




Faculty of Mathematics, Physics, and Informatics
Comenius University, Bratislava

How Much Information About the Future is Needed?

S. Dobrev, R. Kráľovič, D. Pardubská

TR-2007-007



Technical Reports in Informatics

How Much Information About the Future is Needed?*

Stefan Dobrev¹, Rastislav Kráľovič², and Dana Pardubská²

¹ Institute of Mathematics,
Slovak Academy of Sciences

`Stefan.Dobrev@savba.sk`

² Department of Computer Science,
Comenius University, Bratislava, Slovakia
`{kralovic,pardubska}@dcs.fmph.uniba.sk`

Abstract. We propose a new way of characterizing the complexity of online problems. Instead of measuring the degradation of output quality caused by the ignorance of the future we choose to quantify the amount of additional global information needed for an online algorithm to solve the problem optimally. In our model, the algorithm cooperates with an oracle that can see the whole input. We define the advice complexity of the problem to be the minimal number of bits (normalized per input request, and minimized over all algorithm-oracle pairs) communicated between the algorithm and the oracle in order to solve the problem optimally. Hence, the advice complexity measures the amount of problem-relevant information contained in the input.

We introduce two modes of communication between the algorithm and the oracle based on whether the oracle offers an advice spontaneously (helper) or on request (answerer). We analyze the Paging and DiffServ problems in terms of advice complexity and deliver tight bounds in both communication modes.

1 Introduction

The term “online” is used to describe algorithms that operate without the full knowledge of the input: a typical scenario would be a server that must continually process a sequence of requests in the order they arrive. More formally, an online algorithm processing an input sequence of requests $\mathbf{x} = \langle x_1, x_2, \dots, x_n \rangle$ produces an output sequence $\mathbf{y} = \langle y_1, y_2, \dots, y_n \rangle$ in such a way that each y_i is computed as a function of the prefix $\langle x_1, x_2, \dots, x_i \rangle$. On the other hand, an algorithm computing the whole output sequence \mathbf{y} from the entire input sequence \mathbf{x} is termed “offline”. The systematic study of online problems began in the late sixties [16], and has received much attention over the years (see e.g. [3], [6]). The standard measure used for evaluating online algorithms is the competitive ratio [20], [24], i.e. the worst case ratio between the solution quality of the given online algorithm and that of the optimal offline algorithm. The competitive complexity of an online problem is the best competitive ratio attainable by an online algorithm solving the problem. Intuitively, this measure describes the price, in terms of solution quality, that has to be paid for not knowing the whole input from the beginning.

Consider, for example, the well known SKIRENTAL problem [21]: a skier can rent a set of skis for one day for a unit price, or buy them for a fixed price c . However it is only the morning of each day when it becomes clear if the skier wants to continue skiing or

* Supported by APVV-0433-06 and VEGA 1/3106/06. Preliminary version of this paper appeared at SOFSEM 2008.

not. A classical result of the competitive analysis [21] shows that the optimal worst case performance is achieved by first renting the skis for $c - 1$ days, and then buying them, which gives a competitive ratio of $2 - 1/c$. The interpretation of this result is that the input carries a global information relevant to the problem; ignoring this information leads to a degradation of the best possible solution quality by a factor $2 - 1/c$.

In this paper we propose a new way of characterizing the complexity of online problems. The hardness incurred by the online setting comes from the fact that there is some information about the future input that is not available to the algorithm. In our approach we measure the amount of this hidden information. However, the input contains also information that is irrelevant to the problem at hand, and we have to find a way of distilling the problem-relevant information from the input.

Our approach to measure the relevant information is inspired by the communication complexity research. We consider, in addition to the algorithm itself, an oracle that sees the whole input and knows the algorithm. When computing the i -th output y_i , the algorithm not only sees the sequence $\langle x_1, x_2, \dots, x_i \rangle$, but can also communicate with the oracle. We require that the algorithm always computes an optimal solution. The advice complexity of the algorithm is the number of bits communicated between the algorithm and the oracle, normalized per request. The advice complexity of an online problem is the minimum advice complexity over all oracle–algorithm pairs that together solve the problem.

