

# Delayed Information and Action in On-Line Algorithms<sup>1</sup>

Susanne Albers<sup>2</sup>

*Institut für Informatik, Albert-Ludwigs-Universität Freiburg, Georges-Köhler-Allee 79 79110 Freiburg, Germany*  
E-mail: salbers@informatik.uni-freiburg.de

Moses Charikar<sup>3</sup>

*Computer Science Department, Stanford University, Stanford, California 94305*  
E-mail: moses@cs.stanford.edu

and

Michael Mitzenmacher<sup>4</sup>

*Harvard University, 33 Oxford St., Cambridge, Massachusetts 02138*  
E-mail: michaelm@eecs.harvard.edu

Received March 31, 1999

Most on-line analysis assumes that, at each time step, all relevant information up to that time step is available and a decision has an immediate effect. In many on-line problems, however, the time when relevant information is available and the time a decision has an effect may be decoupled. For example, when making an investment, one might not have completely up-to-date information on market prices. Similarly, a buy or sell order might only be executed some time in the future. We introduce and explore natural delayed models for several well-known on-line problems. Our analyses demonstrate the importance of considering timeliness in determining the competitive ratio of an on-line algorithm. For many problems, we demonstrate that there exist algorithms with small competitive ratios even when large delays affect the timeliness of information and the effect of decisions. © 2001 Academic Press

## 1. INTRODUCTION

The theory of on-line algorithms deals with situations where a decision or a series of decisions must be made with limited information, and specifically without knowledge of future events. Implicit in this approach is the idea that the time information becomes available relative to the time decisions take effect can be of paramount importance in algorithm performance. In most on-line analyses, however, the setting chosen for study is the simple one: at each time step, all relevant information up to that time step is available, and a corresponding decision is made.

In many on-line problems the time relevant information is available and the time a decision has an effect are decoupled. This phenomenon arises, for instance, in *investment problems* where one has to decide *whether* and *when* to buy an expensive piece of equipment. An example of such investment problems is the standard on-line ski rental problem. In these investment problems, once a decision is made for buying equipment, it can take some time before the equipment is delivered to the user. For example, it can take a couple of days or weeks to ship a particular model of skis and even months to deliver and install a new machine in a factory. In such cases, a decision to buy equipment has an effect only later in time and the action corresponding to the decision is delayed.

<sup>1</sup> A preliminary version of this paper was presented at the *39th Annual Symposium on Foundations of Computer Science (FOCS)*, 1998.

<sup>2</sup> Part of this work was done while at the Max-Planck-Institut für Informatik, Saarbrücken, Germany.

<sup>3</sup> Supported by a Stanford Graduate Fellowship, an ARO MURI Grant DAAH04-96-1-0007 and NSF Award CCR-9357849, with matching funds from IBM, Schlumberger Foundation, Shell Foundation, and Xerox Corporation.

<sup>4</sup> Supported by NSF CAREER Grant CCR-9983832 and an Alfred P. Sloan Research Fellowship. Most of this work was done while employed at Compaq Systems Research Center.

This can heavily influence the performance of an on-line strategy. An on-line algorithm  $A$  is called  $c$ -competitive if, for all inputs, the cost incurred by  $A$  is at most  $c$  times the cost incurred by an optimal off-line algorithm that knows the entire input in advance. To illustrate the effect of delayed action, we consider the ski rental problem. Skis cost  $r$  dollars to rent per weekend and  $b$  to buy for a season. Suppose an avid skier skis every weekend there is good snow. Whether it is best for him or her to rent or buy skis for the season depends on the number of good ski weekends. If the skier rents  $s$  times before buying, the competitive ratio  $c$  is  $\frac{sr+b}{\min\{(s+1)r, b\}}$ . When  $b/r$  is an integer, an optimal on-line algorithm is to rent skis  $s = (b/r) - 1$  times and then buy; this yields a competitive ratio of  $2 - \frac{r}{b}$ . If skis take  $d \geq 1$  weeks to ship, the analysis is slightly more involved. If a skier decides to rent  $s$  times before buying, we must consider what happens in the intervening  $d$  other weekends before the skis arrive. If  $i$  of the intervening weekends are snowy, then the worst-case ratio between the actual cost and the optimal cost is now

$$\max_{1 \leq i \leq d} \frac{(s+i)r+b}{\min\{(s+i)r, b\}}.$$

It is easily checked that this ratio is maximized at one of the extremes  $i = 1, d$ ; using this, one can easily determine the value of  $s$  that minimizes the competitive ratio.

In the above example, there is a delay between the time a decision is made and when it has an effect. We refer to this as the *delayed action* model. The parameter  $d$  is the maximum delay after which a decision takes effect. For this problem,  $d = 0$  gives us the original problem without delay.

Similarly, there are problems where it is natural to consider information that arrives only after some delay. In this scenario, at time step  $t$  we might have information about the first  $t - d - 1$  time steps only, for some  $d \geq 0$ . This phenomenon arises, for instance, in on-line *financial games* where we have to devise strategies for converting money from one currency to another or for selecting a portfolio in the stock market [17, 18, 20, 33]. Naturally, we might not have access to the very latest exchange rates or stock prices. We refer to this as the *delayed information* model. Again, the case  $d = 0$  corresponds to the original problem without delay.

Related timing problems occur when a *group* of people or agents take decisions. The group might come together only at particular time instances. The actions are again delayed in that they can only occur at specific points in time. For example, in the case of investing in manufacturing machinery, one may only be able to make budget decisions in concert with the rest of an organization at specific budgeting periods. Another example is that of an investment club, where a group of people pool their money together and invest in the stock market. All investment decisions can be made only at club meetings which occur at regular intervals of time, e.g., once a month.

We use the term *delayed models* to loosely describe models where there is this type of discontinuity between the time information is available and the time decisions take effect. Such models are naturally motivated by situations where one has incomplete information about the past or a decision will have a delayed effect on the state of the system. Interestingly, they also often have a natural interpretation in terms of distributed agents acting with limited coordination. In particular, such models correspond nicely to distributed systems where information about the system is updated only after some delay or at specific synchronization points.

*Our contribution.* In this paper, we consider several standard on-line problems and examine their generalizations to delayed models. These generalizations are generally quite natural and lead to interesting insight into the original problem. We note that in this initial exploration of delayed models, we have focused on cases where one can modify the original on-line analysis to analyze the delayed version. We believe that the resulting relative simplicity of many of our results demonstrates the naturalness and utility of this model. We expect, however, that delayed models will prove more difficult than their standard counterparts in many instances.

We briefly describe the remainder of the paper. In Section 1, we study the delayed information model applied to the classical problem of on-line scheduling on parallel machines to minimize the makespan. Here a scheduling algorithm must assign new jobs to processors based on stale load information. Traditional algorithms for on-line scheduling do badly in this scenario. We develop new algorithms for this model and prove almost matching lower bounds. In Section 1, we study the list update problem in the delayed action model and prove nearly tight upper and lower bounds for deterministic on-line

algorithms. We also show that a randomized on-line algorithm can only beat the deterministic lower bound if it uses paid exchanges. In Section 1, we generalize an on-line stock market model introduced in [17] by studying natural delayed models. Finally, in Section 1, we apply the delayed action model to the general class of relaxed metrical task systems [6, 10]. Relaxed task systems are an abstract model for problems where one has to decide when it is appropriate to make expensive configuration changes. This class includes the ski rental problem, page migration [15], file replication [15], network leasing [6], and other problems (see [10]). We extend the results of [6, 10] to apply to relaxed task systems with delayed action, effectively handling the delayed models of an entire general class of problems.

*Related work.* In subsequent sections, we will mention related work relevant to the specific problems we consider. Here, we offer a brief overview of generally relevant related work.

The importance of when information becomes available has been noted previously, especially in the significant body of work on algorithms with lookahead, e.g. [14, 22, 24, 27]. In the case of on-line decision models, however, the possibility of not having up-to-date information is not generally addressed. For load balancing problems, the question has been considered for statistical models [29, 30, 36]; other queueing based models have also been considered [4, 5, 28]. And recently, [7] considered an on-line load balancing setting where tasks gather some information about system behavior before making a choice of processor.

There is also a large body of work on algorithms with distributed agents, who must coordinate their efforts in the face of some cost for communication, e.g. [3, 8, 12, 13]. These models, however, model communication as an instantaneous event, and hence the communication cost does not directly incorporate a notion of time and delay. Another line of research has addressed distributed decision making when the communication among agents is limited, for example by only allowing local communication. Implicitly this allows distant agents to communicate only after a number of communication rounds. The problems investigated include scheduling, load balancing, routing, and general optimization [11, 19, 25, 31, 32].

## 2. SCHEDULING

We consider a classical problem in on-line scheduling. A sequence of jobs  $J_1, J_2, \dots$  must be scheduled on  $m$  identical parallel machines. Whenever a job arrives, the job must be scheduled immediately on one of the machines, without knowledge of any future jobs. Preemption of jobs is not allowed. The goal is to minimize the *makespan*, i.e., the completion time of the last job that finishes.

The problem was first investigated by Graham [21]. He developed the well-known *List* algorithm that always schedules a job on the least loaded machine. Graham's *List* algorithm is  $(2 - \frac{1}{m})$ -competitive. More recently, on-line algorithms that obtain competitive ratios bounded away from 2 have been devised. The currently best known competitive ratio for this problem is 1.923 obtained by Albers [1].

In a setting with delayed information, we do not have the current loads on the processors available to us. When we are presented with the  $i$ th job  $J_i$ , we have the loads on the machines from up to  $d_i + 1$  requests ago. That is, we know the load after the job  $J_{i-d_i-1}$  was placed. (When  $d_i = 0$  always, we have the original problem.) We must decide where to place job  $J_i$  based on this old information. We examine the setting where we have a bound on how old the information is at each stage, i.e.,  $d_i \leq d$ , for some  $d$ . We will refer to the last  $d_i$  jobs whose contribution to the loads is not known as *unknown* jobs and other jobs as *known* jobs.

In this situation, the strategy of placing each job on the processor with the least known load does very badly. In fact the competitive ratio of that strategy can be as bad as  $d + 2 - \frac{d+1}{m}$  (for  $d \leq m - 1$ ). The problem is that this strategy does not take into account the potential effect of unknown jobs. We will devise new algorithms with better competitive ratios for scheduling with delayed information.

