

Static Optimality and Dynamic Search-Optimality in Lists and Trees¹

Avrim Blum,² Shuchi Chawla,² and Adam Kalai³

Abstract. Adaptive data structures form a central topic of on-line algorithms research. The area of Competitive Analysis began with the results of Sleator and Tarjan showing that splay trees achieve static optimality for search trees, and that Move-to-Front is constant competitive for the list update problem [ST1], [ST2]. In a parallel development, powerful algorithms have been developed in Machine Learning for problems of on-line prediction [LW], [FS]. This paper is inspired by the observation made in [BB] that if computational decision-making costs are not considered, then these “weighted experts” techniques from Machine Learning allow one to achieve a $1 + \varepsilon$ ratio against the best *static* object in hindsight for a wide range of data structure problems.

In this paper we give two results. First, we show that for the case of *lists*, we can achieve a $1 + \varepsilon$ ratio with respect to the best static list in hindsight, by a simple *efficient* algorithm. This algorithm can then be combined with existing results to achieve good static and dynamic bounds simultaneously. Second, for trees, we show a (computationally *inefficient*) algorithm that achieves what we call “dynamic search optimality”: dynamic optimality if we allow the *on-line* algorithm to make free rotations after each request. We hope this to be a step towards solving the longstanding open problem of achieving true dynamic optimality for trees.

Key Words. Adaptive data structures, Binary search trees, Competitive Analysis, Experts Analysis.

1. Introduction. Adaptive data structures form one of the central topics of on-line algorithms research, beginning with the results of Sleator and Tarjan that splay trees perform within a constant factor of the best *static* search tree for any request sequence, and that Move-to-Front is constant competitive with respect to the best off-line (adaptive) algorithm for lists [ST1], [ST2]. On-line algorithms for lists have subsequently been well studied, with upper and lower bounds of 1.6 and 1.5, respectively, on the competitive ratio of randomized algorithms [AvSW], [T]. The case of trees appears much harder: it is not known whether splay trees or any other on-line algorithm performs within a constant factor of the optimal off-line adaptive tree algorithm.

An interesting point is that in terms of competing against the best *static* object in hindsight, one can in principle achieve a ratio of $1 + \varepsilon$ for a wide variety of data structure problems, if we ignore time spent computing which update to perform. In particular, Blum and Burch [BB] show the following general result.

¹ This material is based upon work supported under NSF Grants CCR-9732705, CCR-0105488, and IIS-0121678, an NSF Mathematical Sciences Postdoctoral Research Fellowship, and an IBM Graduate Fellowship.

² Computer Science Department, Carnegie Mellon University, Pittsburgh, PA 15213, USA. {avrim,shuchi}@cs.cmu.edu.

³ Department of Mathematics, Massachusetts Institute of Technology, Cambridge, MA 02139, USA. akalai@math.mit.edu.

THEOREM 1.1 (Theorem 4 of [BB]). *Given N on-line algorithms for a Metrical Task System Problem of diameter D , and given $\varepsilon > 0$, one can use the Randomized Weighted Majority algorithm [LW] to achieve on any request sequence σ an expected cost at most $(1 + \varepsilon)L_\sigma + O((1/\varepsilon)D \log N)$, where L_σ is the cost of the best of the N algorithms on σ in hindsight.⁴*

For example, the list-update problem is a Metrical Task System problem with one state for each ordering of the n elements. There are $N = n!$ different lists of n elements, and the diameter of the space (the maximum cost to move between one list and another) is $O(n^2)$. Therefore, viewing each static list as an on-line algorithm, applying Randomized Weighted Majority achieves a $(1 + \varepsilon)$ static ratio with additive constant $O((1/\varepsilon)n^3 \log n)$. For trees, there are “only” $2^{O(n)}$ different search trees on n elements, and the diameter of the space is $O(n)$. So, this algorithm achieves a $1 + \varepsilon$ static ratio with additive constant $O((1/\varepsilon)n^2)$. We call this type of bound (namely, a $1 + \varepsilon$ ratio with additive constant polynomial in n and $1/\varepsilon$) *strong static optimality*. Unfortunately, the algorithm of Theorem 1.1 is horribly computationally inefficient, because it must explicitly maintain a probability distribution over all the component algorithms. Some progress on efficient Markov Chain Monte Carlo simulations for these type of algorithms has been made [CLS], [KV1] but not for the situations of interest to us.

