
Online Optimization

lecture notes, summer term 2023
Prof. Dr. Yann Disser



TECHNISCHE
UNIVERSITÄT
DARMSTADT

Contents

1	Introduction	3
1.1	Offline optimization	4
1.2	Online optimization	4
1.3	The ski rental problem	6
2	The Paging Problem	9
2.1	Optimal offline paging	9
2.2	Lower bound for online paging	10
2.3	Online paging algorithms	11
2.4	Phase partitioning	12
3	Randomized Algorithms	15
3.1	Randomized paging algorithms	15
3.2	Yao's principle	18
3.3	Lower bound for randomized paging	19
4	The List Update Problem	21
4.1	Potential function method (amortized analysis)	21
4.2	Upper bound for move-to-front	24
4.3	Lower bound via averaging	25
4.4	Upper bounds via phase partitioning	26
4.5	Randomized list update algorithms	31
4.6	Lower bound for randomized list update	33
5	Metrical Task Systems*	37
5.1	Lower bound for metrical task systems	38
5.2	Work function algorithm	39
6	The k-Server Problem	43
6.1	Lower bound for the k -server problem	44
6.2	The k -server problem on the line	45
6.3	Balancing algorithm*	48
6.4	General metric spaces	52
7	Primal-Dual Algorithms	53
7.1	The primal-dual method	53
7.2	Randomized rounding	56
7.3	Online bipartite matching	57
8	Online Load Balancing	65
8.1	Identical machines	66
8.2	Related machines	67
8.3	Restricted assignment	69
8.4	Unrelated machines	71

1 Introduction

In classical (*offline*) optimization problems we are given some input data and need to make optimization decisions (i.e., set variables) such that a certain objective is optimized. This framework is too restrictive for problems where parts of the input become available over time (i.e., *online*) and only after some optimization decisions already had to be fixed. For example, consider the problem of controlling a personal elevator: People arrive over time on different floors and request the elevator, but the elevator cannot wait for all requests to arrive before starting to move, at least if we want to achieve a good completion time until all requests have been served.

Online optimization addresses problems with a temporal component as in the elevator problem above. Conceptually, an online algorithm gets its input item by item, one input *event* at a time (e.g., elevator button press). Upon each event, the algorithm immediately needs to respond by making an optimization decision (e.g., move the elevator up/down or wait). Importantly, the algorithm needs to act on the same timeline (“online”) on which events occur, i.e., each decision must only be based on past events and must be independent of the future. The challenge for the online algorithm lies in finding a good balance between efficiently handling events while protecting itself against possible futures. For example, assume the elevator is at the ground floor and receives a request in the fifth floor. Should it start moving immediately, or should it wait a bit to see whether there are passengers that want to go upwards from the ground floor?

In this lecture, we devise “good” online algorithms for various problems. Of course, no online algorithm can always be good, because no matter how it decides to react to a request, its response can always turn out to be a mistake in retrospect. A possible quality measure for an online algorithm would be to consider its expected objective function value with respect to a fixed distribution over inputs. The main issue with this approach lies in how to choose a suitable distribution, and that the fact whether an algorithm is good heavily depends on this choice. Instead, we will use *competitive analysis*, an absolute quality measure for online algorithms based on their worst-case behavior, much like classical running time analysis concentrates on worst-case running times.

Competitive analysis challenges an online algorithm to compete against an *adversary* that counters every action of the algorithm with a worst-possible sequence of future events. For example, if the elevator decides to move away from the ground floor immediately, the adversary may decide to send a request at the ground floor in the next time step, and if the elevator waits, the adversary may not send additional requests. As mentioned above, the adversary can often easily ruin the objective function value of the online algorithm. To make the interaction fairer for the online algorithm, we demand that the adversary present its own *offline* solution once all requests have been revealed, and we compare the performance of the online algorithm with the quality of the offline solution by computing the *competitive ratio* between the two. Of course, this measure is still quite harsh for the online algorithm, since the offline solution of the adversary may use full knowledge of the entire event sequence. However, in many cases, we will be able to give online algorithms with bounded competitive ratios. Lower bounds on the best-possible competitive ratio of an online problem express the loss in solution quality that is due to imperfect information in the online setting, and we will consider an online algorithm to be “good”, if its competitive ratio is best-possible.

In the following, we formally define the notions introduced above and apply competitive analysis to various typical online problems. These online problems naturally arise in many settings:

Server Problems: Requests arrive over time and we control server movements (e.g., elevators).

Data Reorganization: Data is accessed over time and we have to reorganize a data structure to speed up future accesses.

Machine Scheduling: Tasks arrive over time and we have to assign them to machines (workers).

Bin Packing: Items arrive over time and we have to pack them into bins of limited size.

Bipartite Matching: Candidates arrive over time and we have to assign each to a partner.

Trading: Stock prices arrive over time and we have to decide whether to buy or sell stocks.

1.1 Offline optimization

We start by revisiting offline optimization and formally establish the notation we will use in this lecture.

Definition 1.1. An (offline) optimization problem \mathcal{P} is a 4-tuple $(\mathcal{I}, \text{Sol}, c, \text{type})$, where:

- \mathcal{I} is a set of input instances
- Sol is a function mapping each instance $I \in \mathcal{I}$ to a set of feasible solutions with $\mathcal{S} := \bigcup_{I \in \mathcal{I}} \text{Sol}(I)$.
- $c: \mathcal{I} \times \mathcal{S} \rightarrow \mathbb{R}_{\geq 0}$ is the objective function (cost or utility)
- $\text{type} \in \{\min, \max\}$ expresses whether we want to minimize the cost c or maximize the utility c

Most optimization problems we consider are minimization problems, and, for clarity, the following definitions are with respect to a fixed optimization problem $\mathcal{P} = (\mathcal{I}, \text{Sol}, c, \min)$, where c is a cost function. We introduce basic notation to distinguish between the solutions computed by an algorithm ALG and its cost.

Definition 1.2. An (offline) algorithm ALG computes a feasible solution $\text{ALG}[I] \in \text{Sol}(I)$ of cost $\text{ALG}(I) := c(I, \text{ALG}[I])$ for all $I \in \mathcal{I}$.

Definition 1.3. By OPT we denote an (arbitrary but fixed) optimum algorithm with $\text{OPT}(I) = \min_{S \in \text{Sol}(I)} c(I, S)$ for all $I \in \mathcal{I}$.

The definition of a competitive ratio will be similar to the approximation ratio of approximation algorithms, which we recall below.

Definition 1.4. A polynomial-time algorithm ALG is an (asymptotic) ρ -approximation algorithm (with constant $\alpha \geq 0$) if, for all $I \in \mathcal{I}$,

$$\text{ALG}(I) \leq \rho \cdot \text{OPT}(I) (+ \alpha).$$

Definition 1.5. The approximation ratio of a polynomial-time algorithm ALG is the infimum over all ρ , such that ALG is a ρ -approximation algorithm.

1.2 Online optimization

We now make the notion of online algorithms formal. Intuitively, instead of being provided with the entire input at once, online algorithms get a sequence of inputs (requests) and need to react to (i.e., to answer) each request immediately. Formally, we express this interaction as a game in the following sense.

Definition 1.6. An online problem is expressed as a request-answer game $(R, \Sigma, A, \mathcal{C})$, where:

- R is the set of possible requests
- $\Sigma \subseteq \bigcup_{i \in \mathbb{N}} R^i$ is the set of all possible instances
- A is the set of possible answers
- $\mathcal{C} = (c_i)_{i \in \mathbb{N}}$ is a sequence of cost (utility) functions $c_i: R^i \times A^i \rightarrow \mathbb{R}_{\geq 0} \cup \{\infty\}$ that capture the cost (utility) of the answers up to time step i and $c_0 := 0$

For an instance $\sigma = (r_1, \dots, r_n) \in \Sigma$ of length $n \in \mathbb{N}$ we write $\sigma_{\leq i} := (r_1, \dots, r_i)$ for $i \leq n$ and let $\sigma_{\leq 0} := \emptyset$. We call an online problem *finite* if R, A, Σ are all finite.

Our definition of an online algorithm needs to allow for sequential input, but at the same time ensures that requests are answered immediately, i.e., without using knowledge of future requests. The following definitions are with respect to a fixed request-answer game $(R, \Sigma, A, \mathcal{C})$, and we again assume that c_1, c_2, \dots are cost functions that need to be minimized, the definitions for utility maximization instead of cost minimization are analogous.

Definition 1.7. A deterministic online algorithm ALG computes a sequence of functions $(f_i: R^i \rightarrow A)_{i \in \mathbb{N}}$. The solution of ALG for the sequence $\sigma = (r_1, \dots, r_n) \in (\Sigma \cap R^n)$ is $\text{ALG}[\sigma] := (a_1, \dots, a_n) \in A^n$, where $a_i := f_i(\sigma_{\leq i})$, its total cost is $\text{ALG}(\sigma) := c_n(\sigma, \text{ALG}[\sigma])$, and its cost in step $i \in \{1, \dots, n\}$ is $\text{ALG}_i(\sigma) := \text{ALG}(\sigma_{\leq i}) - \text{ALG}(\sigma_{\leq i-1})$.

As a benchmark, we will compare the performance of an online algorithm on an instance to the best possible offline solution we could have gotten if we had known the entire request sequence ahead of time.

Definition 1.8. The (offline) optimum solution $\text{OPT}[\sigma]$ for sequence $\sigma = (r_1, \dots, r_n) \in (\Sigma \cap R^n)$ is fixed such that $\text{OPT}[\sigma] \in \arg \min_{\mathbf{a} \in A^n} c_n(\sigma, \mathbf{a})$. Its total cost is $\text{OPT}(\sigma) = c_n(\sigma, \text{OPT}[\sigma])$. The optimum cost in step $i \in \{1, \dots, n\}$ is defined via $\text{OPT}_i(\sigma) := \text{OPT}(\sigma_{\leq i}) - \text{OPT}(\sigma_{\leq i-1})$.

Notice that, in contrast to the definition of online algorithms, OPT need not be consistent between time steps. In particular, the offline optimum can be completely different for instances $\sigma_{\leq n-1}$ and σ_n .

We are now ready to introduce the notion of the competitive ratio that we will use throughout this lecture to assess the quality of online algorithms. Observe that the definition is unfair in the sense that we do not compare to the best possible online algorithm, but to the best offline solution for each instance independently, even without requiring the solutions to these instances to be consistent with one another.

Definition 1.9. A deterministic online algorithm ALG is ρ -competitive, $\rho \geq 1$ if there is a constant $\alpha \geq 0$, such that for all instances $\sigma \in \Sigma$

$$\text{ALG}(\sigma) \leq \rho \cdot \text{OPT}(\sigma) + \alpha.$$

ALG is *strictly* ρ -competitive if the above holds for $\alpha = 0$.

Note that, in general, we allow ρ to be any function of σ (in particular of the length n of the input) and of any other problem parameters (e.g., number of floors of the building for the elevator).

Definition 1.10. If there is a constant $\rho \in \mathbb{R}$ for which the online algorithm ALG is (strictly) ρ -competitive, we say that ALG is (strictly) competitive. The (strict) competitive ratio of ALG then is the infimum over all $\rho \in \mathbb{R}$, such that ALG is (strictly) ρ -competitive.

Definition 1.11. The (strict) competitive ratio of an online problem is the infimum over all ρ for which a (strictly) ρ -competitive algorithm exists.

Note similarities and differences to the notion of approximation ratios for the offline setting. In both notions, we ask how close we can get to optimal solutions in the worst-case if we have limited resources. For approximation algorithms, the allowed processing time is limited, while online algorithms have limited access to information. While we do not require online algorithms to be efficient, most algorithms we will see will have polynomial running times.

Also note that, while the strict competitive ratio is a guarantee for every possible request sequence, the (non-strict) competitive ratio provides an asymptotic guarantee. In order to show lower bounds on the competitive ratio of an algorithm, we will often use the following characterization.

Proposition 1.12. An online algorithm ALG has competitive ratio at least $\rho > 1$ if and only if there exists an infinite sequence $(\sigma^{(i)} \in \Sigma)_{i \in \mathbb{N}}$ with $\lim_{i \rightarrow \infty} \text{ALG}(\sigma^{(i)}) = \infty$ and

$$\limsup_{i \rightarrow \infty} \frac{\text{ALG}(\sigma^{(i)})}{\text{OPT}(\sigma^{(i)})} \geq \rho.$$

1.3 The ski rental problem

As an example to see the above definitions in action, we consider a typical everyday problem: If we do not know how often (or long) we will use some item, how long should we rent it before deciding to buy? This problem is typically phrased in terms of renting skis, but can trivially be adapted to many similar situations. While, formally, we would have to specify the individual components (request set, answer set, cost functions) of the underlying request answer game, we will usually resort to a more intuitive description of online problems in the following standard form:

SKI RENTAL PROBLEM	
given:	cost 1 to rent skis for one day, cost $B \in \mathbb{N}$ to buy skis
online:	days $\sigma = (d_1, \dots, d_n) \in \{1\}^n$ that skis are needed (only n is unknown)
actions (step i):	rent skis (cost 1) or buy skis (cost B) or use skis bought earlier (cost 0)
objective:	minimize the total cost $\sum_{i=1}^n \text{ALG}_i(\sigma)$

For the sake of the example, we will, for once, also give the formal definition of the problem. Since only the rental duration is unknown, we have a trivial set of requests $R = \{\text{need skis}\}$ that only allows the request “need skis” for every day that skis are needed. Then $\Sigma = \bigcup_{i \in \mathbb{N}} R^i$. The possible answers on each day are $A = \{\text{rent, buy, use}\}$. At this point it may be worrying that we allow “use” without ensuring that we previously chose “buy”. Constraints that are based on the previous answers generally need to be modeled via the cost functions. We let $c_0 := 0$ and define

$$c_i(a_1, \dots, a_i) = c_{i-1}(a_1, \dots, a_{i-1}) + \begin{cases} 1, & \text{if } x = \text{rent,} \\ B, & \text{if } x = \text{buy,} \\ 0, & \text{if } x = \text{use} \wedge \text{buy} \in \bigcup_{j=1}^{i-1} \{a_j\}, \\ \infty, & \text{otherwise.} \end{cases}$$

Note that the only information that an online algorithm learns in step i is that $n \geq i$. Therefore, every (sensible) deterministic online algorithm can be expressed as an algorithm ALG^i that always buys in the i -th step or the algorithm ALG^∞ that never buys. We show that the unique best strategy (in the worst-case) for renting skis is to buy in step B .

Theorem 1.13. ALG^B is strictly $(2 - 1/B)$ -competitive.

Proof. First observe that the optimum cost is given by

$$\text{OPT}(\sigma) = \min\{n, B\}.$$

The cost of ALG^B is

$$\text{ALG}^B(\sigma) = \begin{cases} n & \text{if } n < B, \\ 2B - 1 & \text{else.} \end{cases}$$

Now, if $n < B$, then

$$\text{ALG}^B(\sigma) = n = \text{OPT}(\sigma),$$

and if $n \geq B$, then

$$\text{ALG}^B(\sigma) = 2B - 1 = (2 - 1/B) \cdot \text{OPT}(\sigma).$$

Overall, we get $\text{ALG}^B(\sigma) \leq (2 - 1/B) \cdot \text{OPT}(\sigma)$ as claimed. \square

Theorem 1.14. For every $i \neq B$, ALG^i is not strictly $(2 - 1/B)$ -competitive.

Proof. We want to prove a lower bound on the competitive ratio of ALG^i for $i \neq B$, so we take the perspective of the adversary that reacts with the worst-possible request sequence to the answers of the algorithm. The only freedom the adversary has in the ski rental problem lies in the choice of n . Since $\text{ALG}^\infty(\sigma) \rightarrow_{n \rightarrow \infty} \infty$ while $\text{OPT}(\sigma) \leq \min\{n, B\}$, trivially, ALG^∞ is not competitive (cf. Proposition 1.12). For finite i , we consider the most natural strategy (for the adversary) that stops the request sequence just after ALG^i buys, i.e., just after step i .

If $i \leq B - 1$, we have

$$\text{ALG}^i(\sigma) = i - 1 + B \geq 2i = 2 \cdot \text{OPT}(\sigma).$$

If $i \geq B + 1$, we have

$$\text{ALG}^i(\sigma) = i - 1 + B \geq 2B = 2 \cdot \text{OPT}(\sigma).$$

In both cases $\text{ALG}^i(\sigma) \geq 2 \cdot \text{OPT}(\sigma) > (2 - 1/B) \cdot \text{OPT}(\sigma)$, as claimed. \square

Corollary 1.15. The ski rental problem has a strict competitive ratio of $(2 - 1/B)$.

Proof. By Theorems 1.13 and 1.14 it suffices to show that ALG^B has competitive ratio at least $(2 - 1/B)$. To see this, we consider the adversarial request sequence σ of length B . For this sequence, $\text{OPT}(\sigma) = B$ and

$$\text{ALG}^B(\sigma) = 2B - 1 = (2 - 1/B) \cdot \text{OPT}(\sigma). \quad \square$$



2 The Paging Problem

A typical source of online problems is the dynamic re-organization of data structures in order to improve access times for elements that are accessed often. The paging problem is a basic online problem of this type that models data access in memory hierarchies, where memory layers close to the CPU are very fast but small, while memory layers further away are much slower but also much larger. In the paging problem we consider a simple hierarchy consisting of a slow main memory that contains m memory pages and a fast cache that can only hold $k < m$ pages at the same time (see Figure 2.1). The problem now is to decide online which pages to keep in the cache over time in order to minimize the number of accesses to main memory.

PAGING PROBLEM	
given:	memory pages $P = \{p_1, \dots, p_m\}$, cache size $k \in \mathbb{N}$, cache C with initial contents $C_0 \subset P$, $ C_0 = k$
online:	sequence $\sigma = (r_1, \dots, r_n) \in P^n$ of pages that are accessed in this order
actions (step i):	evict (remove) pages from C (free) load any number $m_i \leq k - C $ of pages from $P \setminus C$ into C (paid) at the end of step i , we must have $r_i \in C$
objective:	minimize the total cost $\sum_{i=1}^n \text{ALG}_i(\sigma) = \sum_{i=1}^n m_i$




Figure 2.1: Illustration of the paging memory model with a fast/small cache C and a slow/large memory P .

Remark 2.1. We can restrict ourselves to *demand paging* algorithms that may only load page r_i in step i , and only evict a page if needed to make room for r_i . We can easily bring any algorithm into this form by delaying all unnecessary evicts and loads for as long as possible. Dealing with demand paging algorithms simplifies the analysis, because we only need to count the number of *page faults*, i.e., the number of times a page in $P \setminus C$ is requested.

2.1 Optimal offline paging

If we know the sequence of pages that will be requested ahead of time, it is natural to use the *longest-forward-distance* (LFD) heuristic for the paging problem: On every page fault evict any page that will not be requested again, and, if no such page exists, evict the page whose next access is furthest in the future. As an example, consider the request sequence $\sigma = (a, a, b, e, f, b, b, c, d, a, b)$ for the cache $C_0 = \{a, b, c, d\}$ of capacity $k = 4$. The first cache fault is caused by $r_4 = e$, and LFD evicts a (forward distance 6) and replaces it by e .




Figure 2.2: Illustration of cases 1 and 3 in the proof of Theorem 2.2. We use ' $x \rightsquigarrow y$ ' to denote that page x is loaded and page y is evicted.

Theorem 2.2 (Belady [1966]). LFD is an optimal (offline) paging algorithm.

Proof. For the sake of contradiction, assume that LFD is not optimal and let $\sigma = (r_1, \dots, r_n)$ be an instance where it is not optimal. Let OPT be an optimal (demand paging) algorithm that maximizes the number of steps i until LFD and OPT make a different decision, i.e., OPT and LFD make the same decisions on $\sigma_{\leq i-1}$ but evict a different page for request r_i . We construct an optimum solution OPT' [σ] that behaves like LFD on $\sigma_{\leq i}$, which then contradicts our choice of OPT and therefore completes the proof. Note that the solution OPT' [σ] that we construct is only demand paging up to step i , but can easily be modified as described in Remark 2.1 to be demand paging for all steps, without changing its behavior in the first i steps.

Let p be the page evicted by LFD in step i and $p^* \neq p$ be the page evicted by OPT in the same step. We define $j^* = \max\{j \in \{i, \dots, n\} : p^* \notin \{r_i, r_{i+1}, \dots, r_j\}\}$ to be the last consecutive step after step i , where p^* is not requested. By definition of LFD, we know that p is not requested in steps $\{i, \dots, j^*\}$. We distinguish three cases (see Figure 2.2):

Case 1: p is evicted by OPT in some step $j \in \{i + 1, \dots, j^*\}$.

Since OPT is demand paging and since p, p^* are not accessed in steps $\{i + 1, \dots, j\}$, OPT has neither p nor p^* in its cache at the end of step j . Hence we get a feasible solution OPT' [σ] from OPT [σ] by evicting p instead of p^* in step i , evicting p^* instead of p in step j and otherwise copying OPT's decisions. Since we did not cause additional page faults, we have $\text{OPT}'(\sigma) \leq \text{OPT}(\sigma)$, hence OPT' [σ] is still optimal.

Case 2: p is not evicted by OPT in steps i, \dots, j^* and $j^* = n$.

Since p, p^* are not accessed anymore after step i , it makes no difference whether we evict p or p^* in step i , and we can simply obtain OPT' [σ] from OPT [σ] by evicting p instead of p^* in step i .

Case 3: p is not evicted by OPT in steps i, \dots, j^* and $j^* < n$.

Then, in step $j^* + 1$, page p^* is requested and, because OPT is demand paging, page p is in OPT's cache but p^* is not. Let p' be the page evicted by OPT in step $j^* + 1$ to make room for p^* . If $p' = p$, then, after step $j^* + 1$, OPT has p^* in its cache, but not p . If we evict page p instead of p^* in step i and otherwise copy OPT's behavior, we have the same cache contents after step $j^* + 1$ and one less page fault, which contradicts optimality of OPT. If $p' \neq p$, then, after step $j^* + 1$, OPT has both p^* and p in its cache. We construct the solution OPT' [σ] by evicting p instead of p^* in step i and loading p instead of p^* in step $j^* + 1$. This solution is feasible (since it has the same cache contents after step $j^* + 1$) and does not cause additional page faults, i.e., $\text{OPT}'(\sigma) = \text{OPT}(\sigma)$. Note that OPT' [σ] is not demand paging, but can be modified to a demand paging solution without affecting its behavior in the first i steps, and without affecting its cost (as described in Remark 2.1). □

2.2 Lower bound for online paging

To derive a lower bound on the competitive ratio of every algorithm, we need to construct an adversarial request sequence $\sigma \in P^n$ for any given online algorithm ALG. A natural candidate is the request sequence

that, in every step, requests an item that is not in ALG's cache. For such a sequence, we trivially have $\text{ALG}(\sigma) = n$. To get a lower bound on the competitive ratio of ALG, we need a sufficiently strong upper bound on the cost of the optimum solution. By Theorem 2.2, we can get such a bound by analyzing the performance of LFD. Observe that our adversarial request sequence works even if $|P| = k + 1$.

Lemma 2.3. The cost of LFD for instances $\sigma \in P^n$ with $|P| = k + 1$ is bounded by

$$\text{LFD}(\sigma) \leq \lceil n/k \rceil.$$

Proof. Assume that LFD has a page fault upon request $r_i \in \sigma$ and evicts page $p \in P$. Since $|P| = k + 1$ and the cache starts with k elements, the next page fault occurs when p is requested the next time. By definition of LFD, all other $k - 1$ pages except r_i in the cache are requested before this happens. Hence, LFD has a page fault only at most once in every k steps. \square

If we choose the length n of our adversarial sequence σ to be a multiple of k , we immediately get $\text{ALG}(\sigma) = n \geq k \cdot \text{LFD}(\sigma) = k \cdot \text{OPT}(\sigma)$. Since our construction works for any online algorithm, this bound holds in general. For $n \rightarrow \infty$, we can use Proposition 1.12 to obtain the following bound.

Theorem 2.4 (Sleator and Tarjan [1985]). The competitive ratio of any deterministic paging algorithm is at least k , even when $|P| = k + 1$.

Proof. Consider any online (demand) paging algorithm ALG for instances with $|P| = k + 1$. Let $\sigma^{(i)} \in P^n$ be the request sequence of length $n_i = ik$ that always requests the unique page not in ALG's cache. Then, $\lim_{i \rightarrow \infty} \text{ALG}(\sigma^{(i)}) = \infty$. By Lemma 2.3 and Theorem 2.2, we have $\text{ALG}(\sigma^{(i)}) = n_i \geq k \cdot \text{LFD}(\sigma^{(i)}) = k \cdot \text{OPT}(\sigma^{(i)})$. It follows that $\limsup_{i \rightarrow \infty} \frac{\text{ALG}(\sigma^{(i)})}{\text{OPT}(\sigma^{(i)})} \geq k$. By Proposition 1.12, it follows that ALG has competitive ratio at least k . \square

2.3 Online paging algorithms

Now that we have established an understanding of the offline optimal algorithm LFD and have a bound on the best possible competitive ratio, we shift our attention to finding good online algorithms. Since we may restrict ourselves to demand paging algorithms, possible candidates only need to differ in the decision which page to evict upon a page fault. We list a few canonical heuristics that respond differently to page faults:

Least-Recently-Used (LRU): evict the page that was least recently accessed

First-In-First-Out (FIFO): evict the page that has been in the cache the longest since it was loaded (or any page in the cache initially)

Last-In-First-Out (LIFO): evict the page that was most recently loaded into the cache (or any page if none was loaded yet)

Least-Frequently-Used (LFU): evict any page that has been least used overall

Flush-When-Full (FWF): at the start and upon a page fault with a full cache, evict all pages (this is not demand paging)

We start by analyzing FIFO.

Theorem 2.5 (Sleator and Tarjan [1985]). FIFO is strictly k -competitive for the paging problem.

Proof. Fix a request sequence $\sigma = (r_1, \dots, r_n) \in P^n$ for which FIFO has a page fault for some request $r_i = p$. By definition, before FIFO can have another page fault caused by another access $r_j = p$ with $j > i$, all k pages that were in the cache at the start of step i must have been evicted as well as p . This means that there must be $k + 2$ total page faults in steps i through j , i.e., $j \geq i + k + 1$. Overall, we get that any set of $k + 1$ consecutive page faults must be caused by distinct elements.

We partition the request sequence σ into phases π_1, \dots, π_N , such that $\sigma = \pi_1 \oplus \dots \oplus \pi_N$, FIFO incurs at most k page faults in each phase, and $(|\pi_N|, \dots, |\pi_2|, |\pi_1|)$ is lexicographically maximized, where $|\pi_j|$ denotes the length of phase π_j (see Figure 2.3). Note that, by definition, FIFO incurs exactly k page faults in all phases except π_1 , and at most k page faults in phase π_1 . We get $\text{FIFO}(\sigma) \leq Nk$.




Figure 2.3: Illustration of the partitioning of the request sequence σ in the proof of Theorem 2.5.

Consider any phase $\pi_j \neq \pi_1$, let z_1, \dots, z_k be the pages that caused k faults during the phase, and let z_0 be the last page that was requested before phase π_j (i.e., the last page requested in phase π_{j-1}). Recall that σ was chosen such that FIFO has at least one page fault. Thus, by definition, phase π_1 contains at least one page fault, since otherwise $(\pi_1 \oplus \pi_2), \pi_3, \dots, \pi_N$ would be a lexicographically larger phase partitioning (with respect to $(|\pi_N|, \dots, |\pi_3|, |\pi_1 \oplus \pi_2|)$). Similarly, the access to z_0 just before phase π_j must have caused a page fault, otherwise z_0 could be included in phase π_j instead. Since page faults can only repeat after $k + 1$ distinct requests, we know that z_0, \dots, z_k must all be distinct. Since z_0 must be in OPT's cache after it was accessed, z_0 is in OPT's cache at the end of phase π_{j-1} . Since k pages other than z_0 are accessed during phase π_j , OPT has at least one page fault during π_j . Since phase π_1 causes at least one page fault for FIFO, there must be at least one page in π_1 that was not in the initial cache C_0 . But this means that OPT also has at least one page fault in phase π_1 . We get $\text{OPT}(\sigma) \geq N$, and thus

$$\text{FIFO}(\sigma) \leq k \cdot \text{OPT}(\sigma). \quad \square$$

Corollary 2.6. The paging problem with cache size k has (strict) competitive ratio k .

2.4 Phase partitioning

Motivated by the proof of Theorem 2.5, we want to establish a proof scheme that can be applied more generally. The main idea behind the proof of Theorem 2.5 was to partition the request sequence σ into phases, with the property that each phase causes at least one page fault for OPT. The partitioning of Theorem 2.5 depended on FIFO's behavior, but we can define an even simpler partitioning with the same property and independent of a specific algorithm. As before, $|\pi_j|$ denotes the length of phase π_j .

Definition 2.7. The k -phase partitioning of a request sequence σ is the partition $\sigma = \pi_1 \oplus \pi_2 \oplus \dots \oplus \pi_N$ for which each phase π_j contains at most k distinct requests and $(|\pi_1|, |\pi_2|, \dots, |\pi_N|)$ is maximized lexicographically.

For example, consider $\sigma = (1, 2, 4, 2, 1, 3, 5, 2, 5, 5, 3, 3, 3, 1, 2, 3, 4)$. Then the 3-phase partitioning is $\pi_1 = (1, 2, 4, 2, 1)$, $\pi_2 = (3, 5, 2, 5, 5, 3, 3, 3)$, $\pi_3 = (1, 2, 3)$, $\pi_4 = (4)$. Note that the partitioning is indeed independent of any algorithm.

In order to prove that a given algorithm ALG is k -competitive as in Theorem 2.5, we need that ALG does not have more than k page faults per phase. This is obviously the case if for every phase π_j it is true that no pages $p \in \pi_j$ are evicted during phase π_j . The following definition gives a slightly weaker sufficient condition in terms of a marking scheme.

Definition 2.8. A *marking algorithm* is an algorithm that never evicts a marked page, where a page is considered marked if it has previously been requested during the current phase of the k -phase partitioning.

Note that FIFO itself no longer falls under the simplified definition of marking algorithms. To see this, consider the partitioned request sequence $\sigma = \pi_1 \oplus \pi_2 = (1, 2, 3) \oplus (4, 2, 1)$ for $k = 3$ and $C_0 = \{4, 5, 6\}$. Upon the last access to 1, the marked page 2 will be evicted by FIFO. On the other hand, FWF obviously is a marking algorithm, since it only evicts pages once a new phase starts.

Proposition 2.9. FIFO is not a marking algorithm.

We can now carry over the idea from the proof of Theorem 2.5 to marking algorithms.

Theorem 2.10 (Torng [1998]). Every marking algorithm is k -competitive for the paging problem with (constant) cache size k .

Proof. Let ALG be any marking algorithm and consider any input $\sigma \in P^n$ with k -phase partitioning $\pi_1 \oplus \dots \oplus \pi_N$. Since ALG is a marking algorithm, it has at most one page fault per distinct element requested in every phase. By definition, every phase contains at most k distinct requests, hence $\text{ALG}(\sigma) \leq Nk$.