We study two variants of the basic scheduling problem. In our first model, we assume that in addition to the loads of the machines from  $d_i + 1$  requests ago, we also know where the last  $d_i$  unknown jobs were placed. This scenario describes, e.g., a centralized scheduling algorithm where the size of every new job is not known to the scheduler immediately on arrival, but is revealed at most  $d$  requests later. In practice, the processing times of jobs often is not known in advance. It is possible to compute accurate estimates on the processing times, but the computation of such estimates (by the scheduler or the user) takes a certain amount of time.

It is simpler to work with a less stringent, but for our purposes equivalent, scenario where we have available a complete history of the process up to  $d_i + 1$  requests ago. In this model, by using specific kinds of deterministic algorithms that do not use the length of the current job in a new scheduling step, we can figure out where the *unknown* jobs were scheduled as follows. Suppose we use a deterministic algorithm that bases its decision on the schedule from  $d + 1$  requests ago, i.e. if  $d_i < d$  pretend that the state seen by the algorithm is the schedule exactly  $d + 1$  requests ago. Because we have complete information about the job history, we can also figure out the complete schedule from  $d + 2$  requests ago,  $d + 3$  requests ago, and so on. Hence we can deduce the state seen by the algorithm while scheduling each of the previous  $d$  jobs and thereby determine where each of the last  $d$  unknown jobs were scheduled.

For this model, we consider an algorithm we call *Delayed List* scheduling, as it generalizes Graham's List algorithm. Let  $w_i$  be the known load on machine  $i$ . (This is the load without the unknown jobs.) Let  $S$  denote the total known load on all the machines; i.e.,  $S = \sum_{i=1}^m w_i$ . Let  $u_i$  be the number of unknown jobs on machine  $i$ . Define the *pseudo-load* on machine  $i$  to be  $u_i + (m - u_i - 1)(w_i/S)$ . The algorithm schedules the new job on the machine which has the lowest pseudo-load. (When  $d = 0$ , the algorithm is exactly the same as List.)

LEMMA 2.1. *When the Delayed List algorithm places the current job on machine  $i$ , the load on machine  $i$  is at most  $1 + u_i + (m - u_i - 1)(w_i/S)$  times the optimal load.*

*Proof.* Let  $x$  be the processing time for the  $j$ th job. Consider what happens if the algorithm tries to place the current job on machine  $i$ . Let  $y$  be the average processing time of the unknown jobs on machine  $i$ . Then  $\ell_i = w_i + u_i \cdot y + x$  will be the new load on machine  $i$ .

The sum of the processing times of all the jobs in the sequence is at least  $S + u_i \cdot y + x$ . Thus  $OPT \geq (S + u_i \cdot y + x)/m$ . Also,  $OPT \geq x$  and  $OPT \geq y$ . Hence

$$\frac{\ell_i}{OPT} \leq \min \left( \frac{w_i + u_i y + x}{x}, \frac{w_i + u_i y + x}{y}, \frac{w_i + u_i y + x}{(S + u_i y + x)/m} \right).$$

We obtain the required bound on  $\frac{\ell_i}{OPT}$  by maximizing the above function over all possible values of  $y$  and  $x$ . Let us maximize over  $y$  first. We wish to compute

$$\max_y \min \left( \frac{w_i + u_i y + x}{x}, \frac{w_i + u_i y + x}{y}, \frac{w_i + u_i y + x}{(S + u_i y + x)/m} \right).$$

Let

$$f_1(x, y) = \frac{w_i + u_i \cdot y + x}{x};$$

$$f_2(x, y) = \frac{w_i + u_i \cdot y + x}{y};$$

$$f_3(x, y) = \frac{w_i + u_i \cdot y + x}{(S + u_i \cdot y + x)/m}.$$

Note that each of the three functions is monotone in  $y$ . We want to find the maximum of the lower envelope (i.e., minimum) of these three monotone curves. This must occur either at an end-point of the interval  $y = 0$  or  $y = \infty$  or at a point where two of the three functions are equal. Further, a point where two functions are equal is a potential maximum only if the value of the third function is greater than the two that are equal.

In fact, our analysis will show that the maximum is achieved when all three functions are equal.

1. Let us first consider the maximum value of the function for end-points of the interval. For  $y = \infty$ , the value of the function is 1. For  $y = 0$ , the value of the function is  $\min((w_i + x)/x, (w_i + x)/((S + x)/m))$ . This is maximized when  $x = (S + x)/m$ . Hence the maximum value is  $1 + (m - 1)(w_i/S)$ .

We now consider the three possible points where two of the functions are equal.

2. Suppose  $f_1(x, y) = f_2(x, y) \leq f_3(x, y)$ . This implies that  $x = y \geq (S + u_i \cdot y + x)/m$ . Hence  $f_1(x, y) = f_2(x, y) = u_i + 1 + w_i/x$ . Our bound is maximized for the smallest possible value of  $x$ . But we also have  $x \geq S/(m - u_i - 1)$ . Hence, the maximum value is  $u_i + 1 + (m - u_i - 1)(w_i/S)$ .

3. Suppose  $f_1(x, y) = f_3(x, y) \leq f_2(x, y)$ . This implies that  $x = (S + u_i \cdot y + x)/m \geq y$ . Hence  $f_1(x, y) = f_3(x, y) = m - ((S - w_i)/x)$ . Our bound is maximized for the largest possible value of  $x$ . But we also have  $x \leq S/(m - u_i - 1)$ . Hence, the maximum value is  $u_i + 1 + (m - u_i - 1)(w_i/S)$ .

4. Suppose  $f_2(x, y) = f_3(x, y) \leq f_1(x, y)$ . This implies that  $y = (S + u_i \cdot y + x)/m \geq x$ . Algebraic manipulation yields  $f_2(x, y) = f_3(x, y) = u_i + (m - u_i)(w_i + x)/(S + x)$ , which is increasing in  $x$  since  $w_i \leq S$ . Our bound is maximized for the largest possible value of  $x$ . But we also have  $x \leq S/(m - u_i - 1)$ . Hence, the maximum value is  $u_i + 1 + (m - u_i - 1)(w_i/S)$ .

In all cases  $\ell_i/OPT \leq 1 + u_i + (m - u_i - 1)(w_i/S)$ . ■

We use the result of Lemma 3 to bound the competitive ratio of the algorithm.

**THEOREM 2.1.** *The Delayed List algorithm is  $(2 + \frac{d-1}{m})$ -competitive.*

*Proof.* The algorithm schedules the current job on the machine  $i$  which has the lowest value of  $c_i = 1 + u_i + (m - u_i - 1)(w_i/S) \geq \frac{\ell_i}{OPT}$ . Now,

$$\sum_{i=1}^m c_i \leq \sum_{i=1}^m \left[ 1 + u_i + (m - 1) \frac{w_i}{S} \right] = m + d + m - 1$$

because  $\sum_{i=1}^m w_i = S$ . Hence there must be some  $c_i$  with value at most  $\frac{2m+d-1}{m} = 2 + \frac{d-1}{m}$ . Thus, the competitive ratio of the algorithm is at most  $2 + \frac{d-1}{m}$ . ■

Theorem 2.1 shows that by spreading out the unknown jobs appropriately, we can achieve a competitive ratio that grows at a “rate” of  $d/m$ . In fact, the analysis in the proof of Lemma 2.1 shows that given  $S, x, u_i$ , and  $w_i$ , one can compute precisely the worst case competitive ratio if the algorithm places the current job on machine  $i$ . This is a function of  $S, x, u_i$ , and  $w_i$ , and an exact expression can be obtained. A more intelligent algorithm would compute this function for each machine and place the current job on the machine that minimizes this function. Indeed, this improves the competitive ratio slightly, although it seems difficult to develop a general bound with a better form than Theorem 2.1. Moreover, the result of Theorem 2.1 is nearly tight, as the following lower bound shows.

**THEOREM 2.2.** *The competitive ratio of any deterministic algorithm for the delayed scheduling problem cannot be smaller than  $2 + \frac{d-2}{m+1}$  when this number is an integer less than or equal to  $m$ .*

*Proof.* Let  $A$  be a deterministic algorithm for the delayed scheduling problem with maximum delay  $d$ . For the lower bound, assume that when  $A$  receives job  $J_i$ , it knows the entire schedule after job  $J_{i-d-1}$  was placed. Suppose  $d = (r - 2)m + r$  for an integer  $r \leq m$ . We will construct a request sequence consisting of  $(r - 1)m + 1$  jobs such that the optimal load is essentially 1, but some machine in  $A$ 's schedule has load  $r$ .

The first  $m - r$  requests are jobs of size 1. The next  $(r - 2)m + r + 1$  jobs have size either 1 or  $\epsilon$ , where  $\epsilon > 0$  can be arbitrarily small. An adversary selects at most  $r$  of these to have size 1 as follows. Since a total of  $(r - 1)m + 1$  jobs are scheduled, there must exist one machine  $x$  to which  $A$  assigns at least  $r$  jobs. Note that while the second group of  $(r - 2)m + r + 1$  jobs is presented,  $A$  does not know the size of any of these jobs. Among the jobs assigned by  $A$  to machine  $x$ , the adversary chooses  $r$  jobs to be of size 1. Thus the on-line makespan is  $r$ . On the other hand, the optimal makespan for this sequence is  $1 + ((r - 2) + 1)\epsilon$ . Choosing  $\epsilon$  arbitrarily small, the competitive ratio is at least  $r = 2 + \frac{d-2}{m+1}$ . ■

We now consider a second variant of the problem and a corresponding algorithm. In this scenario, when we are presented with a job  $J_i$ , we know the loads on the machines from  $d_i + 1$  requests ago, but we do not know the actual schedule or job sizes corresponding to these loads. We assume, however, that each job knows its sequence number  $i$  and the number of jobs already scheduled, or  $i - d_i - 1$ . (Implicitly, the number of scheduled jobs is increasing, so  $i - d_i - 1 \leq k - d_k - 1$  when  $i < k$ .) Our

algorithm will make use of this information in its scheduling decision. This model corresponds to a distributed system where tasks may place themselves on an appropriate server before other tasks reveal their processing times, but through simple shared counters limited information such as the values of  $i$  and  $i - d_i - 1$  is maintained.

