L'' .*

It is clear that if the amortized cost per deque operation in a phase is $f(k)$ then the overall cost for m deque operations on an initial n -node splay tree is $O((m+n)f(m+n))$; see [32, 30, 11]. Using Observation 2.1 we can (and do) restrict our attention to bounding the total length of a sequence of halving path compressions up the left spine of an arbitrary rooted tree. However, we may still lapse into deque terminology. The terms “pop” and “push”, for instance, are basically equivalent to “halving path compression” and “leaf addition.”

Related Work. Restricted forms of path compressions have been studied in a number of situations. The most well known example is Hart and Sharir’s

²By *induced* we mean the following: L is the unique binary tree on the set of left nodes that preserves the ancestor-descendant relationship from the original splay tree.




Figure 2: Top row: the transformation from the splay tree (left nodes black, right nodes white) to L , L' , and L'' . Bottom: the effect of a pop on the splay tree and L'' .

result [16] on the equivalence between *(ababa)*-free Davenport-Schinzel sequences and *generalized* postordered path compressions; both have maximum length $\Theta(n\alpha(n))$. Loeb and Nešetřil [23] and Lucas [25] independently proved that standard postordered path compressions with the so-called *rising roots* condition take linear time. Buchsbaum, Sundar, and Tarjan [6] generalized this result to what they called deque-ordered path compressions, again assuming the rising roots condition. Hart and Sharir [16] have conjectured that the rising roots condition is not essential and that standard postordered path compressions take linear time. The path compressions corresponding to deque operations are similar to the special cases studied earlier. Some differences are that the compressions are halving (not full), and although the compressions are spinal, due to pushes and ejects they are not performed in postorder and do not satisfy the rising roots condition.

3 Notation

We say two sequences are isomorphic if they are the same up to a renaming of symbols. The relations $\sigma \succ \Sigma$ and $\sigma \prec \Sigma$ hold, respectively, when σ is a subsequence of Σ and when σ is isomorphic to a subsequence of Σ . A subsequence of Σ is any sequence derived by deleting symbols from Σ . A sequence Σ is called c -regular if any two occurrences of the same

symbol appear at distance at least c . We denote by $|\Sigma|$ and $\|\Sigma\|$ the length and alphabet size of Σ , respectively. Following Klazar [21] we let $\text{Ex}(\sigma, n)$ be the maximum length of σ -free sequences:

DEFINITION 3.1.

$$\text{Ex}(\sigma, n) = \max\{|\Sigma| \mid \sigma \not\prec \Sigma, \|\Sigma\| = n, \text{ and } \Sigma \text{ is } \|\sigma\|\text{-regular}\}$$

It is known [20] that $\text{Ex}(\sigma, n) = n \cdot 2^{\alpha(n)^{O(1)}}$, where the $O(1)$ depends on the length and alphabet size of σ . Nearly tight bounds on $\text{Ex}(\sigma, n)$ are known [2] when σ is of the form $ababa \dots$. Here $\alpha(n)$ is the inverse-Ackermann function, which can be defined as follows. If $f(n)$ is a strictly decreasing function on the positive integers $f^*(n) = \min\{i : f^{(i)}(n) \leq 2\}$, where $f^{(i+1)}(n) = f(f^{(i)}(n))$ and $f^{(1)}(n) = f(n)$.

Define $\alpha(m, n) = \min\{i \geq 1 : \log^{\overbrace{** \dots *}}^{i-1}(n) \leq 2 + \frac{m}{n}\}$ and $\alpha(n) = \alpha(n, n)$. This definition of α differs from others in the literature [31, 16] by at most a small constant.