Figure 2.4: Illustration of the partitioning of the request sequence σ in the proof of Theorem 2.10.

On the other hand, consider any phase $\pi_j \neq \pi_N$ and let z_j, z_{j+1} be the first requests of phases π_j and π_{j+1} , respectively (see Figure 2.4). Since π is lexicographically maximal, π_j contains exactly k distinct requests and $z_{j+1} \notin \pi_j$. Let π'_j be the shifted phase starting with the second request of π_j and ending with the first request of π_{j+1} , i.e., $(z_j) \oplus \pi'_j := \pi_j \oplus (z_{j+1})$. Then π'_j contains exactly k requests different from z_j . Since OPT must have z_j in its cache at the start of π'_j , it has at least one cache fault during π'_j . Since the positions of the shifted phases for all phases $\pi_j \neq \pi_N$ are mutually disjoint, we have $\text{OPT}(\sigma) \geq N - 1$, and, hence, $\text{ALG}(\sigma) \leq Nk \leq k \cdot \text{OPT}(\sigma) + k$. \square

We immediately obtain an upper bound on the competitive ratio of FWF.

Corollary 2.11 (Karlin et al. [1988]). FWF is k -competitive for the paging problem with cache size k .

Note that FWF is not strictly k -competitive, since it has positive cost even if all requested pages are already in the initial cache C_0 . In particular, Theorem 2.10 does not hold for strict competitiveness.



3 Randomized Algorithms

In Chapter 2 we saw that no deterministic online algorithm can have a better competitive ratio for the paging problem than the cache size k . This lower bound is easy for the adversary to achieve simply by always requesting the exact page the algorithm just evicted. Of course this adversarial strategy relies on perfect knowledge of the behavior of the algorithm in every step. Consequently, there is hope that we can better defend against the adversary if we are less predictable, i.e., if we include random choices into the behavior of our algorithm. Formally, we can express this by defining randomized online algorithms that, for every request sequence, randomly choose an algorithm from the set of all deterministic algorithms \mathcal{A}_{det} according to some distribution, and then deterministically execute the chosen algorithm. Note that, on any finite set of instances, every online algorithm that involves random choices can be expressed in this way (see Isbell [1957]).

Definition 3.1. A (distributional) randomized online algorithm ALG_p is identified with a probability distribution $p: \mathcal{A}_{\text{det}} \rightarrow [0, 1]$ with $\sum_{\text{ALG} \in \mathcal{A}_{\text{det}}} p(\text{ALG}) = 1$.

In order to define the competitive ratio of a randomized online algorithm ALG_p , we need to specify whether or not the adversary may use any knowledge about the outcome of the random choice that ALG_p makes at the start of each run to come up with an adversarial input sequence. Of course, if we allow the adversary to be adaptive to the random choice of ALG_p , then the same lower bound construction for paging still works. In this lecture, we will restrict ourselves to an *oblivious adversary* that must specify a worst-case instance without any knowledge regarding ALG_p 's random choice.

Definition 3.2. A (distributional) randomized online algorithm ALG_p is ρ -competitive if there is a constant $\alpha \geq 0$, such that, for all $\sigma \in \Sigma$,

$$\mathbb{E}[\text{ALG}_p(\sigma)] := \mathbb{E}_{\text{ALG} \sim p}[\text{ALG}(\sigma)] \leq \rho \cdot \text{OPT}(\sigma) + \alpha.$$

ALG_p is *strictly* ρ -competitive, if the above holds for $\alpha = 0$.

Definition 3.3. The (strict) randomized competitive ratios of an algorithm and a problem are defined analogously to the deterministic case (Definitions 1.10 and 1.11).

3.1 Randomized paging algorithms

We can easily define a “randomized algorithm” RAND that prevents the adversarial strategy of requesting the previously evicted page by evicting a page uniformly at random upon each page fault.¹ However, this algorithm does not beat the deterministic bound on the competitive ratio.

Proposition 3.4. The competitive ratio of RAND is at least k , even when $|P| = k + 1$.

¹ Note that RAND is not formally a distributional randomized algorithm, and we take the expectation over the outcome of its random choices in Definition 3.2.

Proof. We use $P = \{p_0, p_1, \dots, p_k\}$, $C_0 = \{p_1, \dots, p_k\}$, and we define the infinite sequence $\sigma = (p_0, p_1, \dots, p_{k-1}, p_0, \dots)$. We consider the request sequence $\sigma_{\leq n}$ for $n \in \mathbb{N}$. For any such sequence, OPT evicts p_k in the first step and $\text{OPT}(\sigma_{\leq n}) = 1$. On the other hand, RAND continues to have page faults until it evicts p_k , assuming n is large enough. Since RAND evicts pages uniformly at random, its probability to evict p_k upon a page fault is $1/k$. Hence, the expected number of page faults until p_k is evicted approaches k as n goes to infinity. We have $\lim_{n \rightarrow \infty} \mathbb{E}[\text{RAND}(\sigma_{\leq n})] = k \cdot \text{OPT}(\sigma_{\leq n})$.

Since $\mathbb{E}[\text{RAND}(\sigma_{\leq n})]$ is bounded independently of n , the above only proves that RAND does not have a strict competitive ratio smaller than k (as in Proposition 1.12). Let $\sigma^{(j)} := (p_j, p_{j+1}, \dots, p_{j-2}, p_j, \dots)$ with $j \in \{0, \dots, k\}$ and all indices modulo $k+1$ from now on. Then $\sigma^{(0)}$ is the sequence we defined above (since $(-2 \bmod (k+1)) \equiv k-1$) and in $\sigma^{(j)}$ the role of p_k in $\sigma^{(0)}$ is played by p_{j-1} . We can make the cost of RAND (and OPT) grow unboundedly, by defining the request sequence $\sigma^{(n,m)} := \sigma_{\leq n}^{(0)} \oplus \sigma_{\leq n}^{(k)} \oplus \sigma_{\leq n}^{(k-1)} \oplus \dots \in P^{mn}$ consisting of $m \in \mathbb{N}$ phases of the form $\sigma_{\leq n}^{(j)}$ for $j \in \{0, \dots, k\}$. As above, OPT has at most one page fault per phase, and hence $\text{OPT}(\sigma^{(n,m)}) \leq m$. If we let n go to infinity for a fixed value of m , then the probability that RAND does not have page p_{j-1} in its cache after a phase of the form $\sigma_{\leq n}^{(j)}$ approaches one. Hence, the number of page faults of RAND approaches k per phase, as above.

Now if RAND is ρ -competitive, then there is a constant $\alpha \geq 0$ such that, for all $n, m \in \mathbb{N}$, we have $\mathbb{E}[\text{RAND}(\sigma^{(n,m)})] \leq \rho \cdot \text{OPT}(\sigma^{(n,m)}) + \alpha$. It follows that (we let n grow much faster than m)

$$\rho \geq \lim_{m \rightarrow \infty} \lim_{n \rightarrow \infty} \frac{\mathbb{E}[\text{RAND}(\sigma^{(n,m)})] - \alpha}{\text{OPT}(\sigma^{(n,m)})} \geq \lim_{m \rightarrow \infty} \frac{km - \alpha}{m} = k. \quad \square$$

We can summarize the approach used in the proof of Proposition 3.4 to obtain a lower bound on the randomized competitive ratio analogously to Proposition 1.12. Note that it would be sufficient to require $\lim_{i \rightarrow \infty} \mathbb{E}[\text{ALG}_p(\sigma^{(i)})] = \infty$ below.

Proposition 3.5. Let $(\sigma^{(i)} \in \Sigma)_{i \in \mathbb{N}}$ be a sequence of request sequences with

$$\lim_{i \rightarrow \infty} \text{OPT}(\sigma^{(i)}) = \infty.$$

Then, every randomized online algorithm ALG_p has a competitive ratio of at least

$$\limsup_{i \rightarrow \infty} \frac{\mathbb{E}[\text{ALG}_p(\sigma^{(i)})]}{\text{OPT}(\sigma^{(i)})}.$$

Observe that the k -phase partitioning of the request sequence σ defined in the beginning of the proof of Proposition 3.4 consists of a single phase. The problem with RAND on this instance is that it evicts pages that were previously requested during the phase, i.e., its problem is that it is not a marking algorithm. We turn RAND into a randomized marking algorithm RMARK , by evicting an *unmarked* page uniformly at random upon each page fault. It turns out that this modification suffices to perform much better than the best deterministic online algorithm.²

Theorem 3.6 (Fiat et al. [1991]). RMARK is strictly $2\mathcal{H}_k$ -competitive, where $\mathcal{H}_k := \sum_{i=1}^k \frac{1}{i} \approx \ln k$ is the k -th harmonic number.

Proof. Fix a request sequence $\sigma \in P^n$ and let its k -phase partitioning be $\pi_1 \oplus \pi_2 \oplus \dots \oplus \pi_N = \sigma$. Consider any phase π_j , call the pages in the cache at the start of the phase *old* and all other pages *new*. Let m_j

² Note that RMARK is not formally a distributional randomized algorithm, and we take the expectation over the outcome of its random choices in Definition 3.2.

denote the number of (distinct) new pages requested in phase π_j . Obviously, each new page causes a page fault for R_{MARK}. We compute the expected number of page faults in phase π_j caused by old pages. Let $r \in \pi_j$ be the ℓ -th (distinct) old page that is requested during phase π_j . At the point where r is requested for the first time during π_j , there are at least $k - m_j$ old pages in the cache, and exactly $\ell - 1$ of them are marked, hence there are at least $k - m_j - (\ell - 1)$ unmarked old pages in the cache. Also, the total number of unmarked old pages, both in and not in the cache, is exactly $k - (\ell - 1)$. Overall, by symmetry, the probability that page r causes a page fault is at most

$$1 - \frac{k - m_j - (\ell - 1)}{k - (\ell - 1)} = \frac{m_j}{k - \ell + 1}.$$

Since at most $k - m_j$ old pages are accessed in phase π_j (fewer than $k - m_j$ only in the last phase), the expected number of page faults of R_{MARK} in phase π_j is at most

$$\begin{aligned} m_j + \sum_{\ell=1}^{k-m_j} \frac{m_j}{k-\ell+1} &= m_j \cdot \left(1 + \frac{1}{k} + \frac{1}{k-1} + \cdots + \frac{1}{m_j+1} \right) \\ &= m_j \cdot (1 + \mathcal{H}_k - \mathcal{H}_{m_j}) \\ &\leq m_j \mathcal{H}_k. \end{aligned}$$

By linearity of expectation, the total expected cost of R_{MARK} becomes

$$\mathbb{E}[\text{R}_{\text{MARK}}(\sigma)] \leq \sum_{j=1}^N m_j \mathcal{H}_k.$$




Figure 3.1: Illustration of the complementary accountings of page faults of OPT over consecutive phases in the proof of Theorem 3.6. The final bound is obtained by averaging over both accountings.

It remains to show a lower bound on $\text{OPT}(\sigma)$. By definition, R_{MARK} has all pages requested in a phase in the cache at the end of the phase. Since exactly k distinct pages are requested in every phase except the last, we have that the number of distinct pages accessed in the two consecutive phases π_{j-1}, π_j , $j > 1$, is exactly $k + m_j$. It follows that OPT has at least m_j page faults in phases π_{j-1} and π_j combined. Also OPT has at least m_1 page faults in the first phase. We get a bound for OPT by separately considering even and odd values of j (see Figure 3.1):

$$\text{OPT}(\sigma) \geq \max \left\{ \sum_{\ell=1}^{\lfloor N/2 \rfloor} m_{2\ell}, \sum_{\ell=0}^{\lfloor (N-1)/2 \rfloor} m_{2\ell+1} \right\} \geq \frac{1}{2} \sum_{\ell=1}^{\lfloor N/2 \rfloor} m_{2\ell} + \frac{1}{2} \sum_{\ell=0}^{\lfloor (N-1)/2 \rfloor} m_{2\ell+1} = \frac{1}{2} \sum_{j=1}^N m_j.$$

Finally, we conclude that

$$\mathbb{E}[\text{R}_{\text{MARK}}(\sigma)] \leq \mathcal{H}_k \sum_{j=1}^N m_j \leq 2\mathcal{H}_k \cdot \text{OPT}(\sigma). \quad \square$$

3.2 Yao's principle

So far, we derived lower bounds for the competitive ratios of online algorithms by taking the perspective of an adversary and constructing a request sequence tailored to a given algorithm, in such a way that we respond to every possible action of the algorithm in the worst possible manner. This is no longer possible for randomized algorithms, since we do not know the outcomes of the algorithm's random choices beforehand. We now have to construct a request sequence for which the algorithm performs badly in expectation over all possible combinations of random choices the algorithm makes, as stated by Proposition 3.5. In general, this is a daunting task.

Yao's principle provides a crucial tool to prove lower bounds by allowing us to restrict ourselves to deterministic algorithms. In a way, Yao's principle allows to trade the randomness in the algorithm for randomness in the input sequence. Instead of designing a single sequence of request sequences with high expected cost for *any fixed randomized* algorithm, we can design a sequence of *distributions* of input sequences with high expected cost for *every deterministic* algorithm. This turns out to be much more manageable in many cases.

Theorem 3.7 (Yao [1977]). Let $(q_i : \Sigma \rightarrow [0, 1])_{i \in \mathbb{N}}$ be a sequence of probability distributions with

$$\lim_{i \rightarrow \infty} \mathbb{E}_{\sigma \sim q_i}[\text{OPT}(\sigma)] = \infty.$$

Then, every randomized online algorithm has a competitive ratio of at least

$$\limsup_{i \rightarrow \infty} \frac{\inf_{\text{ALG} \in \mathcal{A}_{\text{det}}} \mathbb{E}_{\sigma \sim q_i}[\text{ALG}(\sigma)]}{\mathbb{E}_{\sigma \sim q_i}[\text{OPT}(\sigma)]}.$$

Proof. Let ALG_p be the randomized algorithm associated with the probability distribution $p : \mathcal{A}_{\text{det}} \rightarrow [0, 1]$. Assume that ALG_p has competitive ratio ρ , i.e., there is a constant $\alpha \geq 0$ such that, for all $\sigma \in \Sigma$,

$$\mathbb{E}_{\text{ALG} \sim p}[\text{ALG}(\sigma)] \leq \rho \cdot \text{OPT}(\sigma) + \alpha.$$

In particular, for every distribution q_i , $i \in \mathbb{N}$, of the given sequence, we obtain

$$\begin{aligned} \mathbb{E}_{\sigma \sim q_i}[\mathbb{E}_{\text{ALG} \sim p}[\text{ALG}(\sigma)]] &= \sum_{\sigma \in \Sigma} q_i(\sigma) \cdot \mathbb{E}_{\text{ALG} \sim p}[\text{ALG}(\sigma)] \\ &\leq \rho \cdot \sum_{\sigma \in \Sigma} q_i(\sigma) \cdot \text{OPT}(\sigma) + \alpha \cdot \sum_{\sigma \in \Sigma} q_i(\sigma) \\ &= \rho \cdot \mathbb{E}_{\sigma \sim q_i}[\text{OPT}(\sigma)] + \alpha. \end{aligned}$$

Since σ and ALG are drawn independently and since $\text{ALG}(\sigma) \geq 0$, the corresponding expectations commute by Tonelli's theorem. We get

$$\begin{aligned} \rho \cdot \mathbb{E}_{\sigma \sim q_i}[\text{OPT}(\sigma)] + \alpha &\geq \mathbb{E}_{\sigma \sim q_i}[\mathbb{E}_{\text{ALG} \sim p}[\text{ALG}(\sigma)]] \\ &= \mathbb{E}_{\text{ALG} \sim p}[\mathbb{E}_{\sigma \sim q_i}[\text{ALG}(\sigma)]] \\ &\geq \inf_{\text{ALG} \in \mathcal{A}_{\text{det}}} \mathbb{E}_{\sigma \sim q_i}[\text{ALG}(\sigma)], \end{aligned}$$

for every $i \in \mathbb{N}$. Solving for ρ and taking \limsup gives

$$\begin{aligned} \rho &\geq \limsup_{i \rightarrow \infty} \left[\frac{\inf_{\text{ALG} \in \mathcal{A}_{\text{det}}} \mathbb{E}_{\sigma \sim q_i} [\text{ALG}(\sigma)]}{\mathbb{E}_{\sigma \sim q_i} [\text{OPT}(\sigma)]} - \frac{\alpha}{\mathbb{E}_{\sigma \sim q_i} [\text{OPT}(\sigma)]} \right] \\ &\geq \limsup_{i \rightarrow \infty} \frac{\inf_{\text{ALG} \in \mathcal{A}_{\text{det}}} \mathbb{E}_{\sigma \sim q_i} [\text{ALG}(\sigma)]}{\mathbb{E}_{\sigma \sim q_i} [\text{OPT}(\sigma)]} - \limsup_{i \rightarrow \infty} \frac{\alpha}{\mathbb{E}_{\sigma \sim q_i} [\text{OPT}(\sigma)]} \\ &= \limsup_{i \rightarrow \infty} \frac{\inf_{\text{ALG} \in \mathcal{A}_{\text{det}}} \mathbb{E}_{\sigma \sim q_i} [\text{ALG}(\sigma)]}{\mathbb{E}_{\sigma \sim q_i} [\text{OPT}(\sigma)]}, \end{aligned}$$

where we used $\lim_{i \rightarrow \infty} \mathbb{E}_{\sigma \sim q_i} [\text{OPT}(\sigma)] = \infty$. \square

If we only need a bound for the *strict* competitive ratio, we can use a simplified version of Yao's principle (proof: exercise).

Corollary 3.8. Let $q: \Sigma \rightarrow [0, 1]$ be any probability distribution. Then, every randomized online algorithm has a strict competitive ratio of at least

$$\frac{\inf_{\text{ALG} \in \mathcal{A}_{\text{det}}} \mathbb{E}_{\sigma \sim q} [\text{ALG}(\sigma)]}{\mathbb{E}_{\sigma \sim q} [\text{OPT}(\sigma)]}.$$

Remark 3.9. Note that Yao's principle is tight in the sense that we can always find a distribution to prove any valid lower bound on the competitive ratio.

3.3 Lower bound for randomized paging

Equipped with Yao's principle we are now able to derive an (almost) tight lower bound on the randomized competitive ratio of the paging problem.

Theorem 3.10 (Fiat et al. [1991]). The competitive ratio of any randomized paging algorithm is at least \mathcal{H}_k , even when $|P| = k + 1$.

Proof. We fix $P = \{p_0, \dots, p_k\}$ and the initial cache $C_0 = \{p_1, \dots, p_k\}$. We want to apply Yao's principle (Theorem 3.7), i.e., we need to provide a sequence $(q_n)_{n \in \mathbb{N}}$ of probability distributions over input instances. We describe q_n by a randomized strategy to construct a request sequence $\sigma = (r_1, \dots, r_{n'})$ of length $n' \geq n$: We set $r_1 = p_0$ and iteratively choose $r_i, i = 2, 3, \dots$ uniformly at random from $P \setminus \{r_{i-1}\}$, until $i > n$ and the last phase of the k -phase partitioning of (r_1, \dots, r_i) consists of exactly one request. At this point, we terminate with $\sigma = (r_1, \dots, r_{i-1})$. We let $q_n(\sigma)$ be the probability that sequence σ is produced by this procedure.




Figure 3.2: Illustration for the proof of Theorem 3.10. OPT incurs a page fault during each shifted phase π'_j .

Consider the k -phase partitioning $\sigma = \pi_1 \oplus \pi_2 \oplus \dots \oplus \pi_N$ of any sequence $\sigma = (r_1, \dots, r_{n'}) \in P^{n'}$ with $q_n(\sigma) > 0$. Let z_j denote the first request during phase π_j , let $\pi'_1 = (z_1)$ and let $(\pi'_j)_{j \in \{2, \dots, N\}}$ be defined via $(z_{j-1}) \oplus \pi'_j = \pi_{j-1} \oplus (z_j)$ (see Figure 3.2). Since $r_1 \notin C_0$ and since, for $j \geq 2$, phase π'_j contains all

k pages other than z_{j-1} , which must be in the cache at the start of the phase π'_j , every offline algorithm must have a page fault during π'_j for every $j \in \{1, \dots, N\}$. On the other hand, there is exactly one page that is not requested in each phase, therefore the offline algorithm that evicts this page upon the first page fault in a phase never has more than one page fault per phase π_j . Hence, OPT has exactly N page faults. Let X be the random variable that measures the length of a phase π_j . Note that, by construction, $\mathbb{E}_{\sigma \sim q_n}[X]$ is independent of j . **We claim that** $\mathbb{E}_{\sigma \sim q_n}[X] = k\mathcal{H}_k$. With this, and since OPT has one fault per phase, we immediately get

$$\lim_{n \rightarrow \infty} \mathbb{E}_{\sigma \sim q_n}[\text{OPT}(\sigma)] = \lim_{n \rightarrow \infty} \mathbb{E}_{\sigma \sim q_n}[N] = \infty,$$

as required by Theorem 3.7.

Now, consider the behavior of *any* deterministic online (demand paging) algorithm ALG for request sequence σ . At the beginning of step $i \geq 2$, the cache of ALG contains all but one page $r \in P \setminus C$ and the probability that r is the next request is exactly $1/k$ by construction of q_n , because $r_i \in P \setminus \{r_{i-1}\}$ and $r_{i-1} \in C$. The first request r_1 always causes a page fault. Hence, the expected number of page faults of ALG per phase is at least

$$\frac{1}{k} \sum_{l=1}^{\infty} \Pr[X \geq l] = \mathbb{E}_{\sigma \sim q_n}[X]/k.$$

Since this holds for every choice of ALG , and since OPT never has more than one fault per phase, we get

$$\limsup_{n \rightarrow \infty} \frac{\inf_{\text{ALG} \in \mathcal{A}_{\text{det}}} \mathbb{E}_{\sigma \sim q_n}[\text{ALG}(\sigma)]}{\mathbb{E}_{\sigma \sim q_n}[\text{OPT}(\sigma)]} \geq \frac{\mathbb{E}[X]}{k} = \mathcal{H}_k,$$

and we can apply Theorem 3.7 to complete the proof.

It remains to prove our claim that $\mathbb{E}_{\sigma \sim q_n}[X] = k\mathcal{H}_k$. Note that for every request after the first, we draw uniformly at random from k possibilities, and the next phase starts once we have drawn all pages. To compute the expected length of a phase, we need to solve the *coupon collector problem*: “If we draw coupons from an urn with replacement, how many do we have to draw to get each at least once?” Let X_ℓ be the random variable that counts the number of requests during a phase until a new page is requested, once $\ell - 1$ different pages have already been requested. Then $X = \sum_{\ell=1}^{k+1} X_\ell - 1$. By definition, $X_1 = 1$. Consider X_ℓ with $j \geq 2$. If $\ell - 1$ pages have been requested before some request r_i , the probability that r_i is again one of these pages is $(\ell - 2)/k$, since r_i is chosen uniformly at random from $P \setminus \{r_{i-1}\}$. Therefore, X_ℓ is a geometric random variable with success probability $s_\ell = 1 - \frac{\ell-2}{k}$, and

$$\mathbb{E}_{\sigma \sim q_n}[X_{\ell \geq 2}] = \frac{1}{s_\ell} = \frac{k}{k - \ell + 2}.$$

By linearity of expectation, we get

$$\begin{aligned} \mathbb{E}_{\sigma \sim q_n}[X] &= \sum_{\ell=1}^{k+1} \mathbb{E}_{\sigma \sim q_n}[X_\ell] - 1 \\ &= \sum_{\ell=2}^{k+1} \frac{k}{k - \ell + 2} \\ &= k \sum_{\ell=1}^k \frac{1}{\ell} = k\mathcal{H}_k, \end{aligned}$$

which completes the proof. □

Remark 3.11. An \mathcal{H}_k -competitive randomized algorithm was given by McGeoch and Sleator [1991].

4 The List Update Problem

In the previous chapter, we considered the abstract problem of dynamically reorganizing a basic random access data structure that allows fast access to an arbitrary subset of its data. We now introduce a similar problem with an unsorted linked list as our underlying data structure. Essentially, a linked list is managed by keeping a reference to the first element of the list and equipping each element with a reference to its successor in the list. To access an element at position i of a linked list, we need to traverse the first $i - 1$ elements, i.e., accessing items further down the list is more expensive than accessing elements near the front. After extracting an element it may therefore make sense to return it in a position closer to the front in order to speed up further accesses to the same element. To simplify our analysis we also allow to invest extra time to switch the order of any consecutive items. In the standard form, the list update problem is given as follows.¹

LIST UPDATE PROBLEM	
given:	set of (different) items X , $l := X $, stored in a list L_0
online:	sequence $\sigma = (r_1, \dots, r_n) \in X^n$ of items accessed in this order
actions (step i):	transpose any number of consecutive items (paid) move r_i closer to the front of the list (free)
objective:	minimize the total cost $\sum_{i=1}^n \text{ALG}_i(\sigma) = \sum_{i=1}^n (s_i + p_i)$ s_i : number of items in front of r_i at the beginning of step i p_i : number of paid actions in step i

On first glance, it seems hopeless to design a good online algorithm for this problem: How can we know whether an element will be accessed often in the future and should be moved closer to the front? This is where the distinction between measuring *expected* or measuring *worst-case* performance makes all the difference. Of course, if we try to optimize expected cost with respect to a uniform distribution of accesses to elements, there is nothing to be done, since every strategy (that does not needlessly transpose items) has the same quality in expectation. However, in competitive analysis we have to defend against the worst case! For example, if we choose not to move elements at all, the adversary has the obvious strategy of repeatedly accessing the last element of the list. The offline optimum for this instance is very cheap, because it immediately moves the last element to the front, while our algorithm needs to repeatedly traverse the entire list and therefore incurs high cost. An obvious way to prevent this worst-case is to use the *move-to-front* heuristic MTF : In every step, move the accessed element to the first position of the list. To analyze this simple algorithm, we first need to introduce a very important technique of competitive analysis: the *potential function method* (or *amortized analysis*).

4.1 Potential function method (amortized analysis)

The competitive analysis of an online algorithm ALG is easiest if we can bound the cost $\text{ALG}_i(\sigma)$ in each step individually by the corresponding cost $\text{OPT}_i(\sigma)$ of the offline optimum. Often however, it is impossible to guarantee that in *every* step ALG incurs a relatively small cost. In Section 2.4, we introduced the phase partitioning technique in order to bound the costs for different phases separately. But even if we are unable to bound ALG 's cost per step/phase, the algorithm may still be competitive if expensive steps only

¹ We use a cost model where accessing the element at position i incurs cost $i - 1$; this is often referred to as the *reduced cost model* in the literature.

	MTF	MTF _i	Φ _i	ΔΦ _i	a _i	OPT _i	OPT
L ₀ :	$x_5 x_4 x_3 x_2 x_1$		0				$x_5 x_4 x_3 x_2 x_1 : L_0^*$
L ₁ :	$x_1 x_5 x_4 x_3 x_2$	4	4	+4	8	4	$x_5 x_4 x_3 x_2 x_1 : L_1^*$
L ₂ :	$x_2 x_1 x_5 x_4 x_3$	4	6	+2	6	3	$x_5 x_4 x_3 x_2 x_1 : L_2^*$
L ₃ :	$x_3 x_2 x_1 x_5 x_4$	4	6	0	4	2	$x_5 x_4 x_3 x_2 x_1 : L_3^*$
L ₄ :	$x_4 x_3 x_2 x_1 x_5$	4	4	-2	2	1	$x_5 x_4 x_3 x_2 x_1 : L_4^*$
L ₅ :	$x_5 x_4 x_3 x_2 x_1$	4	0	-4	0	0	$x_5 x_4 x_3 x_2 x_1 : L_5^*$

Figure 4.1: A bad example for MTF for $\sigma = (x_1, x_2, x_3, x_4, x_5)$. While it is possible to force $\text{MTF}_i(\sigma) = n - 1 = 4$ and $\text{OPT}_i(\sigma) = 0$, this only happens after $n - 1 = 4$ cheaper steps. Note that after the expensive step both lists become more similar.

occur rarely and only if there are sufficiently many cheap steps to compensate. For example, MTF may have expensive steps in which elements are accessed that OPT has near the front of the list while they are near the back in ALG's list. However, for any element near the front to again reach the back of the list, many other elements need to be accessed by MTF, i.e., there must have been many cheaper steps (see Figure 4.1). We can try to account for this effect by comparing the total cost of ALG with the total cost of OPT, but this is often difficult, because we need to consider any possible sequence of requests.

The principal idea of amortized analysis is to charge cheap steps a little extra cost in order to pay for eventual expensive steps. To understand this, we first consider an intuitive example: Take a student Charlie that still has her parents paying for her meals at the university. Because Charlie does not want to waste her parents' money, she tries to spend as little as possible in the cafeteria. Say she spends only 2€ for lunch each day, but after every ten consecutive days of cafeteria she needs a break and goes to a nice restaurant that costs 13€. Sometime in the middle of the semester, Charlie's parents confront her about seeing her at the restaurant, which gives them the impression that Charlie is abusing their support. Now Charlie needs an easy argument to convince her parents that she is effectively only spending a reasonable amount of at most 3€ a day. Strictly speaking, of course, this is not the case. However, it is the case in an *amortized* sense: At every moment of the semester it is correct that her *average* spending is at most 3€ per lunch. To explain this easily, Charlie uses a potential function argument: She has a piggy bank at home (her *potential function* Φ), and every day she takes 3€ from her parents' money, puts 1€ into the bank, and uses the remaining 2€ to pay for lunch at the cafeteria. Every eleventh day, she takes the 10€ that are in the bank plus 3€ from her parents to pay for the restaurant. With this potential function argument, her parents are quickly convinced that Charlie effectively only spends 3€ per day (i.e., her *amortized cost* a_i is 3€), because she only takes 3€ from her parents money every day. We now generalize this argument formally.

Definition 4.1. A *potential function* is a function $\Phi: \mathbb{N} \times \Sigma \rightarrow \mathbb{R}$ that specifies a potential value $\Phi_i(\sigma) := \Phi(i, \sigma) = \Phi(i, \sigma_{\leq i})$ after every time step $i \in \{0, \dots, n\}$ and for every instance $\sigma \in \Sigma$, $|\sigma| \geq i$. The *amortized cost* in step $i \in \{1, \dots, n\}$ of an online algorithm ALG with respect to Φ and an instance $\sigma \in \mathcal{I}$ is

$$a_i(\sigma) := \text{ALG}_i(\sigma) + \Delta\Phi_i(\sigma) := \text{ALG}_i(\sigma) + \Phi_i(\sigma) - \Phi_{i-1}(\sigma).$$

In the example above, Charlie used the simple potential function $\Phi_i(\sigma) = i - \lfloor i/11 \rfloor \cdot 11$, yielding an amortized cost of $a_i(\sigma) = 2 + i - (i - 1) = 3$ on cheap days and $a_i(\sigma) = 13 + 0 - 10 = 3$ on expensive days.