We provide an algorithm for this scenario called the *Delayed Avoid Heavy* algorithm. We describe what happens when the  $i$ th job  $J_i$  arrives. We say that the machine with the  $k$ th smallest load from known jobs at this time has rank  $k$ . The algorithm uses a constant  $c$  as a parameter; this will be specified later. We never schedule a job on the heaviest  $m/c$  machines. (For convenience, we will assume that  $m/c$  is integral throughout.) Let  $b = m(1 - 1/c)$ , i.e., the number of machines excluding the heaviest  $m/c$ . Let  $f(J_i) = (2i - d_i)$ . The Delayed Avoid Heavy algorithm schedules job  $J_i$  on the machine with rank  $b - (f(J_i) \bmod b)$ .

For the purpose of analysis, we will divide the jobs into groups. Job  $J_i$  is placed in group number  $\lfloor f(J_i)/b \rfloor$ .

LEMMA 2.2. *Two jobs  $J_i$  and  $J_k$  in the same group are assigned to different machines.*

*Proof.* Without loss of generality, assume  $i < k$ . When scheduling  $J_i$ , the algorithm sees the schedule  $S_i$  that results after  $i - d_i$  jobs have been assigned to machines and when scheduling  $J_k$ , the algorithm sees the schedule  $S_k$  that results after  $k - d_k$  jobs have been assigned. As the earlier job  $J_i$  cannot see a more recent schedule than the later, it is the case that  $i - d_i \leq k - d_k$ .

Since  $J_i$  and  $J_k$  are in the same group (say  $g$ ),  $g = \lfloor f(J_i)/b \rfloor = \lfloor f(J_k)/b \rfloor$ . Then  $J_i$  is assigned to the machine  $M_i$  of rank  $b - (f(J_i) \bmod b) = b - (f(J_i) - g \cdot b) = (g + 1)b - (2i - d_i)$  (in schedule  $S_i$ ). Similarly,  $J_k$  is assigned to the machine  $M_k$  of rank  $(g + 1)b - (2k - d_k)$  in schedule  $S_k$ .

Now, schedule  $S_k$  results from schedule  $S_i$  by the scheduling of an additional  $(k - d_k) - (i - d_i)$  jobs. Observe that a machine that has rank  $r$  in a certain schedule  $S$  has rank at least  $r - i$  in the schedule obtained by placing  $i$  additional jobs in  $S$ . Thus, in schedule  $S_k$ , the machine  $M_i$  must have rank at least

$$(g + 1)b - (2i - d_i) - ((k - d_k) - (i - d_i)) \geq (g + 1)b - (k + i - d_k) > (g + 1)b - (2k - d_k).$$

This implies that the machines  $M_i$  and  $M_k$  are distinct. ■

LEMMA 2.3. *The competitive ratio of the Delayed Avoid Heavy algorithm is at most  $2 + \frac{2d}{b} + c$ .*

*Proof.* When job  $J_i$  arrives, we know the loads on all machines except for the contributions to the loads by the last  $d_i$  jobs. Let  $S$  be the set of the last  $d_i$  jobs together with job  $J_i$ . Observe that the  $f$  values of any two jobs in  $S$  can differ by at most  $d + d_i \leq 2d$ . Thus the number of distinct groups that the jobs in  $S$  belong to is at most  $2 + \lfloor \frac{2d}{b} \rfloor \leq 2 + \frac{2d}{b}$ . Since no two jobs in the same group get placed on the same machine, the maximum number of jobs in  $S$  that get placed on the same machine is at most  $2 + \frac{2d}{b}$ , and in particular there are at most  $1 + \frac{2d}{b}$  unknown jobs on the processor that gets  $J_i$ . Let  $w_i$  be the known load on the machine on which job  $J_i$  is placed. Let  $S$  be the total known load on all the machines. Then  $w_i/S \leq c/m$ . If not, then the loads on the heaviest  $m/c$  machines must each be greater than  $Sc/m$ , implying that the total load is greater than  $S$ . This is clearly not possible. Now, Lemma 3 implies that, after  $J_i$  is placed on  $M_i$ , the total load on  $M_i$  is at most  $2 + 2d/b + c$  times the optimal load. Hence the competitive ratio is at most  $2 + \frac{2d}{b} + c$ . ■

Substituting  $b = m(1 - 1/c)$  and optimizing for  $c$ , we get that, for  $c = 1 + \sqrt{2d/m}$ , the competitive ratio of the Delayed Heavy Load algorithm is bounded by  $2 + 2d/m + 2\sqrt{2d/m}$ . It is possible to get slightly better bounds by being a bit more careful in Lemma 3. However, the expressions that result are far from elegant and the improvements are very minor, so we choose to omit them. The main point is that in this more limited model, by again spreading out the unknown jobs appropriately, we can achieve a competitive ratio that grows at a “rate” of about  $2d/m$ .

### 3. LIST UPDATE

The list update problem is a fundamental problem in the theory of on-line algorithms. It consists of maintaining an unsorted list so as to minimize the total cost of accesses on a sequence of requests. Formally, we are given  $n$  items that are stored in an unsorted linear linked list. A list update algorithm receives a sequence of *requests*, where each request specifies one item in the list. To *serve* a request the algorithm must *access* the requested item; i.e., it starts at the front of the list and proceeds linearly through the items until the desired item is found. Serving an access to the item at position  $i$  in the list incurs a cost of  $i$ .

In the standard problem, the list may be updated at any time. More specifically, after each request the accessed item may be moved at no extra cost to any position closer to the front of the list. These exchanges are called *free exchanges*. At any time, two adjacent items in the list may be exchanged with cost 1; these exchanges are called *paid exchanges*. The goal is to serve a sequence of requests so that the total cost is as small as possible.

We investigate a model with delayed action where the free and paid exchanges made by an on-line algorithm in response to a request only take effect some time later. More specifically, we consider a setting where the updates are implemented at the end of a *round*, where every round consists of  $1 + d$  consecutive requests in the request sequence. This setting can also be viewed as a scenario where the on-line algorithm can update the list only at the end of a round. Items requested during the round may be moved closer to the front of the list using free exchanges before the next round. Items not requested in the round can be moved only using paid exchanges. In the following we work with this latter scenario. Note that when  $d = 0$ , we have the original standard problem.

To motivate the delayed model, consider the case where the linked list data structure is a shared object among a number of agents. In this case agents may read the list simultaneously without any problems; however, while the data structure is being updated, it may be necessary for consistency to lock the structure. In this case infrequent updates may provide better overall performance. We may think of the update operations as being batched, in which case the update actions are delayed.

In the following we first concentrate on deterministic on-line algorithms. When analyzing on-line algorithms, we consider two types of *adversaries* that generate a request sequence and serve the generated sequence *off-line*.

- The *standard adversary* may update the list after each request.
- The *limited adversary* can update the list only at the end of each round.

We call a deterministic list update algorithm  $A$   $c$ -competitive against any standard (limited) adversary  $ADV$  if, for *all list lengths*  $n$  and for every request sequence generated by  $ADV$ , the cost incurred by  $A$  is not greater than  $c$  times the cost paid by  $ADV$ .

For the standard list update problem, Sleator and Tarjan [35] showed that the well-known on-line algorithm Move-To-Front (MTF) is 2-competitive. This algorithm moves an item to the front of the list each time it is accessed. This is the best competitive ratio any deterministic on-line algorithm can obtain in the standard model [26].

We now study the problem with delayed action.

**THEOREM 3.1.** *Let  $A$  be a deterministic on-line algorithm for the list update problem with delayed action. If  $A$  is  $c$ -competitive, then  $c \geq 1 + d$ . This lower bound holds for both types of adversaries.*

*Proof.* In each round the adversary issues  $1 + d$  requests to the item that is stored at the last position in  $A$ 's list. Thus, in each round  $A$  incurs a cost of  $(1 + d)n$ .

At the end of each round, the adversary moves the item requested in the next round to the front of the list using paid exchanges. Thus, its cost in each round is at most  $n + d$ . The ratio of the cost incurred by  $A$  to the cost incurred by the adversary is

$$\frac{(1 + d)n}{n + d} = \frac{1 + d}{1 + d/n}$$

and, for large values of  $n$ , this expression can be arbitrarily close to  $1 + d$ . ■

Next we give an adaptation of MTF to the model of delayed action.

**ALGORITHM MTF( $d$ ).** At the end of each round, the algorithm moves the requested items to the front of the list. At the head of the list, for any two items  $i$  and  $j$  requested in the round,  $i$  precedes  $j$  if and only if the last request to  $i$  is more recent than the last request to  $j$ .

The Algorithm MTF( $d$ ) can also be thought of as the algorithm that batches all Move-To-Front operations until an update is allowed.

**THEOREM 3.2.** *The algorithm MTF( $d$ ) is  $(2 + d)$ -competitive. This upper bound holds for both types of adversaries.*

Note that for  $d = 0$  we obtain the upper bound of 2 achieved by the MTF algorithm in the standard list update problem.

*Proof.* We prove the theorem for the standard adversary. For the analysis of MTF( $d$ ) it is convenient to work with a different model for updating the list. In this modified model, an on-line algorithm may move an item accessed in a round only on the last request to the item in that round. Let MTF'( $d$ ) be the algorithm that moves an item to the front of the list whenever it is requested for the last time in a round. Given any request sequence  $\sigma$ , at the end of each round the lists maintained by MTF'( $d$ ) and MTF( $d$ ) are the same. Thus, in each round the cost incurred by MTF( $d$ ) is not higher than the cost incurred by MTF'( $d$ ). We show that the cost incurred by MTF'( $d$ ) is at most  $2 + d$  times the cost incurred by the adversary, for any  $\sigma$ .