In this paper all trees are rooted and the children of any vertex assigned some left-to-right order. The trees we deal with are occasionally restructured with *path compressions*. Let $p(u)$ be the parent of u at some specific time. If $C = (u_1, \dots, u_k)$ is a path,

where $u_{i+1} = p(u_i)$, performing a *total* compression of C means to reassign $p(u_i) = u_k$, for all $1 \leq i \leq k - 2$. A *halving* compression of C sets $p(u_i) = u_{i+2}$ for all odd $i \leq k - 2$. If k is even a halving compression may set $p(u_{k-1}) = p(u_k)$. A specific case of interest is when u_k is the tree root: setting $p(u_{k-1}) = p(u_k)$ causes u_{k-1} to be a tree root as well. We say that C *originates* at u_1 and *terminates* at u_k . A total/halving path compression that does *not* terminate at the tree root is *stunted*. The *length* of a total/halving path compression is the number of vertices whose parent pointers are altered. We never consider compressions with zero length.

A postordering of a tree rooted at v , having children v_1, \dots, v_k from left to right, is the concatenation of the postorderings of the subtrees rooted at v_1, \dots, v_k followed by v . The *spine* of a tree is the path from its root to its leftmost leaf. A path compression is *spinal* if it affects a subpath of the spine; it need not include the leftmost leaf nor the tree root. It is easy to see that a total/halving spinal compression can be postorder preserving. For a total compression (u_1, \dots, u_k) we just prepend u_1, \dots, u_{k-2} to the pre-existing left-right ordering on the children of u_k . For a halving compression we make u_i the new leftmost child of u_{i+1} for all odd i .

In order to use a clean inductive argument we look at a restrictive type of instance called a *spinal compression system*. The initial structure consists of a single path containing a mixture of *essential* nodes and *fluff* nodes. The tree is manipulated through halving spinal path compressions, leaf deletions, and spontaneous compressions. Let us go through each of these in turn. Whenever a fluff node becomes a leaf it is automatically deleted. The leftmost essential node (always a leaf) may be deleted at any time. We are mainly interested in the total length of the “official” path compressions, which are always halving and spinal. Spontaneous compressions are any postorder preserving path compressions, the cost of which we need not account for. Let $\mathcal{R}(n, f, m)$ be the maximum total length of the official compressions on an instance with n essential nodes and f fluff nodes, where at most m of the compressions are stunted. In Section 4 we derive a recursive expression bounding $\mathcal{R}(n, f, m)$ and in Section 5 we relate \mathcal{R} to the time taken by deque operations.

4 Recursive Bounds on Spinal Compression Systems

Consider an initial instance with n essential nodes and f fluff nodes. We first divide up the path into n/B blocks, each containing B essential nodes, where the bottommost node in each block is essential. The

sequence of official compressions is partitioned into n/B epochs, where the j th epoch begins at the first compression when the leftmost leaf belongs to the j th block and ends at the beginning of the $(j + 1)$ th epoch. Let I_j be the set of nodes, not including those in the j th block, that are touched by some compression in the j th epoch. It follows that at the commencement of the j th epoch I_j is a single path; see Figure 3. At this time some members of I_j may




Figure 3:

be *affiliated* with previous epochs. It is always the case that nodes affiliated with any $j' < j$ appear as a connected subpath of the spine. We call an essential node v *exposed* if it has no essential ancestor that either shares an affiliation with v or lies in the same block as v . Let $\hat{I}_j \subseteq I_j$ be the set of essential exposed nodes. We call epoch j *sparse* if $|\hat{I}_j| \log |\hat{I}_j| < |I_j|$ and *dense* otherwise. If epoch j is dense we *affiliate* all nodes in I_j with j . We view I_j as a separate spinal compression system, where \hat{I}_j is the set of essential nodes and all others in $I_j \setminus \hat{I}_j$ are fluff. Notice that a subpath of I_j can be affiliated with j and a previous epoch, say j' . If a compression influences nodes affiliated with both j and j' we charge the cost to the j th spinal compression system and call its effect on the j' th compression system a spontaneous compression. If the j th epoch is sparse we do not affiliate any nodes with j .

