

Dr. Frankenstein's Approach to On-line Algorithms

(Extended abstract)

Yossi Azar * Andrei Z. Broder* Mark S. Manasse*

Abstract

Let $\{A_1, A_2, \dots, A_m\}$ be a set of on-line algorithms for a problem P with input set I . We assume that P can be represented as a metrical task system. Each A_i has a competitive ratio a_i with respect to the optimum off-line algorithm, but only for a *subset* of the possible inputs such that the union of these subsets covers I . Given this setup, we construct a generic deterministic on-line algorithm and a generic randomized on-line algorithm for P that are competitive over *all* possible inputs. We show that their competitive ratios are optimal up to constant factors. Our analysis proceeds via an amusing card game.

*DEC Systems Research Center, 130 Lytton Ave. Palo-Alto, CA 94301.

1 Introduction

A common trick of the trade in algorithm design is to combine several algorithms using round robin execution. The basic idea is that, given a set of m algorithms for a problem P , one can simulate them one at a time in round robin fashion until the fastest of them solves P on the given input. It is easily seen that round robin execution is optimal among deterministic combining algorithms that have no specific knowledge of the problem domain and input. (As an aside, we show in the full version of this paper that randomization helps very little in this context; for any randomized combining scheme, and any $\epsilon > 0$, there is an input set of algorithms, such that the expected cost is greater than $(m - \epsilon)$ times the minimum cost, versus m times the minimum cost for round robin execution.)

For on-line algorithms the situation is more complicated: We are given a set $S = \{A_1, A_2, \dots, A_m\}$ of on-line algorithms (deterministic or randomized – see below) for a problem P with input set I . Each algorithm A_i has a known competitive ratio a_i with respect to the optimum off-line algorithm, but only for a *subset* of the possible inputs, such that the union of these subsets covers I . We assume that P can be represented as a metrical task system (see [2] for definitions). Our goal is to construct an on-line algorithm for P that is competitive over *all* possible inputs.

Again, we are interested in algorithms that have no specific knowledge of the problem domain and input. More precisely, let σ_t be the sequence of requests up to time t . Let $C_i(\sigma_t)$ and $c_i(\sigma_t)$ be the configuration (respectively the cost) associated to A_i serving σ_t . At time t , the configuration associated to the combining algorithm must be one of the $C_i(\sigma_t)$'s. The decision to switch from $C_i(\sigma_{t-1})$ to $C_j(\sigma_t)$ can be based only on the values $c_i(\sigma_{t'})$ for $1 \leq i \leq m$ and $0 \leq t' \leq t$ and on no other domain or input specific information.

If the algorithms A_i and the combining construction are deterministic, then we call the new algorithm *on-line combine* and denote it MIN_S . If the construction uses random bits, we call the algorithm *randomized on-line combine*, denoted RMIN_S . In this later case the A_i 's might be randomized as well.

Fiat *et al.* [3] addressed the question of on-line combine in a context restricted to paging algorithms. Fiat, Rabani, and Ravid [5] considered the general case and showed that constructing a MIN_S algorithm for an arbitrary set of on-line algorithms is equivalent to the *layered graph traversal* problem analyzed by Papadimitriou and Yannakakis [6] and Baeza-Yates, Culberson, and Rawlins [1]. Using the results of these analyses, they obtained a MIN_S algorithm with competitive ratio $O(m \cdot \max_i \{a_i\})$, which is optimal up to a constant factor when all the a_i 's are equal, but not in general. In this paper we completely solve the general case, when the a_i 's are arbitrary. This immediately yields a better competitive ratio for the k -server algorithm of [5].