Apart from its theoretical significance, this measure can be of use in some semi-online scenarios where the input is available, but has to be accessed sequentially by the algorithm. As a motivation example, consider the scenario where a simple device (e.g. a remote robot) is supposed to process a large amount of data (e.g. a series of orders) in an online fashion. The data are stored and sequentially fed to the device from a powerful entity (base station) over a (wireless) communication link. In order to guide the robot in the processing, the base station may pre-process the data and send some additional information together with each data item. However, since communication rapidly depletes the robots battery, the amount of this additional communication should be kept as small as possible.

To return to the original example, the structure of the SKIRENTAL problem is very simple from the advice complexity point of view: the optimal offline solution is either to buy the skis on the first day or rent them during all days. Hence, although the competitive ratio of the SKIRENTAL problem is 2, a single bit of information about the whole input is sufficient for the online algorithm to achieve optimal performance. As we show later, there are other problems, for which this amount of required information is much higher.

We are primarily interested in the relationship between the competitive ratio and the advice complexity. If the competitive ratio measures the price paid for the lack of information about future, the advice complexity quantifies for how much information is this price paid.

Note that there are two ways to achieve trivial upper bounds on advice complexity: (1) the oracle can send, in some compressed way, the whole input to the algorithm, which then can proceed as an optimal offline algorithm, and (2) the oracle can tell the algorithm exactly what to do in each step. However, both these approaches can be far from optimum. In the first case all information about the future input is communicated, although it may not be relevant³. In the second case, the power of the online algorithm is completely ignored.

³ Consider, e.g. the PAGING problem. There may be a long incompressible sequence of requests that do not result in a page fault; the information about the actual requests in this sequence is useless for the algorithm.

Indeed, an online algorithm may be able to process large parts of the input optimally without any advice, requiring only occasional help from the oracle.

In the paper, we define two modes of interaction with the oracle. In the *helper* mode, the algorithm itself cannot activate the oracle; instead, the oracle oversees the progress of the algorithm, and occasionally sends some pieces of advice. In the *answerer* mode the oracle remains passive, and the algorithm may, in any particular step, ask for advice.

To model the impact of the timing of the communication, let the algorithm work in a synchronous setting: in the i -th step, it receives the i -th input request x_i , and possibly some advice a_i , based on which it produces the output y_i . In a manner usual in the synchronous distributed algorithms (see e.g. [26] and references therein) we count the number of bits communicated between the oracle and the algorithm, relying upon the timing mechanism for delimiting both input and advice sequences⁴. We show that these two modes are different, but are related by $B_H(\mathcal{P}) \leq B_A(\mathcal{P}) \leq 0.92 + B_H(\mathcal{P})$ where $B_H(\mathcal{P})$ is the advice complexity of a problem \mathcal{P} in the helper mode, $B_A(\mathcal{P})$ is the complexity in the answerer mode. Moreover, we analyze two well studied online problems from the point of view of advice complexity, obtaining the results shown in Figure 1. To conclude this section we

	competitive ratio	helper	answerer
PAGING	K [28]	(0.1775, 0.2056)	(0.4591, $0.5 + \varepsilon$)
DIFFSERV	≈ 1.281 [11]	$\frac{1}{K}$	$(\frac{\log K}{2K}, \frac{\log K}{K})$

Fig. 1. Communication complexities of some online problems compared with competitive ratio (asymptotics for large K).

note that there has been a significant amount of research devoted to developing alternative complexity measures for online problems. The competitive ratio has been criticized for not being able to distinguish algorithms with quite different behavior on practical instances, and giving too pessimistic bounds [17]. Hence, several modifications of competitive ratio have been proposed, either tailored to some particular problems (e.g. loose competitiveness [31]), or usable in a more general setting. Among the more general models, many forms of *resource augmentation* have been studied (e.g. [9], [19],[25]). The common idea of these approaches is to counterbalance the lack of information about the input by granting more resources to the online algorithm (e.g. by comparing the optimal offline algorithm to an online algorithm that works k -times faster). Another approach was to use a *look-ahead* where the online algorithm is allowed to see some limited number of future requests [2],[5],[19],[29]. The main problem with the look-ahead approach is that a look-ahead of constant size generally does not improve the worst case performance measured by the competitive ratio. Yet another approach is based on not comparing the online algorithms to offline ones, but to other online algorithms instead (e.g. Max/Max ratio [5], relative worst-order ratio [10]; see also [8], [12]). Still another approach is to limit the power of the adversary as e.g. in the access graph model [7, 18], statistical adversary model [27], diffuse adversary model [22], etc.