LEMMA 4.1. *Let $f \geq \sum_j f_j$ and $m \geq \sum_j \bar{m}_j +$*

$\sum_j \tilde{m}_j$.

$$\begin{aligned} \mathcal{R}(n, f, m) &= \sum_{j=1}^{n/B} \mathcal{R}(B, f_j, \tilde{m}_j) \\ &+ \sum_{\substack{\text{epoch } j \\ \text{dense}}} 3\mathcal{R}(|\hat{I}_j|, |\hat{I}_j| \log |\hat{I}_j|, \tilde{m}_j) \\ &+ 2m + n + f \end{aligned}$$

Proof. Consider a compression C occurring in the j th epoch. We can always break C into three parts: (i) a compression C' inside the j th block followed by (ii) an edge e leading from the j th block to some vertex in I_j , followed by (iii) a compression C'' lying entirely within I_j . Notice that if C'' is present it may be either stunted or unstunted. However, if C'' is present then C' (if present) must be an unstunted compression. We call C'' *short* if it has unit cost, i.e., if only one node changes parent. Whether C' is stunted or not its cost is covered in the $\mathcal{R}(B, f_j, \tilde{m}_j)$ term, where f_j is the number of fluffy nodes and \tilde{m}_j the number of compressions that are stunted with respect to block j . Similarly, if the j th epoch is dense the cost of C'' , whether stunted or not, is covered in the term $\mathcal{R}(|\hat{I}_j|, |\hat{I}_j| \log |\hat{I}_j|, \tilde{m}_j)$, where \tilde{m}_j is the number of stunted compressions that terminate in the spinal compression system on I_j . We account for the cost of e in one of two ways. If C'' (and therefore C) is stunted we charge it to the compression; hence the m term. If not, notice that after every unstunted compression C'' the number of nodes whose grandparents are roots or nonexistent increases by at least one. This can obviously happen at most $n + f$ times. Consider the compression C'' when the j th epoch is sparse. If C'' is short we charge its cost to the compression; hence the other m term. In general, C'' will intersect one or more previously spawned spinal compression systems. That is, it can be divided up into pieces $C''_1, C''_2, C''_3, \dots$, where for odd i , C''_i intersects an existing compression system. The costs of C''_1, C''_3, \dots are covered in the $\sum_{j'} \mathcal{R}(|\hat{I}_{j'}|, |\hat{I}_{j'}| \log |\hat{I}_{j'}|, \tilde{m}_{j'})$ terms. By doubling this sum we can account for some of C''_2, C''_4, \dots ; at the very least this includes those with unit cost. Over all of epoch j , the total length of the other mini-compressions of the form C''_i (i even) is $|\hat{I}_j| \log |\hat{I}_j|$; see Sundar [30].³ By tripling the sum $\sum_{j'} \mathcal{R}(|\hat{I}_{j'}|, |\hat{I}_{j'}| \log |\hat{I}_{j'}|, \tilde{m}_{j'})$ we account for the cost of these mini-compressions as well.

³In Sundar's terminology each of these mini-compressions is a k -cascade for some $k \geq 2$.

Lemma 4.1 looks as though it may be effectively vacuous. We are trying to bound the growth of $\mathcal{R}(n, f, m)$ and all Lemma 4.1 can say is that it depends on the magnitude of the $\{\hat{I}_j\}_j$ sets. It does not even suggest a trivial bound on their size. Our strategy is to transcribe the $\{\hat{I}_j\}_j$ sets as a repetition-free sequence \mathcal{S} that avoids a specific forbidden subsequence. It follows from the results of Agarwal et al. [2] that $|\mathcal{S}|$ is nearly linear in the size of its alphabet. By choosing an appropriate block size B we can apply Lemma 4.1 to obtain useful bounds on \mathcal{R} .