In this paper, we show that the problem of combining on-line algorithms is

equivalent, up to a small constant factor, to finding the value of a very simple two-player card game¹. In this game, two identical decks of cards are given to two players. Simplifying slightly, the first player (corresponding to the on-line combining algorithm) places a card face-down on the table. The second player (the adversary) chooses a card from his hand, and turns it face up on the table. The first player then exposes the matching card, either from her hand, in which case no score is recorded, or by showing the card on the table, in which case the second player wins the value of the card. The pair is removed from play, and the players play another round with the reduced decks, until both players run out of cards. A few variations exist, all of which turn out to be equivalent in terms of optimal strategy: the first player can be required to pick the same card to place face down for each round until it is matched by the second player, or not; the second player can be required to arrange the order in which he will play cards in every round before play commences; one can even require the first player to select a schedule in advance for which card she will place face down next, except for those cards that are matched before their turn comes. It turns out that optimal play for all these variants results in the same total or expected score for the second player. In the deterministic case, the total is clearly the value of all the cards, since the second player can inspect the first player's strategy, and play cards in exactly the same order. The randomized case has the following amusing optimal strategy for building a schedule: let each card have a probability proportional to the inverse of its value, and choose a card using that distribution. That card is the last card in the schedule. Repeat, with the chosen card excluded, to find the schedule in reverse order.

Using this card game, we obtain the following results:

Theorem 1 *Let $S = \{A_1, A_2, \dots, A_m\}$ be a set of deterministic on-line algorithms for a metrical task system P with input set I . Assume that each A_i has a competitive ratio a_i with respect to the optimum off-line algorithm for a subset of the possible inputs such that the union of these subsets covers I . Then there exists a deterministic MIN_S algorithm with competitive ratio $O(\sum_{i=1}^m a_i)$, and no deterministic on-line algorithm can do better in general, except for a constant factor.*

The improvement with respect to the previous bounds is relevant when the average of the a_i 's is substantially smaller than their maximum. In particular, our MIN_S algorithm reduces the competitive ratio of the k -server algorithm of [5] by a factor of $k!/2^{O(k)}$. (See section 6.)

For the randomized case we need first to discuss a function that will play an important role in what follows. Let a_1, a_2, \dots be a sequence of positive numbers.

¹Field experiments show that 5-year olds can easily play it...

For any set T of natural numbers we define $f(T)$ by the recurrence

$$\begin{aligned} f(\emptyset) &= 0 \\ f(T) &= \frac{1 + \sum_{i \in T} f(T \setminus \{i\})/a_i}{\sum_{i \in T} 1/a_i}, \quad T \neq \emptyset. \end{aligned} \tag{1}$$

Let $[m]$ stand for the set $\{1, 2, \dots, m\}$. Note that $f([m])$ is a symmetric rational function of a_1, \dots, a_m . In particular

$$\begin{aligned} f(\{1\}) &= a_1, \\ f(\{1, 2\}) &= \frac{a_1^2 + a_2^2 + a_1 a_2}{a_1 + a_2}, \end{aligned}$$

but $f(\{1, 2, 3\})$ has 19 terms, $f(\{1, 2, 3, 4\})$ has 390, and submissions are limited to 10 pages. Nevertheless, we can crudely bound $f(T)$ by

$$H_m \min_{i \in T} a_i \leq f(T) \leq H_m \max_{i \in T} a_i, \tag{2}$$

where $m = |T|$, and H_m is the m 'th harmonic number. Better but more complex bounds will be presented in Section 3.1.

Now we can state our result for randomized on-line combine.

Theorem 2 *Let $S = \{A_1, A_2, \dots, A_m\}$ be a set of deterministic or randomized on-line algorithms for a metrical task system P with input set I . Assume that each A_i has a competitive ratio a_i with respect to the optimum off-line algorithm for a subset of the possible inputs such that the union of these subsets covers I . Then there exists a randomized RMIN_S algorithm with competitive ratio $O(f([m]))$, and no randomized on-line algorithm can do better in general, except for a constant factor.*

Plugging equation (2) into the theorem yields the weak upper bound $O(\log n \cdot \max_i a_i)$ which was obtained in [4].

The remainder of the paper is as follows: in Section 2 we discuss the equivalence between on-line combine and layered graph traversal introduced in [5]; in Section 3 we describe the card game more completely and analyze it, discovering that the value of the game is the mysterious $f(T)$ above; in sections 4 and 5 we use the game analysis to derive lower and upper bounds for the layered graph traversal problem.

2 The layered graph traversal problem

This problem was introduced and analyzed in [1] and [6].