⁴ Alternatively, we might require that both the input requests, and the oracle advices come in a self-delimited form. This would alter our upper bounds by a factor of at most 4, as discussed in the appendix.

Finally, a somewhat similar approach of measuring the complexity of a problem by the amount of additional information needed to solve it has been recently pursued in a different setting by Fraigniaud, Gavoille, Ilcinkas, and Pelc [13], [14], [15].

2 Definitions and Preliminaries

An online algorithm receives the input incrementally, one piece at a time. In response to each input portion, the algorithm has to produce output, not knowing the future input. Formally, an online algorithm is modeled by a *request-answer* game [6]:

Definition 1 Consider an input sequence $\mathbf{x} = \langle x_1, x_2, \dots, x_n \rangle$. An online algorithm \mathcal{A} computes the output sequence $\mathbf{y} = \mathcal{A}(\mathbf{x}) = \langle y_1, y_2, \dots, y_n \rangle$, where $y_i = f(x_1, \dots, x_i)$. The cost of the solution is given by a function $C_{\mathcal{A}}(\mathbf{x}) = \text{COST}(\mathbf{y})$.

In the competitive analysis, the online algorithm \mathcal{A} is compared with an optimal offline algorithm OPT , which knows the whole input in advance (i.e. $\mathbf{y} = f(\mathbf{x})$) and can process it optimally. The standard measure of an algorithm \mathcal{A} is the competitive ratio:

Definition 2 An online algorithm is c -competitive, if for each input sequence \mathbf{x} , $C_{\mathcal{A}}(\mathbf{x}) \leq c \cdot C_{\text{OPT}}(\mathbf{x})$

Let us suppose that the algorithm \mathcal{A} is equipped with an oracle \mathcal{O} , which knows \mathcal{A} , can see the whole input, and can communicate with \mathcal{A} . We shall study pairs $(\mathcal{A}, \mathcal{O})$ such that the algorithm (with the help of the oracle) solves the problem optimally. We are interested in the minimal amount of communication between \mathcal{A} and \mathcal{O} , needed to achieve the optimality.

We distinguish two modes of communication: the *helper mode*, and the *answerer mode*. In the helper mode, the oracle (helper) sends in each step i a binary *advice* string \mathbf{a}_i (possibly empty), thus incurring a communication cost of $|\mathbf{a}_i|$. \mathcal{A} can use this advice, together with the input x_1, \dots, x_i to produce the output y_i .

Definition 3 (Online algorithm with a helper) Consider an online algorithm \mathcal{A} , an input sequence $\mathbf{x} = \langle x_1, x_2, \dots, x_n \rangle$, and a helper sequence $\mathcal{O}(\mathbf{x}) = \langle \mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_n \rangle$ of binary strings \mathbf{a}_i . The online algorithm with helper $(\mathcal{A}, \mathcal{O})$ computes the output sequence $\mathbf{y} = \langle y_1, y_2, \dots, y_n \rangle$, where $y_i = f(x_1, \dots, x_i, \mathbf{a}_1, \dots, \mathbf{a}_i)$. The cost of the solution is $C_{(\mathcal{A}, \mathcal{O})}(\mathbf{x}) = \text{COST}(\mathbf{y})$, and the advice (bit) complexity is

$$B_{(\mathcal{A}, \mathcal{O})}^H(\mathbf{x}) = \sum_{i=1}^n |\mathbf{a}_i|$$

In the answerer mode, on the other hand, the oracle is allowed to send an advice only when asked by the algorithm. However, this advice must be a non-empty string. For the ease of presentation we define the answerer oracle as a sequence of non-empty strings. However, only those strings requested by the algorithm are ever considered.