Our transcription method for \mathcal{S} is very similar to the one used by Hart and Sharir [16]. In Lemma 4.2 we show that $abaabba \not\prec \mathcal{S}$ and $abababa \not\prec \mathcal{S}$. Using bounds on $\text{Ex}(abababa, \cdot)$ from Agarwal, Sharir, and Shor [2] we are able to show, in Lemmas 4.3 and 4.4 that $\mathcal{R}(n, f, m)$ is $O((n + f + m)\alpha^*(n))$.

Before giving the actual transcription method we give a simpler one, point out why it isn't quite what we need, then adjust it. The transcription is based on an evolving labeling of the nodes. A label is simply a sequence of block/epoch indices, listed in *descending* order. Since nodes can belong to several spinal compression systems a node may keep several labels, one for each system. At the commencement of the j th epoch we selected out of I_j a subset \hat{I}_j with several properties, one of which was that \hat{I}_j has at most one node from any block. If the j th epoch is dense we prepend j to the label of each node in \hat{I}_j . The sequence \mathcal{S}' consists of the concatenation of all node labels, where the nodes are given in postorder. An equivalent definition is that \mathcal{S}' begins empty; whenever the leftmost leaf is deleted we append its label to \mathcal{S}' and continue. Besides having the property that $abaabba, abababa \not\prec \mathcal{S}'$, we need \mathcal{S}' to be repetition-free and contain not too many occurrences of any one symbol, say, at most t occurrences. To enforce this, if $|\hat{I}_j| > t$ we simply split it up into $\lceil |\hat{I}_j|/t \rceil$ pieces and treat each piece as a distinct spinal compression system. Thus, the number of systems could exceed the number of blocks/epochs. There may be repetitions in \mathcal{S}' but not too many. Note that the labels of nodes appearing in any block share no symbols in common. Therefore, \mathcal{S} can only contain repetitions at block boundaries. By removing at most $n/B - 1$ symbols from \mathcal{S}' we can eliminate all repetitions. Using the above two observations, let \mathcal{S} be a repetition-free sequence derived from \mathcal{S}' in which no symbol occurs more than t times.

LEMMA 4.2. For $\sigma \in \{abaabba, abababa\}$, $\sigma \not\prec \mathcal{S}'$.

Proof. First note that if the lemma holds for \mathcal{S}'

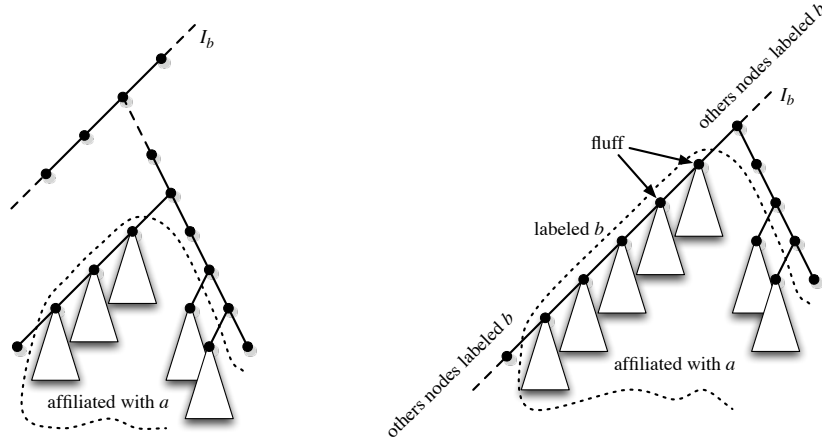


Figure 4:

it holds for \mathcal{S} as well. We show that for any $b > a$, $\text{babba} \not\prec \mathcal{S}'$. This would prove the lemma since, for any $\sigma' \in \{\text{babba}, \text{abaab}\}$ and $\sigma'' \in \{\text{abaabba}, \text{abababa}\}$, $\sigma' \succ \sigma''$. Consider the commencement of the b th epoch. If no nodes affiliated with a appear on the path I_b then all nodes labeled with b occur, in postorder, strictly before or strictly after all nodes labeled with a ; see the left part of Figure 4. In this case $\text{baba} \not\prec \mathcal{S}'$. The more interesting case is when some interval of the nodes in I_b are affiliated with a . Recall that in our transcription method only exposed nodes were labeled and there could be at most one exposed node in each set of nodes with a common affiliation. Thus, only one node affiliated with a can be labeled b . All other nodes labeled b occur strictly before or strictly after all nodes labeled a . Thus, after the appearance of babb in \mathcal{S}' all nodes labeled a would have been deleted. We conclude that $\text{babba} \not\prec \mathcal{S}'$.