A *layered graph* consists exactly of m paths with a common first vertex s , called the *source*, but otherwise vertex disjoint. Thus, the graph can be divided

into layers $L_0 = \{s\}, L_1, L_2, \dots$, such that layer i for $i > 0$ consists of the m vertices that are i edges away from the source on each path. Each edge e , has a certain non-negative length $l(e)$.

An on-line layered graph traversal (LGT) algorithm starts at the source and moves along the edges of the graph. Each time it moves along an edge (in any direction), it pays a cost which is the length of the edge. Its goal is to reach a *target* which is a vertex in the last layer. The lengths of the edges between layer L_{i-1} to L_i are revealed to the algorithm only when a vertex in L_{i-1} is reached for the first time. (The lengths do not change over time.) The target vertex becomes known only when the algorithm reaches a vertex in the next-to-last layer.

Observe that any LGT algorithm must advance one layer at a time either by continuing on its current path, or by backtracking to the source and choosing a different path.

The competitive ratio of the on-line traversal algorithm is the worst case ratio between the distance traveled by the on-line algorithm and the length of the (unique) shortest path from the source to the target.

The optimal deterministic algorithm for this problem has competitive ratio $1 + 2m(1 + \frac{1}{m-1})^{m-1} \approx 2em$ (see [6] and [1]). For randomized algorithms Fiat *et al.* [4] showed that the competitive ratio is $\Theta(\log m)$.

For our purposes we need a slight generalization of this model. We assume that each path P_i has a known associated *waste factor* a_i . For each edge e on P_i , the off-line algorithm pays $l(e)/a_i$, while the on-line algorithm pays $l(e)$ as before. Thus the competitive ratio becomes a function of a_1, \dots, a_m , and the preceding model corresponds to $a_1 = a_2 = \dots = a_m = 1$.

Following [5] we show now that constructing a MIN_S algorithm is equivalent to an algorithm for the modified LGT with the same competitive ratio.

First assume that a (modified) LGT algorithm is given. To construct a MIN_S algorithm, we construct a layered graph which associates a path P_i with each algorithm A_i . We set the waste factors to be the competitive ratios a_1, \dots, a_m and simulate A_1, \dots, A_m on the sequence of requests as follows.

When a request is made, the MIN_S algorithm computes the costs of the edges to the next layer; the cost of the edge on P_i is the cost of serving the request by A_i , as if A_i had been continually simulated from the beginning.

Then, the MIN_S algorithm applies the LGT algorithm in order to decide how to serve the request. If the LGT algorithm continues with the current path P_i , then MIN_S continues to simulate the current algorithm, A_i . On the other hand, if the LGT algorithm backtracks and moves to a vertex v (in the next layer) via another path, P_j , then the MIN_S algorithm switches to the configuration corresponding to v , and A_j becomes the current algorithm. Since we assumed that the underlying problem is a metrical task system, the triangle inequality holds for the cost of switching between configurations; thus the cost of MIN_S is bounded by the cost of LGT. Clearly then, a competitive LGT algorithm yields a MIN_S algorithm with

the same competitive ratio, or better.

Conversely, one can easily use a MIN_S algorithm to construct an LGT algorithm with the same competitive ratio: Let the metrical task system P be the layered graph traversal, where the states correspond to vertices in the graph, with the transition cost between states equal to the total distance in the graph. (It is readily seen that P is well defined.) Let A_1, A_2, \dots, A_m be the m algorithms that correspond to sticking to path P_i and let a_1, \dots, a_m be the waste factors. Clearly, the LGT algorithm that follows MIN_S in the obvious manner has the same competitive ratio.

3 The Guess Game

In this section we define and analyze a certain two player zero-sum game, called the *Guess Game*. Later we will use this analysis to derive upper and lower bounds for the layered graph traversal problem.

One participant is called the *player* and the other is called the *adversary*. Both start with the same set of cards $T = [m] = \{1, \dots, m\}$. The value of card i is $a_i > 0$.