In many situations the best static algorithm may be much worse than an optimal dynamic algorithm. In the case of list update, for instance, the best static algorithm can be up to $\Omega(n)$ times worse than the best dynamic algorithm if the adversary repeatedly accesses every element in succession. On the other hand, in many cases the optimal static algorithm *is* near optimal for dynamic algorithms—for instance, this would be the case if accesses come from some fixed distribution. The possibility of achieving strong optimality then becomes appealing.

In this paper we present two results. First, for the case of lists, we show how to achieve a $1 + \varepsilon$ ratio with respect to the best static list in hindsight by a simple *efficient* algorithm. The algorithm is randomized: the standard lower bound of 2 for deterministic algorithms holds for static ratio as well. We combine this algorithm with the currently best known dynamic algorithm (COMB, see [BEY] for details) to obtain one which has almost the same dynamic competitive ratio but *at the same time* has strong static optimality. Second, for trees, we show a (computationally *inefficient*) algorithm motivated by on-line learning, that achieves what we call “dynamic search optimality”: dynamic optimality if we allow the *on-line* algorithm to make free rotations after each request. We hope this to be a step towards addressing the long-standing open problem of achieving true dynamic optimality for trees. At the very least, this shows that dynamic optimality cannot be proven *impossible* by any argument that neglects the rotation cost of the on-line algorithm.⁵

⁴ The Randomized Weighted Majority algorithm [LW] maintains a probability distribution over the N component algorithms (“experts”). At each time step, the experts are penalized based on their cost at that step, multiplying the weight on an expert incurring cost c by $(1 - \varepsilon')^c$, and then renormalizing, where $\varepsilon' = O(\varepsilon/D)$. The name of the algorithm comes from a setting in which each expert is making a prediction. In that case this procedure can be viewed as taking a weighted vote and then choosing a prediction probabilistically based on the vote totals. It is also called the Hedge algorithm [FS].

⁵ In contrast, the standard $\Omega(\log k)$ lower bound for randomized paging, based on considering a sequence in which pages in $\{1, \dots, k + 1\}$ are requested completely at random, holds even if the on-line algorithm is allowed to move items into and out of its cache between requests for free.

Following our work, Kalai and Vempala [KV2] developed an efficient algorithm for Binary Search Trees that achieves a ratio of $1 + \varepsilon$ with respect to the best static tree in hindsight.

2. Strong Static Optimality for List Update. The List Update problem is to maintain a list of n elements while serving requests of the form $\text{ACCESS}(x)$, where x is an element in the list. The cost of an access for element x is the depth of x in the list. That is, the i th element has a cost of $i - 1$. This is known as the *Partial Cost* model (as opposed to the *Full Cost* model in which the cost of access is i for the i th element). Following an access, the element x may be moved anywhere higher in the list without additional costs. Then other elements may be moved at a cost of one per transposition.

In this section we describe two different algorithms for list update, both based on the same idea of reducing the problem to a two-expert “combining expert advice” problem, and both achieving a $1 + \varepsilon$ competitive ratio with respect to the best static list. We prove the bound for the first algorithm using a theorem of [LW], and prove the bound for the second algorithm from first principles. We refer to the optimal static list as *OPT*.

2.1. *Algorithm A.* Let L be a list of n elements and let σ be a sequence of accesses presented to the algorithm. Algorithm A proceeds as follows:

1. Pick n numbers r_1, \dots, r_n from $[0, 1]$ uniformly at random such that $\sum_{i=1}^n r_i = 1$. (This is easily done by throwing $n - 1$ darts into $[0, 1]$ and letting r_i be the length of the i th interval.) Assign r_i to the i th element in L .

Let $\beta < 1$ be a performance parameter for the algorithm.

2. Assign a weight w_i to each element in the list, with initial value 1. Order elements in the list in increasing order of $r_i w_i$.
3. When an element x at position i is accessed, update $w_i = \beta w_i$.
4. Move element x up in the list, if necessary, to maintain ordering of elements by weight.
5. Go to 3.

THEOREM 2.1. *Algorithm A with $\beta = e^{-\varepsilon}$ is $(1 + \varepsilon)$ -competitive with the best static list.*

In the following analysis, we use L_{xy} to denote the sublist of L that contains only the elements x and y . We use σ_{xy} to denote the subsequence of sequence σ consisting of only the elements x and y .