We assume that MTF'( $d$ ) and the adversary start with the same list. Given an arbitrary request sequence  $\sigma = \sigma(1), \sigma(2), \dots, \sigma(m)$ , let  $t$  denote the point in time *after* the  $t$ th request  $\sigma(t)$  is served. We define a potential function  $\Phi$ . For any time  $t$  and any item  $x$  in the list, let  $r(t, x)$  be the next round in the request sequence in which  $x$  is requested. If  $x$  is still requested in the current round, then  $r(t, x)$  is equal to the current round. Let  $n(t, x)$  be the number of remaining requests to  $x$  in  $r(t, x)$ . We have  $n(t, x) \leq 1 + d$ . An *inversion* is an ordered pair  $(y, x)$  of items such that  $x$  occurs before  $y$  in the adversary's list and after  $y$  in the list maintained by MTF'( $d$ ). At any time the potential  $\Phi$  is the number of inversions  $(y, x)$ , where each inversion is multiplied by  $n(t, x)$ , which can be seen as the weight of inversion  $(y, x)$ .

Consider any request  $\sigma(t)$  and let  $x$  be the item requested. Let  $C_{MTF}(t)$  and  $C_{ADV}(t)$  be the actual costs paid by MTF'( $d$ ) and the adversary during the service of  $\sigma(t)$ . Clearly,  $C_{MTF}(t) \leq C_{ADV}(t) + \text{inv}(t-1, x)$ , where  $\text{inv}(t-1, x)$  is the number of inversions  $(y, x)$  immediately before the request. We show that during the service of  $\sigma(t)$  the potential decreases by  $\text{inv}(t-1, x)$  due to inversions removed or due to inversions whose weights change. If  $x$  is not requested for the last time in the round, then the number of remaining requests to  $x$  in the round decreases by 1; i.e.,  $n(t-1, x) - n(t, x) = 1$  and the weight of each inversion  $(y, x)$  decreases by 1. If  $x$  is requested for the last time in the round,  $n(t, x)$  can increase, i.e.,  $n(t, x) \geq n(t-1, x)$ . However,  $x$  is moved to the front of the list, which implies that all inversions  $(y, x)$  are removed and  $n(t, x)$  does not contribute to the potential. In any case, the potential decreases by  $\text{inv}(t-1, x)$  during the service of  $\sigma(t)$ . If  $x$  is moved to the front of the list, then at most  $C_{ADV}(t)$  new inversions  $(x, z)$  can be created, each of which increases the potential by  $n(t, z) \leq 1 + d$ . Since  $n(t-1, y) = n(t, y)$  for all  $y \neq x$ , we conclude that at any time  $t$ ,

$$\begin{aligned} C_{MTF}(t) + \Delta\Phi &\leq C_{ADV}(t) + (1 + d) \cdot C_{ADV}(t) \\ &\leq (2 + d)C_{ADV}(t). \end{aligned}$$

Finally we have to consider a paid exchange made by the adversary. Each paid exchange can create an inversion, which increases the potential by at most  $1 + d$ , but the adversary has to pay a cost of 1. So again  $C_{MTF}(t) + \Delta\Phi \leq (1 + d)C_{ADV}(t)$ .

Summing over all the steps of  $\sigma$  and noting  $\Phi \geq 0$  yields  $C_{MTF}(\sigma) \leq (2 + d)C_{ADV}(\sigma)$ . ■

It is straightforward to modify the above proof to show the following:

**COROLLARY 3.1.** *If each item is requested at most  $k$  times in a round, then MTF( $d$ ) is  $(1 + k)$ -competitive.*



This corollary shows that if one is attempting to choose a value of  $d$  to balance reading and writing costs, a key parameter to consider is how often items can be requested repeatedly.

Next we consider randomized on-line algorithms and give two lower bounds. None of the randomized on-line algorithms that have been presented so far for the standard list update problem uses paid exchange, see e.g. [2, 34]. We show that such algorithms cannot be better than  $(1 + d)$ -competitive in the setting with delayed action.

**THEOREM 3.3.** *Let  $A$  be a randomized on-line algorithm for the list update problem with delayed action and suppose that  $A$  does not use paid exchanges. If  $A$  is  $c$ -competitive against any oblivious adversary, then  $c \geq 1 + d$ . This lower bound holds for both types of adversaries.*

*Proof.* An adversary constructs a request sequence in *phases*. In each phase the adversary inspects its current list and requests the  $n$  items in ascending order. To each of the  $n$  items, the adversary issues  $1 + d$  consecutive requests, which form a round. At the end of each round, the adversary moves the item requested in the next round to the front of the list using paid exchanges. The adversary needs  $i - 1$  paid exchanges for the item requested in the  $i$ th round because the list items are requested in ascending order during the phase. Thus, the adversary's cost for serving requests in the  $i$ th round is  $i - 1 + 1 + d = i + d$ . Hence, in each phase the adversary incurs a total cost of at most  $\sum_{i=1}^n (i + d) = n(n + 1)/2 + nd$ . The on-line algorithm can only move an item after all the  $1 + d$  requests of a round have been served. Considering the on-line algorithm's list configuration at the beginning of a phase, for  $i = 1, \dots, n$ , the item stored at position  $i$  in the list is requested in exactly one of the rounds. If an item requested in a round is moved closer to the front of the list at the end of a round, this cannot decrease the cost of subsequent rounds in the phase. Thus the online algorithm's cost in a phase is at least  $\sum_{i=1}^n (1 + d)i = (1 + d)n(n + 1)/2$ . The ratio of the cost incurred by  $A$  to the cost incurred by the adversary is at least  $(1 + d)/(1 + 2d/(n + 1))$ . For large values of  $n$ , this can be arbitrarily close to  $(1 + d)$ . ■

If a randomized on-line algorithm uses paid exchanges, our lower bound is slightly weaker.

**THEOREM 3.4.** *Let  $A$  be a randomized on-line algorithm for the list update problem with delayed action and suppose that  $A$  does use paid exchanges. If  $A$  is  $c$ -competitive against any oblivious adversary, then  $c \geq (1 + d)/2$ . This lower bound holds for both types of adversaries.*

*Proof.* We give a probability distribution on request sequences such that the expected cost incurred by any deterministic on-line algorithm is at least  $(1 + d)/2$  times the expected cost incurred by an adversary. The result then follows from Yao's minimax principle [37]. The request sequence is constructed as follows. In each round one of the  $n$  items is chosen uniformly at random; this item is requested  $1 + d$  times. The expected cost incurred by a deterministic on-line algorithm in a round is  $(1 + d)n/2$  whereas the adversary's cost is no more than  $n + d$ . The ratio of the cost incurred by  $A$  to the cost incurred by the adversary is at least  $(1 + d)/(2 + 2d/n)$ . For large values of  $n$ , this can be arbitrarily close to  $(1 + d)/2$ . ■

#### 4. STOCK TRADING

We consider an on-line stock market model studied in [17] based on similar probabilistic models used for stock price fluctuations (see, e.g., [23]). Consider a game where at each step, the price of a stock either increases by a constant factor  $\alpha > 1$  or decreases by a factor  $1/\alpha$ . The game lasts for  $n$  steps, and the price moves up for  $m$  of these steps. At each step, one can invest a fraction  $s$  of one's wealth in the stock and the rest in cash. If the price moves up, the *return* from that step is the factor  $\alpha s + 1 - s$  that the player's wealth increases; if the price moves down, the return  $\frac{s}{\alpha} + 1 - s$  is less than 1. The *total return* is the factor by which the player's wealth increases over the course of the game. Following [17], we say in this setting that the on-line trader plays against an  $(\alpha, m, n)$ -adversary if an adversary determines the price fluctuations subject to the initial constraints.

We review the relevant results from [17]. Let  $R_\alpha(m, n)$  be the optimal on-line return against the  $(\alpha, m, n)$ -adversary. We have boundary conditions  $R_\alpha(n, n) = \alpha^n$  and  $R_\alpha(0, n) = 1$ . As the optimal algorithm obtains a return of  $\alpha^m$  by investing fully whenever the price will go up, studying the on-line

return is sufficient to find the competitive ratio. The return  $R_\alpha(m, n)$  satisfies the recurrence

$$R_\alpha(m, n) = \max_{0 \leq s \leq 1} \min \left\{ (\alpha s + 1 - s)R_\alpha(m - 1, n - 1), \left( \frac{s}{\alpha} + 1 - s \right)R_\alpha(m, n - 1) \right\},$$

and if we define the partial binomial sum  $B(k; n, p) = \sum_{i=0}^k \binom{n}{i} p^i (1 - p)^{n-i}$ , then the solution to the recurrence satisfies

$$R_\alpha^{-1}(m, n) = B \left( n - m - 1; n - 1, \frac{\alpha}{\alpha + 1} \right) + \alpha^{n-2m} B \left( m - 1; n - 1, \frac{\alpha}{\alpha + 1} \right).$$

An interesting consequence is that even if the number of up movements  $m$  is less than the number of down movements, that is  $m < \frac{n}{2}$ , the on-line player can make a profit. In fact this holds true even if  $m = 1$ .

We consider an extension of this model to two delayed models. In the first model, we consider the problem when the player initially sets a fraction  $s$  of his or her wealth to remain invested over the next  $d + 1$  time steps and can only change the investment  $s$  every  $d + 1$  time steps. This model might apply, for example, to an investor who only performs trades at specific or less frequent time intervals and is unwilling to follow every change in the market. Also this model describes scenarios where a group of agents takes decisions at particular time instances, as mentioned in the Introduction. When  $d = 0$ , we have the original model. We call every set of  $d + 1$  steps a *round*. For convenience we let  $r = d + 1$  be the round length below. Without loss of generality we assume that  $n$  is a multiple of  $r$ .

We let  $P_\alpha(r, m, n)$  be the optimal on-line return for a player playing against an  $(\alpha, m, n)$ -adversary who can change its investment only every  $r$  steps. (Of course  $P_\alpha(1, m, n) = R_\alpha(m, n)$ .) For convenience we drop the  $\alpha$  from the notation where the meaning is clear.

Note then that  $P(r, m, n)$  satisfies the following recurrence:

$$P(r, m, n) = \max_s \min_{\substack{i \\ 0 \leq i \leq r, m}} P(r, m - i, n - r)(\alpha^{2i-r} s + 1 - s).$$

That is, for each round, the optimal player chooses the investment  $s$  that maximizes his or her return regardless of the number of up movements the adversary chooses.