The game starts with the adversary putting all his cards face-down on the table in a certain order, that he will be unable to change during the game. Then the player chooses one of his cards and puts it face-down on the table. We call this the *hidden* card. The adversary then turns up the first of his cards and the player has to match it. If the card matches the hidden card (a hit) then the adversary wins the value of the card, the matched pair is discarded, and the player must pick a new card face down. If not (a miss), then the player matches the adversary's card with a card from his hand and the matched pair is discarded without further ado.² Hence, there are m rounds. The value of the game is the sum of the values of the cards that the adversary wins.

Observe that the player pays a_i if and only if he hides card i before he hides any card that comes after i in the adversary's order. In particular the player always pays for the last card in the adversary's order.

Let's assume that that the player selects his algorithm first, and that the adversary is aware of the selection made. If the player's algorithm is deterministic, then the adversary's best strategy is obvious: he chooses the order of his guesses to be the same order as the hidden cards of the player and thus he wins at every round. Hence, the value of the game is exactly $\sum_{i=1}^m a_i$.

In the randomized case the situation is more complicated. The player can choose his hidden cards according to distributions that might depend on the

²Note that cheating at this game requires sleight-of-hand skills not commonly found in 5-year olds; if neither player can surreptitiously modify the cards on the table, no cheating can occur.

history of the game. On the other hand, basic game theory implies that, given the probability distribution on the player strategies, there is a deterministic strategy for the adversary, that is, a fixed order of guesses, that maximizes his profit.

In order to analyze the value of the game we define two other models for players. A *strong* player is a player which, after each miss, is allowed to replace the hidden card by a card which is still in his hand. This, of course, can only help the player and does not increase his expected cost with respect to a standard player. A *weak* player is one that chooses the order of his hidden cards in advance (using random bits) and is not allowed to change this order later in the game. More precisely, the weak player chooses an order for his cards at the beginning of the game and then, whenever there is a hit, he replaces the hidden card by the lowest ordered card which has not been discarded yet. Clearly the expected cost for a weak player is no lower than the expected cost for a standard player.

Let $f(T)$ be defined by equation (1). Our main result in this section is

Theorem 3 *The value of the game with a set of cards T is at least $f(T)$, even against a strong player, and is at most $f(T)$ even against a weak player. Thus the game value is exactly $f(T)$ for all three types of players.*

Proof: We start with the lower bound and assume a strong player. Let $g(T)$ be the value of the game. We have to show that $g(T) \geq f(T)$ for any set T . We use induction on the size of T . If $T = \{i\}$, then $g(T) = a_i = f(T)$ and we are done. For the general case, let p_j be the probability that the player chooses card j as his first hidden card. Now, if the adversary chooses card i to be his first guess, and then chooses the best order for the remaining cards as if the game started with $T \setminus \{i\}$, he can clearly guarantee, even against a strong player, an expected cost of at least $p_i a_i + g(T \setminus \{i\})$. The adversary can choose the i which maximizes this expression. That implies that for all i

$$g(T) \geq p_i a_i + g(T \setminus \{i\}),$$

or

$$\frac{g(T) - g(T \setminus \{i\})}{a_i} \geq p_i.$$

But $\sum_i p_i = 1$, and therefore

$$\sum_i \frac{g(T) - g(T \setminus \{i\})}{a_i} \geq 1$$

or

$$g(T) \geq \frac{1 + \sum_{i \in T} g(T \setminus \{i\})/a_i}{\sum_{i \in T} 1/a_i}$$

Thus $g(T) \geq f(T)$.

We now turn to the upper bound and assume a weak player. Again the proof is by induction on the size of T . The case $T = \{i\}$ is trivial. For the general case, recall that a weak player hides his cards in a fixed order. Assume that the player constructs his order as follows: Among all cards he picks a card with probability inversely proportional to its value. Let the card so chosen be the *last* card in his order. From the remaining cards he picks again a card with probability inversely proportional to its value. Let it be the next-to-last card in his order. And so on. (That is, if after k choices the set of remaining cards is T and $i \in T$, the probability that i becomes the $m - k$ card in the order is $(1/a_i) / \sum_{j \in T} 1/a_j$.)³

Let $h(T)$ be the value of the game when the adversary knows that the player has chosen this particular strategy. Let j be the card chosen by the adversary to be last in his order. Let i be the last card of the player. Note that j is fixed, but i is a random variable.