With our intuition for the potential function Φ as a piggy bank and amortized cost as the money we need to invest in each step, the following sufficient condition for competitiveness should be natural. To be

slightly more general, we allow bounded negative potentials (i.e., debts of at most α) and a nonzero initial potential (i.e., a starting capital β).

Lemma 4.2 (potential function method). Let Φ be a potential function and let $\rho \geq 1$. If there are $\alpha \geq 0$ and $\beta \in \mathbb{R}$, such that $\Phi_i(\sigma) \geq -\alpha$, $\Phi_0(\sigma) \leq \beta$ and $a_i(\sigma) \leq \rho \cdot \text{OPT}_i(\sigma)$ for all $\sigma \in \Sigma$ and $i \leq |\sigma|$, then ALG is ρ -competitive. If $\alpha = \beta = 0$, then ALG is strictly ρ -competitive.

Proof. Fix any instance $\sigma = (r_1, \dots, r_n) \in \Sigma$. Using $\text{OPT}(\sigma) = \sum_{i=1}^n \text{OPT}_i(\sigma)$ and $\text{ALG}(\sigma) = \sum_{i=1}^n \text{ALG}_i(\sigma)$, we obtain

$$\begin{aligned}
 \text{ALG}(\sigma) &= \sum_{i=1}^n \text{ALG}_i(\sigma) \\
 &= \sum_{i=1}^n (a_i(\sigma) - \Phi_i(\sigma) + \Phi_{i-1}(\sigma)) \\
 &= \Phi_0(\sigma) - \Phi_n(\sigma) + \sum_{i=1}^n a_i(\sigma) \\
 &\leq \beta + \alpha + \rho \cdot \sum_{i=1}^n \text{OPT}_i(\sigma) \\
 &= \rho \cdot \text{OPT}(\sigma) + \alpha + \beta,
 \end{aligned}$$

i.e., ALG is ρ -competitive. If $\alpha = 0$ and $\beta = 0$, we get $\text{ALG}(\sigma) \leq \rho \cdot \text{OPT}(\sigma)$, which implies that ALG is strictly ρ -competitive. \square

Sometimes it is useful to separately consider the cost of OPT and of ALG. We can do this by rewriting the condition $a_i(\sigma) \leq \rho \cdot \text{OPT}_i(\sigma)$ of Lemma 4.2 as

$$\Delta\Phi_i(\sigma) = \Phi_i(\sigma) - \Phi_{i-1}(\sigma) \leq \rho \cdot \text{OPT}_i(\sigma) - \text{ALG}_i(\sigma). \quad (4.1)$$

This allows us to separate the contributions of OPT and ALG to the potential function. We can split each step i into a substep where ALG makes its decision and one (before or after) where OPT decides. The potential function may depend on both OPT's and ALG's state/behavior (as we will see), and we can separately regard these contributions. We define Φ'_i to be the intermediate potential function value, after the first and before the second substep of step i . We can then define the contributions $\Delta\Phi_i^{\text{ALG}}$ and $\Delta\Phi_i^{\text{OPT}}$ of ALG and OPT to $\Delta\Phi$, respectively. If ALG acts before OPT, we define $\Delta\Phi_i^{\text{ALG}}(\sigma) := \Phi'_i(\sigma) - \Phi_{i-1}(\sigma)$ and $\Delta\Phi_i^{\text{OPT}}(\sigma) := \Phi_i(\sigma) - \Phi'_i(\sigma)$, and otherwise we exchange these definitions. In either case, $\Delta\Phi_i(\sigma) = \Delta\Phi_i^{\text{ALG}}(\sigma) + \Delta\Phi_i^{\text{OPT}}(\sigma)$, and, using eq. (4.1), we can reformulate Lemma 4.2 as the following corollary. We call this interpretation the *alternating moves approach*.

Corollary 4.3 (potential function method, alternating moves). Let Φ be a potential function and let $\rho \geq 1$. If the following conditions hold for $\alpha \geq 0$ ($\alpha = 0$), $\beta \in \mathbb{R}$ ($\beta = 0$), and for all $\sigma \in \Sigma$ and $i \leq |\sigma|$, then ALG is (strictly) ρ -competitive:

- (i) $\Phi_i(\sigma) \geq -\alpha$,
- (ii) $\Phi_0(\sigma) \leq \beta$,
- (iii) $\Delta\Phi_i^{\text{OPT}}(\sigma) \leq \rho \cdot \text{OPT}_i(\sigma)$,
- (iv) $\Delta\Phi_i^{\text{ALG}}(\sigma) \leq -\text{ALG}_i(\sigma)$.

4.2 Upper bound for move-to-front

We are now ready to analyze M_{TF} . As discussed above (see Figure 4.1), M_{TF} has an expensive access whenever an element in the back of its list is accessed that O_{PT} has close to the front of its list, but this situation can only occur after many cheaper accesses. This characteristic suggests that we may be able to apply the potential function method.

Theorem 4.4 (Sleator and Tarjan [1985]). M_{TF} is strictly 2-competitive for the list update problem.

Proof. Fix a request sequence $\sigma \in \Sigma$ and let L_i, L_i^* denote the lists of ALG and O_{PT} , respectively, at the end of step i . Further let s_i, s_i^* be the number of elements in front of r_i in L_{i-1} and L_{i-1}^* , respectively, at the start of step i . Also, let p_i^* be the number of paid transpositions of O_{PT} in step i . We then have $O_{PT}_i(\sigma) = s_i^* + p_i^*$ and, since M_{TF} does not use paid actions, we have $M_{TF}_i(\sigma) = s_i$. We want to define a potential function that “keeps enough in the bank” to compensate for steps where O_{PT} is significantly cheaper than M_{TF} . We do this by keeping some budget for every inconsistency between L and L^* . More precisely, we set the potential $\Phi_i(\sigma)$ after step i to be the number of inversions (flipped pairs) between L_i and L_i^* . Formally, we define

$$\Phi_i(\sigma) := \left| \{(x, y) \in X^2 : x \text{ before } y \text{ in } L_i \text{ but } y \text{ before } x \text{ in } L_i^*\} \right|.$$

Observe that $\Phi_i(\sigma) \geq \Phi_0(\sigma) = 0$. It remains to show that this potential function does the trick, i.e., that M_{TF} is not paying too much in cheap steps and has sufficient potential to pay for expensive steps.

In order to determine the amortized cost $a_i(\sigma)$, we need to compute the potential difference $\Delta\Phi_i(\sigma) := \Phi_i(\sigma) - \Phi_{i-1}(\sigma)$ in each step i . To do this, we use the alternating moves approach, where we pretend that M_{TF} moves first and O_{PT} only afterwards. This results in an intermediate potential value of $\Phi'_i(\sigma) := \left| \{(x, y) \in X^2 : x \text{ before } y \text{ in } L_i \text{ but } y \text{ before } x \text{ in } L_{i-1}^*\} \right|$ after the move of M_{TF} but before O_{PT} 's move. We can then separately compute $\Delta\Phi_i^{M_{TF}}(\sigma) := \Phi'_i(\sigma) - \Phi_{i-1}(\sigma)$ and $\Delta\Phi_i^{O_{PT}}(\sigma) := \Phi_i(\sigma) - \Phi'_i(\sigma)$ in order to obtain $\Delta\Phi_i(\sigma) = \Delta\Phi_i^{M_{TF}}(\sigma) + \Delta\Phi_i^{O_{PT}}(\sigma)$.

Let X_i be the set of items that are in front of r_i in L_{i-1} , but behind r_i in L_{i-1}^* (indicated as “-” in Figure 4.2). Since M_{TF} moves r_i to the front in L_i , it can only introduce additional flipped pairs between r_i and items that are in front of r_i in L_{i-1}^* . The number of these items is exactly $s_i - |X_i|$ and must, by definition, be at most s_i^* , hence we get

$$s_i - |X_i| \leq s_i^*. \quad (4.2)$$

On the other hand, all flipped pairs between r_i and items in X_i get repaired in L_i by moving r_i to the front. Overall, the change in potential caused by M_{TF} is $\Delta\Phi_i^{M_{TF}}(\sigma) = s_i - 2|X_i|$.



Figure 4.2: Schematic illustration of the potential change during M_{TF} 's move in step i . The items are marked according to whether they contribute positively, negatively, or not at all to $\Delta\Phi_i^{M_{TF}}(\sigma)$ when r_i is moved to the front.

Now consider O_{PT} 's action in step i , after M_{TF} already moved r_i to the front of its list. If O_{PT} moves r_i forward, this is guaranteed to decrease the number of flipped pairs, as r_i is at the very front in M_{TF} 's list. Every paid action, however, can increase the number of flipped pairs by at most one. The change in potential due to O_{PT} 's decisions is therefore bounded by $\Delta\Phi_i^{O_{PT}}(\sigma) \leq p_i^*$. Overall, we obtain

$$\Delta\Phi_i(\sigma) = \Delta\Phi_i^{M_{TF}}(\sigma) + \Delta\Phi_i^{O_{PT}}(\sigma) \leq s_i - 2|X_i| + p_i^*.$$

For the amortized cost this implies

$$\begin{aligned}
a_i(\sigma) &= \text{MTF}_i(\sigma) + \Delta\Phi_i(\sigma) \\
&\leq s_i + s_i - 2|X_i| + p_i^* \\
&\stackrel{(4.2)}{\leq} 2s_i^* + p_i^* \\
&\leq 2 \cdot \text{OPT}_i(\sigma).
\end{aligned} \tag{4.3}$$

We can immediately apply Lemma 4.2 with $\alpha = \beta = 0$ to obtain that MTF is strictly 2-competitive, i.e., MTF does not require more than twice the number of operations as the best offline solution that knows the order in which items will be requested ahead of time. \square

4.3 Lower bound via averaging

Now that we have an upper bound on the competitive ratio of MTF , it is natural to ask whether there are better algorithms than using the simple MTF heuristic. In other words, asymptotically speaking, we are looking for a lower bound on the competitive ratio $\text{ALG}(\sigma)/\text{OPT}(\sigma)$ of the list update problem. This means that, for every possible online algorithm ALG , we need an instance $\sigma \in \Sigma$ together with a lower bound on $\text{ALG}(\sigma)$ and an upper bound on $\text{OPT}(\sigma)$. We will first derive a general upper bound on the cost of the offline optimum and then construct a bad instance for any given online algorithm.

In Section 2.2 we obtained an upper bound by explicitly analyzing the behavior of OPT . Without understanding the structure of optimum offline solutions or even optimum offline algorithms, it is generally difficult to derive good bounds on $\text{OPT}(\sigma)$ that depend only on n but are otherwise independent of the request sequence σ . We use an averaging argument – a powerful technique that is often used to obtain upper bounds on the cost of optimum solutions. The general idea of averaging is the following:

1. Introduce a uniform class \mathcal{A} of offline algorithms with a reasonably simple structure.
2. For any given instance σ , compute the average cost c_{avg} over all algorithms in \mathcal{A} . This average should ideally be independent of σ , since we chose a uniform class of algorithms, and it should be easy to compute, since we chose a class of simple structure.
3. Use the fact that there is at least one algorithm $\text{ALG} \in \mathcal{A}$ with $\text{ALG}(\sigma) \leq c_{\text{avg}}$ and that $\text{OPT}(\sigma) \leq \text{ALG}(\sigma)$ to get the bound $\text{OPT}(\sigma) \leq c_{\text{avg}}$.

To illustrate, we apply this technique to the list update problem to obtain an upper bound on $\text{OPT}(\sigma)$.

Lemma 4.5. The offline optimum for the list update problem is bounded by $\text{OPT}(\sigma) \leq n \cdot \frac{l-1}{2} + \alpha$ for all $\sigma \in \Sigma$ and fixed $\alpha := \alpha(l) \in \mathbb{N}$.

Proof. We consider the simple and uniform class of all algorithms that order the list into one of the $l!$ possibilities using paid actions in the first step, and then never change the order again. By definition, there are $l!$ such algorithms. The maximum cost in step 1 over all these algorithms is $\frac{l(l-1)}{2}$, hence the average setup cost $\alpha(l) < \frac{l(l-1)}{2}$ only depends on the number l of elements in the list.

As the second step of the averaging technique, we compute the average cost c_{avg} over all $l!$ algorithms for any input $\sigma = (r_1, \dots, r_n)$. In every step i , there are exactly $(l-1)!$ algorithms that have r_i at position j , for every $j \in \{1, \dots, l\}$. The average cost $c_{\text{avg},i}$ of step $i \geq 2$ over all algorithms in our class therefore is

$$c_{\text{avg},i} = \frac{1}{l!} \sum_{j=1}^l (j-1)(l-1)! = \frac{(l-1)!}{l!} \cdot \frac{l(l-1)}{2} = \frac{l-1}{2}.$$

Note that we chose our class of algorithms in such a way that $c_{\text{avg},i}$ is independent of both i and σ and is easy to compute. This is the secret behind the power of the averaging technique!

To bound the cost of OPT it now only remains to compute c_{avg} , since at least one algorithm of our class does not exceed the average cost and, hence, neither does OPT . We get

$$\text{OPT}(\sigma) \leq c_{\text{avg}} = \alpha(l) + \sum_{i=2}^n c_{\text{avg},i} = (n-1) \cdot \frac{l-1}{2} + \alpha(l) = n \cdot \frac{l-1}{2} + \alpha'(l).$$

□

Now that we have a general upper bound on $\text{OPT}(\sigma)$, we can sabotage any given online algorithm ALG and obtain the following lower bound on the competitive ratio.

Theorem 4.6 (Karp and Raghavan [1990]). Every deterministic algorithm for the list update problem has competitive ratio at least 2.

Proof. Fix any deterministic algorithm ALG . An obvious request sequence that comes to mind is the sequence that always requests the last element on ALG 's list, which forces a total cost of $\text{ALG}(\sigma) \geq n(l-1) \rightarrow_{n \rightarrow \infty} \infty$ (more than $n(l-1)$ if ALG uses paid exchanges). Together with Lemma 4.5 this yields

$$\frac{\text{ALG}(\sigma)}{\text{OPT}(\sigma)} \geq \frac{n(l-1)}{n(l-1)/2 + \alpha(l)} = \frac{2(l-1)}{(l-1) + 2\alpha(l)/n} \rightarrow_{n \rightarrow \infty} 2.$$

By Proposition 1.12 and since ALG was chosen arbitrarily, this concludes the proof. □

Together with Theorem 4.4, we have shown a tight bound on the strict and non-strict deterministic competitive ratio of the list update problem.

Corollary 4.7. The list update problem has (strict) competitive ratio 2.

4.4 Upper bounds via phase partitioning

In Section 4.2 we employed amortized analysis in order to avoid understanding the evolution of MTF 's list during the course of the algorithm. In this section, we use phase partitioning (cf. Section 2.4) to explicitly analyze this evolution. Importantly, we will first break the general case down to understanding lists of only two elements.

To this end, we let the cost $\text{ALG}_{x,y}(\sigma)$ of an algorithm ALG induced by $(x, y) \in X^2$ be defined as the number of steps in which y is accessed and x precedes y in ALG 's list, i.e., $\text{ALG}_{x,y}(\sigma)$ is the part of the cost for accessing y that is “caused by” x . Observe that, by definition,

$$\text{ALG}(\sigma) = \sum_{(x,y) \in X^2} \text{ALG}_{x,y}(\sigma) + \sum_{i=1}^n p_i. \quad (4.4)$$

In the same vein, for any sequence of items, we obtain the *projection* onto $\{x, y\} \subseteq X$ by removing all occurrences of items in $X \setminus \{x, y\}$. In this way, let $\sigma_{x,y}$ denote the projection of σ onto $\{x, y\} \subseteq X$, let $L_{x,y}$ denote the projection of L onto $\{x, y\} \subseteq X$, etc. Finally, for convenience, we abuse notation and write $\text{ALG}[\sigma_{x,y}]$, $\text{ALG}(\sigma_{x,y})$ to denote the solution produced by ALG and its cost, respectively, when executed on list $L_{x,y}$ with input $\sigma_{x,y}$. See Figure 4.3 for an example.

	MTF[σ]	$\text{MTF}_{x_1x_2}$	$\text{MTF}_{x_2x_1}$	σ	$\sigma_{x_1x_2}$	$\text{MTF}(\sigma_{x_1x_2})$	$\text{MTF}[\sigma_{x_1x_2}]$
L_0	$x_1 x_2 x_3 x_4 x_5$	0	0	x_4			$x_1 x_2$: $L_{x_1x_2,0}$
L_1	$x_4 x_1 x_2 x_3 x_5$	1	0	x_2	x_2	1	
L_2	$x_2 x_4 x_1 x_3 x_5$	0	1	x_1	x_1	1	$x_2 x_1$: $L_{x_1x_2,1}$
L_3	$x_1 x_2 x_4 x_3 x_5$	0	0	x_5			$x_1 x_2$: $L_{x_1x_2,2}$
L_4	$x_5 x_1 x_2 x_4 x_3$	0	0	x_1	x_1	0	
L_5	$x_1 x_5 x_2 x_4 x_3$	1	0	x_2	x_2	1	$x_1 x_2$: $L_{x_1x_2,3}$
L_6	$x_2 x_1 x_5 x_4 x_3$						$x_2 x_1$: $L_{x_1x_2,4}$

Figure 4.3: Illustration of the cost $\text{MTF}_{x_1x_2}(\sigma)$ induced by (x_1, x_2) and the behavior of $\text{MTF}[\sigma_{x_1x_2}]$ on the projection $\sigma_{x_1x_2} = (x_2, x_1, x_1, x_2)$ of the request sequence $\sigma = (x_4, x_2, x_1, x_5, x_1, x_2)$ onto $\{x_1, x_2\}$. Observe that $\text{MTF}_{x_1x_2}(\sigma) + \text{MTF}_{x_2x_1}(\sigma) = \text{MTF}(\sigma_{x_1x_2})$.

Definition 4.8. A (deterministic or randomized) list update algorithm ALG satisfies the *pairwise property* if, for all $\sigma \in \Sigma$ and $x, y \in X$, $\text{ALG}_{xy}(\sigma) + \text{ALG}_{yx}(\sigma) = \text{ALG}(\sigma_{xy})$.

Intuitively, the pairwise property states that the costs induced by any pair of elements are fully captured already in the projected instance. Observe that this need not be the case in general (exercise). We give an alternative characterization of the pairwise property.

Lemma 4.9. If an online algorithm ALG does not use paid exchanges, it satisfies the pairwise property if and only if, for every $\sigma \in \Sigma$ and every $x, y \in X$ the relative order of x and y at the end of $\text{ALG}[\sigma]$ and $\text{ALG}[\sigma_{xy}]$ is the same.

Proof. Consider a shortest possible request sequence $\sigma = (r_1, \dots, r_n) \in \Sigma$ for which x and y are in different relative orders after executing $\text{ALG}[\sigma]$ and $\text{ALG}[\sigma_{xy}]$. Since ALG does not use paid exchanges, the relative order of x and y can only change in steps where one of these two items is accessed. Since σ is shortest possible, we have $r_n \in \{x, y\}$ and x and y are in opposite relative orders for the first time after the last step. Then, $\text{ALG}_{xy}(\sigma) = \text{ALG}_{xy}(\sigma_{xy})$ and $\text{ALG}_{yx}(\sigma) = \text{ALG}_{yx}(\sigma_{xy})$ and thus $\text{ALG}_{xy}(\sigma \oplus (x)) = \text{ALG}_{xy}(\sigma_{xy} \oplus (x))$ but $\text{ALG}_{yx}(\sigma \oplus (x)) \neq \text{ALG}_{yx}(\sigma_{xy} \oplus (x))$. Since $\sigma_{xy} \oplus (x) = (\sigma \oplus (x))_{xy}$, it follows that

$$\begin{aligned}
\text{ALG}_{xy}(\sigma \oplus (x)) + \text{ALG}_{yx}(\sigma \oplus (x)) &\neq \text{ALG}_{xy}((\sigma_{xy} \oplus (x))) + \text{ALG}_{yx}((\sigma_{xy} \oplus (x))) \\
&\stackrel{(4.4)}{=} \text{ALG}((\sigma_{xy} \oplus (x))) \\
&= \text{ALG}((\sigma \oplus (x))_{xy}),
\end{aligned}$$

and the pairwise property is violated for $(\sigma \oplus (x))$.

Conversely, assume that ALG violates the pairwise property and let $\sigma = (r_1, \dots, r_n) \in \Sigma$ be a shortest possible request sequence for which this is the case. By definition, $\text{ALG}_{xy}(\sigma) + \text{ALG}_{yx}(\sigma) \neq \text{ALG}(\sigma_{xy})$ for $x = r_n$ and some $y \in X$. By minimality of σ , and since ALG does not use paid exchanges, x and y are in different relative orders at the beginning of the last step in $\text{ALG}[\sigma]$ and $\text{ALG}[\sigma_{xy}]$. This means that the relative order of x and y is different after executing ALG for $\sigma_{\leq n-1}$ and for its projection onto $\{x, y\}$. \square

Corollary 4.10. MTF satisfies the pairwise property.

Proof. This follows immediately by Lemma 4.9, since the relative order of two items $x, y \in X$ in MTF's list only depends on which of the two elements was accessed most recently (and on L_0). \square

We are now ready to formulate a condition under which it is sufficient to analyze the behavior of an algorithm on lists of length two. Using this approach is often referred to as the *list factoring technique*.

Theorem 4.11 (Bentley and McGeoch [1985]). Let ALG be a (deterministic or randomized) online algorithm that does not use paid exchanges and that satisfies the pairwise property, and let $\rho \geq 1$. Then, ALG is (strictly) ρ -competitive if and only if it is (strictly) ρ -competitive on lists of length two.

Proof. For the non-trivial part of the theorem, we assume that ALG is ρ -competitive on lists of length two for some $\rho \geq 1$, and show that this implies ρ -competitiveness in general. In particular, for a deterministic algorithm ALG, assume there exists $\alpha \geq 0$ such that we have

$$\text{ALG}(\sigma_{xy}) \leq \rho \cdot \text{OPT}(\sigma_{xy}) + \alpha, \quad (4.5)$$

for all $\sigma \in \Sigma$ and $x, y \in X$.

Fix a request sequence $\sigma \in \Sigma$ and two items $x, y \in X$. Consider the offline solution $\text{OFF}[\sigma_{xy}]$ on the projected instance σ_{xy}, L_{xy} that behaves as follows. If $\text{OPT}[\sigma]$ switches the relative order of x and y by using a free action during the k -th access to an item in $\{x, y\}$, then $\text{OFF}[\sigma_{xy}]$ does the same on the projected instance. If $\text{OPT}[\sigma]$ uses a paid exchange to switch x and y between two accesses to items in $\{x, y\}$, then $\text{OFF}[\sigma_{xy}]$ does the same between the corresponding steps in σ_{xy} . By definition, $\text{OFF}[\sigma_{xy}]$ maintains the same relative order of x and y and uses the same number of paid exchanges between x and y as $\text{OPT}[\sigma]$. Hence,

$$\text{OPT}(\sigma_{xy}) \leq \text{OFF}(\sigma_{xy}) = \text{OPT}_{xy}(\sigma) + \text{OPT}_{yx}(\sigma) + p_{xy}^*, \quad (4.6)$$

where p_{xy}^* denotes the number of paid exchanges between x and y in $\text{OPT}[\sigma]$.

Now, since ALG does not use paid exchanges (i.e., $p_i = 0$) and satisfies the pairwise property, we obtain

$$\begin{aligned} \text{ALG}(\sigma) &\stackrel{(4.4)}{=} \sum_{(x,y) \in X^2} \text{ALG}_{xy}(\sigma) \\ &= \sum_{\{x,y\} \subseteq X} (\text{ALG}_{xy}(\sigma) + \text{ALG}_{yx}(\sigma)) \\ &\stackrel{\text{Def. 4.8}}{=} \sum_{\{x,y\} \subseteq X} \text{ALG}(\sigma_{xy}) \\ &\stackrel{(4.5)}{\leq} \sum_{\{x,y\} \subseteq X} (\rho \cdot \text{OPT}(\sigma_{xy}) + \alpha) \\ &\stackrel{(4.6)}{\leq} \rho \sum_{\{x,y\} \subseteq X} (\text{OPT}_{xy}(\sigma) + \text{OPT}_{yx}(\sigma) + p_{xy}^*) + \frac{l(l-1)}{2} \alpha \\ &\stackrel{(4.4)}{=} \rho \cdot \text{OPT}(\sigma) + \alpha'(l), \end{aligned}$$

with $\alpha'(l) \geq 0$. Hence ALG is ρ -competitive on σ , as claimed. If ALG is strictly ρ -competitive on instances of length two, then $\alpha = 0$ and $\alpha'(l) = 0$, thus ALG is strictly ρ -competitive on σ .

For randomized algorithms, the proof is identical if we replace all costs by expected costs. \square

This provides an alternative way to upper bound the competitive ratio of MTF.

Corollary 4.12. M_{TF} is strictly 2-competitive.

Proof. By Corollary 4.10 and Theorem 4.11, it suffices to show that M_{TF} is strictly 2-competitive on lists of length two. To see this, consider a request sequence $\sigma = (r_1, \dots, r_n)$ for the list update problem with $|X| = 2$ items. Let $i_1 < \dots < i_m$ with $m \leq n$ denote the steps in which M_{TF} has non-zero cost, i.e., $M_{TF_{i_j}}(\sigma) = 1$ for $j \in \{1, \dots, m\}$, and $M_{TF}(\sigma) = m$. For convenience, let $i_0 := 0$. Observe that $r_{i_j} \neq r_{i_{j+1}}$ for $j \in \{1, \dots, m-1\}$, and x_{i_1} must be in second position in L_0 . It follows that $\sum_{i=i_j}^{i_{j+1}} OPT_i(\sigma) \geq 1$ for $j \in \{0, \dots, m\}$ (see Figure 4.4).




Figure 4.4: Illustration of the accounting in the proof of Corollary 4.12 for $\sigma = (0, 0, 1, 1, 1, 1, 0, 0, 1, 1, 0)$ and $L_0 = (0, 1)$.

With this, we have

$$\begin{aligned}
 OPT(\sigma) &= \sum_{i=1}^n OPT_i(\sigma) \\
 &\geq \sum_{j=0}^{\lfloor (m-1)/2 \rfloor} \sum_{i=i_{2j}}^{i_{2j+1}} OPT_i(\sigma) \\
 &\geq \lfloor (m-1)/2 \rfloor + 1 \\
 &\geq m/2 \\
 &= \frac{1}{2} M_{TF}(\sigma). \quad \square
 \end{aligned}$$

4.4.1 The timestamp algorithm

The list factoring technique allows to analyze algorithms that are more involved than M_{TF} . For example, consider the algorithm `TIMESTAMP` that is defined as follows. If an item $x \in X$ is accessed for the first time, `TIMESTAMP` does not change the order of the list. Each subsequent time $x \in X$ is accessed, `TIMESTAMP` moves x in front of the first item on its list that was accessed at most once since the previous access to x . We will see that `TIMESTAMP` is a competitive algorithm. To this end, we employ the list factoring technique.

Lemma 4.13. `TIMESTAMP` satisfies the pairwise property.

Proof. We use the characterization of the pairwise property of Lemma 4.9. For the sake of contradiction, let $\sigma \in X^n$ be the shortest request sequence for which $x \in X$ is in front of $y \in X$ after execution of `TIMESTAMP` $[\sigma]$, but not after `TIMESTAMP` $[\sigma_{xy}]$. By minimality of σ , `TIMESTAMP` $[\sigma]$ moves x in front of y in the last step, but `TIMESTAMP` $[\sigma_{xy}]$ keeps y in front of x . Note that the converse is not possible by definition of `TIMESTAMP`.

First observe that $i = \max\{i' < n \mid r_{i'} = x\}$ must exist, otherwise `TIMESTAMP` $[\sigma]$ would not move x forward. Since `TIMESTAMP` $[\sigma_{xy}]$ does not move x forward, there must be steps j and j' with $i < j < j' < n$ and $r_j = r_{j'} = y$, i.e., y was accessed at least twice since step i . At the end of step j' , all items in front of y in

$\text{TIMESTAMP}[\sigma]$ must have been accessed at least twice in steps j, \dots, j' . By induction, any additional item moved in front of some item y in steps $j' + 1, \dots, n$ must have been requested at least twice since step j . In particular, x is not moved in front of y by $\text{TIMESTAMP}[\sigma]$ – a contradiction. \square

To analyze the behavior of an algorithm on lists of length two, the following phase partitioning will be useful. Note that we can extend any request sequence so that it starts and ends with two consecutive accesses to the same element, without affecting $\text{OPT}(\sigma)$.

Definition 4.14. The (unique) *alternating phase partitioning* of a request sequence $\sigma_{xy} \in \{x, y\}^n$ for $L_0 = (x, y)$ is given by $\sigma'_{xy} = \pi_0 \oplus \dots \oplus \pi_N = (x, x) \oplus \sigma_{xy} \oplus (z)^k$, where z denotes the item in front of the list at the end of $\text{OPT}[\sigma_{xy}]$, $k \in \{1, 2\}$ is defined such that the number of consecutive requests to z at the end of σ'_{xy} is even, and each phase $\pi_j = (a_1, a_2, \dots, a_m)$ for $j \in \{0, \dots, N\}$ is defined such that $a_1 \neq a_2 \neq \dots \neq a_{m-1} = a_m$. Note that $\text{OPT}(\sigma'_{xy}) = \text{OPT}(\sigma_{xy})$.

For example, the alternating phase partitioning of $\sigma_{xy} = (x, y, x, x, x, y, x, y, x, x, y, y, y)$ for $L_0 = (x, y)$ is $\pi_0 \oplus \dots \oplus \pi_4 = (x, x) \oplus \sigma_{xy} \oplus (y)$ with phases $\pi_0 = (x, x)$, $\pi_1 = (x, y, x, x)$, $\pi_2 = (x, y, x, y, x, x)$, $\pi_3 = \pi_4 = (y, y)$.

We now use the alternating phase partitioning to analyze TIMESTAMP .

Theorem 4.15 (Albers [1998]). TIMESTAMP is strictly 2-competitive.

Proof. By Theorem 4.11 and Lemma 4.13, it is sufficient to show that TIMESTAMP is strictly 2-competitive on lists of length two. To this end, let $\sigma_{xy} \in \Sigma$ be a request sequence for a list with $X = \{x, y\}$ and consider the alternating phase partitioning $\sigma'_{xy} = \pi_0 \oplus \dots \oplus \pi_N$ of σ_{xy} . Observe that $\text{TIMESTAMP}(\sigma_{xy}) \leq \text{TIMESTAMP}(\sigma'_{xy})$, since y is first moved in the same corresponding steps. Also, $\text{OPT}(\sigma'_{xy}) = \text{OPT}(\sigma_{xy})$, which means that it suffices to show $\text{TIMESTAMP}(\sigma'_{xy}) \leq 2 \text{OPT}(\sigma'_{xy})$. We fix OPT to be the algorithm that changes the order of the items exactly when the second item in the list is accessed and the next request is for the same item (note that we can adapt any solution to be of this form without increasing its cost). We abuse notation and let $\text{TIMESTAMP}(\pi_j)$ and $\text{OPT}(\pi_j)$ denote the costs of TIMESTAMP and OPT , respectively, incurred during phase π_j for $j \in \{0, \dots, N\}$. We claim that $\text{TIMESTAMP}(\pi_j) \leq 2 \cdot \text{OPT}(\pi_j)$ for every phase π_j . This completes the proof, since then

$$\text{TIMESTAMP}(\sigma'_{xy}) = \sum_{j=0}^N \text{TIMESTAMP}(\pi_j) \leq 2 \sum_{j=0}^N \text{OPT}(\pi_j) = 2 \text{OPT}(\sigma'_{xy}).$$

To prove our claim for π_0 , observe that $\text{OPT}(\pi_0) = \text{TIMESTAMP}(\pi_0) = 0$.