Definition 4 (Online algorithm with an answerer) Consider an algorithm \mathcal{A} , an input sequence $\mathbf{x} = \langle x_1, x_2, \dots, x_n \rangle$, and an answerer sequence $\mathcal{O}(\mathbf{x}) = \langle \mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_n \rangle$ of non-empty binary strings \mathbf{a}_i . The online algorithm with answerer $(\mathcal{A}, \mathcal{O})$ computes the output sequence $\mathbf{y} = \langle y_1, y_2, \dots, y_n \rangle$ as follows:

1. in each step i , a query $r_i \in \{0, 1\}$ is generated first as a function of previous inputs and advices, i.e. $r_i = f_r(x_1, \dots, x_i, r_1 \star \mathbf{a}_1, \dots, r_{i-1} \star \mathbf{a}_{i-1})$ ⁵
2. then, the output is computed as $y_i = f(x_1, \dots, x_i, r_1 \star \mathbf{a}_1, \dots, r_i \star \mathbf{a}_i)$
The cost of the solution is $C_{(\mathcal{A}, \mathcal{O})}(\mathbf{x}) = \text{COST}(\mathbf{y})$, and the advice (bit) complexity is

$$B_{(\mathcal{A}, \mathcal{O})}^A(\mathbf{x}) = \sum_{i=1}^n |r_i \star \mathbf{a}_i|$$

As already mentioned, we are interested in the minimal amount of information the algorithm must get from the oracle, in order to be optimal. For an algorithm \mathcal{A} with an oracle (helper or answerer) \mathcal{O} , the communication cost is the worst case bit complexity, amortized per one step:

Definition 5 Consider an online algorithm \mathcal{A} with an oracle \mathcal{O} using communication mode $M \in \{H, A\}$ ⁶. The bit complexity of the algorithm is

$$B_{(\mathcal{A}, \mathcal{O})}^M = \limsup_{n \rightarrow \infty} \max_{|\mathbf{x}|=n} \frac{B_{(\mathcal{A}, \mathcal{O})}^M(\mathbf{x})}{n}$$

The advice complexity of an online problem \mathcal{P} is the minimum bit complexity of an optimal pair $(\mathcal{A}, \mathcal{O})$:

Definition 6 Consider a problem \mathcal{P} . The advice complexity of \mathcal{P} in communication mode $M \in \{H, A\}$ is

$$B_M(\mathcal{P}) = \min_{(\mathcal{A}, \mathcal{O})} B_{(\mathcal{A}, \mathcal{O})}^M$$

where the minimum is taken over all $(\mathcal{A}, \mathcal{O})$ such that $\forall \mathbf{x} : C_{(\mathcal{A}, \mathcal{O})}(\mathbf{x}) = C_{OPT}(\mathbf{x})$

We start analyzing the advice complexity with an immediate observation that the answerer model is more restrictive in the following sense:

Claim 1 For each problem \mathcal{P} , $B_H(\mathcal{P}) \leq B_A(\mathcal{P}) \leq 0.92 + B_H(\mathcal{P})$.

Proof. The left-hand inequality is obvious. Consider an algorithm \mathcal{A} with an answerer. The same algorithm can be used with a helper with the same bit complexity: the helper can simulate \mathcal{A} and locally compute in each step, if the advice is requested. If so, it sends the advice, otherwise it sends an empty string.

For the right-hand part, the inequality $B_A(\mathcal{P}) \leq 1 + B_H(\mathcal{P})$ is easy to see: an algorithm with an answerer can ask a question in every step. The first bit of the answer indicates, if a helper would send a non-empty advice, and the rest is the actual advice as would be given by the helper.

Now we show how to reduce the constant to 0.92. In the first step, the algorithm asks a question and receives, together with the “usual” advice also some initialization information of $o(n)$ bits. First of all, this information allows the algorithm to distinguish two cases:

Case 1: the number of empty strings sent by the helper is less than $n/3$ or more than $2n/3$. The answerer can, in the initial information, encode the positions of the non-empty strings.

⁵ The function “ \star ” is defined $c \star \alpha = \begin{cases} \text{empty string} & \text{if } c = 0 \\ \alpha & \text{otherwise} \end{cases}$

⁶ In the description of communication modes, H stands for helper and A for answerer.