Interestingly, the behavior in this delayed model depends precisely on whether the period length  $r$  is even or odd.

LEMMA 4.1. *For  $r$  even,  $P(r, m, n) = 1$  if  $m \leq n/2$  and  $P(r, m, n) = \alpha^{2m-n}$  if  $m \geq n/2$ .*

*Proof.* If  $m \leq n/2$ , then the adversary can arrange so that each round has at least as many down moves as up moves, and hence no round has a return greater than 1. Of course the player can guarantee a return of 1 by not investing, i.e., choosing  $s = 0$  in each round.

Similarly, if  $m \geq n/2$ , then the player can guarantee a total return of  $\alpha^{2m-n}$  by investing everything each round, i.e., always choosing  $s = 1$ . The adversary can ensure that no greater return is possible by alternating up and down moves on the first  $2(n - m)$  steps. ■

The analysis for  $r$  odd generalizes and makes use of the result from [17] corresponding to the case  $r = 1$  (i.e.,  $d = 0$ ).

LEMMA 4.2. *Let  $N = \frac{n}{r}$  and  $M = m - \lfloor \frac{r}{2} \rfloor \frac{n}{r}$ . For  $r$  odd,  $P(r, m, n) = 1$  if  $m \leq \lfloor \frac{r}{2} \rfloor N$ ,  $P(r, m, n) = \alpha^{2m-n}$  if  $m \geq \lceil \frac{r}{2} \rceil N$ , and  $P(r, m, n) = R_\alpha(M, N)$  otherwise.*

*Proof.* The trivial cases where  $m \leq \lfloor \frac{r}{2} \rfloor n$  or  $m \geq \lceil \frac{r}{2} \rceil n$  handled as in Lemma 4.1.

Otherwise, the problem is more interesting. We first show in this case that  $P(r, m, n) \leq R_\alpha(M, N)$ . Suppose that the adversary announces that in each round there will either be  $\lfloor \frac{r}{2} \rfloor$  or  $\lceil \frac{r}{2} \rceil$  up moves. Then, in total, each round the invested value changes by a factor of  $\alpha$  or  $1/\alpha$ , and there are  $M$  up rounds

out of the  $N$  total rounds. In this case, the problem reduces to the standard case ( $r = 1$ ) from [17]. In particular, the adversary can guarantee a competitive ratio of no more than  $R_\alpha(M, N)$ .

To prove the other direction,  $P(r, m, n) \geq R_\alpha(M, N)$ , we must show that the adversary cannot gain by using any other strategy. We use induction on  $n$ . The base case is trivial.

Now suppose the adversary uses  $\lfloor \frac{r}{2} \rfloor + j$  up moves in the first round. (The cases  $j > 0$  and  $j < 0$  are entirely similar.) By induction, the return for the subsequent rounds is  $R_\alpha(M - j, N - 1)$ . Simple algebraic manipulation (by determining the investor's first investment) yields that the payoff from the first round is

$$\frac{\alpha^{2j-1} - 1}{\alpha - 1} (R_\alpha^{-1}(M - 1, N - 1) - R_\alpha^{-1}(M, N)) + R_\alpha^{-1}(M, N).$$

Hence we have left to show that

$$\left[ \frac{\alpha^{2j-1} - 1}{\alpha - 1} (R_\alpha^{-1}(M - 1, N - 1) - R_\alpha^{-1}(M, N)) + R_\alpha^{-1}(M, N) \right] R_\alpha(M - j, N - 1) \geq R_\alpha(M, N).$$

This is a combinatorial identity that can be checked in a straightforward but quite tedious manner; we spare the reader the details. ■

Next we consider our second delayed model. Suppose that information about trades is continuously updated, but remains  $d$  steps behind. That is, we only know the results from the first trade after the  $(d + 1)$ st trade completes. Again  $d = 0$  corresponds to the original model. Investors can again invest a fraction of their wealth each step. (They may not have accurate knowledge of how much wealth they have, since not all trade results are known. However, investments are possible because only the fraction of the wealth to be invested has to be determined.) This model accounts for situations where one receives updates on prices, but not in real-time. Surprisingly, we can show that there exist money-making schemes for arbitrarily large  $d$  even when there is only 1 up day.

**THEOREM 4.1.** *There exist money-making schemes for  $m = 1$ , regardless of  $n$  and  $d$ .*

*Proof.* Let  $\epsilon_i$  be the investment on the  $i$ th day. We may set  $\epsilon_i = 0$  at any point after the player sees a result which is an up move. It will also be convenient notationally if we define  $\epsilon_i = 0$  for  $i \geq n$ . If the up move is on day  $j$ , then the total return to the player will be

$$(\epsilon_j \alpha + 1 - \epsilon_j) \prod_{i \neq j, i \leq j+d} \left( \frac{\epsilon_i}{\alpha} + 1 - \epsilon_i \right).$$

Note that

$$\left( \frac{\epsilon_a}{\alpha} + 1 - \epsilon_a \right) \left( \frac{\epsilon_b}{\alpha} + 1 - \epsilon_b \right) = \left( \frac{\epsilon_a + \epsilon_b}{\alpha} + 1 - \epsilon_a - \epsilon_b + \epsilon_a \epsilon_b \left( 1 - \frac{1}{\alpha} \right)^2 \right). \quad (1)$$

It will be convenient to assume that the  $\epsilon_i$  will be chosen sufficiently small that we may simplify by removing the nonlinear terms; that is, we proceed as though

$$\left( \frac{\epsilon_a}{\alpha} + 1 - \epsilon_a \right) \left( \frac{\epsilon_b}{\alpha} + 1 - \epsilon_b \right) = \left( \frac{\epsilon_a + \epsilon_b}{\alpha} + 1 - \epsilon_a - \epsilon_b \right).$$

Also,

$$(\epsilon_a \alpha + 1 - \epsilon_a) \left( \frac{\epsilon_b}{\alpha} + 1 - \epsilon_b \right) > 1 \quad \text{if } \epsilon_a > \frac{\epsilon_b}{\alpha(1 - \epsilon_b)}.$$

Hence, the condition

$$(\epsilon_j \alpha + 1 - \epsilon_j) \prod_{i \neq j, i \leq j+d} \left( \frac{\epsilon_i}{\alpha} + 1 - \epsilon_i \right) > 1$$

is satisfied if

$$\epsilon_j > \frac{\sum_{i \neq j, i \leq j+d} \epsilon_i}{\alpha \left( 1 - \sum_{i \neq j, i \leq j+d} \epsilon_i \right)}.$$

This condition is easily satisfied simultaneously for all possible values of  $j$  by choosing  $\epsilon_1$  to be suitably small and having the  $\epsilon_i$  grow geometrically at a suitably small rate (say, less than  $\alpha^{1/d}$ ). Also note the values can easily be chosen so that the effects of the nonlinear terms of Eq. (1) are suitably dominated, justifying our previous simplification. ■

## 5. DELAYED RELAXED TASK SYSTEMS

In this section, we will consider the delayed action model applied to relaxed metrical task systems [6, 10]. An example of a relaxed metrical task system is the ski rental problem described in the Introduction. Another example of a relaxed metrical task system is the  $k$ -page migration problem [10, 15]. For this problem, we wish to keep  $k$  copies of a page available on a network. When a processor wishes to access a page, it requests a copy from a processor holding that page. The communication cost incurred is proportional to the distance between processors. Alternatively, a page copy may migrate from one processor to another at a higher communication cost proportional to the distance between processors. In the delayed model, we assume that the time to transfer a page is nonnegligible, and hence there is a time between when a migration begins and ends during which the old copy serves these requests.

A relaxed metrical task system is associated with a parameter  $D$  and an underlying metrical task system with the same set of configurations. A configuration change in the relaxed task system is  $D$  times more expensive than the corresponding change in the underlying task system. Conveniently, we can demonstrate how to find a competitive algorithm for a relaxed metrical task system in the delayed action model, given a competitive algorithm for the associated metrical task system. Hence we can effectively handle an entire general class of problems, generalizing the work of [6, 10] on relaxed metrical task systems to the setting of delayed actions. We begin by defining a metrical task system [16] and then move on to define relaxed metrical task systems. Here we follow [10].

**DEFINITION 5.1.** A *task system*,  $\mathcal{P}$ , consists of a set of configurations (or states)  $\mathcal{C}$  and a distance function between any two configurations  $C_1, C_2 \in \mathcal{C}$ , denoted  $\text{dist}(C_1, C_2)$ . (This is the *move cost* between the configurations.) The task system consists of a set of requests, called tasks. A task  $r$  is associated with a service cost in each configuration, denoted  $\text{task}(C, r)$  (this is the *task cost*). An algorithm for  $\mathcal{P}$  is associated with a configuration  $C_1$ . Given a request  $r$ , the algorithm serves it by moving to configuration  $C_2$  paying a cost of  $\text{cost}(C_1, C_2, r) = \text{dist}(C_1, C_2) + \text{task}(C_2, r)$ . If the move cost function  $\text{dist}$  forms a metric space over  $\mathcal{C}$ , then the task system is called *metrical*.

**DEFINITION 5.2.** A  $D$ -relaxed task system,  $D\text{-}\mathcal{P}$ , with respect to a task system  $\mathcal{P}$  and some parameter  $D \geq 1/2$ , is the task system with cost, distance, and task functions denoted  $\text{cost}_D$ ,  $\text{dist}_D$ , and  $\text{task}_D$ , respectively.  $\text{dist}_D$  and  $\text{task}_D$  are defined as follows: Given  $C_1, C_2 \in \mathcal{C}$ ,  $\text{dist}_D(C_1, C_2) = D \cdot \text{dist}(C_1, C_2)$ . Given  $C \in \mathcal{C}$  and a task  $r$ ,  $\text{task}_D(C, r) = \min_{C'} \text{dist}(C, C') + \text{task}(C', r)$ .