LEMMA 2.2 (Pairwise Property Lemma [BEY]). *An algorithm ALG satisfies the pairwise property if and only if for every request sequence σ and every pair $\{x, y\} \in L$, the probability that x precedes y in L when ALG serves σ , is the same as the probability that x precedes y in L_{xy} when ALG serves σ_{xy} .*

Let ALG be an algorithm that satisfies the pairwise property. Suppose that for every pair $\{x, y\} \subseteq L$, and for every request sequence σ , $\text{ALG}(\sigma_{xy}) \leq c \cdot \text{OPT}(\sigma_{xy})$ in the Partial Cost model, then ALG is strictly c competitive in both the Partial Cost and Full Cost models.

LEMMA 2.3. *Algorithm A satisfies the pairwise property.*

PROOF. Consider two elements x and y in L . The relative order of x and y in L depends only on which of $w_x r_x$ or $w_y r_y$ is greater. What we need to show is that $\Pr[w_x r_x < w_y r_y]$ is unaffected by any of the other elements in the list or request sequence.

Let $R = 1 - \sum_{i \neq x, y} r_i$. The distribution over the r_i is such that conditioned on R , r_x and r_y are random in $[0, R]$ subject to having a sum of R . So, $\Pr[w_x r_x < w_y r_y] = \Pr[r_x / r_y < w_y / w_x] = w_x / (w_x + w_y)$. Notice that this is independent of R (and everything except w_x and w_y). Therefore, the probability that x precedes y depends only on the number of requests so far to x and y , and thus A has the pairwise property. \square

LEMMA 2.4. *For a list of two elements, algorithm A has a performance ratio of $\log(1/\beta)/(1 - \beta)$ compared with the best static list.*

PROOF. Consider a list L of two elements x and y . If a request accesses the element which is second in the list at the time of access, the algorithm incurs a cost of 1; otherwise, it incurs no cost. There are only two static lists on two elements— (x, y) and (y, x) . We consider these as experts. When element y is accessed, the list (x, y) incurs a cost of 1 (or makes a “mistake”), while (y, x) incurs no cost (so it is “correct” on this request). The reverse happens when element x is accessed.

Notice that if we associate weight w_y with the list (x, y) and w_x with the list (y, x) , then the weight of each expert is $\beta^{(\text{\# of mistakes made by that expert})}$. Algorithm A uses one of the two experts with probability in proportion to its weight because, as noted in the proof of Lemma 2.3, $\Pr[w_x r_x < w_y r_y] = w_y / (w_x + w_y)$. This is the same as the Randomized Weighted Majority algorithm of [LW] for the case of two experts. Therefore, following [LW], we get

$$E[A(\sigma_{xy})] \leq \frac{OPT(\sigma_{xy}) \log(1/\beta) + \log 2}{1 - \beta}. \quad \square$$

PROOF OF THEOREM 2.1. First, notice that algorithm A does not incur any movement costs other than ordering the list initially after choosing the weights. This is because the relative order of elements changes only when one of them is accessed, and, at that time, the list is reordered just by moving up the accessed element. The initial movement cost is at most $n(n - 1)/2$.

Using the *List Factoring Lemma* [BEY] with Lemmas 2.3 and 2.4, we obtain the competitive ratio of algorithm A with respect to the cost incurred by OPT as

$$E[A(\sigma)] \leq \frac{OPT(\sigma) \log(1/\beta) + O(n(n - 1))}{1 - \beta}.$$

Putting $\beta = e^{-\varepsilon}$, we get a competitive ratio of $(1 + \varepsilon)$ for A with an additive cost of at most $O(n(n - 1)(1 + 1/\varepsilon))$. \square

2.2. *Algorithm B.* Algorithm B is a simpler variant of algorithm A. The basic difference between the two is that here we pick the initial random numbers iid, and we update

additively instead of multiplicatively. The algorithm proceeds as follows:

1. Pick n numbers r_1, \dots, r_n from $[0, 1/\varepsilon]$ independently uniformly at random. Assign r_i to the i th element in L .
2. Assign a weight w_i to each element in the list, which tracks the number of accesses to that element. Initially, all $w_i = 0$. Order elements in the list in decreasing order of $r_i + w_i$.
3. When an element x at position i is accessed, increase w_i by 1.
4. Move element x up if necessary to maintain ordering by weights.
5. Go to step 3.

THEOREM 2.5. *Algorithm B is $(1 + O(\varepsilon))$ -competitive with the optimal static list.*

The proof of Theorem 2.5 closely follows the proof of Theorem 2.1. We first note that algorithm B also satisfies the pairwise property, because the probability that x is above y in the list is just $\Pr[r_x + w_x > r_y + w_y]$ which does not depend on any of the other elements. The following lemma analyzes performance of algorithm B on a two element list.