Lemma 4.3 incorporates the recursive characterization of \mathcal{R} from Lemma 4.1 and the (abababa) -freeness of \mathcal{S} established in Lemma 4.2.

LEMMA 4.3.

$$\begin{aligned} \mathcal{R}(n, f, m) &= \sum_{j=1}^{n/B} \mathcal{R}(B, f_j, \tilde{m}_j) \\ &+ 3 \sum_{i=1}^q \mathcal{R}(l_i, l_i \log l_i, 2l_i \log^2 l_i) \\ &+ 2m + n + f \end{aligned}$$

where

$$\begin{aligned} f &\geq \sum_j f_j, & m &\geq \sum_j \tilde{m}_j, & l_i &\leq t, \\ q &\leq \frac{\sum_i l_i}{t} + \frac{n}{B}, & \sum_i l_i &\leq \text{Ex}(\text{abababa}, q) \end{aligned}$$

Proof. Here $\{l_i\}_i$ represent the sizes of the spawned spinal compression systems. They would correspond with the epochs/blocks if we did not artificially break them up to guarantee that each $l_i \leq t$. It follows that the number of spawned systems is at most $q = (\sum_i l_i)/t + n/B$. By Lemma 4.2 we have $\sum_i l_i \leq \text{Ex}(\text{abababa}, q)$. The term $\mathcal{R}(l_i, l_i \log l_i, 2l_i \log^2 l_i)$ reflects the cost of the i th spawned compression system, with l_i essential nodes and at most $l_i \log l_i$ fluff nodes. By [30] the number of stunted compressions (with greater than unit cost) terminating in this system is at most $(l_i + l_i \log l_i) \log(l_i + l_i \log l_i) \leq 2l_i \log^2 l_i$.

LEMMA 4.4. $\mathcal{R}(n, f, m) = O((n + f + m)\alpha^*(n))$

Proof. Let $\min\{\text{Ex}(\text{abaabba}, z), \text{Ex}(\text{abababa}, z)\} = z \cdot \beta(z)$. From the bounds established by Hart and Sharir [16] and Agarwal et al. [2] we only know that $\beta(z) = \Omega(\alpha(z))$ and $O(\alpha(z))^{\alpha(z)}$.⁴ We apply

⁴Klazar's and Valtr [22] claimed that $\text{Ex}(\text{abaabba}, n) = \Theta(n\alpha(n))$ could be had by fiddling with Hart and Sharir's proof [16]. This claim was later retracted [21].

Lemma 4.3 with $B = t = \beta^2(n)$:

$$\begin{aligned}
\sum_i l_i &\leq \min\{\text{Ex}(abababa, q), \text{Ex}(abaabba, q)\} \\
&\leq \min\left\{\text{Ex}\left(abababa, \frac{\sum_j l_j}{t} + \frac{n}{B}\right), \right. \\
&\quad \left. \text{Ex}\left(abaabba, \frac{\sum_j l_j}{t} + \frac{n}{B}\right)\right\} \\
&\leq \left(\frac{\sum_i l_i}{t} + \frac{n}{B}\right)\beta(n) \\
&\leq \frac{\sum_i l_i}{\beta(n)} + \frac{n}{\beta(n)}
\end{aligned}$$

Thus $\sum_j l_j \leq n/(\beta(n) - 1)$. Assume inductively that $\mathcal{R}(n, f, m) \leq c(n+f+m)\gamma(n)$, where $\gamma(n)$ is the minimum i such that $(\beta^2)^{(i)}(n)$ is less than some large enough constant. (We need this constant threshold since $\beta^2(n)$ is not necessarily decreasing for very small n .)