Now consider a phase π_j with $j \in \{1, \dots, N\}$ and assume, without loss of generality, that π_{j-1} ends with a request to x . Note that by definition of both TIMESTAMP and OPT , the item x is the first item in both OPT 's and TIMESTAMP 's list at the start of phase π_j . There are four cases, depending on whether π_j starts with x or y and whether it ends with x or y . The following table lists the costs incurred by TIMESTAMP and OPT in each case, where $k \in \mathbb{N} \cup \{0\}$. Observe that in each case $\text{TIMESTAMP}(\pi_j) \leq 2 \cdot \text{OPT}(\pi_j)$, as claimed.

form of phase π_j	$\text{TIMESTAMP}(\pi_j)$ ($k > 0$)	$\text{TIMESTAMP}(\pi_j)$ ($k = 0$)	$\text{OPT}(\pi_j)$
$x(yx)^k x$	$2k - 1$	0	k
$x(yx)^k y y$	$2k$	2	$k + 1$
$yx(yx)^k x$	$2k + 1$	1	$k + 1$
$(yx)^k y y$	$2k$	2	$k + 1$

\square

4.5 Randomized list update algorithms

In Theorem 4.6, we have seen that no deterministic list update algorithm can be better than 2-competitive. The proof used a request sequence that depends on the behavior of the online algorithm by repeatedly accessing the item currently at the end of the list. Similarly to the paging problem (cf. Section 3.1), we may try to use randomization in order to be less predictable and improve on MTF 's performance on instances of this kind. Recall that we assumed the adversary to be oblivious, i.e., it does not know the outcomes of random choices of our algorithm when constructing its worst-case instance.

A natural way to randomize MTF would be to move accessed items to the front independently with probability $1/2$. It turns out that RMTF does not improve on the performance of MTF at all.²

Proposition 4.16. The competitive ratio of RMTF is at least 2.

Proof. Let $L_0 = (x, y)$ and consider the adversarial request sequence $\sigma_{\leq n}^{(y)} = (y, y, \dots, y)$ of length $n \in \mathbb{N}$. The offline optimum solution for this instance is $\text{MTF}[\sigma_{\leq n}^{(y)}]$ and has cost

$$\text{OPT}(\sigma_{\leq n}^{(y)}) = \text{MTF}(\sigma_{\leq n}^{(y)}) = 1.$$

Let I be the random variable indicating the step in which RMTF moves y to the front. Then

$$\mathbb{E}[\text{RMTF}(\sigma_{\leq n}^{(y)})] = \sum_{i=1}^n \Pr[I \geq i] = \sum_{i=1}^n \frac{1}{2^{i-1}} = \frac{1 - 1/2^n}{1 - 1/2} = 2 - \frac{1}{2^{n-1}} \rightarrow_{n \rightarrow \infty} 2 \cdot \text{OPT}(\sigma_{\leq n}^{(y)}).$$

In order to apply Proposition 3.5, we need to ensure that $\lim_{n \rightarrow \infty} \text{OPT}(\sigma) = \infty$. We can achieve that similarly to the proof of Proposition 3.4, by using the request sequence $\sigma_{\leq n}^{(y)} \oplus \sigma_{\leq n}^{(x)} \oplus \sigma_{\leq n}^{(y)} \oplus \dots$, where $\sigma_{\leq n}^{(x)} = (x, x, x, \dots) \in X^n$. \square

Intuitively, RMTF sometimes takes too long before it moves an item forward. We introduce an algorithm BIT that ensures that items are moved forward every second time. To this end, BIT initially sets a bit $b(x) \in \{0, 1\}$ randomly and with equal probability for every item $x \in X$. Whenever an item $x \in X$ is accessed, BIT flips the value of $b(x)$ and then moves x to the front of the list if $b(x) = 1$. We show that BIT indeed improves on the competitive ratio of MTF .

Theorem 4.17 (Reingold et al. [1994]). BIT is strictly 1.75-competitive.

Proof. Observe that, during execution of BIT for a request sequence σ , the relative order of two elements only depends on σ_{xy} (and the initial choice of b). In other words, BIT satisfies the pairwise property (by Lemma 4.9). Since BIT does not use paid exchanges, we can therefore use the list factoring technique (Theorem 4.11) to restrict our analysis to request sequences σ_{xy} for lists of length two. As in the proof of Theorem 4.15, we fix OPT to be the algorithm that changes the order of the items exactly when the second item is accessed and the next request is for the same item.

To this end, let $\sigma'_{xy} = \pi_0 \oplus \dots \oplus \pi_N$ be the alternating phase partitioning of σ_{xy} . Note that $\text{BIT}(\sigma'_{xy}) \geq \text{BIT}(\sigma_{xy})$ and $\text{OPT}(\sigma'_{xy}) = \text{OPT}(\sigma_{xy})$, which means that it suffices to show $\text{BIT}(\sigma'_{xy}) \leq 1.75 \cdot \text{OPT}(\sigma'_{xy})$. As in the proof of Theorem 4.15, it suffices to show that $\mathbb{E}[\text{BIT}(\pi_j)] \leq \frac{7}{4} \cdot \text{OPT}(\pi_j)$ for every phase π_j . By definition of the alternating phase partitioning, the only item accessed during phase π_0 is already at the front of L_0 , thus $\text{OPT}(\pi_0) = \text{BIT}(\pi_0) = 0$.

² Note that RMTF is not formally a distributional randomized algorithm, and we take the expectation over the outcome of its random choices in Definition 3.2.

Now, consider any phase π_j with $j \in \{1, \dots, N\}$. Let $x \in X$ be the item last requested in phase π_{j-1} and $y \in X$ be the other item. Observe that at the start of phase π_j , the item x is in front in both OPT 's and BIT 's list, and the values of $b(x)$ and $b(y)$ are both independently 0 or 1 with probability $1/2$. Let $\mathbb{E}[\text{BIT}_{xx}(\sigma')]$ denote the expected cost of BIT for the subsequence σ' , assuming that that x was requested in the two previous steps. Then, using the abbreviation $abc \dots \equiv (a, b, c, \dots)$, we have

$$\begin{aligned}\mathbb{E}[\text{BIT}_{xx}(xy)] &= 1, \\ \mathbb{E}[\text{BIT}_{xx}(yx)] &= \mathbb{E}[\text{BIT}_{xx}(yy)] = 1 + \frac{1}{2} = \frac{3}{2}.\end{aligned}$$

Let $\mathbb{E}[\text{BIT}_{xyx}(\sigma')]$ denote the expected cost of BIT for the request sequence σ' , assuming that the previously requested items were x, y, x in this order, and let $\mathbb{E}[\text{BIT}_{yxy}(\sigma')]$ be defined analogously. Observe that the probability that y is in front of BIT 's list after the subsequence (x, y, x) is $1/4$, independently of the earlier requests, and symmetrically for x and (y, x, y) . Hence,

$$\begin{aligned}\mathbb{E}[\text{BIT}_{xyx}(y)] &= \mathbb{E}[\text{BIT}_{yxy}(x)] = \frac{3}{4}, \\ \mathbb{E}[\text{BIT}_{xyx}(x)] &= \mathbb{E}[\text{BIT}_{yxy}(y)] = \frac{1}{4}, \\ \mathbb{E}[\text{BIT}_{xyx}(yy)] &= \mathbb{E}[\text{BIT}_{xyx}(y)] + \mathbb{E}[\text{BIT}_{yxy}(y)] = \frac{3}{4} + \frac{1}{4} = 1, \\ \mathbb{E}[\text{BIT}_{xyx}((yx)^k)] &= 2k \cdot \frac{3}{4} = \frac{3}{2}k, \\ \mathbb{E}[\text{BIT}_{xx}((yx)^k)] &= \frac{3}{2} + \frac{3}{2}(k-1) = \frac{3}{2}k.\end{aligned}$$

With this, we have the following costs, depending on the first and last request of π_j , assuming that the previous phase ended with two requests for $x \in X$. In particular, we have $\mathbb{E}[\text{BIT}(\pi_j)] \leq \frac{7}{4} \cdot \text{OPT}(\pi_j)$ in each case.

form of phase π_j	$\mathbb{E}[\text{BIT}(\pi_j)]$ ($k > 0$)	$\mathbb{E}[\text{BIT}(\pi_j)]$ ($k = 0$)	$\text{OPT}(\pi_j)$
$x(yx)^k x$	$0 + \frac{3}{2}k + \frac{1}{4} = \frac{1}{4} + \frac{3}{2}k$	0	k
$x(yx)^k yy$	$0 + \frac{3}{2}k + 1 = 1 + \frac{3}{2}k$	$1 + \frac{1}{2} = \frac{3}{2}$	$k + 1$
$yx(yx)^k x$	$\frac{3}{2} + \frac{3}{2}k + \frac{1}{4} = \frac{7}{4} + \frac{3}{2}k$	$\frac{3}{2} + \frac{1}{4} = \frac{7}{4}$	$k + 1$
$(yx)^k yy$	$\frac{3}{2} + \frac{3}{2}(k-1) + 1 = 1 + \frac{3}{2}k$	$\frac{3}{2}$	$k + 1$

□

If we compare the performance of BIT to the performance of TIMESTAMP for each possible phase in terms of competitive ratios, we observe that both algorithms have their worst-cases in different types of phases:

form of phase π_j	$\mathbb{E}[\text{BIT}(\pi_j)]/\text{OPT}(\pi_j)$				$\text{TIMESTAMP}(\pi_j)/\text{OPT}(\pi_j)$			
	$k = 0$	$k = 1$	$k = 2$	$k \geq 2$	$k = 0$	$k = 1$	$k = 2$	$k \geq 2$
$x(yx)^k x$	–	1.75	1.63	$\in [1.5, 1.63]$	–	1	1.5	$\in [1.5, \mathbf{2}]$
$x(yx)^k yy$	1.5	1.25	1.33	$\in [1.33, 1.5]$	2	1	1.33	$\in [1.33, \mathbf{2}]$
$yx(yx)^k x$	1.75	1.63	1.58	$\in [1.5, 1.58]$	1	1.50	1.67	$\in [1.67, \mathbf{2}]$
$(yx)^k yy$	1.5	1.25	1.33	$\in [1.33, 1.5]$	2	1	1.33	$\in [1.33, \mathbf{2}]$

Linearity of expectation allows us to interpolate linearly between the performance of BIT and TIMESTAMP . We define the randomized algorithm COMB as a superposition of BIT and TIMESTAMP . More precisely, COMB initially selects BIT with probability $4/5$ and TIMESTAMP with probability $1/5$.

Theorem 4.18 (Albers et al. [1995]). COMB is strictly 1.6-competitive.

Proof. By Lemma 4.13 and since BIT satisfies the pairwise property, any superposition of BIT and TIMESTAMP , including COMB , has the pairwise property. This means that, as before, we can use list factoring (Theorem 4.11) and restrict our analysis to phases π_j of the alternating phase partitioning of request sequences on lists of length two. Because of

$$\mathbb{E}[\text{COMB}(\pi_j)] = \mathbb{E}\left[\frac{4}{5}\mathbb{E}[\text{BIT}(\pi_j)] + \frac{1}{5}\text{TIMESTAMP}(\pi_j)\right] = \frac{4}{5}\mathbb{E}[\text{BIT}(\pi_j)] + \frac{1}{5}\text{TIMESTAMP}(\pi_j),$$

we obtain the values $\mathbb{E}[\text{COMB}(\pi_j)]/\text{OPT}(\pi_j)$ as a linear combination of the corresponding values for BIT and TIMESTAMP . For the first phase π_0 , we have

$$\mathbb{E}[\text{COMB}(\pi_0)] = \left(\frac{4}{5} \cdot \mathbb{E}[\text{BIT}(\pi_0)] + \frac{1}{5} \cdot \text{TIMESTAMP}(\pi_0)\right) = 0 = \text{OPT}(\pi_0).$$

Regarding the other phases, the following table contains all linear combinations of the above tables for BIT and TIMESTAMP (assuming that the previous phase ended with two requests for $x \in X$). Observe that the competitive ratio is indeed bounded by 1.6, and that this bound is tight in many cases.

form of phase π_j	$\mathbb{E}[\text{COMB}(\pi_j)]/\text{OPT}(\pi_j)$			
	$k = 0$	$k = 1$	$k = 2$	$k \geq 2$
$x(yx)^k x$	–	1.6	1.6	1.6
$x(yx)^k y y$	1.6	1.2	1.33	$\in [1.33, 1.6]$
$yx(yx)^k x$	1.6	1.6	1.6	1.6
$(yx)^k y y$	1.6	1.2	1.33	$\in [1.33, 1.6]$

□

4.6 Lower bound for randomized list update

Now that we established a good upper bound for randomized list update in the previous section, we derive a lower bound using Yao's principle (cf. Section 3.2).

Theorem 4.19 (Teia [1993]). Every randomized list update algorithm has competitive ratio at least 1.5.

Proof. In order to apply Yao's principle, we define a sequence $(q_N : \Sigma \rightarrow [0, 1])_{N \in \mathbb{N}}$ of probability distributions for a set of items X with $l := |X|$. For each $N \in \mathbb{N}$, we describe q_N by a randomized construction of a request sequence $\sigma = \pi_1 \oplus \dots \oplus \pi_N \in X^n$ consisting of N phases. The construction of the phases uses a simple offline algorithm OFF that moves a requested item to the front of the list if and only if this item is requested again later in the same phase. Every phase π_j , $j \in \{1, \dots, N\}$ is defined in the following way.

We let L'_i denote OFF 's list *at the beginning of step i* . Let i_j denote the first step of phase π_j and set $i_{N+1} = n + 1$. Observe that, by definition of OFF , the order of the list L'_{i_j} at the start of phase π_j only depends on the previous phases. This allows us to base the construction of phase π_j on L'_{i_j} in the following way. In phase π_j we simply request all items one by one in the order given by L'_{i_j} . Each item is either requested a single time or three times in a row, and both cases occur with probability $1/2$. We use phase partitioning to analyze the cost in each phase. By definition of OFF , its cost in phase π_j is given by exactly

$$\text{OFF}(\pi_j) := \sum_{i=1}^l (i-1) = l(l-1)/2. \quad (4.7)$$

Now fix any deterministic online algorithm $\text{ALG} \in \mathcal{A}_{\text{det}}$ and let L_i denote ALG 's list *at the beginning of step i* . We assume that ALG does not use paid exchanges for now and argue later why we may do this. We use amortized analysis to bound the cost of the algorithm in each phase. To that end, consider the potential function $\Phi_i(\sigma) := \sum_{\{x,y\} \subseteq X} \Phi_i^{xy}$ *at the beginning of step i* (analogous to the one of Section 4.2), with

$$\Phi_i^{yx} \equiv \Phi_i^{xy} := \begin{cases} 0 & \text{if } x \text{ and } y \text{ are in the same relative order in } L_i \text{ and } L'_i, \\ 1 & \text{else.} \end{cases}$$

Let $\Delta\Phi(\pi_j) := \Phi_{i_{j+1}}(\sigma) - \Phi_{i_j}(\sigma)$ denote the change in potential during phase j , and let $\Delta\Phi(\sigma) := \sum_{j=1}^N \Delta\Phi(\pi_j) = \Phi_{N+1} - \Phi_{i_1}$ denote the total change in potential. Observe that $\Phi_1(\sigma) = 0$ and $0 \leq \Phi_i(\sigma) \leq l(l-1)/2$ throughout. Therefore,

$$0 \leq \Delta\Phi(\sigma) \leq l(l-1)/2. \quad (4.8)$$

We claim that $\mathbb{E}_{\sigma \sim q_N}[\Delta\Phi(\pi_j) + \text{ALG}(\pi_j)] \geq \frac{3}{4}l(l-1)$. Application of Yao's principle (Theorem 3.7) then completes the proof, since

$$\begin{aligned} \limsup_{N \rightarrow \infty} \frac{\inf_{\text{ALG} \in \mathcal{A}_{\text{det}}} \mathbb{E}_{\sigma \sim q_N}[\text{ALG}(\sigma)]}{\mathbb{E}_{\sigma \sim q_N}[\text{OPT}(\sigma)]} &\geq \limsup_{N \rightarrow \infty} \frac{\inf_{\text{ALG} \in \mathcal{A}_{\text{det}}} \mathbb{E}_{\sigma \sim q_N}[\sum_{j=1}^N \text{ALG}(\pi_j)]}{\mathbb{E}_{\sigma \sim q_N}[\text{OFF}(\sigma)]} \\ &\geq \limsup_{N \rightarrow \infty} \frac{\sum_{j=1}^N \left(\frac{3}{4}l(l-1) - \inf_{\text{ALG} \in \mathcal{A}_{\text{det}}} \mathbb{E}_{\sigma \sim q_N}[\Delta\Phi(\pi_j)] \right)}{\sum_{j=1}^N \mathbb{E}_{\sigma \sim q_N}[\text{OFF}(\pi_j)]} \\ &\stackrel{(4.7)}{=} \limsup_{N \rightarrow \infty} \frac{\frac{3}{4}Nl(l-1) - \sup_{\text{ALG} \in \mathcal{A}_{\text{det}}} \mathbb{E}_{\sigma \sim q_N}[\Delta\Phi(\sigma)]}{\frac{1}{2}Nl(l-1)} \\ &\stackrel{(4.8)}{\geq} \frac{3}{2} - \limsup_{N \rightarrow \infty} \frac{1}{N} \\ &= \frac{3}{2}, \end{aligned}$$

and, by construction, $\lim_{N \rightarrow \infty} \mathbb{E}_{\sigma \sim q_N}[\text{OPT}(\sigma)] = \infty$.

It remains to show our claim. We use list factoring. Let $\text{ALG}_{xy}(\pi_j)$ denote the number of requests to y during phase π_j in which x is in front of y in ALG 's list, i.e., $\text{ALG}_{xy}(\pi_j)$ is the cost for ALG of requests to y "caused" by x during phase π_j . With this, since ALG does not use paid actions, the cost of ALG in phase j can be expressed as

$$\text{ALG}(\pi_j) := \sum_{\{x,y\} \subseteq X} (\text{ALG}_{xy}(\pi_j) + \text{ALG}_{yx}(\pi_j)). \quad (4.9)$$

Fix any phase π_j and any two items $x, y \in X$ such that x is in front of y in the list L'_{i_j} at the start of π_j , i.e., in particular, x is requested before y in phase π_j . Let $\Delta\Phi^{xy}(\pi_j) = \Phi_{i_{j+1}}^{xy} - \Phi_{i_j}^{xy} \in \{-1, 0, 1\}$ denote the change of Φ^{xy} during phase π_j . It suffices to show $\mathbb{E}[\Delta\Phi^{xy}(\pi_j) + \text{ALG}_{xy}(\pi_j) + \text{ALG}_{yx}(\pi_j)] \geq 3/2$, since then, by linearity of expectation,

$$\mathbb{E}[\Delta\Phi(\pi_j) + \text{ALG}(\pi_j)] \stackrel{(4.9)}{=} \sum_{\{x,y\} \subseteq X} \left(\mathbb{E}[\Delta\Phi^{xy}(\pi_j) + \text{ALG}_{xy}(\pi_j) + \text{ALG}_{yx}(\pi_j)] \right) \geq \frac{3}{4}l(l-1).$$

Let i_x, i_y denote the steps in which x , respectively y , are accessed for the first time during phase π_j . With this, define $\Delta_x := \Phi_{i_y}^{xy} - \Phi_{i_j}^{xy}$ and $\Delta_y := \Phi_{i_{j+1}}^{xy} - \Phi_{i_x}^{xy}$, so that $\Delta\Phi^{xy}(\pi_j) = \Delta_x + \Delta_y$ (see Figure 4.5). Note that, because ALG does not use paid exchanges, $\Phi_{i_j}^{xy} = \Phi_{i_x}^{xy}$. We distinguish two cases.




Figure 4.5: Illustration of the notation in the proof of Theorem 4.19 for phase π_j . Note that only steps where x (resp. y) are accessed can contribute to Δ_x (resp. Δ_y).

Case 1: x is in front of y in L_{i_y} , i.e., $\Phi_{i_y}^{xy} = 0$.

Note that the relative order of x and y in L'_{i_y} is still the same as in L'_{i_j} , in particular, x is in front of y in L'_{i_y} as well, hence $\Phi_{i_y}^{xy} = 0$. If $\Phi_{i_j}^{xy} = 1$, then y is in front of x in L_{i_x} and $\text{ALG}_{yx}(\pi_j) \geq 1$. This means that, for either value of $\Phi_{i_j}^{xy}$, we have

$$\Delta_x + \text{ALG}_{yx}(\pi_j) \geq 0. \quad (4.10)$$

Case 1.1: ALG moves y in front of x in step i_y .

If y is accessed once during phase π_j , we have $\Delta_y + \text{ALG}_{xy}(\pi_j) = 1 + 1 = 2$, otherwise we have $\Delta_y + \text{ALG}_{xy}(\pi_j) = 0 + 1 = 1$. Since both cases occur with probability $1/2$, using eq. (4.10) we have $\mathbb{E}[\Delta\Phi^{xy}(\pi_j) + \text{ALG}_{xy}(\pi_j) + \text{ALG}_{yx}(\pi_j)] \geq \mathbb{E}[\Delta_y + \text{ALG}_{xy}(\pi_j)] = 3/2$. Note that we need amortized analysis here, since for $\Phi_{i_j} = 0$ we have $\mathbb{E}[\text{ALG}_{xy}(\pi_j) + \text{ALG}_{yx}(\pi_j)] = 0 + 1 < 3/2$.

Case 1.2: ALG leaves x in front of y in step i_y .

If y is accessed once during phase π_j , we have $\Delta_y + \text{ALG}_{xy}(\pi_j) = 0 + 1 = 1$, otherwise we have $\Delta_y + \text{ALG}_{xy}(\pi_j) \geq 0 + 2 = 2$. Since both cases occur with probability $1/2$, using eq. (4.10) we have $\mathbb{E}[\Delta\Phi^{xy}(\pi_j) + \text{ALG}_{xy}(\pi_j) + \text{ALG}_{yx}(\pi_j)] \geq \mathbb{E}[\Delta_y + \text{ALG}_{xy}(\pi_j)] = 3/2$.

Case 2: x is behind y in L_{i_y} , i.e., $\Phi_{i_y}^{xy} = 1$.

Note that x is in front of y in L'_{i_y} by definition of σ , hence $\Phi_{i_y}^{xy} = 1$. We have $\text{ALG}_{xy}(\pi_j) = 0$. Since ALG does not use paid exchanges, y must already have been ahead of x in L_{i_j} . This implies $\Phi_{i_j}^{xy} = 1$ and thus $\Delta_x = 0$.

Because of $\Delta_x = 0$ and because the relative order of x and y in OFF's list does not change until step i_y , the same must be true for the order in ALG's list, since moving x in front of y cannot be undone by ALG without paid exchanges before y is accessed. If x is accessed once, we have $\text{ALG}_{yx}(\pi_j) = 1$; if x is accessed three times, we have $\text{ALG}_{yx}(\pi_j) = 3$. Since both cases occur with probability $1/2$, we have $\mathbb{E}[\text{ALG}_{yx}(\pi_j)] = 2$.

Now, if y is accessed once, we have $\Delta_y = 1 - 1 = 0$. If y is accessed three times, OFF moves y ahead of x so that the relative order of x and y becomes the same in both lists. Note again that ALG cannot move x ahead of y without paid exchanges. Thus $\Delta_y = 0 - 1 = -1$ in this case. Since both cases occur with equal probabilities, we obtain $\mathbb{E}[\Delta_y] = -1/2$.

Overall, we have $\mathbb{E}[\Delta\Phi^{xy}(\pi_j) + \text{ALG}_{xy}(\pi_j) + \text{ALG}_{yx}(\pi_j)] = 0 - 1/2 + 0 + 2 = 3/2$.

Observe that the above proof estimates the expected cost of ALG by only taking into account the relative positions of pairs of elements. In particular, we do not use in any way that the pairwise orders of the items must be consistent, i.e., a total order. This means that the proof applies even for algorithms that do not need to maintain an ordered list, but only an (arbitrary) relation between pairs of elements. For this kind of algorithms, paid exchanges can be eliminated easily, by deferring each exchange that moves item x ahead of item y to the next time x is requested, without increasing the total cost. Thus our assumption that ALG does not use paid exchanges is without loss of generality. Note that if we require ALG to maintain a total order, this transformation may not be possible, since the order of paid actions cannot be changed. \square

Remark 4.20. Note the gap between the bounds of Theorems 4.18 and 4.19. Closing this gap is a prominent open problem in online optimization. A lower bound of 1.6 is known for algorithms fulfilling the pairwise property, in particular for COMB, see Ambühl et al. [2013].



5 Metrical Task Systems*

In the previous chapters, we have designed competitive online algorithms that were tailored for specific online problems. Ideally, we would like to generalize these algorithms to be applicable to a wide variety of settings, i.e., to design competitive algorithms for a large class of online problems. Of course, larger classes generally only allow for worse competitive ratios, since we empower the adversary by allowing additional instances. In particular, we have to limit our scope if we want online algorithms with bounded competitive ratios.

Abstractly, all problems we have studied so far have in common that the underlying system (e.g., the cache, the list, ...) is in some *state* after every step, and that the cost in each step only depends on the newest request, the current state (e.g., the position of the requested item on the list), and the state transition in the step (e.g., number of transpositions required to change the order of the list). In contrast, the cost in a step of a request-answer game can depend arbitrarily of previous actions and requests. We define the abstract framework of a *metrical task system* that captures the structure outlined above, with the additional requirement that the costs of state transitions define a metric.

Definition 5.1. A *metric space* \mathcal{M} is a pair (S, d) , where S is a set of points and $d : S \times S \rightarrow \mathbb{R}_{\geq 0}$ is a metric distance function that satisfies (for pairwise distinct $s, s', s'' \in S$):

1. $d(s, s) = 0$ (reflexivity)
2. $d(s, s') > 0$ (positivity)
3. $d(s, s') = d(s', s)$ (symmetry)
4. $d(s, s') + d(s', s'') \geq d(s, s'')$ (triangle inequality)

A metrical task system consists of a set of states with (metric) transition costs, as well as a cost associated with responding to a request (or *task*) in a given state.

Definition 5.2. A *metrical task system* is a pair $(\mathcal{M}, \mathcal{R})$, where $\mathcal{M} = (S, \tau)$ is a metric space and \mathcal{R} is a set of *tasks* $r : S \rightarrow \mathbb{R}_{\geq 0} \cup \{\infty\}$. We say that S are *states* and τ are the *transition costs* between states, and we set $\tau_{\min} := \min_{s \neq s'} \tau(s, s')$ and $\tau_{\max} := \max_{s, s' \in S} \tau(s, s')$.

A metrical task system naturally induces the online problem where tasks arrive online:

METRICAL TASK SYSTEM	
given:	metrical task system $(\mathcal{M} = (S, \tau), \mathcal{R})$, initial state $s_0 \in S$
online:	sequence $\sigma = (r_1, \dots, r_n) \in \mathcal{R}^n$ of tasks requested in this order
actions (step i):	transition from state s_{i-1} to state s_i (cost: $\tau(s_{i-1}, s_i)$)
objective:	minimize the total cost $\sum_{i=1}^n \text{ALG}_i(\sigma) = \sum_{i=1}^n [\tau(s_{i-1}, s_i) + r_i(s_i)]$

Example 5.3. We can model the paging problem by setting $S = \binom{P}{k}$ to be every possible cache configuration, letting $\tau(C, C') = |C' \setminus C|$ be the number of loads needed to transition from cache $C \in S$ to cache $C' \in S$, and modeling a request for a page $p_i \in P$ in step i by the task

$$r_i(C) = \begin{cases} 0 & \text{if } p_i \in C \\ \infty & \text{otherwise,} \end{cases}$$

which forces $p_i \in C$ at the end of step i .

5.1 Lower bound for metrical task systems

We start our analysis of the abstract class of online problems induced by metrical task systems by giving a general lower bound for the competitive ratio of any online algorithm.

Theorem 5.4 (Borodin et al. [1992]). No deterministic online algorithm for metrical task systems with N states has a competitive ratio smaller than $(2N - 1)$.

Proof. In order to derive a lower bound on the best possible competitive ratio for any fixed metric space (S, τ) of size $N > 1$, we need to adopt the perspective of the adversary. An obvious candidate for an adversarial request sequence $\sigma = (r_1, \dots, r_n)$ for a given algorithm ALG presents tasks such that ALG happens to be in the worst possible state for each task. Introducing some parameter $\varepsilon > 0$, we can achieve this by setting

$$r_i(s) = \begin{cases} \varepsilon, & \text{if } s = s_{i-1}, \\ 0, & \text{else,} \end{cases}$$

where s_{i-1} is the state of ALG at the start of step i .

Assume that ALG changes its state $n' \leq n$ times overall, and let $(s_0 = s_{i_0}), s_{i_1}, \dots, s_{i_{n'}}$ be the states that ALG transitions to, in this order. Then

$$\text{ALG}(\sigma) = (n - n')\varepsilon + \sum_{j=1}^{n'} \tau(s_{i_{j-1}}, s_{i_j}). \quad (5.1)$$

Observe that, by positivity of τ and ε , we have $\lim_{n \rightarrow \infty} \text{ALG}(\sigma) = \infty$. If $\lim_{n \rightarrow \infty} \text{OPT}(\sigma)$ is bounded, then ALG has an unbounded competitive ratio. In the following, we can therefore assume that $\lim_{n \rightarrow \infty} \text{OPT}(\sigma) = \infty$.

We bound the cost of OPT using the averaging technique. We use a class of algorithms \mathcal{A} that depends on the behavior of ALG . Specifically, \mathcal{A} consists of $(2N - 1)$ algorithms, such that, after every step, exactly two algorithms of \mathcal{A} are in each state *not* occupied by ALG and exactly one algorithm is in the same state as ALG , while at most one algorithm in \mathcal{A} switches state in every step. Observe that we can achieve this by switching one algorithm from state s_i to state s_{i-1} whenever ALG transitions from s_{i-1} to s_i . Let $B_0 = B_0(\mathcal{M}) > 0$ denote the total (constant) setup cost required in step 1 to distribute the algorithms in \mathcal{A} over starting states according to the system we just described.