LEMMA 2.6. *For a list of two elements, algorithm B has a competitive ratio of $1 + O(\varepsilon)$ with the optimal static list.*

PROOF. Consider a list L of two elements x and y . The probability that x precedes y after w_x accesses to x and w_y to y , is given by $\Pr[r_x + w_x > r_y + w_y] = \Pr[r_x - r_y > w_y - w_x]$. Let $\Delta = w_y - w_x$. Since both r_x and r_y are chosen uniformly at random from $[0, 1/\varepsilon]$, this probability $p(\Delta) = (1 - \varepsilon\Delta)^2/2$ if $\Delta > 0$, and $p(\Delta) = 1 - (1 + \varepsilon\Delta)^2/2$ otherwise. When $\Delta > 1/\varepsilon$ or $\Delta < -1/\varepsilon$, $p(\Delta) = 0$ or 1 , respectively.

As before, we consider the two lists (x, y) and (y, x) as experts. The probability of picking expert (x, y) then becomes $p(\Delta)$. $p(\Delta)$ takes on $2/\varepsilon + 1$ distinct values corresponding to $\Delta \in \{-1/\varepsilon, \dots, 1/\varepsilon\}$. As the algorithm proceeds, this probability either moves one step up or one step down depending on whether (x, y) or (y, x) makes a mistake, respectively.

The algorithm starts at $\Delta = 0$. As it moves through various values of Δ , we can couple every movement from $\Delta = x$ to $\Delta = x + 1$ with a movement from $\Delta = x + 1$ to $\Delta = x$, with at most $2/\varepsilon$ movements left uncoupled. Notice that this coupling has the property that both experts make exactly one mistake among the two moves ((x, y) while moving up and (y, x) while moving down).

Algorithm B on the other hand, makes an expected number of mistakes equal to $p(x) + 1 - p(x + 1)$. For $x \geq 0$, $p(x) + 1 - p(x + 1) = 1 + \varepsilon - \varepsilon^2(x + \frac{1}{2}) < 1 + \varepsilon$. For $x < 0$, $p(x) + 1 - p(x + 1) = 1 + \varepsilon + \varepsilon^2(x + \frac{1}{2}) < 1 + 2\varepsilon$.

The total number of mistakes made by algorithm B on uncoupled steps are at most $O(1/\varepsilon)$. Putting these together, we get a competitive ratio of $1 + O(\varepsilon)$ with the best static list with an additive term of $O(1/\varepsilon)$. \square

PROOF OF THEOREM 2.5. The result follows, as before, from the *List Factoring Lemma* and Lemma 2.6. \square

Notice that both algorithms A and B can be thought of as randomized versions of the frequency count algorithm.

3. Combining Statically and Dynamically Optimal Algorithms. In the previous section we described two algorithms which are variants of the Experts algorithm and are $(1 + \varepsilon)$ optimal with respect to the best static list. In this section we reapply the experts technique, this time using algorithm A (or B) and algorithm COMB [BEY] as the two experts. There are some issues that need to be addressed, but through this we can achieve both $(1 + \varepsilon)$ static optimality and $(1.6 + \varepsilon)$ dynamic optimality simultaneously.

The key idea is to use one of the experts for some number of accesses, and then probabilistically decide to switch the expert or stay with the same algorithm. However, in making the probabilistic decision, we do not want to calculate the weights of the two experts explicitly, as that would require running the two algorithms simultaneously, defeating the very purpose of this algorithm.⁶ In order to circumvent this difficulty, we use a variant of the Exp3 algorithm proposed by Auer et al. [ACB⁺]. This algorithm was proposed for the multi-armed bandit problem in which a gambler must decide to play on one of K slot machines, but gets to see the profit of only the machine that he is playing on.

Burch [B] extends the Exp3 algorithm to the case where there is a cost d for switching between the different experts. We will use this to deal with the fact that when we switch algorithms, we may have a different internal state (probability distribution over lists) than had we used the other algorithm from the start. We now describe the two algorithms Hedge and Hedge-Bandit as given in [B].

3.1. Algorithm Hedge. Algorithm Hedge uses K algorithms (experts), choosing one for each time step and obtaining a gain equal to the gain of the chosen expert in that time period. Expert i obtains a gain $x_i(t)$ in time step t with $x_i(t) \in [0, 1]$. Define $G_i(t) = \sum_{j < t} x_i(j)$, the total gain obtained by expert i up to time step t . Let $\eta > 0$ be a performance parameter. The algorithm proceeds as follows:

In every time step the algorithm chooses an expert i_t according to the distribution $p_i = e^{\eta G_i(t-1)} / \sum_j e^{\eta G_j(t-1)}$, and updates the gains according to the obtained vector $x(t)$.