Consider an algorithm for a task system  $\mathcal{P}$ . Suppose the algorithm starts out in configuration  $C_0$ . It receives a sequence of requests  $r_1, r_2, \dots$ . When request  $r_i$  is received, the algorithm is in configuration  $C_{i-1}$ . The algorithm first moves to configuration  $C_i$  and then services request  $r_i$  from this configuration. The cost of the configuration change is  $\text{dist}(C_{i-1}, C_i)$  and the request service cost is  $\text{task}(C_i, r_i)$ . In the delayed action model, we distinguish between the *real* state of the algorithm and the *desired* state of the algorithm. The algorithm should be in the desired configuration  $C_i$  when it is just about to service request  $r_i$ . However, state changes may not be instantaneous, but occur only after a certain delay. Hence,

the algorithm's state may not be  $C_i$ , but some earlier state  $C_{i-d_i}$ , where  $d_i$  is some delay parameter. Thus, the algorithm must service the request  $C_i$  from state  $C_{i-d_i}$ . The request service cost is therefore  $\text{task}(C_{i-d_i}, r_i)$ . Eventually, the algorithm's real state will go through the same sequence of states as the desired state, i.e.,  $C_0, C_1, C_2, \dots$ . Thus, we can think of the configuration change cost as  $\text{dist}(C_{i-1}, C_i)$ , even though the configuration change may not occur right away. We will assume that the delay is bounded by  $d$ , i.e.,  $d_i \leq d$  for some  $d$ . Note that the case  $d = 0$  gives us the original task system. We consider algorithms for task systems in the delayed action model and determine their competitive ratio as a function of the maximum delay  $d$ . For the analysis, we will assume that the adversary does not have any delay associated with its configuration changes.

For an arbitrary metrical task system  $\mathcal{P}$ , the delayed action model may not be meaningful. In fact, there are task systems  $\mathcal{P}$  such that, in the delayed action model, it is impossible to have a finite competitive ratio even for delay  $d = 1$ , even though there is an algorithm with finite competitive ratio for  $d = 0$ . For example, this could happen in the case of *forcing* task systems, where the request service costs are either 0 or  $\infty$ . For relaxed task systems, however, the delayed action model is meaningful, as we now show.

### 5.1. Cost Analysis for Delayed Relaxed Task Systems

Let  $\mathcal{P}$  be a metrical task system. Let  $\text{task}(C, r)$  be the cost of servicing request  $r$  from configuration  $C$  in  $\mathcal{P}$ . Let  $C_{\min}(C, r)$  denote any configuration  $C'$  which minimizes  $\text{dist}(C, C') + \text{task}(C', r)$ . Let  $\text{task}_D(C, r)$  be the cost of servicing request  $r$  from configuration  $C$  in  $D$ - $\mathcal{P}$ . Then  $\text{task}_D(C, r) = \text{dist}(C, C') + \text{task}(C', r)$ , where  $C' = C_{\min}(C, r)$ .

Consider an algorithm for  $D$ - $\mathcal{P}$ . The total cost in servicing a sequence of requests  $r_1, r_2, \dots, r_n$  by moving through the sequence of states  $C_0, C_1, C_2, \dots, C_n$  is

$$\sum_{i=1}^n \text{dist}_D(C_{i-1}, C_i) + \sum_{i=1}^n \text{task}_D(C_i, r_i) = D \sum_{i=1}^n \text{dist}(C_{i-1}, C_i) + \sum_{i=1}^n (\text{dist}(C_i, C'_i) + \text{task}(C'_i, r_i)),$$

where  $C'_i = C_{\min}(C_i, r_i)$ .

On the other hand, the cost of servicing the request sequence in the delayed model is

$$\begin{aligned} \sum_{i=1}^n \text{dist}_D(C_{i-1}, C_i) + \sum_{i=1}^n \text{task}_D(C_{i-d_i}, r_i) &\leq D \sum_{i=1}^n \text{dist}(C_{i-1}, C_i) + \sum_{i=1}^n (\text{dist}(C_{i-d_i}, C'_i) + \text{task}(C'_i, r_i)) \\ &\leq D \sum_{i=1}^n \text{dist}(C_{i-1}, C_i) + \sum_{i=1}^n \text{dist}(C_{i-d_i}, C_i) + \sum_{i=1}^n (\text{dist}(C_i, C'_i) + \text{task}(C'_i, r_i)) \leq D \sum_{i=1}^n \text{dist}(C_{i-1}, C_i) \\ &\quad + \sum_{i=1}^n \sum_{j=i-d_i+1}^i \text{dist}(C_{j-1}, C_j) + \sum_{i=1}^n (\text{dist}(C_i, C'_i) + \text{task}(C'_i, r_i)) \leq (D + d) \sum_{i=1}^n \text{dist}(C_{i-1}, C_i) \\ &\quad + \sum_{i=1}^n (\text{dist}(C_i, C'_i) + \text{task}(C'_i, r_i)). \end{aligned}$$

Thus for the purpose of analysis, we can think of the delayed model as being equivalent to the model without delay where the cost of moving from configuration  $C_1$  to  $C_2$  is  $(D + d)\text{dist}(C_1, C_2)$  and the request service cost is the same as before. The cost estimate we get using this approximation is an upper bound on the actual cost incurred by the algorithm in the delayed model. On the other hand, since we compare with an adversary that does not face delays, the cost for the adversary is the same as for the relaxed task system without delays. This considerably simplifies the analysis. In particular, this means that if we use the same algorithm for the delayed model as for the original relaxed task system, the cost increases by at most a factor of  $(1 + \frac{d}{D})$ . Hence if  $A$  is a  $c$ -competitive algorithm for the relaxed task system without delays, then  $A$  is a  $c(1 + \frac{d}{D})$ -competitive algorithm for the relaxed task system in the delayed model.

Since the results of [6, 10] show how to turn competitive algorithms for metrical task system into competitive algorithms for relaxed metrical task systems, we now have a means of turning competitive algorithms for metrical task system into competitive algorithms for relaxed metrical task system in the delayed model. The above observation shows that the competitive ratio we achieve for the delayed model is at most a factor of  $(1 + \frac{d}{D})$  times the competitive ratio for the original relaxed task system.

In fact, it is possible to improve on this observation and get better competitive ratios by modifying the algorithm and/or the analysis of [6, 10] to tailor them to the delayed model. We demonstrate this below. Our results generalize the algorithms of [6, 10]; in fact, when  $d = 0$ , our arguments reduce to theirs.

## 5.2. Randomized Algorithm

Let  $A$  be a  $c$ -competitive algorithm for  $\mathcal{P}$ , and let  $D \geq 1/2$ . We give a randomized algorithm Delayed  $D$ -Alg that is competitive against adaptive on-line adversaries for  $D$ - $\mathcal{P}$  in the delayed model. The algorithm is exactly the same as the algorithm in [6] for relaxed task systems.

**ALGORITHM DELAYED  $D$ -Alg.** Algorithm Delayed  $D$ -Alg simulates a version of algorithm  $A$ . At all times, the configuration of Delayed  $D$ -Alg is equal to that of the simulated version of  $A$ . Upon receiving a request  $r$ , with probability  $\frac{1}{2D}$ , feed  $A$  with new request  $r$ , and change the configuration to the new configuration of  $A$ . With probability  $1 - \frac{1}{2D}$ , the algorithm stays in the same configuration and serves the request from there.

**THEOREM 5.1.** *Let  $\mathcal{P}$  be a metrical task system, and let  $A$  be  $c$ -competitive for  $\mathcal{P}$  against adaptive on-line adversaries. Algorithm Delayed  $D$ -Alg is  $(3 + \frac{d-1}{D})c$ -competitive for  $D$ - $\mathcal{P}$  with delay  $d$ , against adaptive on-line adversaries, for  $D \geq 1/2$ .*

The proof is a modification of the proof of Theorem 4.1 in [6]. The definition of relaxed task system we use is from [10]. This is slightly more general than the definition of relaxed task systems used in [6]. However, the proof of Theorem 4.1 in [6] can be easily modified to work for the more general definition [9]. We briefly indicate the modifications in the proof of [6] to obtain the above theorem for relaxed task systems with delay. We will use the same notation as in [6]; we refer the reader to [6] for definitions.

The potential function used is

$$\Phi(h_n, A_n) = (3D + d - 1) \cdot \overline{\text{Up}}(\hat{h}_n, A_n),$$

where  $\overline{\text{Up}}$  is defined by

$$\overline{\text{Up}}(\hat{h}_n, A_n) = \min_{\bar{A}} \{\text{Up}(\hat{h}_n, \bar{A}) + c \cdot \text{dist}(\bar{A}, A_n)\}.$$

When the adversary changes configuration from  $A_n$  to  $A_{n+1}$ , the change in potential is bounded by

$$\Delta\Phi \leq \left(3 + \frac{d-1}{D}\right) \cdot c \cdot \text{dist}_D(A_n, A_{n+1}).$$

The expected cost of algorithm Delayed  $D$ -Alg on receiving request  $r_{n+1}$  is bounded by

$$\mathbb{E}(\text{Cost}_{\text{Del } D\text{-Alg}}(h_n, r_{n+1})) \leq \frac{3D + d - 1}{2D} \cdot \mathbb{E}(\text{Cost}_{\text{Alg}}(\hat{h}_n, r_{n+1})).$$

This then allows us to prove that

$$\mathbb{E}(\Delta\Phi) \leq \left(3 + \frac{d-1}{D}\right) \cdot c \cdot \text{task}_D(A_{n+1}, r_{n+1}) - \mathbb{E}(\text{Cost}_{\text{Del } D\text{-Alg}}(h_n, r_{n+1})).$$

## 5.3. Deterministic Algorithm

For any deterministic algorithm  $A$ , request sequence  $\sigma$ , and request  $r$ , let  $\text{cost}_A(\sigma, r)$  (or  $\text{cost}_A(r)$  when  $\sigma$  follows from the context) be the cost incurred by  $A$  while servicing  $r$  from the configuration reached by previously servicing  $\sigma$ . Also, let  $\text{cost}_A(\sigma)$  be the total cost of  $A$  on  $\sigma$ . Assuming that  $A$  is  $c$ -competitive for  $\mathcal{P}$ , we define the competitive algorithm Delayed  $D$ -DALg for  $D$ - $\mathcal{P}$  as follows. (The algorithm is a modification of the algorithm  $D$ -DALg in [10] for relaxed task systems.)