- If $i = j$, an event whose probability is proportional to $1/a_j$, then the player has to pay a_j in the last round. Furthermore, the distribution used by the weak player with respect to the set of remaining cards (that is, $T \setminus \{j\}$) is exactly the same as if he started the game with the set $T \setminus \{j\}$. Hence in this case, the player's expected cost is at most $a_j + h(T \setminus \{j\})$ even if the adversary plays optimally on the remaining cards.
- If $i \neq j$, the player will never have to pay a_i and again his distribution on the remaining cards is exactly as if he started the game with $T \setminus \{i\}$, so his cost is at most $h(T \setminus \{i\})$.

This implies that

$$h(T) \leq \frac{1/a_j}{\sum_{i \in T} 1/a_i} (a_j + h(T \setminus \{a_j\})) + \frac{\sum_{i \in T \setminus \{j\}} h(T \setminus \{i\})/a_i}{\sum_{i \in T} 1/a_i}$$

or

$$h(T) \leq \frac{1 + \sum_{i \in T} h(T \setminus \{i\})/a_i}{\sum_{i \in T} 1/a_i}$$

That is, $h(T) \leq f(T)$.

We conclude that $h(T) = f(T) = g(T)$ and thus the value of the game is exactly $f(T)$ for all three types of players. \square

3.1 Properties of $f(T)$

In this section we discuss some of the interesting properties of $f(T)$.

Let's return to the weak player's strategy as described in Theorem 3. Let $P(i, R)$ for $i \in R \subset T$ be the probability that the player chooses card i the last

³Note that this strategy is not the same as choosing the sequence from first to last with probabilities proportional to a_i .

among the cards in R (which means that, in the player's hiding order, card i will be the first among the cards in R .) We claim that $P(i, R)$ does not depend on the values of the cards in $T \setminus R$. Indeed, call the cards in R , *red*. We can think that when the player builds his order, he first decides, with suitable probability, *whether* to pick a red card from the remaining cards, and if so, he then decides, with suitable probability, *which* red card to pick. Clearly the order among the red cards depends only on the values of the red cards.

Let $\pi_1, \pi_2, \dots, \pi_m$ be the adversary's order. As we have already observed, the player pays a_i if and only if he hides card i before he hides any card that comes after i in the adversary's order. That implies that the probability that the weak player pays a_{π_i} is exactly, $P(\pi_i, \{\pi_i, \pi_{i+1}, \dots, \pi_m\})$.

But the proof of Theorem 3 implies that the weak player's strategy as described is optimal, and therefore game theoretical considerations imply that the order chosen by the adversary is irrelevant – the expected value of the game is the same. It follows that

$$f([m]) = \sum_i^m a_{\pi_i} P(\pi_i, \{\pi_i, \pi_{i+1}, \dots, \pi_m\}) \quad (3)$$

for any permutation π ! In particular,

$$f([m]) = \sum_i^m a_i P(i, \{i, i+1, \dots, m\}). \quad (4)$$

In this form, it is rather hard to see that $f([m])$ is symmetric in the a_i 's, since the i 'th term in the sum depends only on a_i, a_{i+1}, \dots, a_m . We also don't know of any direct proof that shows that (4) is a solution of (1).

Unfortunately, the alternate expression is not computationally easier, since $P(i, R)$ does not seem to have a simple closed form. It can be computed with the formula

$$P(i, R) = \frac{\sum_{j \in R \setminus \{i\}} P(i, R \setminus \{j\})/a_j}{\sum_{j \in R} 1/a_j} \quad (5)$$

Similar considerations lead to

Theorem 4 *Let $T = [m]$. Without loss of generality assume that $a_1 \leq a_2 \leq \dots \leq a_m$. Then*

$$\sum_{i=1}^m \frac{a_i}{i} \leq \sum_{i=1}^m \frac{a_i^2}{a_1 + \dots + a_i} \leq f(T) \leq \sum_{i=1}^m \frac{a_i^2}{a_i + \dots + a_m} \leq \sum_{i=1}^m \frac{a_{m+1-i}}{i}.$$

Proof: Omitted. \square

4 The lower bound

Let r be an arbitrary positive number. Given a_1, \dots, a_m , we show that an adversary can construct a layered graph such that the cost of the off-line (modified) LGT algorithm is r while the cost of any on-line LGT is at least $rv([m])$, where $v([m])$ is the value of the Guess game on m cards with values a_1, \dots, a_m . (If the on-line LGT algorithm is deterministic (resp. randomized) then so is the player; and the value of $v([m])$ changes accordingly: $v([m]) = \sum_1^m a_i$ in the deterministic case and $v([m]) = f([m])$ in the randomized case.)