3.2. Algorithm Hedge-Bandit. Algorithm Hedge-Bandit is similar to Hedge except that it does not get to see the entire gain vector $x(t)$, but only gets to see the gain of the expert that it chooses at a particular time step. In each time step, the algorithm simulates Hedge to obtain a distribution that it uses to select the expert for that time step. It then returns a “fake” gain vector $\hat{x}(t)$ to Hedge based on the gain $x(t)$ that it actually observed, in order to get the next distribution.

To avoid spending too much in switching between experts, this algorithm runs each expert for s steps before making the next decision. It obtains a gain of at most s in a single time segment, and loses at most d in switching from one expert to another. To

⁶ If we only care about access cost, and not the running time of the algorithms, then this is not a problem.

use this algorithm for list update, in every time step, we return a gain of $1 - i/n$, if the access cost for an element is i . We assume that the algorithm runs for T time steps. The following describes the behavior of the algorithm in each time segment:

For time segments $t = 1, 2, \dots, T/s$ do:

- Get $\hat{p}_i(t)$ from Hedge. Let $p_i(t) = (1 - \gamma)\hat{p}_i(t) + \gamma/K$.
- Select i_t according to $p_i(t)$.
- Use expert i_t for s steps observing a total gain of $x_{i_t}(t)$. Send gains $\hat{x}_j(t)$ to Hedge, where $\hat{x}_j(t) = x_{i_t}(t)/p_{i_t}(t)$ if $j = i_t$, and 0 otherwise.

THEOREM 3.1 (Theorem 5.3 of [B]). *The expected gain of Hedge-Bandit is at least*

$$G^* - (1 - \gamma)\frac{T}{s}d - \frac{1 - \gamma}{\gamma}sK \ln K - (e - 1)\gamma G^*,$$

where G^* is the largest total actual gain acquired by any single expert, and where $\eta = e^{\gamma/sn}$.

For the List Update problem we have $K = 2$. We also have $d = O(n^2)$ for reasons described below. Reverting back to our loss model and choosing appropriate values for s and γ , we obtain the following result:

THEOREM 3.2. *Algorithm Hedge-Bandit using algorithms A and COMB as experts is $(1 + 2\varepsilon)$ statically optimal and $1.6(1 + \varepsilon)$ dynamically optimal in the Full Cost model, with an additive term of $O(n^5/\varepsilon^2)$ for the List Update problem, if it makes a switching decision every en^3/ε steps and uses $\gamma = \varepsilon/en$.*

PROOF. From Theorem 3.1,

$$G_{\text{HB}} \geq G^* - (1 - \gamma)\frac{T}{s}d - \frac{1 - \gamma}{\gamma}sK \ln K - (e - 1)\gamma G^*.$$

Putting $L_{\text{HB}} = n(T - G_{\text{HB}})$, $L^* = n(T - G^*)$ and using $G^* < T$, we obtain the following:

$$\begin{aligned} L_{\text{HB}} &\leq n \left\{ T - G^* + (1 - \gamma)\frac{T}{s}d + \frac{1 - \gamma}{\gamma}sK \ln K + (e - 1)\gamma G^* \right\} \\ &= L^* + (1 - \gamma)\frac{T}{s}dn + \frac{1 - \gamma}{\gamma}nsK \ln K + (e - 1)\gamma nT - (e - 1)\gamma L^*. \end{aligned}$$

However, $L^* \geq T$ in the full cost model (we always incur a cost of at least 1). So we have

$$L_{\text{HB}} \leq L^* \left\{ 1 + (1 - \gamma)\frac{dn}{s} + (e - 1)\gamma(n - 1) \right\} + \frac{1 - \gamma}{\gamma}nsK \ln K.$$

Using the values $\gamma = \varepsilon/en$, $s = en^3/\varepsilon$, $d = n^2$, and $K = 2$ gives us the desired result:

$$L_{\text{HB}} \leq (1 + \varepsilon)L^* + 8.9\frac{n^5}{\varepsilon^2}. \quad \square$$

We omitted a few details in the above description of the algorithm. When we switch from one algorithm to the other at the end of a phase, we need to know the state of the other algorithm at the end of that phase. It is not obvious that this is calculable, as we want to avoid running the other algorithm explicitly. However, it can be argued that extra overhead incurred in starting the algorithms in a “wrong” state is at most $O(n^2)$. Alternatively, the state can be easily determined without running the algorithms explicitly, by keeping access counts for each element in the case of algorithm A, and keeping timestamps for each element in the case of COMB. These variables can be updated in constant time per access. Unfortunately, although using these extra variables would help us determine the next state of the list, it is not clear how to use it to calculate the losses made by the experts, which necessitates the use of the Hedge-Bandit algorithm.