To compute the average cost $c_{\text{avg}}(\sigma)$ of the algorithms in \mathcal{A} , observe that the algorithms in state s_{i-1} after step i have non-zero cost ε , while all other algorithms have cost 0. If ALG changes state away from s_{i-1} in step i , two algorithms in \mathcal{A} have cost ε , otherwise only one algorithm has cost ε . Also, whenever ALG transitions between states, so does a single algorithm in \mathcal{A} , otherwise there are no state transitions. Since ALG transitions exactly n' times between states, we get

$$c_{\text{avg}}(\sigma) = \frac{1}{2N - 1} \left[(n - n')\varepsilon + \sum_{j=1}^{n'} \tau(s_{i_j}, s_{i_{j-1}}) + 2n'\varepsilon + B_0 \right].$$

With eq. (5.1) this gives

$$\text{OPT}(\sigma) \leq c_{\text{avg}}(\sigma) = \frac{2n'\varepsilon + \text{ALG}(\sigma) + B_0}{2N - 1}. \quad (5.2)$$

With $\text{ALG}(\sigma) \geq \sum_{j=1}^{n'} \tau(s_{i_{j-1}}, s_{i_j}) \geq n' \tau_{\min}$, this yields

$$(2\varepsilon/\tau_{\min} + 1) \cdot \text{ALG}(\sigma) \geq 2n'\varepsilon + \text{ALG}(\sigma) \stackrel{(5.2)}{\geq} (2N - 1) \cdot \text{OPT}(\sigma) - B_0.$$

Observe that $\lim_{n \rightarrow \infty} \text{ALG}(\sigma) = \lim_{n \rightarrow \infty} \text{OPT}(\sigma) = \infty$. We obtain

$$\frac{\text{ALG}(\sigma)}{\text{OPT}(\sigma)} \geq \frac{2N - 1}{1 + 2\varepsilon/\tau_{\min}} - \frac{B_0}{(1 + 2\varepsilon/\tau_{\min}) \cdot \text{OPT}(\sigma)} \xrightarrow{\varepsilon \rightarrow 0} (2N - 1) - B_0/\text{OPT}(\sigma) \xrightarrow{n \rightarrow \infty} 2N - 1.$$

By Proposition 1.12, this implies that ALG cannot be better than $(2N - 1)$ -competitive for all $\varepsilon > 0$. \square

Remark 5.5. Observe that this lower bound implies that if we solve the paging problem with an algorithm for metrical task systems applied to the metrical task system of Example 5.3, we can only provide very weak bounds on the competitive ratio, since, generally, $N = \binom{|P|}{k} \gg k$.

5.2 Work function algorithm

We now design an online algorithm W_{FA} for metrical task systems using the paradigm of a *work function algorithm*. The principal idea of this approach is to, in every step, explicitly compute the final state of an optimum solution for the request sequence so far, and try to stay close to this state. Obviously, we need a way to compute optimum solutions to use this scheme.

Let $w_i(s)$ be the optimum offline cost required to first process the request sequence $\sigma_{\leq i} = (r_1, \dots, r_i)$ and then transition to state $s \in S$, and, in particular, $\text{OPT}(\sigma) = \min_{s \in S} w_n(s)$. We can compute this cost via dynamic programming:

$$\begin{aligned} w_0(s) &= \tau(s_0, s), \\ w_i(s) &= \min_{x \in S} \{w_{i-1}(x) + r_i(x) + \tau(x, s)\}. \end{aligned} \quad (5.3)$$

Ideally, we want W_{FA} to always be in a state s that minimizes $w_i(s)$. However, if $\tau(s_{i-1}, s)$ is large and $w_i(s_{i-1})$ is not much worse than $w_i(s)$, it may not be worth it to transition to s . Instead, we let W_{FA} move to a state

$$s_i \in \arg \min_{x \in S} \{w_i(x) + \tau(s_{i-1}, x)\}. \quad (5.4)$$

In fact, this choice leaves some freedom. The following lemma states that we can always select s_i such that there is an offline optimum for first serving the request sequence $\sigma_{\leq i}$ and then moving to s_i , such that the last request is already served in state s_i . We define W_{FA} such that it selects s_i accordingly.

Lemma 5.6. W_{FA} can select s_i such that

$$w_i(s_i) = w_{i-1}(s_i) + r_i(s_i).$$

Proof. Take any state $s \in \arg \min_{x \in S} \{w_i(x) + \tau(s_{i-1}, x)\}$ and let $x' \in \arg \min_{x \in S} \{w_{i-1}(x) + r_i(x) + \tau(x, s)\}$ be a state that minimizes eq. (5.3) for s , i.e.,

$$w_i(s) = w_{i-1}(x') + r_i(x') + \tau(x', s). \quad (5.5)$$

We claim that we can set $s_i = x'$. To show this, we need to show that (i) $x' \in \arg \min_{x \in S} \{w_i(x) + \tau(s_{i-1}, x)\}$ and (ii) $w_i(x') = w_{i-1}(x') + r_i(x')$.

To see (i), we observe that, by definition of w_i (eq. (5.3)), we have $w_i(x') \leq w_{i-1}(x') + r_i(x')$, and thus

$$\begin{aligned}
 w_i(x') + \tau(s_{i-1}, x') &\stackrel{(5.3)}{\leq} w_{i-1}(x') + r_i(x') + \tau(s_{i-1}, x') \\
 &\stackrel{(5.5)}{=} w_i(s) - \tau(x', s) + \tau(s_{i-1}, x') \\
 &\stackrel{\Delta\text{-ineq.}}{\leq} w_i(s) + \tau(s_{i-1}, s).
 \end{aligned} \tag{5.6}$$

To see (ii), we observe that, by definition of s , we have $w_i(s) + \tau(s_{i-1}, s) \leq w_i(x') + \tau(s_{i-1}, x')$, and hence eq. (5.6) holds with equality. But then, all steps in the derivation of eq. (5.6) also hold with equality, and, in particular, the first step yields

$$w_i(x') = w_{i-1}(x') + r_i(x'). \quad \square$$

Recall that the proof of Theorem 5.4 relied on a class of algorithms \mathcal{A} that depends on the online algorithm. In particular, the proof used averaging, i.e., the fact that $\text{OPT}(\sigma)$ is bounded by the average cost over the $(2N - 1)$ algorithms in \mathcal{A} . To get a tight upper bound for W_{FA} , we compare the cost incurred by W_{FA} to the *total* cost of the algorithms in \mathcal{A} , which is obviously at most $(2N - 1)$ times the average cost. Since, after step i , \mathcal{A} has exactly two algorithms in every state $s \neq s_i$ and exactly one algorithm in state s_i , we can compute a lower bound B_i on the total cost up to step i by pretending that every algorithm in \mathcal{A} chose an optimal way of reaching its final state. We get

$$B_i = w_i(s_i) + 2 \sum_{s \neq s_i} w_i(s). \tag{5.7}$$

We can now show that the cost of W_{FA} in step i is bounded from above by the total cost in step i of all algorithms in \mathcal{A} .

Lemma 5.7. For all $i \in \{1, \dots, n\}$ we have $W_{\text{FA}_i}(\sigma) \leq B_i - B_{i-1}$.

Proof. The cost of W_{FA} in step i is defined as

$$W_{\text{FA}_i}(\sigma) = \tau(s_{i-1}, s_i) + r_i(s_i).$$

We separately bound both terms. For $\tau(s_{i-1}, s_i)$, we use the definition of s_i to obtain

$$\begin{aligned}
 w_i(s_i) + \tau(s_{i-1}, s_i) &\stackrel{(5.4)}{=} \min_{x \in S} \{w_i(x) + \tau(s_{i-1}, x)\} \\
 &\leq w_i(s_{i-1}) + \tau(s_{i-1}, s_{i-1}) \\
 &= w_i(s_{i-1}),
 \end{aligned}$$

which implies

$$\tau(s_{i-1}, s_i) \leq w_i(s_{i-1}) - w_i(s_i).$$

For $r_i(s_i)$, by Lemma 5.6,

$$r_i(s_i) = w_i(s_i) - w_{i-1}(s_i).$$

Together, we get

$$W_{\text{FA}_i}(\sigma) \leq w_i(s_{i-1}) - w_{i-1}(s_i). \quad (5.8)$$

On the other hand, using $w_i(s) \geq w_{i-1}(s)$, we have

$$\begin{aligned} B_i - B_{i-1} &= 2 \sum_{s \neq s_i} w_i(s) + w_i(s_i) - 2 \sum_{s \neq s_{i-1}} w_{i-1}(s) - w_{i-1}(s_{i-1}) \\ &= 2 \sum_{s \notin \{s_i, s_{i-1}\}} (w_i(s) - w_{i-1}(s)) + 2w_i(s_{i-1}) + w_i(s_i) - 2w_{i-1}(s_i) - w_{i-1}(s_{i-1}) \\ &\geq (2w_i(s_{i-1}) - w_{i-1}(s_{i-1})) + (w_i(s_i) - 2w_{i-1}(s_i)) \\ &\geq w_i(s_{i-1}) - w_{i-1}(s_i) \\ &\stackrel{(5.8)}{\geq} W_{\text{FA}_i}(\sigma), \end{aligned}$$

as claimed. \square

Lemma (5.7) implies that W_{FA} is not more expensive than the total cost of all algorithms in the class \mathcal{A} . This means that on the adversarial request sequence that we constructed for the lower bound, W_{FA} is only a factor of at most $(2N - 1)$ away from the bound on the optimal offline cost we constructed via averaging. We show that this holds for all request sequences and all bounds for $\text{OPT}(\sigma)$.

Theorem 5.8 (Borodin et al. [1992]). W_{FA} is $(2N - 1)$ -competitive for metrical task systems with N states.

Proof. Using Lemma 5.7, and setting $\alpha = \alpha(\mathcal{M}) := (2N - 1) \max_{x, y \in S} \tau(x, y) - B_0(\mathcal{M})$, we obtain

$$\begin{aligned} W_{\text{FA}}(\sigma) &= \sum_{i=1}^n W_{\text{FA}_i}(\sigma) \\ &\stackrel{\text{Lem. 5.7}}{\leq} B_n - B_0 \\ &\stackrel{(5.7)}{\leq} (2N - 1) \max_{s \in S} w_n(s) - B_0 \\ &\stackrel{(5.3)}{=} (2N - 1) \max_{s \in S} \min_{x \in S} (w_{n-1}(x) + r_n(x) + \tau(x, s)) - B_0 \\ &\leq (2N - 1) \min_{x \in S} (w_{n-1}(x) + r_n(x)) + (2N - 1) \max_{x, y \in S} \tau(x, y) - B_0 \\ &\leq (2N - 1) \min_{x, s \in S} (w_{n-1}(x) + r_n(x) + \tau(x, s)) + \alpha \\ &\stackrel{(5.3)}{=} (2N - 1) \min_{s \in S} w_n(s) + \alpha \\ &= (2N - 1) \cdot \text{OPT}(\sigma) + \alpha \end{aligned}$$

\square



6 The k -Server Problem

In the last chapter, we considered metrical task systems as a broad, abstract class of online problems that feature a system state with metric transition costs between states. Aside from this restriction, the costs of a request in a metrical task system may vary arbitrarily between states. As a result, the competitive ratio for metrical task systems depends on the total number of states of the system, which only allows for very weak bounds in general. In this chapter, we consider a subclass of metrical task systems with a more restrictive cost structure that, aside from metric transition costs, only allows to distinguish between permitted and forbidden states (cost 0 or ∞). More specifically we assume that the system state is comprised of the individual states of k servers, and a request is given by a state that must be occupied by at least one of the servers. Because of the geometric interpretation, we speak of points of the metric space instead of states.

k -SERVER PROBLEM	
given:	metric space (P, d) , initial server locations $S_0 = (S_{0,1}, \dots, S_{0,k}) \in P^k$
online:	sequence $\sigma = (r_1, \dots, r_n) \in P^n$ of points
actions:	for every $j \in \{1, \dots, k\}$: move server j to $S_{i,j}$ (cost: $d(S_{i-1,j}, S_{i,j})$) (step i) afterwards there must be a server j with $S_{i,j} = r_i$
objective:	minimize the total cost $\sum_{i=1}^n \text{ALG}_i(\sigma) = \sum_{i=1}^n \sum_{j=1}^k d(S_{i-1,j}, S_{i,j})$

We observe that k -server problems are a restriction of metrical task systems that still covers the paging problem (for example).

Remark 6.1. The k -server problem can be formulated as a metrical task system where S_i is the system state after step i and a request for point r_i is translated into a cost function that assigns S_i a cost of 0 if there is a server j with $S_{i,j} = r_i$ and ∞ otherwise. With this formulation, Theorem 5.8 gives us an exponential upper bound of $(2N^k - 1)$ on the competitive ratio for the k -server problem for metric spaces with N points. In the following, we derive better bounds tailored to the k -server problem.

Example 6.2. The k -server problem is still abstract enough to encompass the (k -)paging problem. We can see this by letting the set of points P be identical to the set of pages and interpreting the server locations as the set of pages that are in the cache. To obtain cost 1 for every page load, we simply set the distance between any pair of points/pages to 1.

Remark 6.3. Recall that for the paging problem we could restrict ourselves to demand paging algorithms that only evict/load (at most) one page at a time, and only upon a page fault. Very similarly, for the k -server problem, we can restrict ourselves to *lazy* algorithms that, in step i , only move the server that serves request r_i . We can transform any algorithm to be lazy as follows: If a server s moves without serving a request, we instead keep the server stationary and only pretend that the server moved, and remember that we still need to execute the move later. Once server s needs to move to serve a request, we first execute all remembered moves, which brings the server to its supposed location, and then we execute the move to serve the request. The cost of the modified algorithm may be smaller than that of the original algorithm, but not larger.

A straight-forward lazy algorithm for the k -server problem is the GREEDY algorithm that always moves the server that is closest to the requested point with respect to the distance function d . To see that this algorithm is not competitive, consider a metric space with three points $P = \{a, b, c\}$ where $d(a, b) < \min\{d(a, c), d(b, c)\}$, and two servers initially at $S_0 = (c, a)$ (see Figure 6.1). For the request sequence $\sigma = (b, a, b, a, b, a, \dots) \in P^n$, GREEDY keeps one of the servers at c forever and incurs a cost of $\text{GREEDY}(\sigma) = n \cdot d(a, b)$, while $\text{OPT}(\sigma) = d(c, b)$ is independent of n .




Figure 6.1: Linear embedding of the bad example for GREEDY with two servers initially at $S_0 = (a, c)$ and the request sequence $\sigma = (b, a, b, a, \dots)$.

6.1 Lower bound for the k -server problem

The k -server problem becomes trivial for finite metric spaces with $|P| = k$ once servers have been distributed. As soon as P contains more than k points, the lower bound of k for paging (Thm. 2.4) carries over, by Example 6.2. We now show that this is the case independently of the metric space, i.e., also beyond paging.

Theorem 6.4 (Manasse et al. [1990]). No deterministic online k -server algorithm has a competitive ratio lower than k on **any** metric space (P, d) with at least $k + 1$ points.

Proof. For any fixed (lazy) online algorithm ALG and any metric space (P, d) with at least $k + 1$ points, we construct an adversarial request sequence $\sigma = (r_1, \dots, r_n) \in P^n$ for ALG. Without loss of generality, we may assume that P consists of exactly $k + 1$ points, since we can simply ignore any additional point. A natural idea is to choose start with servers distributed over k distinct points of P , and let σ always request the unique point not occupied by ALG. Since ALG is lazy, in step $i < n$, it moves exactly one server, and, by construction of σ , this move is from r_{i+1} to r_i . Let r_{n+1} denote the position of the server that serves request r_n at the start of step n . Then,

$$\text{ALG}(\sigma) = \sum_{i=1}^n d(r_{i+1}, r_i), \quad (6.1)$$

and, in particular, by positivity of d , we have $\lim_{n \rightarrow \infty} \text{ALG}(\sigma) = \infty$.

We derive a bound on $\text{OPT}(\sigma)$ using the averaging technique. For this, we need to define a class \mathcal{A} of simple algorithms, i.e., simple rules which server to select for serving request r_i . If there is a server at r_i , we obviously use this server. Otherwise, r_i is the only unoccupied point. The most natural candidate to serve r_i is the server at the closest point to r_i , but this gives the GREEDY algorithm, which we already showed to be bad. Another server that stands out is the server at r_{i-1} (for $i > 1$). In order to define a class of algorithms that follow the simple rule of serving each request for an unoccupied point r_i with the server at r_{i-1} , we define an algorithm for every possible set S_1 of (distinct) positions of all servers at the end of step 1. Of course, we must have $r_1 = S_{1,j}$ for some $j \in \{1, \dots, k\}$, therefore, since $|P| = k + 1$, there are exactly $\binom{|P \setminus \{r_1\}|}{k-1} = k$ different starting configurations (we do not distinguish individual servers). Since no two servers initially share a position, the average cost to put the algorithms in their configuration in step 1 is a constant $\alpha = \alpha(P, d) \leq \max_{p, q \in P} d(p, q)$. Note that the algorithms in \mathcal{A} are not lazy (in step 1), but could be transformed to be lazy.

We claim that no two algorithms can ever be in the same configuration. This is true initially for the request sequence introduced above. If it is the case until step $i - 1$, then, at the start of step $i \geq 2$, there is exactly one algorithm in \mathcal{A} that does not have a server at r_i already, and all algorithms have a server at r_{i-1} . The only algorithm that needs to move a server in step i (from r_{i-1} to r_i) is the only one without a server at position r_{i-1} after step i . Our claim follows by induction.

Since exactly one algorithm needs to move in every step, the total cost of the algorithms in \mathcal{A} is $k\alpha + \sum_{i=2}^n d(r_{i-1}, r_i)$. The optimum cost is bounded by the average cost $c_{\text{avg}} = \alpha + \frac{1}{k} \sum_{i=2}^n d(r_{i-1}, r_i)$. With eq. (6.1), and by symmetry and positivity of d , we get

$$\frac{\text{ALG}(\sigma)}{\text{OPT}(\sigma)} \geq \frac{\sum_{i=1}^n d(r_i, r_{i+1})}{\alpha + \frac{1}{k} \sum_{i=2}^n d(r_{i-1}, r_i)} \geq \frac{k}{1 + k\alpha / \sum_{i=2}^n d(r_{i-1}, r_i)} \xrightarrow{n \rightarrow \infty} k.$$

Using $\lim_{n \rightarrow \infty} \text{ALG}(\sigma) = \infty$, Proposition 1.12 thus excludes any competitive ratio below k . \square

Based on this lower bound and matching upper bounds for some special cases that we will see below, Manasse, McGeoch and Sleator posed their famous “ k -server conjecture”:

Conjecture 6.5 (Manasse et al. [1990]). *There is a deterministic k -competitive k -server algorithm.*

Proving or disproving this conjecture for general metric spaces is an important open problem in online optimization. In the following sections, we prove the k -server conjecture for specific metric spaces.

6.2 The k -server problem on the line

We first consider the simple metric space (P, d) of the (Euclidean) line, i.e., with $P = \mathbb{R}$ and where d is the Euclidean distance between points. Recall the bad instance for the GREEDY algorithm we discussed above (Figure 6.1). Since this example is already embedded on the line, this means that we need to be more careful about how to decide which server to move. If $d(b, c) \gg d(a, b)$, the first requests must be served by the server at a , otherwise we are not competitive on short sequences. But at some point we have to involve the server at c to avoid an unbounded cost compared to the optimum as instances grow.

One (non-lazy) idea how to improve GREEDY’s behavior in the example of Figure 6.1 would be to move both servers towards a request at b and let the server serve the request that first reaches it. The algorithm Dc (double cover) generalizes this idea to any number of servers by moving a closest server towards request r_i from both sides if such a server exists. We call the two servers closest to r_i from either side its neighbors, and if there are multiple closest servers on one side, we pick only one of them arbitrarily as the neighbor of r_i . Then, in step i , Dc moves the (at most two) neighbors of r_i towards r_i with the same speed until (at least) one of them reaches r_i . Observe that Dc behaves as intended on the example of Figure 6.1: The first requests are served by the server starting at a , and at some point the server that started at c reaches b and stays there. Eventually, the servers are located at a and b and do not need to move anymore.

Theorem 6.6 (Chrobak et al. [1991]). Dc is k -competitive for the k -server problem on the line if $\sum_{j \neq j'} d(S_{0,j}, S_{0,j'})$ is bounded by a constant.

Proof. We analyze Dc using the potential function method. Recall that the principal idea of defining a potential function is to allow investing additional cost in cheap steps in order to pay for expensive steps. One reason why a k -server request on the line can be costly for Dc, in terms of relative cost compared to OPT, is that the server positions of Dc and OPT are very different. We can measure this by the cost $M_i(\sigma)$ of a minimum (distance) matching between the server positions of Dc and OPT at the end of step i . Note that this matching matches the j -th server from the left of Dc to the j -th server from the left of OPT for $j \in \{1, \dots, k\}$. A reason why a request can be costly in terms of absolute cost is that Dc’s servers are far apart. We can measure this by the sum $\Sigma_i(\sigma) := \sum_{j \neq j'} d(S_{i,j}, S_{i,j'})$ of pairwise distances of Dc’s servers at the end of step i . With this intuition, we define a potential function of the form

$$\Phi_i(\sigma) = a_1 M_i(\sigma) + a_2 \Sigma_i(\sigma),$$

with some parameters $a_1, a_2 \geq 0$.

We want to use the alternating moves approach (Corollary 4.3), i.e., we pretend that OPT moves before Dc in step i and let $\Phi'_i(\sigma)$ denote the intermediate potential function value after OPT moved. By Corollary 4.3, to show that Dc is ρ -competitive, it is sufficient to show that (i) $\Phi_i(\sigma) \geq 0$, (ii) there is a constant $\beta \geq 0$ independent of σ with $\Phi_0(\sigma) \leq \beta$, (iii) $\Delta \Phi_i^{\text{OPT}} := \Phi'_i(\sigma) - \Phi_{i-1}(\sigma) \leq \rho \cdot \text{OPT}_i(\sigma)$, and (iv) $\Delta \Phi_i^{\text{Dc}} := \Phi_i(\sigma) - \Phi'_i(\sigma) \leq -\text{Dc}_i(\sigma)$.




Figure 6.2: Illustration of the proof of Theorem 6.6, Case 2. If OPT's server at r_i (crossed out) is not matched to a neighboring server (wlog to its right), there is a matching of at most the same cost, where it is (dashed).

Properties (i) and (ii) hold as long as $S_0 \in P^k$ is fixed, because $\Phi_i \geq 0$ and $\Phi_0 \leq \beta(S_0) := \sum_{j \neq j'} d(S_{0,j}, S_{0,j'})$.

For Property (iii), we may assume that OPT is lazy, i.e., that it moves at most one server in every step i . Assume that OPT moves a server by $\text{OPT}_i(\sigma) = d \geq 0$. This move may increase $M_i(\sigma)$ by at most d , even if we maintain the same (not necessarily minimum) matching. Then Property (iii) holds if $\Delta\Phi_i^{\text{OPT}} \leq a_1 d \leq \rho \cdot \text{OPT}_i(\sigma) = \rho d$, i.e., if $a_1 \leq \rho$. For Property (iv), we distinguish two cases in step i .

Case 1: Dc moves a single server by $d > 0$. Then, by definition of Dc, this server must lie on the boundary of the convex hull of all server positions, and the requested point must lie outside of the convex hull. This means that the server moves away from all other servers and $\Sigma_i(\sigma)$ increases by exactly $(k-1)d$. Also, there is a minimum matching that matches the server of Dc at r_i after the move to the server at r_i in OPT. This must also have been a minimum matching before the move, since there are no servers closer to r_i in Dc. Therefore $M_i(\sigma)$ is decreased by exactly d . Property (iv) holds in this case if $\Delta\Phi_i^{\text{Dc}} = a_2(k-1)d - a_1 d \leq -\text{Dc}_i(\sigma) = -d$, i.e., if $a_1 \geq a_2(k-1) + 1$.

Case 2: Dc moves two servers by $d \geq 0$ each. Then, by definition of Dc, two servers $j, j' \in \{1, \dots, k\}$ move towards each other with no other servers between them. This means that the contribution of the pair (j, j') to Σ_i decreases by exactly $2d$, and the changes in distance from j to every other server cancels with the change in distance from j' . During the move, there is a minimum matching that matches j or j' to the server of OPT at r_i , since there are no servers of Dc between j and j' (see Figure 6.2). Consequently, $M_i(\sigma)$ does not increase during the move. Property (iv) holds in this case if $\Delta\Phi_i^{\text{Dc}} \leq -a_2 2d \leq -2d = -\text{Dc}_i(\sigma)$, i.e., if $a_2 \geq 1$.

Overall, all properties hold if

$$\rho \geq a_1 \geq a_2(k-1) + 1 \geq (k-1) + 1 = k,$$

and we get equality for $a_1 = \rho = k$ and $a_2 = 1$, i.e., for the potential $\Phi_i(\sigma) = kM_i(\sigma) + \Sigma_i(\sigma)$. \square

6.2.1 Double Cover on trees

There is a direct generalization for Dc from the real line to the metric space induced by a tree $T = (V, E)$ with edge lengths, where the distance between two vertices is defined by the metric closure (i.e., the shortest path distance). Let $v \in V$ be a vertex of T , consider a server j located at a vertex v_j , and let P be the unique v_j - v -path in T . If there is no other server located along $P \setminus \{v_j\}$ and if there is no server $j' \in \{1, \dots, j-1\}$ located at v_j , we call j a *neighbor* of v . We generalize Dc by continuously moving *all* neighbors of v towards each request for a vertex $v \in V$, until the first server reaches v (see Figure 6.3). Note that the set of neighbors can become smaller over time. Also note that, since the vertices of T are the points of the metric space, we cannot continuously move servers along edges. Similarly to the transformation into a lazy algorithm, we can avoid this issue by executing all movements of a server only virtually, until the server actually serves a request at a vertex.




Figure 6.3: Server movements over time (left to right) in Dc for trees when the central node is requested. Rectangles correspond to servers. Note that the set of neighbors (thick) changes over time.

Theorem 6.7 (Chrobak and Larmore [1991b]). Dc is k -competitive for the k -server problem on a tree if $\sum_{j \neq j'} d(S_{0,j}, S_{0,j'})$ is bounded by a constant.

Proof. We use the same potential as in the proof of Theorem 6.6, i.e.,

$$\Phi_i(\sigma) = kM_i(\sigma) + \Sigma_i(\sigma).$$

As before, we can still assume OPT to be lazy, i.e., to move a single server by d in step i , and therefore $\Delta\Phi_i^{\text{OPT}} \leq kd = k \cdot \text{OPT}_i(\sigma)$.

Now subdivide Dc's move in step i into time periods $(t_1, t_2], (t_2, t_3], \dots, (t_{N-1}, t_N]$, such that the set of servers that move stays the same during each period. Consider a period $(t_j, t_{j+1}]$ where m_j servers move by $d_j = t_{j+1} - t_j$ each. Observe that, during the whole period, there is a minimum matching where one of the moving servers is matched to OPT's server at the requested node $r_i \in P$. This means that M_i can increase by at most $(m_j - 1)d_j - d_j = m_j d_j - 2d_j$, which causes an increase in potential by at most $km_j d_j - 2kd_j$. Each stationary server gets closer to at least $m_j - 1$ other servers and further away from at most one other server. The contribution to Σ_i of all these changes in pairwise distance between one of the $k - m_j$ stationary servers and a moving server sum to

$$-(k - m_j)((m_j - 1)d_j - d_j) = -(k - m_j)(m_j d_j - 2d_j).$$

Finally, the pairwise distances between moving servers each decrease by exactly $2d_j$, which changes Σ_i by $-\frac{m_j(m_j-1)}{2} \cdot 2d_j = -m_j(d_j m_j - d_j)$.

Overall, the change in potential during one period is at most

$$km_j d_j - 2kd_j - (k - m_j)(m_j d_j - 2d_j) - m_j(d_j m_j - d_j) = -m_j d_j,$$

which corresponds exactly to the total distance traveled in the period. Summed over all periods in step i , we get

$$\Delta\Phi_i^{\text{Dc}} = -\sum_{j=1}^{N-1} m_j d_j = -\text{Dc}_i(\sigma).$$

Together with our bound on $\Delta\Phi_i^{\text{OPT}}$ this gives an amortized cost of

$$a_i(\sigma) = \text{Dc}_i(\sigma) + \Delta\Phi_i(\sigma) = k \cdot \text{OPT}_i(\sigma).$$

We also have $\Phi_i \geq 0$ and $\Phi_0 \leq \frac{k(k-1)}{2} \max_{v, v' \in V} d(v, v')$, which is constant if the tree is fixed (and finite). We can therefore apply the potential function method (Corollary 4.3) to prove the theorem. \square

Remark 6.8. Note that the (k) -paging problem can be modeled as the k -server problem on a star-shaped tree with one vertex per memory page and one server per cache slot (see Figure 6.4). The positions of the servers correspond to the pages in the cache, evicting a page corresponds to moving the respective server to the center node, and loading a page corresponds to moving a server to the respective node. Note that in our model, costs are split between loading and evicting, but this makes no difference for demand paging algorithms, provided that the cache is full initially. On the star associated with a paging instance, D_C becomes equivalent to the paging algorithm F_{WF} . This means that k -competitiveness of D_C on trees implies k -competitiveness of F_{WF} for paging. It is worth emphasizing that the k -server framework captures the paging problem tightly enough to allow tight bounds to be derived via the much more general setting.




Figure 6.4: Illustration of the (k) -paging problem expressed as a k -server problem.

6.2.2 Double Cover on finite metric spaces

We can further generalize D_C to any finite metric space with N points, by interpreting such a space as a complete graph $G = (V, E)$ with a vertex for every point of the space and edge lengths that correspond to the distances between every pair of points. We can apply our adaptation of D_C for trees by restricting ourselves to a minimum spanning tree T of G . This means that we forget all edges not in T and serve all requests by only moving along T . By definition of minimum spanning trees, every edge e of $G \setminus T$ closes a cycle C in T , and e must be a longest edge of C . If e has length d , this means that $C \setminus \{e\}$ has length at most $(N - 1)d$, i.e., there is a detour in T for e that has length at most $(N - 1)d$. Applying this argument multiple times yields that every shortest path of length d' in G has a detour of length at most $(N - 1)d'$ in T .

Let OPT_{tree} be the offline optimum cost when restricted to moving along edges of T . Then, by Theorem 6.7, there is $\alpha \in \mathbb{R}$ with

$$D_C(\sigma) \leq k \cdot \text{OPT}_{\text{tree}}(\sigma) + \alpha \leq k(N - 1) \cdot \text{OPT}(\sigma) + \alpha,$$

and we obtain the following.

Theorem 6.9 (Chrobak and Larmore [1991b]). D_C is $(Nk - k)$ -competitive for the k -server problem on any finite metric space with N points if $\sum_{j \neq j'} d(S_{0,j}, S_{0,j'})$ is bounded by a constant.

Note that, in general, this is much better than the competitive ratio of $2N^k - 1$ that follows from the fact that the k -server problem is a metrical task system.