The graph consists of m paths and $m + 1$ layers. Each path starts with an edge with finite positive length, followed by a number of zero length edges, followed by a (practically) infinite length edge, except for the path to the target, which does not contain the infinite edge. Path P_i starts with an edge with length ra_i and has waste factor a_i . Each path has its infinite edge starting on a different layer, and for $j = 1, \dots, m - 1$, every layer L_j has an infinite edge out.

Thus, the on-line algorithm has no reason to visit the same path twice, and whenever it paid the first edge on the path, it can be presumed that it will not backtrack before reaching the infinite edge, or the target, since it costs nothing to advance and return on the zero length edges.

Path P_i corresponds to card i in the game. The on-line algorithm starting path i corresponds to the player hiding card i . An infinite edge on P_i between layer j and $j + 1$ corresponds to the adversary guessing card i at round j in the game. With these correspondences, it can be easily verified (see the example below) that a (randomized) strategy for the LGT algorithm immediately translates into a strategy for the (randomized) standard player in the Guess game.

Given the player's strategy, the adversary starts by choosing an order on the paths corresponding to his optimal order of guesses in the Guess game. Let this order be $\pi_1, \pi_2, \dots, \pi_m$. The adversary completes the construction as follows:

- The target is on path π_m at layer m .
- For $i = 1, \dots, m - 1$ the path π_i gets its infinite edge between layers i and $i + 1$.

For instance if the algorithm starts on path π_3 then it pays ra_{π_3} but will not pay the first edges on π_1 and π_2 . This corresponds exactly to the player in the Guess game that hides π_3 at the first turn, and hence pays a_{π_3} but does not pay a_{π_1} or a_{π_2} .

Clearly the cost of the LGT algorithm is at least r times greater than the cost of the player. ("At least" because the LGT algorithm also pays on the way back to source.) Hence we conclude that the cost of the on-line LGT algorithm is at least $rv([m])$ while the cost of the off-line algorithm is only r . (The path to the target has length ra_m and waste factor a_m .) This concludes the proof of the lower bound.

5 The upper bound

In the full paper, we present the complete proof of the theorem for arbitrary algorithms. The following proof assumes that the algorithms A_1, \dots, A_m are deterministic. The randomized proof is quite similar, but it requires a lot of summing and averaging, which obscures the essential ideas.

We are given a layered graph traversal problem: the graph consists of m paths P_1, \dots, P_m with waste factors a_1, \dots, a_m . We show how to construct an LGT algorithm using a strategy for the Guess game.

Let $l_{i,j}$ be the distance from the source to layer j on path P_i . Let $s_j = \min_i l_{i,j}/a_i$, that is, s_j is the minimum cost that the off-line algorithm must pay to get to layer j . Let j' be the minimum j such that $s_j > 0$. Let $s = s_{j'}$.

We now partition the layers into *strata*. All layers j such that

$$s2^{k-1} \leq s_j < s2^k,$$

belong to stratum k . All layers j with $j < j'$ belong to stratum 0.

Consider a layer j in stratum k . If $l_{i,j}/a_i > s2^k$ call path i *blocked at j* . Notice that on each layer there must be some path which is not blocked. (If all the paths have $l_{i,j}/a_i > s2^k$ then we just started a new stratum, and the “blocked” notion is redefined.) For stratum 0 we call a path blocked if $l_{i,j}/a_j > 0$.