4. Search Trees. An on-line binary search tree algorithm is *dynamically optimal* if its total cost (sum of search cost and number of rotations) is never more than a constant times the total cost of the best off-line algorithm, for any sequence of accesses. Much work has gone into an attempt to prove the dynamic optimality of splay trees since Sleator and Tarjan made their conjecture in [ST2]. We are unable to prove the dynamic optimality of splay trees or any other on-line search tree algorithm.

Instead, we show that it is possible to have *dynamic search-optimality*, a property much weaker than dynamic optimality. In general the total cost of an algorithm is its search cost, the sum of the depths of the accesses, plus its rotation cost, the number of rotations made. We say that our algorithm has dynamic search-optimality because its search cost is at most a constant times the *total cost* of any off-line algorithm. While our algorithm is exponentially slow, it is based on simple principles.

There are two main difficulties in achieving dynamic optimality. First, the algorithm has to decide which nodes to keep near the root.⁷ Secondly, for full dynamic optimality it has to be able to get these nodes near the root without using too many rotations. We show that the first difficulty is not insurmountable. Equivalently, dynamic optimality cannot be proven impossible by any argument that neglects the rotation cost of the on-line algorithm.

There is no especially strong evidence suggesting that any BST algorithm is dynamically optimal. A natural approach to disproving this would be to present a set of sequences, and argue that no on-line algorithm can handle all these sequences in a dynamically optimal manner. The simplest form of this argument would be that the on-line algorithm cannot “guess” which node comes next and therefore has too many nodes to keep near the root. However, the existence of an algorithm with dynamic search optimality implies that this type of argument is not possible.

Wilber [W] made some progress by proving lower bounds for off-line algorithms. In particular, he has shown that a random sequence of accesses in $\{1, 2, \dots, n\}$ costs

⁷ Since rotations are free, a natural idea is at each step, to choose the tree of the optimal off-line algorithm so far. This is not as simple as it sounds, because off-line optimality is ambiguous. For example, suppose you start with a seven node “line” tree of depth 7. After several accesses to node 1, the deepest node, any optimal off-line algorithm must have brought it to the root. There are 132 equally optimal ways to do this, all using six rotations.

an expected $\Omega(\log n)$ per access for off-line algorithms. This is a necessary condition for dynamic optimality to be possible, because any on-line algorithm pays an expected $\theta(\log n)$ just in search cost (not counting rotations) since the average depth of a node is $\theta(\log n)$.

We go slightly farther in analyzing the cost of random sequences. We show that the number of sequences with off-line cost k is less than 2^{12k} for any k . The fact that there are only $2^{O(k)}$ sequences of cost k is also a necessary condition for dynamic optimality, for information-theoretic reasons.⁸ Essentially, we give a way to describe off-line rotations in 12 bits per rotation, even though there are, in general, $n - 1$ possible rotations one can perform.

4.1. Structural Preliminaries. We assume that the nodes in the tree are simply the numbers $1, 2, \dots, n$, and let m be the length of the access sequence. We further assume that all algorithms, on- and off-line, begin with the same fixed tree, say, rooted at n and having depth n . In this section we prove the following. The constant 12 is not really important: we are mostly interested in the fact that it is $2^{O(k)}$.

THEOREM 4.1. *The number of access sequences having optimal off-line cost k is at most 2^{12k} , for all $k \geq 0$, regardless of n or m .*

PROOF. We use an information-theoretic argument based on the fact that one can concisely describe any sequence having optimal off-line cost k . We argue that it is possible to describe any access sequence via the trees used in the optimal off-line algorithm. First, as several people have observed, we may assume that the off-line algorithm rotates the next node to be accessed to the root before each access, and we only lose a factor of 2 in the optimal cost. The reason is that given any off-line algorithm, we can modify it by making it rotate a node to the root immediately before accessing it, and then reverse these rotations to move the node back to where it was. If the node was at depth d , we have paid $2(d - 1)$ in rotations (and incurred no additional search cost), whereas the former algorithm paid only $d - 1$.