6.3 Balancing algorithm*

For arbitrary metric spaces, one of the most natural online algorithms is the algorithm BAL that greedily balances the distances moved by the servers. Let $D_j^i = \sum_{i'=1}^i d(S_{i'-1,j}, S_{i',j})$ denote the total distance

that server j moved until the end of step i . Then BAL serves the request r_i in step i with any server $j \in \arg \min_{j' \in \{1, \dots, k\}} \{D_{j'}^{i-1} + d(S_{i-1, j'}, r_i)\}$ that is closest to r_i among servers minimizing $D_{j'}^{i-1} + d(S_{i-1, j'}, r_i)$. In particular, if there is a server j at r_i , this server is selected, since $d(S_{i-1, j}, r_i) = 0$ and $D_j^{i-1} \leq D_{j'}^{i-1} + d(S_{i-1, j'}, r_i)$ for every $j' \in \{1, \dots, k\}$, because either $D_j^{i-1} = 0$ or because of the fact that j was selected earlier to move to r_i . In particular, if servers start at distinct locations, they always remain at distinct locations. In the following, we assume that this is the case.

Proposition 6.10. For every $k \geq 2$, BAL is not competitive for the k -server problem on metric spaces of size $N \geq k + 2$.

Proof. Consider the example in Figure 6.5 with $k = 2$ servers and $N = 4$ points. We define the adversarial input sequence $\sigma = (a, b, c, d, a, b, c, d, \dots) \in P^n$. For this sequence $\text{BAL}(\sigma) = n$, since one server serves all requests at $\{a, d\}$ and the other serves the requests at $\{b, c\}$. On the other hand $\text{OPT}(\sigma) = 1 + (n - 2)\varepsilon$ for $n \geq 3$, since one server serves all requests at $\{a, b\}$ and the other serves the requests at $\{c, d\}$. We immediately get

$$\frac{\text{BAL}(\sigma)}{\text{OPT}(\sigma)} = \frac{n}{1 + (n - 1)\varepsilon} \xrightarrow{\varepsilon \rightarrow 0} n.$$

Note that the example can be extended to larger k, N simply by adding additional points and servers far away from $\{a, b, c, d\}$. \square




Figure 6.5: Adversarial instance of Proposition 6.10 with points $P = \{a, b, c, d\}$ and initial server locations $S_0 = (c, d)$. The distances obey $d(a, b) = d(c, d) = \varepsilon \ll 1$.

Remark 6.11. In the following, we want to show that BAL is k -competitive on metric spaces with $N = k + 1$ points. Since it does not make sense for the adversary to make requests that do not incur a cost for BAL, we may assume that in every step the unique unoccupied location of BAL is requested. Let q_i be the unique unoccupied point at the end of step $i \geq 0$, i.e., $q_{i-1} := r_i$ for all $i \geq 1$. Similarly, let q_i^* be an unoccupied point after step i in $\text{OPT}[\sigma]$. For every point $p \in P$ that is occupied by a server $j \in \{1, \dots, k\}$ in BAL at the end of step i , we define $D_p^i := D_j^i$ to be the total distance travelled by the server at p . Finally, let $S_{i, j}^*$ denote the location of server $j \in \{1, \dots, k\}$ in $\text{OPT}[\sigma]$ at the end of step i .

We want to apply the potential function method with the alternating moves approach (Corollary 4.3), pretending that OPT moves before BAL in every step. We introduce a rather complicated potential function, which will turn out to work well later:

$$\Phi_i(\sigma) := \begin{cases} kD_{q_i^*}^i - \sum_{j=1}^k D_j^i & \text{if } q_i^* \neq q_i = r_{i+1}, \\ k(D_{r_i}^i - d(r_i, r_{i+1})) - \sum_{j=1}^k D_j^i & \text{if } q_i^* = q_i = r_{i+1}. \end{cases} \quad (6.2)$$

This potential function distinguishes between whether OPT and BAL occupy the same points ($q_i^* = q_i$) or not ($q_i^* \neq q_i$). At the beginning of step i , BAL has point $q_{i-1} = r_i$ unoccupied. After OPT moved, it must have a server at $q_{i-1} = r_i$, i.e., we are in the first case of the potential function. Since BAL's servers did not move yet, the intermediate potential value is

$$\Phi_i'(\sigma) := kD_{q_i^*}^{i-1} - \sum_{j=1}^k D_j^{i-1}.$$

In order to apply Corollary 4.3, we have to ensure that conditions (i)-(iv) are satisfied. For condition (i), it is sufficient to show that there is a constant $\alpha(\mathcal{M}) \geq 0$ such that for all $j \in \{1, \dots, k\}$ we have $kD_j^i - \sum_{j'=1}^k D_{j'}^i \geq -\alpha(\mathcal{M})$, where $\alpha(\mathcal{M})$ only depends on the metric space $\mathcal{M} = (P, d)$. Then $\Phi_i(\sigma) \geq -kd_{\max}(\mathcal{M}) - \alpha(\mathcal{M})$ where $d_{\max}(\mathcal{M}) := \max_{p, p' \in P} d(p, p')$ only depends on \mathcal{M} , and condition (i) holds. To see this, we may individually bound differences of the form $|D_{j'}^i - D_j^i|$ for $j, j' \in \{1, \dots, k\}$, because we need an upper bound for

$$\sum_{j'=1}^k D_{j'}^i - kD_j^i = \sum_{j'=1}^k (D_{j'}^i - D_j^i) \leq \sum_{j'=1}^k |D_{j'}^i - D_j^i|.$$

Claim 6.12. For all $j, j' \in \{1, \dots, k\}$ we have $|D_{j'}^i - D_j^i| \leq d(S_{i,j}, S_{i,j'})$.

With this claim, we can set $\alpha(\mathcal{M}) := k \cdot d_{\max}(\mathcal{M})$ to fulfill condition (i), since then $\sum_{j'=1}^k D_{j'}^i - kD_j^i \leq \sum_{j'=1}^k d(S_{i,j}, S_{i,j'}) \leq k \cdot d_{\max}(\mathcal{M})$ and, thus, $\Phi_i(\sigma) \geq -2kd_{\max}(\mathcal{M})$.

Proof of Claim (6.12). We prove the claim by induction on i . The induction base for $i = 0$ holds, since $|D_{j'}^0 - D_j^0| = 0 \leq d(S_{i,j}, S_{i,j'})$. For the inductive step, assume that $|D_{j'}^{i-1} - D_j^{i-1}| \leq d(S_{i-1,j}, S_{i-1,j'})$ holds. The claim trivially holds if j, j' do not move in step i . Otherwise, without loss of generality, assume that j moves in step i , i.e., j serves the request at $r_i = S_{i,j}$. We need to show that

$$|D_{j'}^{i-1} - (D_j^{i-1} + d(S_{i-1,j}, r_i))| \leq d(r_i, S_{i,j'}), \quad (6.3)$$

since then

$$|D_{j'}^i - D_j^i| = |D_{j'}^{i-1} - (D_j^{i-1} + d(S_{i-1,j}, r_i))| \leq d(r_i, S_{i,j'}) = d(r_i, S_{i-1,j'}).$$

First, by induction and triangle inequality, we have

$$D_{j'}^{i-1} - D_j^{i-1} - d(S_{i-1,j}, r_i) \stackrel{\text{ind.}}{\leq} d(S_{i-1,j'}, S_{i-1,j}) - d(S_{i-1,j}, r_i) \stackrel{\Delta\text{-ineq.}}{\leq} d(r_i, S_{i-1,j'}). \quad (6.4)$$

Now, by definition of BAL, since BAL selected server j to serve r_i , we have $D_j^{i-1} + d(S_{i-1,j}, r_i) \leq D_{j'}^{i-1} + d(S_{i-1,j'}, r_i)$, and thus

$$D_j^{i-1} + d(S_{i-1,j}, r_i) - D_{j'}^{i-1} \stackrel{\text{BAL}}{\leq} D_{j'}^{i-1} + d(S_{i-1,j'}, r_i) - D_{j'}^{i-1} = d(S_{i-1,j'}, r_i). \quad (6.5)$$

Together, eqs. (6.4) and (6.5) imply eq. (6.3). \square

For condition (ii) of Corollary 4.3, we simply observe that $\Phi_0 \leq 0$.

For condition (iii), we need that $\Delta\Phi_i^{\text{OPT}} := \Phi'_i - \Phi_{i-1} \leq k \cdot \text{OPT}_i(\sigma)$. If OPT does not move in step i , then $q_{i-1}^* = q_i^* \notin \{r_{i-1}, r_i\}$, and therefore $\Delta\Phi_i^{\text{OPT}} = 0$. Otherwise, by definition, OPT moves a server from $q_i^* \neq r_i$ to $r_i = q_{i-1}^*$, and $\text{OPT}_i(\sigma) = d(q_i^*, r_i)$. Because $q_{i-1}^* = r_i$, the potential Φ_{i-1} is given by the second case of eq. (6.2) and we get

$$\begin{aligned} \Delta\Phi_i^{\text{OPT}} &= \Phi'_i - \Phi_{i-1} \\ &= (kD_{q_i^*}^{i-1} - \sum_{j=1}^k D_j^{i-1}) - (k(D_{r_{i-1}}^{i-1} - d(r_{i-1}, r_i)) - \sum_{j=1}^k D_j^{i-1}) \\ &= k(D_{q_i^*}^{i-1} - D_{r_{i-1}}^{i-1} + d(r_{i-1}, r_i)). \end{aligned}$$

If $q_i^* = r_{i-1}$, it immediately follows that $\Delta\Phi_i^{\text{OPT}} = kd(r_{i-1}, r_i) = k \cdot \text{OPT}_i(\sigma)$. Otherwise, in step $i-1$, BAL moved a server from $r_i \neq q_i^*$ to $r_{i-1} \neq q_i^*$. This means that there is a server at q_i^* that does not move in step $i-1$. We have

$$\begin{aligned} D_{q_i^*}^{i-1} &= D_{q_i^*}^{i-2} \\ D_{r_{i-1}}^{i-1} &= D_{r_i}^{i-2} + d(r_i, r_{i-1}). \end{aligned}$$

By Claim 6.12, we have $D_{q_i^*}^{i-2} - D_{r_i}^{i-2} \leq d(q_i^*, r_i)$, and, therefore,

$$\begin{aligned} \Delta\Phi_i^{\text{OPT}} &= k(D_{q_i^*}^{i-1} - D_{r_{i-1}}^{i-1} + d(r_{i-1}, r_i)) \\ &= k(D_{q_i^*}^{i-2} - D_{r_i}^{i-2}) \\ &\leq kd(r_i, q_i^*) \\ &= k\text{OPT}_i(\sigma), \end{aligned}$$

as required by condition (iii).

Finally, condition (iv) of Corollary 4.3 demands that $\Delta\Phi_i^{\text{BAL}} := \Phi_i - \Phi'_i \leq -\text{BAL}_i(\sigma)$, where $\text{BAL}_i(\sigma) = d(r_{i+1}, r_i)$. If $q_i^* \neq q_i = r_{i+1}$, since also $q_i^* \neq r_i$, BAL has a server at q_i^* that does not move during step i , i.e., $D_{q_i^*}^i = D_{q_i^*}^{i-1}$. The potential Φ_i is given by the first case of eq. (6.2) and we get

$$\begin{aligned} \Delta\Phi_i^{\text{BAL}} &= (kD_{q_i^*}^i - \sum_{j=1}^k D_j^i) - (kD_{q_i^*}^{i-1} - \sum_{j=1}^k D_j^{i-1}) \\ &= \sum_{j=1}^k D_j^{i-1} - \sum_{j=1}^k D_j^i \\ &= -d(r_{i+1}, r_i). \end{aligned}$$

Otherwise, we have $q_i^* = q_i = r_{i+1}$, and $D_{r_i}^i - D_{q_i^*}^{i-1} = D_{r_i}^i - D_{r_{i+1}}^{i-1} = d(r_i, r_{i+1})$. The potential Φ_i is given by the second case of eq. (6.2) and we get

$$\begin{aligned} \Delta\Phi_i^{\text{BAL}} &= (k(D_{r_i}^i - d(r_i, r_{i+1})) - \sum_{j=1}^k D_j^i) - (kD_{q_i^*}^{i-1} - \sum_{j=1}^k D_j^{i-1}) \\ &= k(D_{r_i}^i - D_{q_i^*}^{i-1} - d(r_i, r_{i+1})) - d(r_{i+1}, r_i) \\ &= -d(r_{i+1}, r_i). \end{aligned}$$

In both case condition (iv) is fulfilled. Using conditions (i)-(iv), we can apply Corollary 4.3 and obtain the following result.

Theorem 6.13 (Manasse et al. [1990]). BAL is k -competitive for the k -server problem on finite metric spaces of size $N = k + 1$ if the initial server locations are distinct.

Remark 6.14. Irani and Rubinfeld [1991] showed that minimizing $D_j^{i-1} + 2d(S_{i-1,j}, r_i)$ yields a 10-competitive 2-server algorithm, and Chrobak and Larmore [1991a] showed that this algorithm has competitive ratio at least 6. Kleinberg [1994] showed that minimizing $D_s^{i-1} + f(d(S_{i-1,s}, r_i))$ for any function $f: \mathbb{R} \rightarrow \mathbb{R}$ does not yield a 2-server algorithm with competitive ratio smaller than $\frac{5+\sqrt{7}}{2} \approx 3.82$.

6.4 General metric spaces

We now briefly outline additional results for general metric spaces that rely on work function approaches, i.e., on online algorithms that take the optimum behavior into account for their decisions.

For 2 servers, Manasse et al. [1990] define the algorithm RES (*residue*) that, in step i , serves the request at r_i with the server that served the last request (at r_{i-1}) if and only if it can afford to do so while staying 2-competitive, even if it later needs to undo the move from r_{i-1} to r_i and instead move the other server to r_i via r_{i-1} . Formally, we can express this condition via

$$2\text{OPT}(\sigma_{\leq i}) - \text{RES}(\sigma_{\leq i-1}) \geq 2d(r_{i-1}, r_i) + d(S_{i-1,1}, S_{i-1,2}),$$

where the left hand side is the so-called *residue*, i.e., the cost that RES can cause in step i while still staying 2-competitive.

Theorem 6.15 (Manasse et al. [1990]). RES is 2-competitive for the online 2-server problem.

For arbitrary $k \in \mathbb{N}$ we can define a work function algorithm much like for metrical task systems. Assuming that servers start in distinct locations, we use the fact that there is an optimum solution where servers always remain at distinct locations. We define the work function $w_i(C)$ that maps a configuration $C \subset P$, $|C| = k$ of server locations to the offline optimum cost for serving all requests in $\sigma_{\leq i}$ and then positioning the servers on the points in C . Let $d(C, C')$ denote the total distance of a minimum distance matching between the points in C and C' , and let $C_0 := \{S_{0,j} : j \in \{1, \dots, k\}\}$ denote the initial configuration of servers. We get

$$\begin{aligned} w_0(C) &= d(C_0, C) \\ w_i(C) &= \min_{p \in C} \{w_{i-1}((C \setminus \{p\}) \cup \{r_i\}) + d(r_i, p)\}. \end{aligned}$$

With this, the straight-forward work function algorithm WFA serves the request r_i in step i with a server

$$j \in \arg \min_{j \in \{1, \dots, k\}} \{w(C_{i-1} \setminus \{S_{i-1,j}\} \cup \{r_i\}) + d(S_{i-1,j}, r_i)\},$$

where C_i denotes the server configuration after step i . We state the following theorem without proof.

Theorem 6.16 (Koutsoupias and Papadimitriou [1995]). WFA is $(2k - 1)$ -competitive for the online k -server problem.

7 Primal-Dual Algorithms

All underlying offline problems we have discussed in the previous chapters can be expressed in terms of an integer linear program (ILP), i.e., in the form

$$(ILP) \quad \begin{array}{l} \min \quad \mathbf{c}^\top \mathbf{x} \\ \text{s.t.} \quad \mathbf{Ax} \geq \mathbf{b} \\ \mathbf{x} \in \mathbb{N}^n \end{array}$$

where $A \in \mathbb{R}^{m \times n}$, $\mathbf{b} \in \mathbb{R}^m$ and $\mathbf{c} \in \mathbb{R}^n$ are input parameters, and $\mathbf{x} \in \mathbb{N}^n$ are the solution variables.

For example, we can express the ski rental problem with n days and buying price B by introducing a variable x that indicates whether we decide to buy skis, and variables z_i that indicate whether we rent skis on day i . We obtain the following ILP:

$$(SR) \quad \begin{array}{l} \min \quad B \cdot x + \sum_{i=1}^n z_i \\ \text{s.t.} \quad x + z_j \geq 1, \quad \forall j \in \{1, \dots, n\} \\ x, z_i \in \{0, 1\}, \quad \forall i \in \{1, \dots, n\} \end{array}$$

We can turn the problem of finding a good solution to an ILP to an online problem by letting variables and constraints appear over time, while requiring monotonicity in the sense that variables that have previously been set may not be decreased.

ONLINE ILP PROBLEM	
given:	linear program (ILP) $(A_0, \mathbf{b}_0, \mathbf{c}_0)$, initial solution \mathbf{x}_0
online:	sequence $\sigma = ((A_1, \mathbf{b}_1, \mathbf{c}_1), (A_2, \mathbf{b}_2, \mathbf{c}_2), \dots, (A_n, \mathbf{b}_n, \mathbf{c}_n))$ of ILPs, s.t. $A_i, \mathbf{b}_i, \mathbf{c}_i$ extend $A_{i-1}, \mathbf{b}_{i-1}, \mathbf{c}_{i-1}$ by adding rows & columns
actions: (step i)	extend solution \mathbf{x}_{i-1} to a new feasible solution \mathbf{x}_i for $A_i, \mathbf{b}_i, \mathbf{c}_i$ by setting all new variables and by increasing or keeping old variables
objective:	minimize $\mathbf{c}_n^\top \mathbf{x}_n$

This view on online problems will allow us to employ powerful linear programming methods to develop and analyze online algorithms.

7.1 The primal-dual method

Consider the LP-relaxation (P) of an integer linear program (ILP) that we obtain by dropping the requirement that variables need to be integral (below). Including fractional solutions allows us to tap into fundamental results from LP duality, which can be used to bound the quality of a solution to (P) by considering the dual (D). Intuitively, the dual LP computes the best lower bound on the objective function value of (P) that we can obtain by a linear combination of its inequalities, where y_j is the weight of the j -th inequality in this combination. The constraints of the dual ensure that the contribution of no primal variable to this combination of inequalities exceeds its contribution to the primal objective, i.e., the linear

combination lower bounds the objective function. This immediately implies *weak duality*: the fact that every feasible solution to (D) gives a lower bound for the optimum value of (P).

(P)

(D)

$$\begin{array}{ll} \min & \sum_{i=1}^n c_i x_i \\ \text{s.t.} & \sum_{i=1}^n a_{ij} x_i \geq b_j, \quad \forall j \in \{1, \dots, m\} \\ & x_i \geq 0, \quad \forall i \in \{1, \dots, n\} \end{array}$$

$$\begin{array}{ll} \max & \sum_{j=1}^m b_j y_j \\ \text{s.t.} & \sum_{j=1}^m a_{ij} y_j \leq c_i, \quad \forall i \in \{1, \dots, n\} \\ & y_j \geq 0, \quad \forall j \in \{1, \dots, m\} \end{array}$$

Theorem 7.1 (weak duality). Let \mathbf{x}, \mathbf{y} be feasible solutions for the dual linear programs (P) and (D). Then,

$$\sum_{i=1}^n c_i x_i \geq \sum_{j=1}^m b_j y_j.$$

Proof. We have

$$\sum_{i=1}^n c_i x_i \stackrel{\text{(D)}}{\geq}_{x_i \geq 0} \sum_{i=1}^n \sum_{j=1}^m a_{ij} y_j x_i = \sum_{j=1}^m y_j \sum_{i=1}^n a_{ij} x_i \stackrel{\text{(P)}}{\geq}_{y_j \geq 0} \sum_{j=1}^m b_j y_j. \quad \square$$

In particular, weak duality allows us to prove near-optimality of a solution by providing a feasible dual solution with similar objective function value.

Corollary 7.2. Let \mathbf{x}, \mathbf{y} be feasible solutions for the dual linear programs (P) and (D) with $\sum_{i=1}^n c_i x_i \leq \alpha \sum_{j=1}^m b_j y_j$ for some $\alpha \geq 1$. Then \mathbf{x}, \mathbf{y} are α -approximations of the optimum for (P) and (D), respectively.

It is not clear that we can generally expect to get a good bound on the optimal objective function value only by a linear combination of inequalities. *Strong duality* provides the fundamental guarantee that we can always get a tight bound this way. This very powerful result is at the core of LP duality theory. For its proof we refer to any introductory class on linear optimization.

Theorem 7.3 (strong duality). A linear program (P) is feasible and bounded if and only if its dual (D) is feasible and bounded. In this case, the optimum values of (P) and (D) coincide.

Another fundamental result of duality theory is the principle of *complementary slackness*. This principle states that a pair of feasible solutions of (P) and (D) is optimal if and only if it has the property that a variable of (P)/(D) is non-zero if the corresponding constraint of (D)/(P) is tight and vice-versa. We slightly extend this principle by including approximately optimal solutions with approximately tight constraints for non-zero variables.

Theorem 7.4 (approximate complementary slackness). Let \mathbf{x}, \mathbf{y} be feasible solutions for the dual linear programs (P) and (D), such that the following conditions hold for $\alpha, \beta \geq 1$:

(i) primal slackness: If $x_i > 0$ for $i \in \{1, \dots, n\}$, then $\sum_{j=1}^m a_{ij} y_j \geq c_i / \alpha$.

(ii) dual slackness: If $y_j > 0$ for $j \in \{1, \dots, m\}$, then $\sum_{i=1}^n a_{ij} x_i \leq \beta b_j$.

Then, $\sum_{i=1}^n c_i x_i \leq \alpha \beta \sum_{j=1}^m b_j y_j$ and hence \mathbf{x}, \mathbf{y} are $\alpha\beta$ -approximations for (P) and (D), respectively.

Proof. We have

$$\begin{aligned}
\sum_{i=1}^n c_i x_i &\stackrel{(i)}{\leq} \sum_{i=1}^n \alpha x_i \sum_{j=1}^m a_{ij} y_j \\
&= \alpha \sum_{j=1}^m y_j \sum_{i=1}^n a_{ij} x_i \\
&\stackrel{(ii)}{\leq} \alpha \beta \sum_{j=1}^m b_j y_j.
\end{aligned}$$

The statement thus follows by Corollary 7.2. \square

With these structural results in mind, we are ready to state the *primal-dual method*. The idea of *primal-dual approximation algorithms* is to compute not only a primal solution, but to simultaneously compute a feasible dual solution. If we can ensure that the ratio between the value of both solutions is bounded, weak duality (Corollary 7.2) implies a bounded approximation ratio for the primal solution. We use the same method to devise primal-dual online algorithms that maintain a pair of feasible solutions, one for the primal and one for the dual LP, while ensuring that the ratio between the value of both solutions stays bounded.

Theorem 7.5 (primal-dual method). Let $\rho \geq 1$ and ALG be an online algorithm for the online ILP problem that, in every step i , computes feasible solutions x_i, y_i for the LP relaxation of (A_i, b_i, c_i) and its dual (A_i^\top, c_i, b_i) , such that x_i is integral and $c_i^\top x_i \leq \rho b_i^\top y_i$. Then ALG is strictly ρ -competitive.

In order to illustrate the primal-dual method on a simple example, we apply it to the ski rental problem. To do this, we first need to formulate the LP-relaxation (SRP) of the ski rental ILP (SR) and its dual (SRD).

(SRP)

$$\begin{array}{ll}
\min & B \cdot x + \sum_{i=1}^n z_i \\
\text{s.t.} & x + z_j \geq 1, \quad \forall j \in \{1, \dots, n\} \\
& x, z_i \geq 0, \quad \forall i \in \{1, \dots, n\}
\end{array}$$

(SRD)

$$\begin{array}{ll}
\max & \sum_{j=1}^n y_j \\
\text{s.t.} & \sum_{j=1}^n y_j \leq B \\
& 0 \leq y_j \leq 1, \quad \forall j \in \{1, \dots, n\}
\end{array}$$

In every online step of the ILP version of the ski rental problem, we are presented with an additional variable z_j and an additional constraint $x + z_j \geq 1$. We need to satisfy this constraint by either setting $x = 1$ (i.e., by buying skis) or by setting $z_j = 1$ (i.e., by renting skis). At the same time, we want to set y_j in order to maintain a feasible dual solution with large objective value. The simplest approach is to set y_j as large as possible without violating dual constraints, i.e., such that $\sum_{j=1}^n y_j = B$ or $y_j = 1$, whichever happens first when increasing y_j from 0. We can let the dual solution guide our primal decisions by setting $x = 1$ if the corresponding dual constraint $\sum_{j=1}^n y_j \leq B$ becomes tight, and setting $z_j = 1$ otherwise (in that case, $y_j = 1$).

Observe that this algorithm is equivalent to the algorithm ALG^B that always buys on day B . We already saw in Section 1.3 that ALG^B is strictly $(2 - 1/B)$ -competitive. We now derive a slightly weaker bound using LP duality theory.

Theorem 7.6. ALG^B is strictly 2-competitive.

Proof. We use the primal-dual method (Theorem 7.5). By definition, the LP-interpretation of ALG^B maintains integral feasible solutions of (SRP) and (SRD). It remains to show that $B \cdot x + \sum_{i=1}^n z_i \leq 2 \sum_{j=1}^n y_j$. We can easily see this using the approximate complementary slackness conditions (Theorem 7.4). For

primal slackness, observe that if $x > 0$, then, by integrality of x , we have $x = 1$ and thus $\sum_{j=1}^n y_j = B$ must hold, by definition of ALG^B . Also, if $z_i > 0$, then $z_i = 1$ and thus we have $y_i = 1$, again by definition of ALG^B . For dual slackness, we simply observe that $x + z_j \leq 2$, i.e., the primal constraints are always tight up to a factor of 2. By Theorem 7.4, we get $B \cdot x + \sum_{i=1}^n z_i \leq 2 \sum_{j=1}^n y_j$, i.e., the primal solution is a 2-approximation. \square

7.2 Randomized rounding

The primal-dual method as stated in Theorem 7.5 requires that, in every step i , the computed solution x_i of the LP relaxation must be integral. This requirement is very restrictive, and it is often easier to formulate a primal-dual algorithm that gradually increases variables, e.g., the variable x in the ski rental problem. Of course, the fractional solution that is computed this way is not a feasible solution to the original ILP. However, we can obtain a randomized feasible solution by interpreting fractional variable values (smaller 1) as probabilities of setting the corresponding variable to 1. This way, we can hope to retain the solution quality of the fractional solution by introducing randomization. This approach is the so-called *randomized rounding* technique.

For example, consider again the ski rental problem, this time allowing fractional solutions. On day i , the additional primal variable z_i is presented, together with the primal constraint $x + z_i \geq 1$, which corresponds to a new dual variable y_i . We define a natural primal-dual algorithm PD that computes a fractional solution in every step. First, as before, if $x < 1$, we set $y_i = 1$ and otherwise we set $y_i = 0$. Then, in order to satisfy $x + z_i \geq 1$, we set $z_i = 1 - x$. Finally, PD increases x by some amount $\alpha(x + \beta)$ to a maximum of 1, where $\alpha, \beta \in \mathbb{R}$ will be determined later. Since x grows monotonically over time, this ensures that the primal constraints always stay satisfied.

Theorem 7.7. For $\alpha = 1/B$ and $\beta = 1/[(1 + 1/B)^B - 1]$, PD computes a $(1 + \beta)$ -approximate, fractional solution for the ski rental problem.

Proof. We show that PD computes a good fractional solution using the primal-dual method (without requiring integrality). To do this, we need to show that the primal and dual solutions computed in every step are feasible and that the primal objective function value remains bounded by $\rho \geq 1$ times the dual objective function value (analogously to Theorem 7.5).

We already argued that the primal solution is feasible, because the constraints $x + z_i \geq 1$ remain satisfied by definition of PD . To show dual feasibility, we need to ensure $\sum_{j=1}^n y_j \leq B$, i.e., $x \geq 1$ after at most B days. Let $x_i \geq 0$ denote the increment of x in step i , i.e., $x = \sum_{i=1}^n x_i$ at the end, $x_1 = \alpha\beta$, and

$$x_i = \alpha \left(\sum_{j=1}^{i-1} x_j + \beta \right) = \alpha x_{i-1} + \alpha \left(\sum_{j=1}^{i-2} x_j + \beta \right) = (1 + \alpha)x_{i-1} = (1 + \alpha)^{i-1} \alpha\beta,$$

for $i \in \{2, \dots, n\}$. After B days we have

$$x = \sum_{j=1}^B x_j = \sum_{j=0}^{B-1} (1 + \alpha)^j \alpha\beta = \alpha\beta \frac{(1 + \alpha)^B - 1}{(1 + \alpha) - 1} = \beta \cdot [(1 + \alpha)^B - 1].$$

To ensure that $x = 1$ after B days, and have dual feasibility, we may thus set

$$\beta := 1/[(1 + \alpha)^B - 1].$$

In order to bound the ratio between the total primal and dual objective function value in the end, we can bound the ratio between the change in primal and dual objective function values in every step. Consider

step i . We may assume $x = \sum_{j=1}^{i-1} x_j < 1$ at the start of step i , otherwise the objective function values do not change, since z_i, y_i are set to 0 and x remains 1. But then y_i is set to 1, which increases the dual objective by 1. On the other hand, the primal objective changes by

$$Bx_i + z_i = B\alpha\left(\sum_{j=1}^{i-1} x_j + \beta\right) + 1 - \sum_{j=1}^{i-1} x_j.$$

We can set $\alpha = 1/B$ in order to eliminate the dependence on x_1, \dots, x_{i-1} . We get a constant change in primal objective of exactly $1 + \beta$ times the change in dual objective in every step. Consequently, the ratio between primal and dual objective is always bounded by $1 + \beta$. This concludes the proof by weak duality (Theorem 7.1). \square

Based on the fractional solution computed by P_D , we can define a randomized algorithm R_{PD} using randomized rounding. To do this, we interpret the fractional contribution x_i to x in step i as the probability of buying skis on day i and the final value of x as the probability of buying skis at all. We can achieve these probabilities by initially choosing a threshold $\tau \in [0, 1]$ uniformly at random and buying skis as soon as x exceeds τ . The expected cost for buying skis then becomes $B \cdot x = B \cdot \sum_{i=1}^n x_i$. The probability of renting skis on the i -th day is $1 - \sum_{j=1}^{i-1} x_j = z_i$, thus the total expected cost of the rounded solution is

$$\mathbb{E}_{\tau \sim [0,1]}[R_{PD}(\sigma)] = Bx + \sum_{i=1}^n z_i,$$

which coincides with the cost of the fractional solution produced by P_D . Since P_D computes a $(1 + \beta)$ approximation, even to the fractional optimum, we obtain the following.