Now we are ready to describe the on-line LGT algorithm. For stratum 0, the algorithm follows a 0-length path until it blocks, then switches arbitrarily to another 0-length path, and so on, until all paths have strictly positive lengths and stratum 1 starts. Notice that in general, once the algorithm has reached layer $j - 1$, it can compute s_j at no cost.

Once it gets to the first layer of stratum k , the on-line algorithm gets to the first layer of the next non-empty stratum k' this way: it follows a path until it blocks (with respect to stratum k), then it returns to the source and follows another path not yet blocked, and so on, until it arrives on the first layer of stratum k' . The crux of the algorithm is how to choose the next path to try.

The idea is that on the first layer of stratum $k > 0$, the algorithm starts playing a Guess game. As in the lower bound proof, path P_i corresponds to card i in the game, the on-line algorithm trying path i corresponds to the player hiding card i , and a blocked path P_i corresponds to the adversary guessing card i at a certain round in the game.

The difference is that now the adversary might guess (and miss) some cards even before any card is hidden – this corresponds to paths that are already blocked with respect to stratum k on the first layer of the stratum, and the adversary might guess several cards at once, against a single hidden card – this corresponds to paths that block on the same layer. Of course, both these maneuvers work to the advantage of the player.

Notice that card i on stratum k costs at most $s2^k a_i$. Taking into account backtracking, the total cost of the on-line algorithm on stratum k is bounded by $2s2^k v([m])$ where $v([m])$ is the value of the Guess game on m cards with values a_1, \dots, a_m .

Assume that the target belongs to stratum t . Then the total cost of the off-line algorithm is at least $s2^{t-1}$ while the total cost of the on-line algorithm is at most

$$2s(1 + 2 + \dots + 2^t)v([m]) < s2^{t+2}v([m]).$$

Hence the competitive ratio is at most $8v([m])$: that is, $8 \sum_{i=1}^m a_i$ for the deterministic case, and $8f([m])$ for the randomized case.

6 Application to the k server problem

The k -server algorithm of [5] is based on recursive calls to the MIN_S operation with i^2 algorithms whose competitive ratio can be divided into i groups, each of size i . The competitive ratios of the algorithms in the same group are about the same but the ratios differ greatly among groups. More precisely, the sum of the competitive ratios of all the algorithms is dominated by the sum in one group. Thus the average competitive ratio is $\Theta(i)$ times smaller than the maximum one. The MIN_S originally used in [5] has competitive ratio $O(m \cdot \max_i \{a_i\})$, while our algorithm has competitive ratio $O(\sum_{i=1}^m a_i)$. Hence, our algorithm saves a $\Theta(i)$ factor in each recursive call and this results in a $k!/2^{O(k)}$ overall savings factor. Unfortunately the competitive ratio of the modified algorithm is still exponential in k , namely $O((k!)^2 2^{O(k)})$.

References

- [1] R. Baeza-Yates, J. Culberson and G. Rawlins, “Searching in the plane,” to appear in *Information and Computation*.
- [2] A. Borodin, N. Linial, and M. Saks. “An Optimal Online Algorithm for Metrical Task Systems” *Proceedings of the 19th Annual ACM Symposium on Theory of Computing*, 1987, pp. 373–382.
- [3] A. Fiat, R. Karp, M. Luby, L. McGeoch, D. Sleator, and N. Young. “Competitive paging algorithms,” *Journal of Algorithms*, 12(1991), pp. 685-699.
- [4] A. Fiat, D. Foster, H. Karloff, Y. Rabani, Y. Ravid and S. Vishwanathan, “Competitive algorithms for layered graph traversal,” *Proceedings of the 32nd IEEE Symposium on Foundations of Computer Science*, 1991, pp. 288-297.

- [5] A. Fiat, Y. Rabani and Y. Ravid, “Competitive k -server algorithms,” *Proceedings of the 31st IEEE Symposium on Foundations of Computer Science*, 1990, pp. 454-463.
- [6] C. Papadimitriou and M. Yannakakis, “Shortest paths without a map,” *Proceedings of the 16th ICALP*, 1989, pp. 610-620.
- [7] D. Sleator and R. Tarjan, “Amortized efficiency of list update and paging rules,” *Communications of the ACM*, 23(1985), pp. 202–208.