Then, the algorithm asks only in those steps, in which the helper would send a non-empty advice. To encode the positions of non-empty strings means to encode a binary string of length n with at most $n/3$ ones. There are

$$\sum_{i=0}^{n/3} \binom{n}{i} \leq \frac{n}{3} \cdot \binom{n}{n/3} \approx \frac{n}{3} \cdot \left(\frac{3}{2} \sqrt[3]{2}\right)^n \cdot c \cdot \frac{1}{\sqrt{n}}$$

such strings so that asymptotically

$$\log \frac{n}{3} \cdot \left(\frac{3}{2} \sqrt[3]{2}\right)^n \cdot c \cdot \frac{1}{\sqrt{n}} \approx n \log \left(\frac{3}{2} \sqrt[3]{2}\right) \approx 0.9182n$$

bits are needed to encode the string. Hence the overhead is at most 0.92 bits per request.

Case 2: *the number of empty strings sent by the helper is between $n/3$ and $2n/3$*

For a given i , let μ_i be a binary string of length i , such that it is the least frequently occurring string among the advices of length i given by the helper⁷. The initial information includes the sequence $\mu_1, \mu_2, \dots, \mu_{\log n}$, which is $\log^2 n$ bits overall. The algorithm asks a question in every step, and receives the same answer as the helper would give with the exception that instead of empty string, the string μ_1 is sent, and instead of μ_i , μ_{i+1} is sent. Obviously, since the string $\mu_{\log n}$ is the least frequent of all n strings of length $\log n$, it is never used by the helper, and so does not need to be re-mapped.

To bound the communication overhead, let there be x empty string sent by the helper. These incur x additional bits in the answerer model. From the remaining $n-x$ strings, there may be at most $(n-x)/2^i$ occurrences of μ_i , thus incurring at most $(n-x)/2^i$ additional bits. Obviously, the worst case is when all non-empty strings are of length 1, in which case we get the overhead $x + (n-x)/2$. Since $x < 2n/3$, the overhead is at most $5n/6 \approx 0.83n$. \square

In the lower bound arguments, we shall use the notion of a *communication pattern*. Informally, a communication pattern is the entire information that the algorithm receives from the oracle. Since the algorithms are deterministic, the number of different communication patterns gives the number of different behaviors of the algorithm on a given input.

Definition 7 (Communication pattern – helper) *Consider an algorithm with helper. The communication pattern is defined as the sequence of advices given at each particular step, i.e. $\langle \mathbf{a}_1, \dots, \mathbf{a}_n \rangle$, where \mathbf{a}_i is a, possibly empty, binary string.*

Obviously, the input and communication pattern completely determine the behavior of the algorithm.

Lemma 1 *Consider an algorithm with helper, and let the input sequence be of length $n+1$. For a fixed s , consider only communication patterns in which the helper sends in total at most s bits over all $n+1$ advices. The number X of distinct communication patterns with this property is at most*

$$\log X \leq s \left(\log(1 + \alpha) + 1 + \frac{1}{\ln 2} \right) + \frac{1}{2} \left[\log \left(1 + \frac{1}{\alpha} \right) + \log s \right] + c$$

where $\alpha = \frac{n}{s} > 1$, and c is some constant.

⁷ If there is more than one such string, the choice is arbitrary but fixed.

Proof. For a particular number of bits a , a communication pattern is a string of a bits, distributed among $n + 1$ time slots. Hence there are

$$X = \sum_{a=0}^s 2^a \binom{a+n}{a} \leq s 2^s \binom{s+n}{s}$$

different communication patterns that use at most s bits. For the rest of the proof let us suppose that $s < n$. Using Claim 3 we get

$$\begin{aligned} \log X &\leq \log s + s + \log \binom{s+n}{s} \leq \dots \\ &\leq s + (s+n) \log(s+n) - s \log s - n \log n + \frac{1}{2} [\log(s+n) + \log s - \log n] + c \end{aligned} \quad (1)$$

for some constant c . Let, for $x > 0$, $f(x) = x \log x$, and the derivatives $f'(x) = \frac{1}{\ln 2} + \log x$, $f''(x) = \frac{1}{x \ln 2}$. Since $f(x)$ is convex, it holds

$$(s+n) \log(s+n) - n \log n = f(s+n) - f(n) \leq s \cdot f'(s+n) = s \left(\log(s+n) + \frac{1}{\ln 2} \right) \quad (2)$$