Now it suffices to describe the sequence of trees, because the accesses will just be their roots. To do this, we describe the set of rotations performed from each tree to the next, which we show how to do in at most $6r$ bits if there are r rotations. This implies that there are at most 2^{6k} possible sequences of cost k because there are at most 2^{6k} descriptions of length $6k$. We then get 2^{12k} because of the factor of 2 lost in our initial assumption.

What remains is to describe a set of r rotations using $6r$ bits. Like Lucas [L], we think of a rotation as an edge rotation which changes a single edge from either left to right or right to left. Of course, the nodes adjacent to an edge may change. However, based on

⁸ Given an on-line algorithm, such as splaying, any access sequence can be described by the location of each node in its corresponding tree. This can be described using three symbols (left, right, and stop) and has length proportional to the search cost. There is no need to describe the rotations performed by the on-line algorithm, since those can be determined from the access sequence. So, for any on-line algorithm, there are at most 3^k sequences costing less than k .

our assumption that the next access is rotated to the root, it is not difficult to see that all the edges on its path to the root must be rotated at least once.

Lucas argues that without further loss of generality we may assume that the set of edges rotated by an optimal off-line algorithm form a single connected component that includes the root and the next node to be accessed. Briefly, this is because any rotations of edges not in this connected component could easily be delayed (using lazy programming) until they are in such a connected component. Their delay will not affect the search cost, since these rotations cannot affect the depth of the next request, nor does their delay affect the rotation costs.

Next, observe that regardless of the order of the rotations, we can completely describe the result of the rotations in $6w$ bits if there are w edges. First, we describe the subset of edges that were rotated one or more times. Since this is a rooted subtree, we can describe this using the symbols left (00), right (01), and up (1) to form a cycle that traverses each edge twice, using a total of $3w$ bits. Next we describe their position in the resulting tree. In the resulting tree these edges will still be a rooted subtree, so we can describe them also with $3w$ bits. First, note that the set of nodes in this subtree does not change even as the positions of the edges do. Secondly, notice that the shape of this subtree completely determines the positions of all the nodes, because this is a binary search tree. Finally, note that the off-line algorithm has to perform at least w rotations.

Thus, we can describe the optimal sequence of trees (and thus the access sequence) in $6r$ bits if it performs r rotations. Since we lose a factor of 2 due to our first assumption, this proves the theorem. \square

4.2. Dynamic Search-Optimality. In this section we consider probability distributions. A tree can be thought of as predicting the next access, where it predicts nodes closer to the root with higher probability. From our structural result, we see:

COROLLARY 4.2. *There is a probability distribution over arbitrary sequences of accesses that assigns probability at least 2^{-13k} to an access sequence of optimal off-line cost k .*

PROOF. Choose a cost k according to the distribution $1/2^k$, then pick a random sequence of that cost. By Theorem 4.1, there are at most 2^{12k} sequences of cost k , so the probability on any given sequence is at most 2^{-13k} . \square

It is easy to convert a binary search tree into a probability distribution p over elements such that $p(j) \geq 3^{-\text{depth}(j)}$. Simply choose j by beginning at the root, going left, right, or stopping, each with probability $1/3$ (when possible). It is also possible to convert a probability distribution over elements into a tree.

OBSERVATION 4.1. *For a probability distribution p over individual accesses, we can create a binary search tree such that $\text{depth}(a) \leq 1 - \log p(a)$ for any node a .*

PROOF. For the root, choose the first i such that $\sum_1^{i-1} p(i) \leq \frac{1}{2}$ and $\sum_{i+1}^n p(i) \leq \frac{1}{2}$. Recurse on the numbers less than i (normalizing p) to create the left subtree, and the

numbers greater than i for the right subtree. It is easy to see that the total probability of any subtree rooted at depth d is at most $1/2^{d-1}$ so a node of probability $p(i)$ cannot be deeper than $-\log p(i)$. \square

We can combine these two ideas to make an on-line algorithm.