$$\begin{aligned}
\mathcal{R}(n, f, m) &\leq \sum_{j=1}^{n/B} \mathcal{R}(B, f_j, \tilde{m}_j) \\
&\quad + \sum_{i=1}^q 3\mathcal{R}(l_i, l_i \log l_i, 2l_i \log^2 l_i) + 2m + n + f \\
&\leq c(n+f + \sum_j \tilde{m}_j)\gamma(\beta(n)^2) \\
&\quad + 3c\left(\frac{n}{\beta(n)-1}\right)\left(1 + \log(\beta^2(n))\right) \\
&\quad + 2\log^2(\beta^2(n))\gamma(\beta^2(n)) + 2m + n + f \\
&\leq c(n+f + \sum_j \tilde{m}_j)(\gamma(n) - 1) \\
&\quad + o(n) + 2m + n + f \\
&\leq c(n+f+m)\gamma(n) \quad \{\text{for } c \text{ sufficiently large}\}
\end{aligned}$$

It is easy to see that $\beta^2(\beta^2(n)) < \alpha(n)$ for n sufficiently large, from which it follows that $\gamma(n) = O(\alpha^*(n))$.

5 Upper Bounds on Deque Operations

THEOREM 5.1. *Starting with an n -node splay tree, the time for m deque operations is $O((m+n)\alpha^*(m+n))$.*

Proof. We reduce the deque problem to a set of spinal compression systems. The main difference between these systems and the deque problem as a whole is that spinal compression systems do not allow general insertion/deletion of leaves and the initial tree is always a single path.

We impose a linear order on all nodes that ever appear in the splay tree. A pushed node (injected node) precedes (proceeds) all nodes that are currently in or were in the splay tree in the past. As in the general reduction from Section 2 we divide the access sequence into phases. At the beginning of each phase the nodes are separated into equally sized left and right sets, say m' in each set. We partition the left set into blocks of size $B = \beta(m')^2$. Let the j th period begin at the first pop of an element in the j th block and end when all of the j th block has been popped. Note that the pop that begins the j th period could delete *any* element from the j th block, not necessarily the first. Also note that periods are not necessarily disjoint. For example, just after the j th period begins we could push B elements and then perform a pop. This would have the effect of starting the j' th period without ending the j th period, where $j' > j$. Only one period is *active*, namely, the one whose block contains the leftmost leaf in the tree. All other periods are *on hold*.

Let J_j be the *set* of nodes that participate in a compression in the active part of the j th period, excluding those that lie in block j . Clearly the nodes in J_j lie on one path when the j th period begins. However, they are not necessarily contiguous. For example, if the j th period is put on hold, compressions from later periods could evict non- J_j nodes from the path between two J_j nodes. If we put ourselves in one period's point of view we can assume without loss of generality that J_j *does* form a contiguous segment since, by definition, this period never "sees" any non- J_j nodes that would expose its misbelief.

Let $\mathcal{D}(m')$ be the maximum cost of deque operations (i.e., the cost of push, pop, and eject operations that affect the left set) when the total number of nodes is m' . After the j th period begins the j th block functions as a mini-deque structure and its total cost at most $\mathcal{D}(B)$. Until the j th period begins the block- j nodes may participate in compressions, the total cost of which is B . (Whenever a node is evicted from the left spine of the tree induced by block- j nodes it can only be seen again after the j th period begins.) The cost of compressions inside the $\{J_j\}_j$ sets can be expressed in terms of the \mathcal{R} function. Again, at the commencement of the j th period we identify the exposed nodes $\hat{J}_j \subseteq J_j$. We do nothing if it is sparse ($|\hat{J}_j| \log |\hat{J}_j| < |J_j|$) and if it is dense, we affiliate all nodes in J_j with j and spawn a separate spinal compression system on J_j . The proofs of Lemmas 4.1 and