Combining (1) and (2), we get

$$\log X \leq s \left(1 + \log(s+n) + \frac{1}{\ln 2} - \log s \right) + \frac{1}{2} [\log(s+n) + \log s - \log n] + c$$

Denote $\alpha = \frac{n}{s} > 1$, then the previous equation becomes

$$\log X \leq s \left(\log(1+\alpha) + 1 + \frac{1}{\ln 2} \right) + \frac{1}{2} \left[\log \left(1 + \frac{1}{\alpha} \right) + \log s \right] + c$$

□

The situation in the answerer mode is slightly more complicated due to the fact that answers are delivered only when requested.

Definition 8 (Communication pattern – answerer) For each execution of an algorithm with q queries to the answerer, the communication pattern is the sequence $\langle \mathbf{a}_1, \dots, \mathbf{a}_q \rangle$ of non-empty answers.

The behavior of the algorithm is clearly completely determined by the input, the communication pattern and a mapping that assigns to each \mathbf{a}_i the step j_i in which the answer was delivered (see Figure 2).

However, this mapping bears no relevant information: for a given input and communication pattern, the algorithm always receives identical answers, and hence it also asks identical questions. We have the following:

Claim 2 The behavior of an algorithm with answerer is completely determined by its input and communication pattern.

Lemma 2 Consider an algorithm with answerer. For a fixed q , and $s \geq q$, consider only communication patterns, in which the algorithm ask q questions, and s is the total number of bits in all answers. Then there are

$$X = \frac{1}{3} (2^{2s+1} + 1)$$




Fig. 2. Communication pattern mapped to input.

different communication patterns with this property⁸.

Proof. Since each answer is non-empty, there are

$$2^a \cdot \binom{a-1}{q-1}$$

possible communication patterns with exactly a bits. So there are

$$X = 1 + \sum_{q=1}^s \sum_{a=q}^s 2^a \cdot \binom{a-1}{q-1} = 1 + \sum_{a=1}^s 2^a \cdot \sum_{q=0}^{a-1} \binom{a-1}{q} = 1 + \sum_{a=1}^s 2^{2a-1} = \frac{1}{3} (2^{2s+1} + 1)$$

possible communication patterns. □

In the rest of the paper we assume that the algorithm knows the length of the input. Indeed, it is always possible to alter the oracle in such a way that it sends the length of the input⁹ in the first step. Since there are $O(\log n)$ additional bits sent, the normalized contribution to one request is $O(\log n/n)$ which is asymptotically zero.

3 Paging

Paging and its many variants belong to the classical online problems. The virtual memory of a computer is divided into logical pages. At any time K logical pages can reside in the physical memory. A paging algorithm is the part of the operating system responsible for maintaining the physical memory. If a program requests access to a logical page that is not currently in the physical memory, a *page fault* interrupt occurs and the paging algorithm has to transfer the requested page into physical memory, possibly replacing another one. Formally, we define the paging problem as follows:

Definition 9 (Paging Problem) *The input is a sequence of integers (logical pages) $\mathbf{x} = \langle x_1, x_2, \dots, x_n \rangle$, $x_i > 0$. The algorithm maintains a buffer (physical memory) $B = \{b_1, \dots, b_K\}$ of K integers. Upon receiving an input x_i , if $x_i \in B$, $y_i = 0$. If $x_i \notin B$ a page fault is generated, and the algorithm has to find some victim b_j , i.e. $B := B \setminus \{b_j\} \cup \{x_i\}$, and $y_i = b_j$. The cost of the solution is the number of faults, i.e. $COST(\mathbf{y}) = |\{y_i : y_i > 0\}|$.*

⁸ Note that the formula does not depend on q .