THEOREM 4.3. *For any probability distribution p over access sequences, we can create an on-line algorithm with search cost at most $m - \log p(a_1 a_2 \cdots a_m)$ for every access sequence $a_1 a_2 \cdots a_m$.*

PROOF. The distribution p can be thought of as predicting the next access from the previous accesses. In particular, the conditional probability of the next access given the previous accesses is

$$p_i(a_i) = p(a_i \mid a_1 a_2 \cdots a_{i-1}) = \frac{\sum_{b_{i+1}, \dots, b_m} p(a_1 \cdots a_{i-1} a_i b_{i+1} \cdots b_m)}{\sum_{b_i, \dots, b_m} p(a_1 \cdots a_{i-1} b_i b_{i+1} \cdots b_m)}.$$

We can write p as a product of the conditional distributions of access i , i.e.,

$$p(a_1 a_2 \cdots a_m) = \prod_1^m p_i(a_i).$$

Our on-line algorithm works as follows. For the i th access, we have enough information to compute p_i . We then convert p_i into a tree by the method of Observation 4.1. Thus, the depth of access a_i in this tree will be no more than $1 - \log p_i(a_i)$. Our total search over m accesses is at most

$$\sum_1^m 1 - \log p_i(a_i) = m - \log p(a_1 a_2 \cdots a_m). \quad \square$$

COROLLARY 4.4. *There is an on-line algorithm that has dynamic search-optimality. In particular, on any access sequence, its search cost is at most 14 times the optimal off-line total cost.*

PROOF. We use the probability distribution of Corollary 4.2 in combination with Theorem 4.3 to get a total cost of m plus 13 times the optimal off-line cost. However, m is no larger than the optimal off-line cost. \square

5. Conclusions and Open Problems. In this paper we have presented two results: an efficient algorithm achieving strong static optimality for the list-update problem, and an inefficient algorithm achieving dynamic search-optimality for trees. Recently Kalai and Vempala [KV2] gave an efficient algorithm achieving strong static optimality for Binary Search Trees. Several natural open problems are: Can dynamic search-optimality be achieved efficiently for trees? and, of course, Can we achieve true dynamic optimality (efficiently or inefficiently) for search trees?

Acknowledgments. We thank Santosh Vempala for a number of helpful discussions.

References

- [ACB⁺] Peter Auer, Nicolò Cesa-Bianchi, Yoav Freund, and Robert Schapire. Gambling in a rigged casino: the adversarial multiarmed bandit problem. In *Proceedings of the 36th Annual Symposium on Foundations of Computer Science*, pages 322–331, November 1995.
- [AvSW] Susanne Albers, Bernhard von Stengel, and Ralph Werchner. A combined BIT and TIMESTAMP algorithm for the list update problem. *Information Processing Letters*, 56:135–139, 1995.
- [B] Carl Burch. Machine learning in metrical task systems and other on-line problems. Technical Report CMU-CS-00-135, Carnegie Mellon University, May 2000.
- [BB] Avrim Blum and Carl Burch. On-line learning and the metrical task system problem. *Machine Learning*, 39(1):35–58, April 2000.
- [BEY] Allan Borodin and Ran El-Yaniv. *Online Computation and Competitive Analysis*. Cambridge University Press, Cambridge, 1998. <http://www.cup.org/Titles/56/0521563925.html>.
- [CLS] Deepak Chawla, Lin Li, and Stephen Scott. Efficiently approximating weighted sums with exponentially many terms. In *Proceedings of the Fourteenth Annual Conference on Computational Learning Theory*, pages 82–98, July 2001.
- [FS] Yoav Freund and Robert Schapire. Game theory, on-line prediction and boosting. In *Proceedings of the 9th Conference on Computational Learning Theory*, pages 325–332, 1996.
- [KV1] Adam Kalai and Santosh Vempala. Efficient algorithms for universal portfolios. In *Proceedings of the 41st Annual Symposium on the Foundations of Computer Science*, pages 486–491, November 2000.
- [KV2] Adam Kalai and Santosh Vempala. Follow the leader for online algorithms. Unpublished manuscript.
- [L] Joan M. Lucas. Canonical forms for competitive binary search tree algorithms. Technical Report DCS-TR-250, Rutgers University, 1988.
- [LW] Nicolò Littlestone and Manfred K. Warmuth. The weighted majority algorithm. *Information and Computation*, 108:212–261, 1994.
- [ST1] Daniel Sleator and Robert E. Tarjan. Amortized efficiency of list update and paging rules. *Communications of the ACM*, 28:202–208, 1985.
- [ST2] Daniel D. Sleator and Robert Endre Tarjan. Self-adjusting binary search trees. *Journal of the ACM*, 32:652–686, 1985.
- [T] Boris Teia. A lower bound for randomized list update algorithms. *Information Processing Letters*, 47:5–9, 1993.
- [W] Robert Wilber. Lower bounds for accessing binary search trees with rotations. *SIAM Journal on Computing*, 18:56–67, 1989.