Theorem 7.8. R_{PD} is a strictly $(1 + \beta)$ -competitive randomized online algorithm for ski rental, with $\beta = 1/[(1 + 1/B)^B - 1]$.

Remark 7.9. For $B \rightarrow \infty$, we have $\beta = \frac{1}{e-1}$ and the resulting randomized competitive ratio of R_{PD} of $\frac{e}{e-1}$ is best possible for the ski rental problem (see exercises).

7.3 Online bipartite matching

An important example of a primal-dual online algorithm is the `RANKING` algorithm for the online bipartite matching problem. In this problem, the set L of vertices (the “left-hand side”) of a bipartite graph $G = (L \cup R, E)$ is given and the vertices of R arrive online, together with their incident edges to vertices in L . Our goal is to match each vertex of R to at most one unmatched neighbor in L as it arrives, such as to maximize the size of the final matching. Formally, we define matchings as follows.

Definition 7.10. A *matching* in a graph $G = (V, E)$ is a set $M \subseteq E$ of mutually disjoint edges. A *maximum matching* is a matching of maximum cardinality, and a *maximal matching* is a matching that is no proper subset of any other matching.

ONLINE BIPARTITE MATCHING PROBLEM	
given:	set of left-hand vertices L of an unknown bipartite graph $G = (L \cup R, E)$
online:	sequence $\sigma = (v_1, v_2, \dots, v_n) \in R^n$ of right-hand vertices, each v_i given by the set $L_i \subseteq L$ of neighbors of v_i in G
actions (step i):	match v_i to an unmatched vertex in L_i , or leave it unmatched
objective:	maximize the size of the matching




Figure 7.1: Illustration of the adversarial construction in the proof of Theorem 7.13. The matching $\text{ALG}[\sigma]$ consists of the thick edges, while $\text{OPT}[\sigma]$ consists of the other edges.

Note that the online bipartite matching problem is a *maximization* problem. Accordingly, we need to adapt the definition of competitiveness. The definition of randomized competitiveness and the (randomized) competitive ratio are changed analogously.

Definition 7.11. A deterministic online algorithm ALG is ρ -competitive for a maximization problem if there is a constant $\alpha \geq 0$, such that for all instances $\sigma \in \Sigma$

$$\text{ALG}(\sigma) \geq \frac{1}{\rho} \cdot \text{OPT}(\sigma) - \alpha.$$

The statement of the primal-dual method needs to be adapted as well.

Theorem 7.12 (primal-dual method (maximization)). Let $\rho \geq 1$ and ALG be an online algorithm for maximization version of the online ILP problem that, in every step i , computes feasible solutions $\mathbf{x}_i, \mathbf{y}_i$ for the LP relaxation of $(A_i, \mathbf{b}_i, \mathbf{c}_i)$ and its dual $(A_i^\top, \mathbf{c}_i, \mathbf{b}_i)$, such that \mathbf{x}_i is integral and $\mathbf{c}_i^\top \mathbf{x}_i \geq (1/\rho) \mathbf{b}_i^\top \mathbf{y}_i$. Then ALG is strictly ρ -competitive.

7.3.1 Deterministic algorithms

We first give a simple lower bound on the best deterministically possible competitive ratio.

Theorem 7.13. Every deterministic online matching algorithm has a competitive ratio of at least 2.

Proof. Consider an adversarial request sequence $\sigma = (v_1, v_2)$ with two requests for the online bipartite matching problem, where v_1 has an edge to two vertices $u_1, u_2 \in L$ and v_2 has a single edge to the vertex in $\{u_1, u_2\}$ that was matched to v_1 in the first step. In this instance, we obviously have $\text{OPT}(\sigma) = 2$ and $\text{ALG}(\sigma) = 1$ for any online algorithm ALG . We can arbitrarily extend this example to a request sequence $\sigma = (v_1, v_2, \dots, v_{2n})$ and a set $L = \{u_1, \dots, u_{2n}\}$ for any $n \in \mathbb{N}$, by, for every $i \in \{1, \dots, n\}$ letting v_{2i-1} have an edge to both u_{2i-1}, u_{2i} , and letting v_{2i} have an edge to the single matched vertex from $\{u_{2i-1}, u_{2i}\}$ (see Figure 7.1). Then $\text{OPT}(\sigma) = 2n$, while $\text{ALG}(\sigma) \leq n$, and thus $\lim_{n \rightarrow \infty} \frac{\text{OPT}(\sigma)}{\text{ALG}(\sigma) - \alpha} \geq 2$ for any $\alpha \in \mathbb{R}$. Note that our construction works for any algorithm ALG that deterministically decides whether to match v_{2i-1} to u_{2i-1} or u_{2i} for every $i \in \{1, \dots, n\}$. We thus get a lower bound of 2 on the competitive ratio of any deterministic online algorithm, analogously to Proposition 1.12. \square

We can achieve this competitive ratio easily, using the straight forward **GREEDY** algorithm that starts by ordering L arbitrarily and then matches every incoming right-hand vertex to the first unmatched neighbor in this order.

Theorem 7.14. **GREEDY** is strictly 2-competitive for the online bipartite matching problem.

Proof. Observe that GREEDY always computes a maximal matching M_{GREEDY} : We need to show that no edge $e = \{u, v\} \in E \setminus M_{\text{GREEDY}}$ can be added to M_{GREEDY} to obtain a larger matching. Either $v \in R$ is matched by GREEDY or $u \in L$ was already matched when v appeared, otherwise GREEDY would have matched u and v when v appeared. Thus $M_{\text{GREEDY}} \cup \{e\}$ is not a matching.

It is now sufficient to show that every maximal matching M has at least half the size of the maximum matching $\text{OPT}[\sigma]$. To see this, we compare the number of matched vertices in M and $\text{OPT}[\sigma]$. Every matching edge $e \in \text{OPT}[\sigma]$ is responsible for two matched vertices in $\text{OPT}[\sigma]$. On the other hand, since M is maximal, at least one endpoint of e is matched in M . Overall, the number of matched vertices is $2|\text{OPT}(\sigma)|$ in $\text{OPT}[\sigma]$ versus at least $|\text{OPT}(\sigma)|$ in M . \square

7.3.2 Ranking algorithm

We can avoid the simple lower bound construction of Theorem 7.13 by making the GREEDY algorithm less predictable. The easiest way to do this is to start with a random order of L and then to deterministically execute GREEDY for this order. The RANKING algorithm is an alternate formulation of this exact algorithm: For every $u \in L$ pick a random number $g_u \in [0, 1]$ uniformly at random, and then, in every step i , match v_i to the unmatched neighbor $u \in L$ with the lowest value g_u . We will interpret g_u as the *greed* of u , which means that RANKING matches v_i to its least greedy (i.e., most “generous”) unmatched neighbor.¹

We analyze RANKING using the primal-dual method. To do this, we first need to formulate the primal and dual linear programs. For convenience, we write $(u, v) \in E$ as a shorthand for $\{u, v\} \in E, u \in L, v \in R$. With this, let x_{uv} be a variable that indicates whether the edge $(u, v) \in E$ is part of the matching ($x_{uv} = 1$), or not ($x_{uv} = 0$). By allowing fractional values of x_{uv} we get the linear program relaxation (BMP) with its dual (BMD).

(BMP)	(BMD)
$\begin{aligned} \max \quad & \sum_{(u,v) \in E} x_{uv} \\ \text{s.t.} \quad & \sum_{v: (u,v) \in E} x_{uv} \leq 1, \quad \forall u \in L \\ & \sum_{u: (u,v) \in E} x_{uv} \leq 1 \quad \forall v \in R \\ & x_{uv} \geq 0, \quad \forall (u, v) \in E \end{aligned}$	$\begin{aligned} \min \quad & \sum_{u \in L} \alpha_u + \sum_{v \in R} \beta_v \\ \text{s.t.} \quad & \alpha_u + \beta_v \geq 1 \quad \forall (u, v) \in E \\ & \alpha_u \geq 0 \quad \forall u \in L \\ & \beta_v \geq 0 \quad \forall v \in R \end{aligned}$

In order to turn RANKING into a primal-dual algorithm, we additionally need to maintain a dual solution in every step. Whenever RANKING matches two vertices $u \in L, v \in R$, the primal objective increases by $x_{uv} = 1$. The primal-dual method for maximization problems (Theorem 7.12) requires that, at the same time, the dual increases by at most ρ , where $\rho \geq 1$ is the competitive ratio we want to prove. We can achieve this by distributing a value of ρ between α_u and β_v . We do this according to the greed of u , by setting

$$\alpha_u = \rho \cdot h(g_u) \cdot \sum_{(u,v) \in E} x_{uv}, \quad \beta_v = \rho \cdot \sum_{(u,v) \in E} (1 - h(g_u)) \cdot x_{uv},$$

where $h: [0, 1] \rightarrow [0, 1]$ is some monotonically increasing function with $h(1) = 1$ that we will determine later. Note that as long as $u \in L$ (respectively, $v \in R$) are unmatched, we have $\alpha_u = 0$ (respectively, $\beta_v = 0$).

We interpret the dual variables α, β as some distribution of the dual objective value over the vertices in L and R , respectively. Each vertex $u \in L$ that is being matched gets a value of ρ that it distributes between itself and its matching partner according to its greed g_u : It keeps $\rho \cdot h(g_u)$ to itself and gives $\rho \cdot (1 - h(g_u))$

¹ Note we may safely ignore the event that a vertex has greed 1 or that two vertices have the same greed, since the combined probability of any such event occurring is 0.

to its partner. Whenever a new vertex $v \in R$ appears, it chooses the most generous available neighbor $u \in L$ as its matching partner in order to get the highest possible value.

Fix any edge $(u, v) \in E$ and assume we run **RANKING**, ignoring the vertex $u \in L$ and using the same greed values $g_{u'}$ for all $u' \in L \setminus \{u\}$. We denote this modified algorithm by $\text{RANKING}_{\bar{u}}$. We are interested in the dual utility value $\bar{\beta}_{uv}$ that v gets during the execution of $\text{RANKING}_{\bar{u}}$. Equivalently, we can ask for the greed value \bar{g}_{uv} of v 's matching partner (if v gets matched) in the execution of $\text{RANKING}_{\bar{u}}$. Formally, we define

$$\bar{g}_{uv} = \begin{cases} g_{u'} & \text{if } v \text{ gets matched to } u' \in L \text{ by } \text{RANKING}_{\bar{u}}, \\ 1 & \text{if } v \text{ remains unmatched.} \end{cases}$$

and $\bar{\beta}_{uv} := \rho \cdot (1 - h(\bar{g}_{uv}))$.

The following *dominance* lemma establishes that vertices $u \in L$ of smaller greed are always preferred.

Lemma 7.15. If $g_u < \bar{g}_{uv}$ for some edge $(u, v) \in E$, then u gets matched by **RANKING**.

Proof. Consider the online step of the execution of **RANKING** where $v \in R$ appears. Assume that u is not matched yet, otherwise the lemma is trivial. Then, the execution of **RANKING** so far is identical to the execution of $\text{RANKING}_{\bar{u}}$. By definition, $\text{RANKING}_{\bar{u}}$ matches v to the neighbor in $L \setminus \{u\}$ with lowest greed \bar{g}_{uv} if it exists. Since $\{u, v\} \in E$, $g_u < \bar{g}_{uv}$ and u is not matched yet, u is the unmatched neighbor of v in L with the lowest greed. Therefore **RANKING** matches v and u . \square

We can also show a *monotonicity* lemma that states that the value a vertex $v \in R$ receives can only increase as we add vertices to L .

Lemma 7.16. For all $(u, v) \in E$, we have $\beta_v \geq \bar{\beta}_{uv}$.

Proof. First observe that the set of unmatched vertices in $L \setminus \{u\}$ after every step of **RANKING** forms a superset of the set of unmatched vertices in $L \setminus \{u\}$ after the same step of $\text{RANKING}_{\bar{u}}$. We can see this by induction over the number of online steps: Assume that when $v' \in R$ appears, the set of unmatched neighbors $N_{v'}$ of v' in $L \setminus \{u\}$ in **RANKING** is a superset of the same set $\bar{N}_{v'}$ in $\text{RANKING}_{\bar{u}}$. If **RANKING** matches v' to a vertex $u' \in N_{v'}$, then this vertex has the lowest greed among all vertices in $N_{v'} \supseteq \bar{N}_{v'}$. If $u' \in \bar{N}_{v'}$, then $\text{RANKING}_{\bar{u}}$ must therefore match the same vertex. Hence the unmatched vertices of **RANKING** in $L \setminus \{u\}$ remains a superset of those of $\text{RANKING}_{\bar{u}}$.

Now when v arrives, by the observation above, we have $N_v \supseteq \bar{N}_v$. Now \bar{N}_v contains a vertex of greed \bar{g}_{uv} if $\bar{g}_{uv} < 1$, by definition of \bar{g}_{uv} . Hence, if $\bar{g}_{uv} < 1$, **RANKING** matches v to a vertex u'' of greed $g_{u''} \leq \bar{g}_{uv}$. By monotonicity of h , we get $\beta_v = \rho \cdot (1 - h(g_{u''})) \geq \rho \cdot (1 - h(\bar{g}_{uv})) = \bar{\beta}_{uv}$, as claimed. If $\bar{g}_{uv} = 1$, then $h(1) = 1$ implies that $\bar{\beta}_{uv} = 0 \leq \beta_v$. \square

We are now ready to prove competitiveness of **RANKING**.

Theorem 7.17 (Karp et al. [1990]). **RANKING** is strictly $\frac{e}{e-1}$ -competitive for online bipartite matching.

Proof. By definition of **RANKING**, we always compute a feasible and integral primal solution (a matching). Since every increase in primal objective increases the dual objective by exactly ρ , we maintain a ratio of ρ between primal and dual objective. The final ingredient of the primal-dual method is to use dual feasibility and apply weak duality in order to obtain a guarantee on the approximation quality of the computed primal solution. Unfortunately, we cannot guarantee the computed dual solution to always be feasible!

Recall that for a bounded randomized competitive ratio we only need a guarantee on the *expected* approximation quality of the computed solution rather than for every possible outcome. Since the ratio between primal and dual objective is ρ in every execution of RANKING, the same is true for the ratio between the expected primal and dual objective. By linearity of expectation, we get

$$\rho = \frac{\mathbb{E}[\sum_{u \in L} \alpha_u + \sum_{v \in R} \beta_v]}{\mathbb{E}[\sum_{(u,v) \in E} x_{uv}]} = \frac{\sum_{u \in L} \mathbb{E}[\alpha_u] + \sum_{v \in R} \mathbb{E}[\beta_v]}{\sum_{(u,v) \in E} \mathbb{E}[x_{uv}]}.$$

We can therefore apply the primal-dual method using the expected primal solution $\mathbb{E}[x]$ and the expected dual solution $\mathbb{E}[\alpha], \mathbb{E}[\beta]$. Since the feasible region of the primal LP is a convex polytope, the expected solution $\mathbb{E}[x]$ is feasible. The expected primal solution $\mathbb{E}[x]$ may be fractional, but it is only important that every possible computed solution is integral. In order to complete the primal-dual argument we need to apply weak duality. To do this, it remains to show that the expected dual solution is feasible, i.e., that $\mathbb{E}[\alpha_u] + \mathbb{E}[\beta_v] \geq 1$ for all $(u, v) \in E$. It then follows, by weak duality (Theorem 7.1), that the optimal fractional solution is bracketed between the primal and dual objective values, which are separated by a ratio of ρ .

To bound $\mathbb{E}[\alpha_u]$, we fix the values $g_{u'}$ for every $u' \in L \setminus \{u\}$ to any values in $[0, 1]$, and only consider the expectation with respect to choosing g_u uniformly at random from $[0, 1]$. Recall that if u is matched by RANKING, then $\alpha_u = \rho \cdot h(g_u)$, otherwise $\alpha_u = 0$. By Lemma 7.15, we know that u is guaranteed to be matched when $g_u < \bar{g}_{uv}$. We also know that \bar{g}_{uv} is independent of g_u . We thus have

$$\mathbb{E}_{g_u \sim [0,1]}[\alpha_u] \geq \rho \cdot \int_0^{\bar{g}_{uv}} h(g) dg =: \rho[H(\bar{g}_{uv}) - H(0)],$$

where H is the antiderivative of h . On the other hand, since $\bar{\beta}_{uv}$ is independent of g_u , by Lemma 7.16 we have

$$\mathbb{E}_{g_u \sim [0,1]}[\beta_v] \geq \bar{\beta}_{uv} = \rho \cdot (1 - h(\bar{g}_{uv})).$$

Together, the expected dual solution is feasible if, for every choice of $g_{u'}$ for every $u' \in L \setminus \{u\}$, we have

$$H(\bar{g}_{uv}) - H(0) + 1 - h(\bar{g}_{uv}) \geq 1/\rho. \quad (7.1)$$

In order to get rid of the dependence on \bar{g}_{uv} , we choose $h: [0, 1] \rightarrow [0, 1]$ to be of the form $h(g) = ce^g$. The condition $h(1) = 1$ then immediately implies $c = 1/e$. With this, the smallest possible competitive ratio that guarantees dual feasibility (eq. (7.1)) is $\rho = \frac{e}{e-1}$. We can apply Theorem 7.12 to obtain the following. \square

7.3.3 Randomized Lower Bound*

We now derive a tight lower bound on the randomized competitive ratio of the online bipartite matching problem using Yao's principle for maximization problems:

Theorem 7.18 (Yao [1977], maximization). Let $(q_n: \Sigma \rightarrow [0, 1])_{n \in \mathbb{N}}$ be a sequence of probability distributions with

$$\lim_{n \rightarrow \infty} \mathbb{E}_{\sigma \sim q_n}[\text{OPT}(\sigma)] = \infty.$$

Then, every randomized online algorithm for a maximization problem has a competitive ratio of at least

$$\limsup_{n \rightarrow \infty} \frac{\mathbb{E}_{\sigma \sim q_n}[\text{OPT}(\sigma)]}{\sup_{\text{ALG} \in \mathcal{A}_{\text{det}}} \mathbb{E}_{\sigma \sim q_n}[\text{ALG}(\sigma)]}.$$

With this, we show the following bound.

Theorem 7.19 (Karp et al. [1990]). The competitive ratio of any randomized online bipartite matching algorithm is at least $\frac{e}{e-1}$.

Proof. To make use of Theorem 7.18, we need to define a sequence $(q_n)_{n \in \mathbb{N}}$ of probability distributions over input sequences. We define $q_n : \Sigma \rightarrow [0, 1]$ to be a distribution over request sequences of length n that results from the following process. We use a set L containing n vertices, and say that all vertices of L are *active* initially. In every step i , the appearing vertex $v_i \in R$ has an edge to every active vertex in L , and at the end of step i a random active vertex becomes inactive. Equivalently, we set $L = \{u_1, \dots, u_n\}$ and initially choose a random ordering $u_{j_1}, u_{j_2}, \dots, u_{j_n}$ of the vertices in L . Then v_i is adjacent to the set of vertices $\{u_{j_i}, u_{j_{i+1}}, \dots, u_{j_n}\}$ of size $n - i + 1$.

By definition, it is possible to match v_i to u_{j_i} in every step i , thus $\text{OPT}(\sigma) = n$ for every possible input sequence σ produced by this process. In particular, we have $\mathbb{E}_{\sigma \sim q_n}[\text{OPT}(\sigma)] = n$ and thus $\lim_{n \rightarrow \infty} \mathbb{E}_{\sigma \sim q_n}[\text{OPT}(\sigma)] = \infty$, as required by Yao's principle.

Now consider the options for an online algorithm in step i . If v_i still has unmatched active neighbors, it is never a mistake to match v_i to one of them. By symmetry of the process, all unmatched active neighbors are completely identical and there is no reason to choose one of them over the others. Therefore, no online algorithm has a better choice than picking a matching partner for v_i uniformly at random from amongst its unmatched active neighbors. We obtain that the corresponding algorithm RAND is a best-possible online algorithm for the specific distribution q_n of instances that we constructed.

Using Yao's principle, we obtain that no online algorithm has competitive ratio below

$$\limsup_{n \rightarrow \infty} \frac{\mathbb{E}_{\sigma \sim q_n}[\text{OPT}(\sigma)]}{\sup_{\text{ALG} \in \mathcal{A}_{\text{det}}} \mathbb{E}_{\sigma \sim q_n}[\text{ALG}(\sigma)]} \geq \limsup_{n \rightarrow \infty} \frac{n}{\mathbb{E}_{\sigma \sim q_n}[\text{RAND}(\sigma)]}. \quad (7.2)$$

It remains to bound $\lim_{n \rightarrow \infty} \mathbb{E}_{\sigma \sim q_n}[\text{RAND}(\sigma)]$. Let $a_i = n - i + 1$ denote the number of active vertices in L at the start of step i , and let $x_i \leq a_i$ denote the number of unmatched active vertices in L at the start of step i . Let further $\Delta a_i := a_i - a_{i+1}$ and $\Delta x_i := x_i - x_{i+1}$ denote the change in a_i and x_i during step i . We have $\Delta a_i = 1$ by definition of the random process. While $x > 0$, if u_{j_i} is unmatched at the start of step i and v_i does not get matched to u_{j_i} , then $\Delta x_i = 2$, and in all other cases $\Delta x_i = 1$. The probability that u_{j_i} is still unmatched, i.e., that u_{j_i} is one of the x_i unmatched active vertices, is $\frac{x_i}{a_i}$, since, by symmetry, every active vertex has the same probability of being matched. The probability that u_{j_i} is not matched to v_i , given that it is one of the x_i unmatched active vertices, is $\frac{x_i - 1}{x_i}$. We have

$$\mathbb{E}_{\sigma \sim q_n}[\Delta x_i] = 1 + \frac{x_i}{a_i} \frac{x_i - 1}{x_i} = 1 + \frac{x_i - 1}{a_i},$$

and, since $\Delta a_i = 1$,

$$\frac{\mathbb{E}_{\sigma \sim q_n}[\Delta x_i]}{\mathbb{E}_{\sigma \sim q_n}[\Delta a_i]} = 1 + \frac{x_i - 1}{a_i}.$$

For $n \rightarrow \infty$ we can approximate the solution to this difference equation, i.e., the distribution over x_i 's and a_i 's that satisfies it, with high probability² by the solution to the ordinary differential equation

$$\frac{dx}{da} = 1 + \frac{x - 1}{a}.$$

² With x, a approximating x_i, a_i with "high probability", we mean that, for every $\varepsilon > 0$, the probability that the function $x(a)$ has error at most ε goes to 1 as n goes to infinity.

The only solution to this differential equation is $x = ca + a \ln a + 1$ for any $c \in \mathbb{R}$. In order to determine c , we observe that $x_1 = a_1 = n$, thus

$$c = 1 - \ln n - \frac{1}{n} = \frac{n-1}{n} - \ln n.$$

We get that with high probability

$$x = 1 + a\left(\frac{n-1}{n} + \ln \frac{a}{n}\right).$$

Consider the step i where $x_i = 1$, i.e., one vertex in L remains active and unmatched. At this point, in the next step this vertex will still be matched, but then all unmatched vertices in L are inactive and the matching does not grow further. For $x = 1$ we get

$$\ln \frac{a}{n} = -\frac{n-1}{n}$$

or

$$a = ne^{-\frac{n-1}{n}}.$$

For $n \rightarrow \infty$, once $x = 1$, the expected number of remaining active vertices in L gets arbitrarily close to $a = n/e$. Since only one of them will still be matched, the total size of the matching computed by RAND is given by

$$\mathbb{E}_{\sigma \sim q_n}[\text{RAND}(\sigma)] = n - \frac{n}{e} + 1 = \frac{e-1}{e}n + 1.$$

With eq. (7.2) we obtain the claimed lower bound. □



8 Online Load Balancing

Scheduling problems have an inherent time component that makes them particularly well-suited to be studied in an online setting. In scheduling, we are given *jobs* that need to be processed by a set of machines. Each machine can run some jobs better (i.e., faster) than others, which is usually expressed in terms of machine-dependent *processing times* for each job. Scheduling problems have been studied in a multitude of variants, e.g., differing in

- whether each job needs to be executed en bloc on one machine, or can be interrupted (*preempted*) and resumed on the same or a different machine (*migration*), or even can be parallelized;
- whether jobs have a *release time* when they become available, and/or a *deadline* when they need to be completed;
- whether jobs have *precedence constraints* that demand a partial order in which jobs need to be processed;
- whether the latest completion time (*makespan*), the (*weighted*) *sum of completion times*, or other measures need to be minimized.

We focus on the *load balancing* problem, where each job needs to be processed without interruption on a single machine and we want to minimize the makespan, i.e., the maximum *load* of a machine. In the online variant, we assume that jobs arrive one after the other at time 0 and need to immediately be assigned to a machine. Importantly, jobs *do not* arrive over time, but rather before the actual processing of the computed schedule begins.

We now formally specify the online load balancing problem. Note that it is customary to identify machines with natural numbers $\{1, \dots, m\}$ and jobs with $\{1, \dots, n\}$ to allow for simpler notation, even though each job is formally given by a vector of its processing times on every machine.

ONLINE LOAD BALANCING PROBLEM	
given:	machines $1, \dots, m$
online:	sequence $\sigma = (1, \dots, n)$ of jobs, each job j comes with processing times $p_{ij} > 0$ on every machine $i \in \{1, \dots, m\}$
actions (for job j):	assign job j to some machine i_j ; we set $J_i(\sigma) := \{j \in \sigma : i_j = i\}$
objective:	minimize the makespan $\max_{i \in \{1, \dots, m\}} L_i(\sigma)$, where $L_i(\sigma) := \sum_{j \in J_i(\sigma)} p_{ij}$ is the load of machine i

In terms of this notation, $C_j(\sigma) := L_{i_j}(\sigma_{\leq j})$ is the completion time of job j , and $S_j(\sigma) := L_{i_j}(\sigma_{\leq j-1})$ is its starting time. See Figure 8.1 for an example.




Figure 8.1: Example of a load balancing instance with $m = 4$ machines and $n = 6$ jobs. The unique optimum solution with maximum load 4.5 is depicted on the right.

8.1 Identical machines

We first consider online load balancing with *identical machines*, where we demand that the processing time of a job does not depend on the machine the job is processed on, i.e., $p_{ij} = p_j$. A very natural online algorithm for this setting is *Graham's list scheduling algorithm* (LIST): Schedule each job j on a machine i that currently has minimum load $L_i(\sigma_{\leq j-1})$. This algorithm was proposed by Graham [1969] and is often considered to be the first approximation algorithm, though it predates the theory of NP-hardness.

Theorem 8.1 (Graham [1969]). LIST has strict competitive ratio $2 - 1/m$ for the online load balancing problem with m identical machines.

Proof. Fix an instance given by sequence $\sigma = (1, \dots, n)$ and associated processing times $p_j > 0$ for $j \in \sigma$. Since OPT needs to schedule every job, we trivially have $\text{OPT}(\sigma) \geq \max_{j \in \sigma} p_j$. We get another simple lower bound by observing that the total load $L = \sum_{i=1}^m L_i(\sigma) = \sum_{i=1}^m \sum_{j \in J_i(\sigma)} p_j = \sum_{j \in \sigma} p_j$ is the same for any schedule and, therefore,

$$\text{OPT}(\sigma) \geq L/m = \sum_{j \in \sigma} p_j/m.$$

Now take any job $j \in \sigma$ that has maximum completion time $C_j(\sigma) = \max_{j' \in \sigma} C_{j'}(\sigma)$. By definition of LIST, machine i_j had minimum load at the starting time $S_j(\sigma)$ of job j , i.e., all machines are busy until at least $S_j(\sigma)$, i.e., $\frac{1}{m} \sum_{j'=1}^{j-1} p_{j'} \geq S_j(\sigma)$. Hence,

$$\begin{aligned} \text{LIST}(\sigma) &= C_j(\sigma) \\ &= S_j(\sigma) + p_j \\ &\leq \sum_{j' \in \sigma \setminus \{j\}} p_{j'}/m + p_j \\ &= \sum_{j' \in \sigma} p_{j'}/m + (1 - 1/m)p_j \\ &\leq \sum_{j' \in \sigma} p_{j'}/m + (1 - 1/m) \max_{j' \in \sigma} p_{j'} \\ &\leq (2 - 1/m) \cdot \text{OPT}(\sigma). \end{aligned}$$




Figure 8.2: Lower bound construction for LIST with $m(m-1)$ jobs with unit processing times followed by one job with processing time m .

It remains to show a tight lower bound on the strict competitive ratio of LIST. Consider a sequence $\sigma = (1, \dots, n)$ with $n = m(m-1) + 1$ and $p_{j \neq n} = 1, p_n = m$ (see Figure 8.2). Since LIST attempts to keep the load on all machines balanced, until step $n-1$ it will produce load $L_i(\sigma_{\leq n-1}) = m-1$ on all machines i . It immediately follows that $\text{LIST}(\sigma) = 2m-1$. On the other hand, we can achieve the same

load on all machines by assigning job n to its own machine and distributing all other jobs evenly. It follows that $\text{OPT}(\sigma) = m$, and thus $\text{LIST}(\sigma) \geq (2 - 1/m)\text{OPT}(\sigma)$, which establishes a matching lower bound. \square

Proposition 8.2. LIST is a best-possible online algorithm for $m = 2$.

Proof. Consider the instance σ with processing times 1, 1, 2 and let ALG be an online algorithm with best-possible competitive ratio. Since any online algorithm that schedules the two first jobs on the same machine only has competitive ratio $2 > 2 - 1/m = 3/2$ for the instance $\sigma_{\leq 2}$, we can conclude that ALG schedules the first two jobs on different machines. But then, ALG has makespan 3 for instance σ , while $\text{OPT}(\sigma) = 2$. We get $\text{ALG}(\sigma) \geq 3\text{OPT}(\sigma)/2 = (2 - 1/m) \cdot \text{OPT}(\sigma)$. \square

Remark 8.3. The best known algorithm for an arbitrary number of identical machines was presented by Fleischer and Wahl [2000] and is 1.9201-competitive. The best known lower bound of 1.88 was given by Rudin III and Chandrasekaran [2003].

8.2 Related machines

Load balancing with *related machines* allows different processing times on each machine by assigning each machine i a speed s_i , $s_1 \leq s_2 \leq \dots \leq s_m$, while demanding $p_{ij} = p_j/s_i$ for all jobs j , where $p_j > 0$. This allows machines to be different as long as this difference (in speed) remains consistent between jobs. If we want to apply LIST to this setting, it makes a difference whether we choose a machine with minimum load *before* scheduling the next job on this machine (PREGREEDY) or *after* scheduling the next job on it (POSTGREEDY). If we have two machines with very different speeds, PREGREEDY is arbitrarily bad (proportional to the ratio between the speeds) already for two identical jobs, since the first job may be put on the slow machine. On the other hand, POSTGREEDY has a logarithmically bounded competitive ratio (without proof).

Theorem 8.4 (upper bound: Cho and Sahni [1988], lower bound: Aspnes et al. [1997]). POSTGREEDY has competitive ratio $\Theta(\log n)$ for the online load balancing problem on related machines.