⁹ in self-delimited form to distinguish it from the possible advice

It is a well known fact [28] that there is a K -competitive paging algorithm, and that K is the best attainable competitive ratio by any deterministic online algorithm. The optimal offline algorithm is due to [4]. Let us consider the advice complexity of this problem for both helper and answerer modes. We prove that for the helper mode the complexity is between 0.1775 and 0.2056, and for the answerer mode the complexity is between 0.4591 and $0.5 + \varepsilon$ bits per request. Let us first analyze the helper mode. We start with a simple algorithm that uses one bit per request:

Lemma 3 *Consider the PAGING problem. There is an algorithm \mathcal{A} with a helper \mathcal{O} , such that \mathcal{O} sends an advice of exactly one bit each step.*

Proof. Consider an input sequence \mathbf{x} , and an optimal offline algorithm OPT processing it. In each step of OPT , call a page currently in the buffer *active*, if it will be requested again, before OPT replaces it by some other page. We design \mathcal{A} such that in each step i , the set of OPT 's active pages will be in B , and \mathcal{A} will maintain with each page an *active* flag identifying this subset. If \mathcal{A} gets an input x_i that causes a page fault, some passive page is replaced by x_i . Moreover, \mathcal{A} gets with each input also one bit from the helper telling whether x_i is active for OPT . Since the set of active pages is the same for OPT and \mathcal{A} , it is immediate that \mathcal{A} generates the same sequence of page faults. \square

Now we are going to further reduce the advice complexity. The algorithm will still receive the required one bit for every input, however, it is possible to encode the bits in a more efficient way using larger strings as advice:

Lemma 4 *For r large enough, the helper can communicate a binary string of length αr using r bits over a period of αr steps, where $\alpha \approx 4.863876183$.*

Proof. Let $f(x) = (1+x)\log(1+x) - x\log x - x + 1$. $f(x)$ is decreasing, and has exactly one positive root $x_0 \approx 4.863876183$. For any $\varepsilon > 0$, let $\alpha = x_0 - \varepsilon$.

There are

$$X = 2^r \binom{\alpha r + r - 1}{r} = 2^r \frac{\alpha}{\alpha + 1} \binom{r(\alpha + 1)}{r}$$

possible tuples $\langle \mathbf{a}_1, \dots, \mathbf{a}_{\alpha r} \rangle$ of advices containing r bits in total. Using Claim 3, we get

$$X \geq 2^r \frac{(1+\alpha)^{r(1+\alpha)}}{\alpha^{r\alpha}} \frac{1}{1.05^2} \frac{\alpha}{\alpha+1} \sqrt{\frac{1+\alpha}{2\pi r\alpha}}$$

We show that $\lim_{r \rightarrow \infty} (\log X - \alpha r) = \infty$ which means that for large enough r , $X \geq 2^{\alpha r}$.

After some calculations we get

$$\log X - \alpha r \geq r((1+\alpha)\log(1+\alpha) - \alpha \log \alpha - \alpha + 1) + O(\log r) > r + O(\log r)$$

thus finishing the proof. \square

Theorem 1 $B_H(\text{PAGING}(K)) \leq \frac{1}{\alpha}$, where $\alpha \approx 4.863876183$.

Proof. Divide the input into frames of length $\alpha \log n$. The helper sends the first $\alpha \log n$ bits at the beginning. During i th frame, $\log n$ bits are used to communicate the string of $\alpha \log n$ for the next frame using Lemma 4. Overall, there are $\alpha \log n + \frac{n}{\alpha}$ bits. \square

On the lower bound side, we can prove the following:

Theorem 2 For every fixed K , there is a constant $\alpha_K < 20.742$ such that $B_H(\text{PAGING}(2K)) \geq \frac{1}{\alpha_K}$. Moreover, α_K is a decreasing function in K and $\lim_{K \rightarrow \infty} \alpha_K \approx 5.632423693$

Proof. We shall consider a particular subset of input sequences $\mathbf{x} = \{x_k\}_{k=1}^{K(2+3i)}$ for some i . Each input sequence consists of the sequence $S_0 = \langle 1, 2, \dots, 2K \rangle$ followed by i frames, each of length $3K$, where the j th frame has the form $D_j \cdot \langle d_j \rangle \cdot S_j$. The first part of each frame, D_j is of length $K - 1$ and contains unused pages that generate page faults: $D_j = \langle (j+1)K + 1, \dots, (j+2)K - 1 \rangle$. The next request $d_j = (j+2)K$ is again an unused page. The last part, S_j is a sequence of length $2K$ consisting of any subsequence¹⁰ of $S_{j-1} \cdot D_j$ of length $2K - 1$, followed by d_j .