4.3 can easily be adapted to show that:

$$\mathcal{D}(m') \leq \frac{m'}{B} \mathcal{D}(B) + O(m' + \sum_{j=1}^{m'/B} \mathcal{R}(|\hat{J}_j|, |\hat{J}_j| \log |\hat{J}_j|, 2|\hat{J}_j| \log^2 |\hat{J}_j|))$$

where $\sum_j |\hat{J}_j| \leq \text{Ex}(abababa, m'/B) = (m'/B)\beta(m'/B) \leq m'/\beta(m')$. A simple proof by induction (along the lines of Lemma 4.4) shows that $\mathcal{D}(m') = O(m'\alpha^*(m'))$.

6 Conclusion

It's fair to say that previous analyses of splay trees could be characterized as using potential functions [29, 14], counting arguments [32, 26, 30, 11] or a complex synthesis of the two [9, 8]. Although these techniques have been wildly successful in proving (or nearly proving) the *corollaries* of dynamic optimality, they have had no impact on the dynamic optimality question itself. The reason for this is probably the fact that so little is known about the behavior of the optimal (offline) dynamic binary search tree; see [24, 5]. It is worth noting that the $O(\log \log n)$ -competitiveness of tango trees and their variants [10, 33, 15] was established *not* by comparing them against an optimal search tree (the standard way to establish competitiveness results) but indirectly, by way of one of Wilber's lower bounds on any binary search tree [34].

The strategy taken in this paper is quite different from previous work and is clearly general enough to be applied to other open problems concerning splay trees. By transcribing the rotations performed by the splay tree into a Davenport-Schinzel sequence avoiding *abaabba* and *abababa* we were able to use an off-the-shelf result of Agarwal et al. [2] and avoid an unmanageable bookkeeping problem. One direction for future research is to study the relationship between splay operations, various transcription methods, and various forbidden subsequences. Of particular interest are those forbidden subsequences whose extremal function is linear since it is these that would be fit for finally settling the deque, split, and traversal conjectures. Adamec et al. [1] and Klazar and Valtr [22] have exhibited a large family of linear forbidden subsequences.

References

[1] R. Adamec, M. Klazar, and P. Valtr. Generalized Davenport-Schinzel sequences with linear upper bound. *Discrete Math.*, 108(1-3):219–229, 1992.

[2] P. Agarwal, M. Sharir, and P. Shor. Sharp upper and lower bounds on the length of general Davenport-Schinzel sequences. *J. Combinatorial Theory, Series A*, 52, 1989.

[3] S. Albers. Improved randomized on-line algorithms for the list update problem. *SIAM J. Comput.*, 27(3):682–693, 1998.

[4] S. Albers, B. von Stengel, and R. Werchner. A combined BIT and TIMESTAMP algorithm for the list update problem. *Info. Proc. Lett.*, 56, 1995.

[5] A. Blum, S. Chawla, and A. Kalai. Static optimality and dynamic search-optimality in lists and trees. *Algorithmica*, 36, 2003.

[6] A. Buchsbaum, R. Sundar, and R. E. Tarjan. Data-structural bootstrapping, linear path compression, and catenable heap-ordered double-ended queues. *SIAM J. Comput.*, 24, 1995.

[7] R. Chaudhuri and H. Höft. Splaying a search tree in preorder takes linear time. *International Journal of Mathematics and Mathematical Sciences*, 13(1):186–186, 1990.

[8] R. Cole. On the dynamic finger conjecture for splay trees. II. The proof. *SIAM J. Comput.*, 30(1):44–85, 2000.

[9] R. Cole, B. Mishra, J. Schmidt, and A. Siegel. On the dynamic finger conjecture for splay trees. I. Splay sorting $\log n$ -block sequences. *SIAM J. Comput.*, 30(1):1–43, 2000.