We develop a better online algorithm that uses a typical design scheme: Assume that we know an upper bound $\Lambda \geq \text{OPT}(\sigma)$ and are able to design a procedure that, given Λ , produces a solution of cost below $\rho \cdot \Lambda$ for some constant ρ . Now make a guess Λ for the value of $\text{OPT}(\sigma)$ and try to use our procedure. We either get a solution of cost below $\rho\Lambda$, or we realize that our guess was incorrect. We can then begin with $\Lambda = \Lambda_0 \leq \text{OPT}(\sigma)$ and restart from the beginning with $\Lambda_i = 2\Lambda_{i-1}$ every time our procedure fails. For every restart, we ignore all choices made so far (but not their costs), and pretend that the input sequence arrives again online from the beginning, even though we could do better since we know part of the sequence already.

Since $\Lambda_0 > 0$, the procedure eventually succeeds for Λ_k , $k \in \mathbb{N}$, at the latest when $\Lambda_k \geq \text{OPT}(\sigma)$. If $k = 0$, we computed a solution of cost at most $\rho\Lambda_0 \leq \rho \cdot \text{OPT}(\sigma)$. Otherwise, we know that $\Lambda_{k-1} < \text{OPT}(\sigma)$. Even summing up the costs for every run independently (see Figure 8.3), i.e., scheduling the same job multiple times, we get $\text{ALG}(\sigma) \leq \sum_{i=0}^k \rho\Lambda_i = \rho \sum_{i=0}^k 2^i \Lambda_0 = \rho \frac{2^{k+1}-1}{2-1} \Lambda_0 < \rho 2^{k+1} \Lambda_0$. Using $2^{k-1} \Lambda_0 = \Lambda_{k-1} < \text{OPT}(\sigma)$ we get $\text{ALG}(\sigma) < 4\rho \cdot \text{OPT}(\sigma)$.

Lemma 8.5 (doubling strategy). If there is an online algorithm that computes a solution of cost at most $\rho\Lambda$ for given $\Lambda \geq \text{OPT}(\sigma)$ and we can compute some bound $0 < \Lambda_0 \leq \text{OPT}(\sigma)$ in the first step, then we can devise a strictly 4ρ -competitive online algorithm.




Figure 8.3: Illustration of the doubling strategy for scheduling. If our subroutine is unable to fit all jobs into a time window of width $\rho\Lambda_i$, it restarts from the beginning with a wider window of width $2\Lambda_{i+1}$. The shading of jobs corresponds to the smallest time window in which they could be scheduled. Observe that jobs are generally scheduled in multiple time windows and thus contribute multiple times to the overall cost.

We introduce the algorithm `SLOWFIT` that uses the doubling strategy (for $\rho = 2$) with the following subroutine for a given bound $\Lambda \geq \text{OPT}(\sigma)$: Put every job j on the slowest machine i where the resulting load does not exceed 2Λ (the *slowest* machine where the job fits), i.e., $i = \min\{i' \in \{1, \dots, m\} : L_{i'}(\sigma_{\leq j-1}) + \frac{p_j}{s_{i'}} \leq 2\Lambda\}$. The following lemma guarantees that if $\Lambda \geq \text{OPT}(\sigma)$, such a machine must always exist. If this is not the case, then `SLOWFIT` can safely abort the subroutine and restart it for $\Lambda' = 2\Lambda$.

Lemma 8.6. Given $\Lambda \geq \text{OPT}(\sigma)$ the subroutine of `SLOWFIT` computes a solution of cost at most 2Λ .

Proof. We need to show that the subroutine of `SLOWFIT` always finds a machine where to schedule a job j without exceeding a load of 2Λ . For the sake of contradiction, assume this is not the case, i.e., there is a job j for which $L_i(\sigma_{\leq j-1}) + \frac{p_j}{s_i} > 2\Lambda$ for all machines $i \in \{1, \dots, m\}$. Let $L_i^*(\sigma)$ denote `OPT`'s load on machine i after processing the entire request sequence σ . Since $\text{OPT}(\sigma) \leq \Lambda$ and $p_j > 0$, we know that $\sum_{i=1}^m \Lambda s_i \geq \sum_{i=1}^m L_i^*(\sigma) \cdot s_i = \sum_{j'=1}^n p_{j'} > \sum_{j'=1}^{j-1} p_{j'}$. This means that there must be a machine with load less than Λ at the start of step j , since otherwise $\sum_{j'=1}^{j-1} p_{j'} = \sum_{i=1}^m L_i(\sigma_{\leq j-1}) \cdot s_i \geq \sum_{i=1}^m \Lambda s_i$. We may thus define $f = \max\{i \in \{1, \dots, m\} : L_i(\sigma_{\leq j-1}) \leq \Lambda\}$ to be the fastest machine that is not overloaded (see Figure 8.4).




Figure 8.4: Sketch of the proof of Lemma 8.6. The machine f is the fastest machine with load at most Λ . Since $\text{OPT}(\sigma) \leq \Lambda$, `SLOWFIT` is processing at least one job on machines $\{f + 1, \dots, m\}$ that `OPT` is not. But this job would fit on machine f , which is a contradiction with the fact that `SLOWFIT` puts each job on the slowest machine where it fits.

If $f = m$, then

$$L_m(\sigma_{\leq j-1}) + p_j/s_m \leq \Lambda + \text{OPT}(\sigma) \leq 2\Lambda,$$

which contradicts the fact that the subroutine of `SLOWFIT` doesn't find a machine for job j . Hence, we can assume $f < m$.

Since all machines $i > f$ have load $L_i(\sigma_{\leq j-1}) > \Lambda$, we get

$$\begin{aligned} \sum_{i=f+1}^m L_i(\sigma_{\leq j-1}) \cdot s_i &> \sum_{i=f+1}^m \Lambda s_i \\ &\geq \sum_{i=f+1}^m \text{OPT}(\sigma) \cdot s_i \\ &\geq \sum_{i=f+1}^m L_i^*(\sigma) \cdot s_i, \end{aligned}$$

i.e., the total processing times of jobs in $\sigma_{\leq j-1}$ processed on machines $\{f+1, \dots, m\}$ is strictly larger for SLOWFIT than for OPT. This means that there must be a job $j' < j$ that SLOWFIT assigns to a machine $i_{j'} \in \{f+1, \dots, m\}$ and that OPT assigns to a slower machine $i^* \in \{1, \dots, f\}$. Then, $p_{j'}/s_{i^*} \leq \text{OPT}(\sigma) \leq \Lambda$. Consider the moment when SLOWFIT assigned job j' to machine $i_{j'}$. At this point, we had

$$L_f(\sigma_{\leq j'-1}) + \frac{p_{j'}}{s_f} \leq L_f(\sigma_{\leq j-1}) + \frac{p_{j'}}{s_{i^*}} \leq 2\Lambda.$$

But then, SLOWFIT could have assigned job j' to machine f , which contradicts the fact that j' was assigned to a machine $i_{j'} > f$. \square

Corollary 8.7 (Aspnes et al. [1997]). SLOWFIT is strictly 8-competitive for the online load balancing problem with related machines.

Proof. The claim follows from Lemmas 8.5 and 8.6 if we initially set $\Lambda_0 = p_1/s_m \leq \text{OPT}(\sigma)$. \square

Remark 8.8. Berman et al. [2000] presented the currently best known deterministic algorithm for online load balancing on related machines which is 5.828-competitive as well as the best known lower bound of 2.438.

8.3 Restricted assignment

So far, we assumed that the difference in processing times on different machines is consistent between jobs, i.e., that we can order machines from slow to fast. We now want to investigate how much harder online load balancing becomes without this property. To do this, we consider load balancing with *restricted assignment*, where every job has a fixed processing time but cannot be processed on all machines, i.e., $p_{ij} \in \{p_j, \infty\}$ for all machines i and jobs j . This setting turns out to be substantially harder.

Theorem 8.9 (Azar et al. [1995]). Every deterministic online algorithm for load balancing with restricted assignment has competitive ratio at least $\log_2 m$.

Proof. Let ALG be any fixed online algorithm. We take $k \in \mathbb{N}$ and construct an adversarial request sequence $\sigma = (1, 2, \dots, m-1)$ for $m = 2^k$ machines and with $p_{ij} \in \{1, \infty\}$ for all jobs $j \in \{1, \dots, m-1\}$ (see Figure 8.5). If we let $M_j = \{i \in \{1, \dots, m\} : p_{ij} = 1\}$ denote the machines that can run job j , we need to specify M_j for all $j \in \{1, \dots, m-1\}$ to fully describe the instance. For $j \in \{1, \dots, m/2\}$ we set $M_j = \{j, j+m/2\}$. By symmetry, we can assume that ALG schedules each job $j \in \{1, \dots, m/2\}$ on machine j , otherwise we rename machines to ensure this is the case. Similarly, we let $M_{j'+m/2} = \{j', j'+m/4\}$ for $j' \in \{1, \dots, m/4\}$ and assume that ALG schedules each job $j \in \{m/2+1, \dots, m/2+m/4\}$

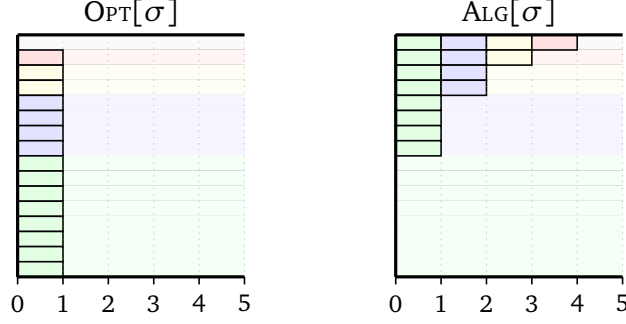


Figure 8.5: Illustration of the proof of Theorem 8.9. Jobs arrive in the order green, blue, yellow, red. Every job can be scheduled by two machines, but not two jobs of the same color can be processed by the same machine.

on machine $j - m/2$, and so on. Eventually, for all $x \in \{1, \dots, k\}$ we have $M_{j'+X_x} = \{j', j' + m/2^x\}$ for $j' \in \{1, \dots, m/2^x\}$ and $X_x := \sum_{l=1}^{x-1} \frac{m}{2^l}$, and job $j' + X_x$ is scheduled on machine j' . The number of jobs scheduled on machine 1 is $k = \log_2 m$, hence $\text{ALG}(\sigma) = \log_2 m$.

On the other hand, we can put a single job on every machine by assigning job $j' + X_x$ to machine $j' + m/2^x$ for every $x \in \{1, \dots, k\}$ and every $j' \in \{1, \dots, m/2^x\}$. Hence $\text{OPT}(\sigma) = 1$ and $\text{ALG}(\sigma) \geq \log_2 m \cdot \text{OPT}(\sigma)$, thus the strict competitive ratio of ALG is at least $\log_2 m$.

Since $\text{OPT}(\sigma)$ scales with the processing times, we can apply Proposition 1.12 to obtain the claimed bound. \square

We claim that the simple greedy algorithm (GREEDY) that puts every job j on the least loaded machine i with $p_{ij} \neq \infty$ is almost optimal in terms of competitive ratio.

Theorem 8.10 (Azar et al. [1995]). GREEDY is strictly $(\lceil \log_2 m \rceil + 1)$ -competitive for online load balancing with restricted assignment.

Proof. We consider an input sequence σ and let $R_{\lambda,i} := \max\{0, L_i(\sigma) - \lambda \text{OPT}(\sigma)\}$ denote the load on machine i after time $\lambda \text{OPT}(\sigma)$ for $\lambda \in \mathbb{N}$, and $R_\lambda := \sum_{i=1}^m R_{\lambda,i}$. Note that, by definition, $R_0 = \sum_{j=1}^n p_j$. **We claim: $R_{\lambda+1} \leq R_\lambda/2$ for all $\lambda \in \mathbb{N}$.** With this claim, for all $\lambda \in \mathbb{N}$, we have $R_\lambda \leq R_0/2^\lambda$ and thus

$$\text{OPT}(\sigma) \geq \frac{1}{m} \sum_{j=1}^n p_j = \frac{1}{m} R_0 \geq \frac{2^\lambda}{m} R_\lambda,$$

and, in particular,

$$\text{OPT}(\sigma) \geq \frac{2^{\lceil \log_2 m \rceil}}{m} R_{\lceil \log_2 m \rceil} \geq R_{\lceil \log_2 m \rceil}.$$

But this means that even if the total load $R_{\lceil \log_2 m \rceil}$ after time $\lceil \log_2 m \rceil \cdot \text{OPT}(\sigma)$ goes on a single machine, GREEDY must finish before time $(\lceil \log_2 m \rceil + 1) \cdot \text{OPT}(\sigma)$, which proves the theorem.

It remains to establish our claim that $R_{\lambda+1} \leq R_\lambda/2$ for all $\lambda \in \mathbb{N}$. Consult Figure 8.6 with the following. We show $w_\lambda := R_\lambda - R_{\lambda+1} \geq R_{\lambda+1}$, which immediately implies $R_\lambda = w_\lambda + R_{\lambda+1} \geq 2R_{\lambda+1}$, as claimed. We define *layer* λ to be the time period $[\lambda \text{OPT}(\sigma), (\lambda + 1) \text{OPT}(\sigma))$ and let $w_{\lambda,i} := R_{\lambda,i} - R_{\lambda+1,i} \leq \text{OPT}(\sigma)$ be the load on machine i in layer λ . Let $r_{\lambda,j} := \max\{0, \min\{p_j, C_j - \lambda \text{OPT}(\sigma)\}\}$ denote the remaining processing time of job j at the start of layer λ . We further let J_i^* be the set of jobs scheduled on machine i by OPT and $R_{\lambda,i}^* := \sum_{j \in J_i^*} r_{\lambda,j} \leq \text{OPT}(\sigma)$. It is sufficient to show $w_{\lambda,i} \geq R_{\lambda+1,i}^*$ for all machines i , because then

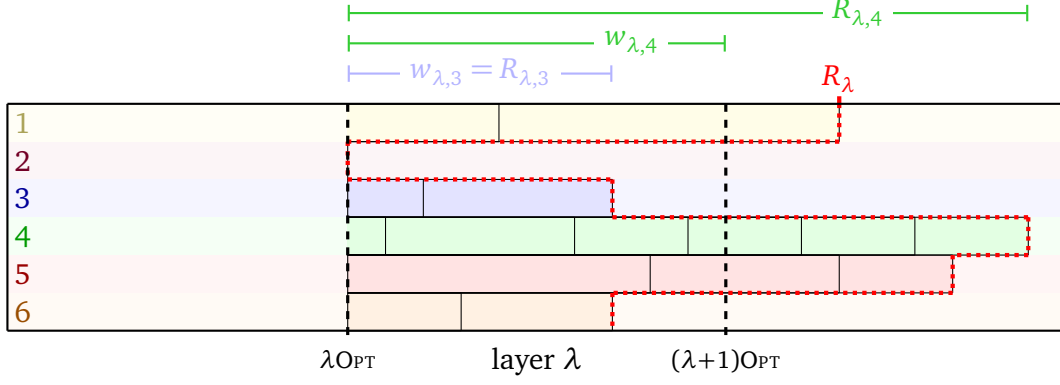


Figure 8.6: Illustration of the definitions in the proof of Theorem 8.10.

$w_\lambda = \sum_{i=1}^m w_{\lambda,i} \geq \sum_{i=1}^m R_{\lambda+1,i}^* = R_{\lambda+1}$. This is obviously true if $w_{\lambda,i} = \text{OPT}(\sigma)$, because $R_{\lambda+1,i}^* \leq \text{OPT}(\sigma)$, or if $R_{\lambda+1,i}^* = 0$.

Now take a machine i with $w_{\lambda,i} < \text{OPT}(\sigma)$ and $R_{\lambda+1,i}^* > 0$. The latter allows us to take any job $j^* \in J_i^*$ that is still unfinished after layer λ , i.e., with $r_{\lambda+1,j^*} > 0$. This job is not scheduled on machine i , because $w_{\lambda,i} < \text{OPT}(\sigma)$ and thus $w_{\lambda',i} = 0$ for all $\lambda' > \lambda$. By definition of GREEDY, the starting time of j^* is $C_{j^*} - p_{j^*} \leq \lambda \text{OPT}(\sigma) + w_{\lambda,i}$, because otherwise GREEDY would have put j^* on machine i . It follows that

$$\begin{aligned}
R_{\lambda+1,i}^* &= r_{\lambda+1,j^*} + \sum_{j \in J_i^* \setminus \{j^*\}} r_{\lambda+1,j} \\
&\leq r_{\lambda+1,j^*} + \text{OPT}(\sigma) - p_{j^*} \\
&\leq C_{j^*} - (\lambda + 1)\text{OPT}(\sigma) + \text{OPT}(\sigma) - p_{j^*} \\
&\leq w_{\lambda,i}.
\end{aligned}$$

□

8.4 Unrelated machines

Finally, we turn to a setting where each job can have arbitrary processing times on the different machines. This is often referred to as the *unrelated machines* setting. We present an elegant algorithm PDLB for this setting that relies on the doubling technique (Lemma 8.5) together with a primal-dual approach for the necessary subroutine.

Recall that Lemma 8.5 requires a subroutine that has an input parameter Λ and must always compute a solution of bounded quality when provided with a valid bound $\Lambda \geq \text{OPT}(\sigma)$. The subroutine of PDLB is based on an LP-relaxation (LBP) for the problem of distributing the maximum number of jobs on the machines while not exceeding load Λ (see below). In the underlying integer program, the variable $x_{ij} \in \{0, 1\}$ indicates whether job j is scheduled on machine i . Since, for $\Lambda \geq \text{OPT}(\sigma)$, we know that $\text{OPT}[\sigma]$ schedules each job j on a machine i with $p_{ij} \leq \Lambda$, we can restrict our variables x_{ij} accordingly. To that end, we define $S := \{(i, j) \in \{1, \dots, m\} \times \{1, \dots, n\} : p_{ij} \leq \Lambda\}$.

Observe that the integer program has a solution of value n (i.e., all jobs can be scheduled) if and only if $\Lambda \geq \text{OPT}(\sigma)$. Consequently, if $\Lambda \geq \text{OPT}(\sigma)$, the optimum solution of the LP-relaxation (LBP) must have value n (the converse may not be true). By weak duality (Theorem 7.1), the dual (LBD) can therefore not have a solution of value below n if $\Lambda \geq \text{OPT}(\sigma)$.

(LBP)

$$\begin{array}{ll}
\max & \sum_{(i,j) \in S} x_{ij} \\
\text{s.t.} & \sum_{j:(i,j) \in S} p_{ij} x_{ij} \leq \Lambda, \quad \forall i \in \{1, \dots, m\} \\
& \sum_{i:(i,j) \in S} x_{ij} = 1, \quad \forall j \in \{1, \dots, n\} \\
& x_{ij} \geq 0 \quad \forall (i,j) \in S
\end{array}$$

(LBD)

$$\begin{array}{ll}
\min & \Lambda \sum_{i=1}^m y_i + \sum_{j=1}^n z_j \\
\text{s.t.} & z_j + p_{ij} y_i \geq 1, \quad \forall (i,j) \in S \\
& y_i \geq 0 \quad \forall i \in \{1, \dots, m\} \\
& z_j \in \mathbb{R} \quad \forall j \in \{1, \dots, n\}
\end{array}$$

The idea behind the subroutine of PDLB is a bit different from the primal-dual technique in Chapter 7. We again maintain a feasible dual solution that informs primal decisions, but we do not use it to bound the quality of our primal solution. In fact, the subroutine always schedules all jobs and may not maintain a feasible primal solution at all. Instead, we use the dual solution as a means to detect when to abort: If the dual objective value ever falls below n , we may conclude that $\Lambda < \text{OPT}(\sigma)$ and abort. The subroutine is specified formally below.

Algorithm 1: PDLB subroutine

```

for  $i \in \{1, \dots, m\}$  do
   $y_i \leftarrow 1/(2m)$ 
foreach job  $j$  arriving online do
  if  $\min_{i \in \{1, \dots, m\}} p_{ij} > \Lambda$  or  $\max_{i \in \{1, \dots, m\}} y_i > 1$  then
    abort
   $i \leftarrow \arg \min_{i': (i', j) \in S} p_{i'j} y_{i'}$ 
   $x_{ij} \leftarrow 1$ 
   $z_j \leftarrow 1 - p_{ij} y_i$ 
   $y_i \leftarrow y_i (1 + p_{ij}/2\Lambda)$ 

```

We now show that the subroutine behaves as desired.

Lemma 8.11. Given $\Lambda \geq \text{OPT}(\sigma)$, the subroutine of PDLB computes a solution of load at most $\Lambda \cdot \mathcal{O}(\log m)$.

Proof. We first show that, if the subroutine does not abort, it computes a feasible solution where each machine has load at most $\Lambda \cdot \mathcal{O}(\log m)$. Note that we do not claim that the subroutine computes a feasible primal solution, but allow it to violate $\sum_{j=1}^n p_{ij} x_{ij} \leq \Lambda$.

In case the subroutine does not abort, we argue that the produced dual solution fulfills $y_i \leq 3/2$ for all $i \in \{1, \dots, m\}$, by induction over the number of jobs assigned to machine i . This is clear if no jobs get assigned, since the initial value $y_{i,0}$ of y_i is $1/(2m) < 3/2$. Now assume that at least one job is assigned to machine i and consider the step in which the last job j gets assigned. Since the algorithm does not abort, we have $y_i \leq 1$ at the beginning of this step, and by choice of i with $(i, j) \in S$ we have $p_{ij} \leq \Lambda$ (note that such an i must exist, since the algorithm did not abort). This means that y_i is updated to $y_i(1 + p_{ij}/2\Lambda) \leq 3/2$ and the claim is true.

Now, recall that $J_i(\sigma) \subseteq \{1, \dots, n\}$ denotes the set of jobs assigned to machine i . Since $y_{i,0} = 1/(2m)$ and we update y_i each time a job gets assigned to machine i , we have

$$y_i = \frac{1}{2m} \prod_{j \in J_i(\sigma)} (1 + p_{ij}/2\Lambda) \stackrel{p_{ij} \leq \Lambda}{\geq} \frac{1}{2m} \prod_{j \in J_i(\sigma)} \left(\frac{3}{2}\right)^{p_{ij}/\Lambda},$$

where we used $1 + x/2 \geq (3/2)^x$ for $x \in [0, 1]$.

Combining this with $y_i \leq 3/2$ and taking logarithms, we get

$$\log \frac{3}{2} \geq \log \frac{1}{2m} + \sum_{j \in J_i(\sigma)} \frac{p_{ij}}{\Lambda} \log \frac{3}{2},$$

and thus

$$L_i(\sigma) = \sum_{j \in J_i(\sigma)} p_{ij} \leq \Lambda \left(\frac{\log(3/2) + \log(2m)}{\log(3/2)} \right) = \Lambda \cdot \frac{\log(3m)}{\log(3/2)} = \Lambda \cdot \mathcal{O}(\log(m)),$$

as claimed.

It remains to show that if the subroutine of P_{DLB} aborts, then our guess of $\text{OPT}(\sigma)$ must have been too small, i.e., $\Lambda < \text{OPT}(\sigma)$. First observe that if the algorithm aborts because $\min_{i \in \{1, \dots, m\}} p_{ij} > \Lambda$, we clearly have $\text{OPT}(\sigma) > \Lambda$, since scheduling job j alone already produces a load of more than Λ . Secondly, observe that if $\Lambda \geq \text{OPT}(\sigma)$, the primal LP has a solution of value n , since all jobs can be scheduled.

We now assume that the algorithm aborts because $\max_{i \in \{1, \dots, m\}} y_i > 1$ at the end of some iteration j^* . It suffices to show that in this case the dual optimum has value less than j^* at this point. By weak duality (Theorem 7.1), it then follows that the primal optimum must also have value less than j^* , which implies that not all jobs in $\{1, \dots, j^*\}$ can be scheduled within load Λ , and, in particular, $\Lambda < \text{OPT}$.

To this end, observe that the algorithm maintains a dually feasible solution: Since the value of z_j is set to $1 - p_{ij}y_i$ with i being the machine that minimizes $p_{ij}y_i$, the constraints involving z_j are satisfied in step j , and, since y_i is non-decreasing, they remain satisfied later on.

To bound the value of the dual optimum, consider a step j where the algorithm assigns job j to some machine i . This increases the value of y_i by $\Delta_{ij} := y_i - y'_i = p_{ij}y'_i/2\Lambda$, where y'_i is the value of y_i before the update. Overall, the dual objective function increases by exactly $(1 - p_{ij}y'_i) + p_{ij}y'_i/2 = 1 - \Lambda \cdot \Delta_{ij}$ due to the increase of z_j and y_i . This means that the dual objective function value at the end of iteration j^* can be written as

$$\begin{aligned} \Lambda \sum_{i=1}^m y_i + \sum_{j=1}^{j^*} z_j &= \Lambda \sum_{i=1}^m y_{i,0} + \sum_{i=1}^m \sum_{j \in J_i(\sigma_{\leq j^*})} (1 - \Lambda \cdot \Delta_{ij}) \\ &= \Lambda \sum_{i=1}^m y_{i,0} + j^* - \Lambda \cdot \sum_{i=1}^m (y_i - y_{i,0}) \\ &= 2\Lambda \sum_{i=1}^m \frac{1}{2m} + j^* - \Lambda \sum_{i=1}^m y_i \\ &= \Lambda + j^* - \Lambda \sum_{i=1}^m y_i. \end{aligned}$$

By assumption, the algorithm aborts because there exists a machine i with $y_i > 1$. But, with $y_i \geq 0$, this implies that the objective function value of the current, feasible dual solution y, z is less than j^* , and therefore the same is true for the optimum primal solution. We can therefore conclude that $\Lambda < \text{OPT}(\sigma)$. \square

Using Lemma 8.11, we can now apply the doubling technique (with $\Lambda_0 = \min_i p_{i1}$) and obtain the following result.

Theorem 8.12 (Buchbinder and Naor [2011]). P_{DLB} is strictly $\mathcal{O}(\log m)$ -competitive for load balancing.

Remark 8.13. The bound in Theorem 8.12 was first proven by Aspnes et al. [1997] using a different approach. The algorithm P_{DLB} is a primal-dual interpretation of one of their algorithms.



Bibliography

- S. Albers. Improved randomized on-line algorithms for the list update problem. *SIAM Journal on Computing*, 27(3):682–693, 1998.
- S. Albers, B. von Stengel, and R. Werchner. A combined BIT and TIMESTAMP algorithm for the list update problem. volume 56, pages 135–139, 1995.
- C. Ambühl, B. Gärtner, and B. von Stengel. Optimal lower bounds for projective list update algorithms. *ACM Transactions on Algorithms*, 9(4):31:1–31:18, 2013.
- J. Aspnes, Y. Azar, A. Fiat, S. Plotkin, and O. Waarts. On-line routing of virtual circuits with applications to load balancing and machine scheduling. *Journal of the ACM*, 44(3):486–504, 1997.
- Y. Azar, J. Naor, and R. Rom. The competitiveness of on-line assignments. *Journal of Algorithms*, 18(2):221–237, 1995.
- L. A. Belady. A study of replacement algorithms for a virtual-storage computer. *IBM Systems Journal*, 5(2):78–101, 1966.
- J. L. Bentley and C. C. McGeoch. Amortized analyses of self-organizing sequential search heuristics. *Communications of the ACM*, 28(4):404–411, 1985.
- P. Berman, M. Charikar, and M. Karpinski. On-line load balancing for related machines. *Journal of Algorithms*, 35(1):108–121, 2000.
- A. Borodin, N. Linial, and M. Saks. An optimal online algorithm for metrical task systems. *Journal of the ACM*, 39:745–763, 1992.
- N. Buchbinder and J. Naor. Fair online load balancing. *Journal of Scheduling*, 16(1):117–127, 2011.
- Y. Cho and S. Sahni. Bounds for list schedules on uniform processors. *SIAM Journal on Computing*, 9:91–103, 1988.
- M. Chrobak and L. L. Larmore. On fast algorithms for two servers. *Journal of Algorithms*, 12(4):607–614, 1991a.
- M. Chrobak and L. L. Larmore. An optimal on-line algorithm for k servers on trees. *SIAM Journal on Computing*, 20(1):144–148, 1991b.
- M. Chrobak, H. Karloff, T. Payne, and S. Vishwnathan. New results on server problems. *SIAM Journal on Discrete Mathematics*, 4(2):172–181, 1991.
- A. Fiat, R. M. Karp, M. Luby, L. A. McGeoch, D. D. Sleator, and N. E. Young. Competitive paging algorithms. *Journal of Algorithms*, 12(4):685–699, 1991.
- R. Fleischer and M. Wahl. Online scheduling revisited. *Journal of Scheduling*, 3:343–353, 2000.
- R. L. Graham. Bounds on multiprocessing timing anomalies. *SIAM Journal on Applied Mathematics*, 17(2):416–429, 1969.
- S. Irani. Two results on the list update problem. *Information Processing Letters*, 38(6):301–306, 1991.
- S. Irani and R. Rubinfeld. A competitive 2-server algorithm. *Information Processing Letters*, 39(2):85–91, 1991.
- J. R. Isbell. Finitary games. *Contributions to the Theory of Games*, III:79–96, 1957.
- A. R. Karlin, M. S. Manasse, L. Rudolph, and D. D. Sleator. Competitive snoopy caching. *Algorithmica*, 3(1-4):79–119, 1988.
- R. Karp and P. Raghavan, 1990. Personal communication reported by Irani [1991].

-
- R. M. Karp, U. V. Vazirani, and V. V. Vazirani. An optimal algorithm for on-line bipartite matching. In *Proceedings of the 22nd annual ACM symposium on Theory of computing (STOC)*, pages 352–358, 1990.
- J. M. Kleinberg. A lower bound for two-server balancing algorithms. *Information Processing Letters*, 52(1): 39–43, 1994.
- E. Koutsoupias and C. H. Papadimitriou. On the k-server conjecture. *Journal of the ACM*, 42(5):971–983, 1995.
- M. S. Manasse, L. A. McGeoch, and D. D. Sleator. Competitive algorithms for server problems. *Journal of Algorithms*, 11(2):208–230, 1990.
- L. A. McGeoch and D. D. Sleator. A strongly competitive randomized paging algorithm. *Algorithmica*, 6 (1-6):816–825, 1991.
- N. Reingold, J. Westbrook, and D. D. Sleator. Randomized competitive algorithms for the list update problem. *Algorithmica*, 11(1):15–32, 1994.
- J. F. Rudin III and R. Chandrasekaran. Improved bounds for the online scheduling problem. *SIAM Journal on Computing*, 32:717–735, 2003.
- D. D. Sleator and R. E. Tarjan. Amortized efficiency of list update and paging rules. *Communications of the ACM*, 28(2):202–208, 1985.
- B. Teia. A lower bound for randomized list update algorithms. *Information Processing Letters*, 47(1):5–9, 1993.
- E. Torng. A unified analysis of paging and caching. *Algorithmica*, 20(2):175–200, 1998.
- A. C.-C. Yao. Probabilistic computations: Toward a unified measure of complexity. In *Proceedings of the 18th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 222–227, 1977.