**ALGORITHM DELAYED D-DAI<sub>g</sub>.** Algorithm Delayed  $D$ -DAI<sub>g</sub> simulates  $2D$  copies  $A_1 \dots A_{2D}$  of  $A$ . Let  $\beta = 2 + \sqrt{1 + d/D}$ . The configuration of Delayed  $D$ -DAI<sub>g</sub> is always the same as that of  $A_1$ . When given a new request  $r$ , the algorithm gives it to one of the  $A_i$  according to the following rule:

- if there exists  $i \geq 2$  such that  $\text{cost}_{A_i}(r) \geq \frac{1}{\beta c} \text{cost}_{A_1}(r)$ ,  $r$  is given to  $A_i$  (i.e., the simulated configuration of  $A_i$  is updated). Then Delayed  $D$ -DAI<sub>g</sub> services  $r$  remotely, without changing its configuration.
- otherwise,  $r$  is given to  $A_1$ . Then Delayed  $D$ -DAI<sub>g</sub> services  $r$  and moves to the new configuration of  $A_1$ .

**THEOREM 5.2.** *Let  $\mathcal{P}$  be a metrical task system and let  $A$  be a  $c$ -competitive deterministic algorithm for  $\mathcal{P}$ . Then algorithm Delayed  $D$ -DAI<sub>g</sub> is  $\beta^2 c^2$ -competitive for the  $D$ -relaxed task system  $D\text{-}\mathcal{P}$ .*

*Proof.* The proof is a modification of the proof of Theorem 2.1 in [10]. It consists of two steps. First, we show that the sum of the costs of algorithms  $A_1 \dots A_{2D}$  is within a factor  $2c$  from the optimal off-line cost of servicing the requests in  $D\text{-}\mathcal{P}$ . Then we show that the cost of Delayed  $D$ -DAI<sub>g</sub> is within a factor  $(\beta^2/2)c$  from the above sum. The result will follow. For brevity, we will often refer to the algorithm Delayed  $D$ -DAI<sub>g</sub> as simply  $D$ -DAI<sub>g</sub>.

The following lemma is proved in [10, Lemma 2.1].

**LEMMA 5.1.** *Let  $\sigma$  be a request sequence, and let  $\sigma_1 \dots \sigma_{2D}$  be (possibly empty) subsequences of  $\sigma$  such that each request from  $\sigma$  appears in exactly one  $\sigma_i$ . Also, let  $A$  be a  $c$ -competitive algorithm for  $\mathcal{P}$  and let  $\text{cost}_{\text{Adv}}(\sigma)$  be the optimal off-line cost of servicing  $\sigma$  in  $D\text{-}\mathcal{P}$ . Then*

$$\sum_{i=1}^{2D} \text{cost}_A(\sigma_i) \leq 2c \cdot \text{cost}_{\text{Adv}}(\sigma).$$

The next lemma is analogous to Lemma 2.2 in [10].

**LEMMA 5.2.** *Let  $\sigma_i$  be a sequence of requests given to  $A_i$  while running  $D$ -DAI<sub>g</sub> on  $\sigma$ . Then*

$$\text{cost}_{D\text{-DAI}_g}(\sigma) \leq \frac{\beta^2}{2} c \sum_{i=1}^{2D} \text{cost}_{A_i}(\sigma_i).$$

*Proof.* We may split  $\text{cost}_{D\text{-DAI}_g}(\sigma)$  into  $\text{cost}_{D\text{-DAI}_g}^S(\sigma)$  (the cost of servicing requests) and  $\text{cost}_{D\text{-DAI}_g}^M(\sigma)$  (the cost of moving between configurations).

We analyze the cost incurred by  $D$ -DAI<sub>g</sub> to service a request  $r$ . If  $r$  is given to  $A_i$ , the cost of servicing  $r$  from the current configuration of  $D$ -DAI<sub>g</sub> is at most  $\beta c$  times  $\text{cost}_{A_i}(r)$ . Hence, we can bound the total cost of servicing requests by  $\beta c \sum_{i=1}^{2D} \text{cost}_{A_i}(\sigma_i)$ .

Therefore, it is sufficient to bound  $\text{cost}_{D\text{-DAI}_g}^M(\sigma) = (D + d) \cdot \text{cost}_{A_1}(\sigma_1)$  in terms of  $\sum_{i=1}^{2D} \text{cost}_{A_i}(\sigma_i)$ . To this end, consider algorithms  $A'_i$  which simulate  $A_i$  on  $\sigma_i$ , but also service all requests from  $\sigma_1$  in the following way: whenever  $r \in \sigma_1$  appears,  $A'_i$  moves from its current configuration  $C$  to  $C' = C_{\min}(C, r)$ , services  $r$  and moves back to  $C$ , paying  $\text{cost}_{A'_i}(r) := 2 \cdot \text{dist}(C, C') + \text{task}(C', r) \leq 2 \cdot (\text{dist}(C, C') + \text{task}(C', r)) \leq 2 \cdot \text{cost}_{A_i}(r)$ . As  $r$  was given to  $A_1$ , we know that

$$\text{cost}_{A_i}(r) \leq \frac{1}{\beta c} \text{cost}_{A_1}(r)$$

which implies

$$\text{cost}_{A'_i}(r) \leq \frac{2}{\beta c} \text{cost}_{A_1}(r).$$

Hence the total cost of  $A'_i$  (denoted by  $\text{cost}_{A'_i}(\sigma_1)$ ) is bounded by

$$\text{cost}_{A_i}(\sigma_i) + \sum_{r \in \sigma_1} \text{cost}_{A'_i}(r) \leq \text{cost}_{A_i}(\sigma_i) + \frac{2}{\beta c} \sum_{r \in \sigma_1} \text{cost}_{A_1}(r) = \text{cost}_{A_i}(\sigma_i) + \frac{2}{\beta c} \text{cost}_{A_1}(\sigma_1).$$

On the other hand, the algorithm  $A_1$  is  $c$ -competitive, so  $\text{cost}_{A_1}(\sigma_1) \leq c \cdot \text{cost}_{A'_1}(\sigma_1)$ . Hence

$$\frac{1}{c} \text{cost}_{A_1}(\sigma_1) \leq \text{cost}_{A'_1}(\sigma_1) \leq \text{cost}_{A_i}(\sigma_i) + \frac{2}{\beta c} \text{cost}_{A_1}(\sigma_1)$$

and thus  $\text{cost}_{A_1}(\sigma_1) \leq \frac{\beta c}{(\beta-2)} \cdot \text{cost}_{A_i}(\sigma_i)$ . Now we can bound the moving cost as follows:

$$\begin{aligned} \text{cost}_{D\text{-DAlg}}^M(\sigma) &= (D+d) \cdot \text{cost}_{A_1}(\sigma_1) \\ &\leq \frac{1}{2} \left(1 + \frac{d}{D}\right) \sum_{i=1}^{2D} \text{cost}_{A_1}(\sigma_1) \\ &\leq \frac{\beta c}{2(\beta-2)} \left(1 + \frac{d}{D}\right) \sum_{i=1}^{2D} \text{cost}_{A_i}(\sigma_i). \\ \text{cost}_{D\text{-DAlg}}(\sigma_1) &= \text{cost}_{D\text{-DAlg}}^S(\sigma) + \text{cost}_{D\text{-DAlg}}^M(\sigma) \\ &\leq \left(\beta + \frac{\beta}{2(\beta-2)} \left(1 + \frac{d}{D}\right)\right) c \sum_{i=1}^{2D} \text{cost}_{A_i}(\sigma_i) \\ &= \frac{\beta^2}{2} c \sum_{i=1}^{2D} \text{cost}_{A_i}(\sigma_i). \end{aligned}$$

■

Theorem 5.2 follows from Lemmas 5.1 and 5.2. ■

#### 5.4. Other Results

Similar to the results in [10], we can get slightly better competitive ratios for *monotonic* task systems (defined below), as well as randomized algorithms against oblivious adversaries.

##### *Monotonic Task Systems*

**DEFINITION 5.3.** A *monotonic task system* is a forcing task system with a monotonicity property between configurations as follows. A configuration  $C$  is said to be dominated by  $C'$  if for all tasks for which  $C$  is allowable so is  $C'$ . A forcing task system is *monotonic* if for every pair of configurations  $C_1, C_2$  there exists a configuration  $C$  dominating both, and for every configuration  $C'_1$  dominated by  $C_1$ ,  $\text{dist}(C_1, C) \leq \text{dist}(C'_1, C_2)$ .

A better ratio of  $\gamma^2 c^2$  (where  $\gamma = 1 + \sqrt{1 + \frac{d}{D}}$ ) may be obtained when the underlying task system  $\square$  is *monotonic*. An example of a monotonic task system is the Steiner tree problem. The corresponding relaxed version is the page replication problem. Another example is the generalized Steiner tree problem; the relaxed version is the network leasing problem.

To get the better bound, we use a modified version of Delayed  $D$ -DAlg, which now simulates  $D$  algorithms  $A_1 \dots A_D$  and gives a requests  $r$  to  $A_i$  for which  $\text{cost}_{A_i}(v) \geq \frac{1}{\gamma c} \text{cost}_{A_1}(r)$  (if such an algorithm exists) or to  $A_1$  otherwise. The analysis is similar to that in [10].

##### *Randomized Algorithm against Oblivious Adversary*

One can define a randomized version of  $D$ -DAlg, called  $D$ -RALg, which is  $(3 + \frac{d}{D})c$ -competitive against an oblivious adversary. For monotonic task systems it is  $(2 + \frac{d}{D})c$ -competitive. The algorithm is exactly the same as the randomized algorithm for relaxed task systems given in [10]; hence the same name. The algorithm  $D$ -RALg simulates  $2D$  algorithms  $A_1 \dots A_{2D}$  ( $D$  algorithms in the monotonic case). At the beginning it chooses one of them at random (say  $A_i$ ) and then always keeps the same configuration as  $A_i$ . The requests are always given to the algorithm which incurs the *highest* cost. The following lemma bounds the expected cost of the algorithm  $D$ -RALg.