[10] E. Demaine, D. Harmon, J. Iacono, and M. Pătraşcu. Dynamic optimality — almost. In *Proceedings 45th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 484–490, 2004.

[11] A. Elmasry. On the sequential access theorem and deque conjecture for splay trees. *Theoretical Computer Science*, 314, 2004.

[12] M. L. Fredman. On the efficiency of pairing heaps and related data structures. *J. ACM*, 46(4):473–501, 1999.

[13] M. L. Fredman, R. Sedgewick, D. D. Sleator, and R. E. Tarjan. The pairing heap: a new form of self-adjusting heap. *Algorithmica*, 1(1):111–129, 1986.

[14] G. F. Georgakopoulos. Splay trees: A reweighing lemma and a proof of competitiveness vs. dynamic balanced trees. *J. Algor.*, 51, 2004.

[15] G. F. Georgakopoulos. How to splay for $\log \log n$ -competitiveness. In *Proc. 4th Int'l Workshop on Experimental and Efficient Algorithms*, volume 3503, pages 570–579, 2005.

[16] S. Hart and M. Sharir. Nonlinearity of Davenport-Schinzel sequences and of generalized path compression schemes. *Combinatorica*, 6(2):151–177, 1986.

[17] Iacono. Key-independent optimality. *Algorithmica*, 42, 2005.

[18] J. Iacono. Improved upper bounds for pairing heaps. In *Scandinavian Workshop on Algorithm Theory (SWAT, LNCS 1851)*, pages 32–43, 2000.

[19] John Iacono. Alternatives to splay trees with $o(\log n)$ worst-case access times. In *SODA*, pages 516–522, 2001.

- [20] M. Klazar. A general upper bound in extremal theory of sequences. *Comment. Math. Univ. Carolin.*, 33(4):737–746, 1992.
- [21] M. Klazar. Generalized Davenport-Schinzel sequences: results, problems, and applications. *Integers*, 2:A11, 39 pp. (electronic), 2002.
- [22] M. Klazar and P. Valtr. Generalized Davenport-Schinzel sequences. *Combinatorica*, 14(4):463–476, 1994.
- [23] M. Loebl and J. Nešetřil. Linearity and unprovability of set union problem strategies. I. Linearity of strong postorder. *J. Algor.*, 23(2):207–220, 1997.
- [24] J. M. Lucas. Canonical forms for competitive binary search tree algorithms. Technical Report DCS-TR-250, Rutgers University, 1988.
- [25] Joan M. Lucas. Postorder disjoint set union is linear. *SIAM J. on Computing*, 19(5):868–882, October 1990.
- [26] Joan M. Lucas. On the competitiveness of splay trees: Relations to the union-find problem. In Lyle A. McGeoch and Daniel D. Sleator, editors, *On-line Algorithms*, volume 7 of *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, pages 95–124, 1991.
- [27] S. Pettie. Towards a final analysis of pairing heaps. In *Proceedings 46th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 174–183, 2005.
- [28] D. D. Sleator and R. E. Tarjan. Amortized efficiency of list update and paging rules. *Comm. ACM*, 28, 1985.
- [29] D. D. Sleator and R. E. Tarjan. Self-adjusting binary search trees. *J. ACM*, 32(3):652–686, 1985.
- [30] R. Sundar. On the deque conjecture for the splay algorithm. *Combinatorica*, 12(1):95–124, 1992.
- [31] R. E. Tarjan. Efficiency of a good but not linear set merging algorithm. *J. ACM*, 22:215–225, 1975.
- [32] R.E. Tarjan. Sequential access in play trees takes linear time. *Combinatorica*, 5(4):367–378, 1985.
- [33] C. Wang, J. Derryberry, and D. Sleator. $O(\log \log n)$ -competitive dynamic binary search trees. In *Proc. 17th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2006.
- [34] R. Wilber. Lower bounds for accessing binary search trees with rotations. *SIAM J. Comput.*, 18(1):56–67, 1989.