LEMMA 5.3. *The expected cost of  $D$ -RALg is at most  $(\frac{3}{2} + \frac{d}{2D}) \sum_{i=1}^{2D} \text{cost}_{A_i}(\sigma_i)$ .*

*Proof.* Suppose the algorithm  $D$ -RALg simulates  $A_i$ . We claim that the cost of  $D$ -RALg is at most  $(D + d) \cdot \text{cost}_{A_i}(\sigma_i) + \sum_{j=1}^{2D} \text{cost}_{A_j}(\sigma_j)$ . We will split the total cost of  $D$ -RALg into two parts: the cost of moving between configurations and the cost of servicing requests.

The total movement cost of  $D$ -RALg is at most  $(D + d) \cdot \text{cost}_{A_i}(\sigma_i)$ . Whenever  $A_i$  changes configuration from  $C$  to  $C'$ ,  $D$ -RALg also changes configuration from  $C$  to  $C'$ . The cost incurred by  $A_i$  is  $\text{dist}(C, C')$  and that by  $D$ -RALg is  $(D + d) \cdot \text{dist}(C, C')$ .

Consider now the total request service cost of  $D$ -RALg. Suppose a request  $r$  is given to  $A_j$ . The cost incurred by  $D$ -RALg to service this request is  $\text{cost}_{A_i}(r) \leq \text{cost}_{A_j}(r)$ . Hence the total request service cost of  $D$ -RALg is bounded by  $\sum_{j=1}^{2D} \text{cost}_{A_j}(\sigma_j)$ .

This proves the claim that if  $D$ -RALg chooses  $A_i$ , the cost incurred by it is at most  $(D + d) \cdot \text{cost}_{A_i}(\sigma_i) + \sum_{i=1}^{2D} \text{cost}_{A_i}(\sigma_i)$ . Since  $i$  is chosen uniformly and at random from the set  $1, \dots, 2D$ , the expected cost of  $D$ -RALg is at most  $\frac{D+d}{2D} \sum_{i=1}^{2D} \text{cost}_{A_i}(\sigma_i) + \sum_{i=1}^{2D} \text{cost}_{A_i}(\sigma_i)$  which is  $(\frac{3}{2} + \frac{d}{2D}) \sum_{i=1}^{2D} \text{cost}_{A_i}(\sigma_i)$ . ■

Combining Lemma 5.3 with Lemma 5.1, the expected competitive ratio of  $D$ -RALg is bounded by  $(3 + \frac{d}{D})c$ .

In the case of monotonic task systems, the algorithm  $D$ -RALg simulates  $D$  algorithms  $A_1, \dots, A_D$ . It is easy to modify the proof of Lemma 5.3 to show that in this case, the expected cost of  $D$ -RALg is bounded by  $(2 + \frac{d}{D}) \sum_{i=1}^D \text{cost}_{A_i}(\sigma_i)$ . Combining this with Lemma 18, the expected competitive ratio of  $D$ -RALg is bounded by  $(2 + \frac{d}{D})c$ .

## 6. CONCLUSION

We have considered the effects of delayed action and delayed information for a variety of on-line problems, including the general class of problems corresponding to relaxed metrical task systems. Our results demonstrate that in many cases appropriate algorithms can deal gracefully with delay to the extent that the competitive ratio grows slowly as the delay increases. We believe that examining delayed situations, besides yielding interesting problems, gives more insight into these on-line problems. In particular, by studying delay one learns more about the underlying model and how reasonable it appears as well as how robust suggested algorithms are for handling slightly different situations.

Further directions to pursue include studying the effects of delay on more challenging on-line problems, such as the  $k$ -server problem. Also, determining how to introduce notions of delay in more general models of on-line problems may yield interesting results.

## REFERENCES

1. Albers, S. (1999), Better bounds for online scheduling, *SIAM J. Comput.* **29**, 459–473.
2. Albers, S., von Stengel, B., and Werchner, R. (1995), A combined BIT and TIMESTAMP algorithm for the list update problem, *Inform. Process. Lett.* **56**, 135–139.
3. Alon, N., Kalai, G., Ricklin, M., and Stockmeyer, L. (1992), Lower bounds on the competitive ratio for mobile user tracking and distributed job scheduling, in “Proc. 33rd Ann. Symp. on the Foundations of Computer Science,” pp. 334–343.
4. Altman, E., and Nain, P. (1992), Closed loop control with delayed information, in “Proceedings of the ACM Sigmetrics Conference on Measurement and Modeling of Computer Systems, Newport, Rhode Island,” pp. 193–204, Assoc. Comput. Mach., New York.
5. Altman, E., and Stidham, S. (1996), Optimality of monotonic policies for two-action Markovian decision processes, with applications to control of queues with delayed information, “Queueing Systems: Theory and Applications.”
6. Awerbuch, B., Azar, Y., and Bartal, Y. (1996), On-line generalized Steiner problem, in “Proc. 7th Ann. ACM-SIAM Symp. on Discrete Algorithms,” pp. 68–74.
7. Awerbuch, B., Azar, Y., Fiat, A., and Leighton, T. (1996), Making commitments in the face of uncertainty: How to pick a winner almost every time, in “Proc. 28th Ann. ACM Symp. on Theory of Computing,” pp. 519–530.
8. Awerbuch, B., Bartal, Y., and Fiat, A. (1993), Competitive distributed file allocation, in “Proc. 25 ACM Symp. on Theory of Computing,” pp. 164–173.
9. Bartal, Y. (1997), personal communication.

10. Bartal, Y., Charikar, M., and Indyk, P. (1997), On page migration and other relaxed task systems, in "Proc. 8th Ann. ACM-SIAM Symp. on Discrete Algorithms."
11. Bartal, Y., Byers, J., and Raz, D. (1997), Global optimization using local information with applications to flow control, in "Proc. 38th Ann. Symp. on Foundations of Computer Science," pp. 303–312.
12. Bartal, Y., Fiat, A., and Rabani, Y. (1992), Competitive algorithms for distributed data management, in "Proc. 24th Ann. ACM Symp. on the Theory of Computing," pp. 39–49.
13. Bartal, Y., and Rosen, A. (1997), The distributed  $k$ -server problem—A competitive distributed translator for  $k$ -server algorithms, *J. Algorithms* **23**, 241–264.
14. Ben-David, S., and Borodin, A. (1994), A new measure for the study of on-line algorithms, *Algorithmica* **11**, 73–91.
15. Black, D. L., and Sleator, D. D. (1989), "Competitive Algorithms for Replication and Migration Problems," Technical Report CMU-CS-89-201, Department of Computer Science, Carnegie-Mellon University.
16. Borodin, A., Linial, N., and Saks, M. (1992), An optimal on-line algorithm for metrical task systems, *J. Assoc. Comput. Mach.* **39**, 745–763.
17. Chou, A., Cooperstock, J., El-Yaniv, R., Klugerman, M., and Leighton, T. (1995), The statistical adversary allows optimal money-making trading schemes, in "Proc. 6th Ann. ACM-SIAM Symp. on Discrete Algorithms," pp. 467–476.
18. Cover, T. M. (1991), Universal portfolios, *Math. Financ.* **1**, 1–29.
19. Deng, X., and Papadimitriou, C. H. (1992), Competitive distributed decision-making, in "Proc. 12th IFIP Congress," pp. 350–356.
20. El-Yaniv, R., Fiat, A., Karp, R., and Turpin, G. (1992), Competitive analysis of financial games, in "Proc. 33rd Ann. Symp. on Foundations of Computer Science," pp. 327–333.
21. Graham, R. L. (1966), Bounds for certain multi-processing anomalies, *Bell Systems Technol. J.* **45**, 1563–1581.
22. Halldórsson, M. M., and Szegedy, M. (1992), Lower bounds for on-line graph coloring, in "Proc. 3rd Ann. ACM-SIAM Symp. on Discrete Algorithms," pp. 211–216.
23. Hull, J. C. (1993), "Options, Futures, and Other Derivative Securities," 2nd ed., Prentice-Hall, New York.
24. Irani, S. (1994), Coloring inductive graphs on-line, *Algorithmica* **11**, 53–62.
25. Irani, S., and Rabani, Y. (1993), On the value of information in coordination games, in "Proc. 34th Ann. Symp. on Foundations of Computer Science," pp. 12–21.
26. Karp, R., and Raghavan, P. From a personal communication cited in [34].
27. Koutsoupias, E., and Papadimitriou, C. H. (1994), Beyond competitive analysis, in "Proc. 35th Ann. Symp. on Foundations of Computer Science," pp. 394–400.
28. Kuri, J., and Kumar, A. (1995), Optimal control of arrivals to queues with delayed queue length information, *IEEE Trans. Automat. Control* **40**, pp. 1444–1450.
29. Mitzenmacher, M. (1997), How useful is old information? in "Proc. 16th Ann. ACM Symp. on Principles of Distributed Computing," pp. 83–91.
30. Mirchandaney, R., Towsley, D., and Stankovic, J. A. (1989), Analysis effects of delays on load sharing, *IEEE Trans. Comput.* **38**, 1513–1525.
31. Papadimitriou, C. H., and Yannakakis, M. (1991), On the value of information in distributed decision making, in "Proc. 25th ACM Symp. on Principles of Distributed Computing," pp. 61–64.
32. Papadimitriou, C. H., and Yannakakis, M. (1993), Linear programming without the matrix, in "Proc. 25th ACM Symp. on Theory of Computing," pp. 121–129.
33. Raghavan, P. (1991), "A Statistical Adversary for On-line Algorithms," On-Line Algorithms DIMACS Series in Discrete Mathematics and Theoretical Computer Science, pp. 79–83.
34. Reingold, N., Westbrook, J., and Sleator, D. D. (1994), Randomized competitive algorithms for the list update problem, *Algorithmica* **11**, 15–32.
35. Sleator, D. D., and Tarjan, R. E. (1985), Amortized efficiency of list update and paging rules, *Commun. Assoc. Comput. Mach.* **28**, 202–208.
36. Towsley, D., and Mirchandaney, R. (1988), The effect of communication delays on the performance of load balancing policies in distributed systems, in "Proc. Second International MCPR Workshop," pp. 213–226.
37. Yao, A. C.-C. (1977), Probabilistic computations: Towards a unified measure of complexity, in "Proc. 17th Ann. Symp. on Foundations of Computer Science," pp. 222–227.