Fig. 3. An example of first two frames for $K = 3$, i.e. with buffer of size 6. The arrows indicate which pages are replaced during faults.

Clearly, since D_j, d_j contain values that have never been used before, any algorithm must generate at least K page faults in each frame. Moreover, there is an algorithm that generates exactly K page faults in each frame as follows. Suppose that after $3Kj$ steps the buffer contains exactly elements from S_j . The following $2K$ steps do not generate any page faults. The next K steps generate a page fault every request, and because from among the $3K - 1$ elements of $S_j \cdot D_{j+1}$ only $2K - 1$ are used in S_{j+1} , the buffer always contains at least one element not in S_{j+1} . Hence, after $3K(j+1)$ steps the buffer contains elements from S_{j+1} .

From the above reasoning it follows that no optimal algorithm can generate a page fault in S_j , which means that at the beginning of S_j , the content of the buffer of any optimal algorithm is uniquely determined.

Now we show that for two different inputs, the communication patterns of an optimal algorithm must be different. Consider two executions of an optimal algorithm with the same communication patterns. We claim that the inputs must have been the same, too. By contradiction, let j be the first frame in which the inputs differ. By the construction of the input sequence, the first two components (D_j, d_j) of the frame are identical. Hence, after $3Kj$ steps (i.e. at the beginning of S_j) the executions are in the same state. However, since the algorithm is optimal, the buffer of both executions contains exactly S_j . Hence the j th frame is identical in both inputs – a contradiction.

For each S_j , there are $\binom{3K-1}{2K-1}$ different possible S_{j+1} , hence there are

$$Y = \binom{3K-1}{2K-1}^i = \left[\frac{2}{3} \binom{3K}{K} \right]^i$$

different inputs of length $2K + 3Ki$, and thus

¹⁰ note that S_j is always an increasing sequence

$$\log Y \geq i \left[K(3 \log 3 - 2) - \frac{\log K}{2} - c_1 \right] \quad \text{where } c_1 = \log(1.05^2 \cdot \sqrt{3\pi}) \approx 1.759 \quad (3)$$

Since any optimal algorithm needs a different communication pattern for each input, there must be at least Y different communication patterns. However, using Lemma 1, we get that there are at most X different communication patterns of length $n + 1$ using at most s bits, where

$$\log X \leq s \left(\log(1 + \alpha) + 1 + \frac{1}{\ln 2} \right) + \frac{1}{2} \left[\log \left(1 + \frac{1}{\alpha} \right) + \log s \right] + c \quad (4)$$

Using (3), and the fact that

$$i = \frac{n + 1 - 2K}{3K} = \frac{\alpha s}{3K} - \frac{2K - 1}{3K}$$

we get

$$\log Y \geq \left(\frac{\alpha s}{3K} - \frac{2K - 1}{3K} \right) \left[K(3 \log 3 - 2) - \frac{\log K}{2} - c_1 \right] \quad (5)$$

Combining (4) and (5), we get

$$\log X - \log Y \leq s \left[\log(1 + \alpha) + 1 + \frac{1}{\ln 2} - \alpha \left(\frac{3 \log 3 - 2}{3} - \frac{\log K}{6K} - \frac{c_1}{3K} \right) \right] + O(\log s) \quad (6)$$

Since $\lim_{n \rightarrow \infty} (\log X - \log Y) \geq 0$, it must hold (possibly for large enough α)

$$F(\alpha, \kappa) := \log(1 + \alpha) + C - \alpha \kappa \geq 0 \quad (7)$$

where

$$C = 1 + \frac{1}{\ln 2} \approx 2.443$$

$$\kappa = \frac{3 \log 3 - 2}{3} - \frac{\log K}{6K} - \frac{c_1}{3K} \approx 0.918 - \frac{\log K}{6K} - \frac{0.586}{K}$$

Consider a fixed K (and hence κ) and examine the function $F_\kappa(\alpha) = F(\alpha, \kappa)$. We have $F_\kappa(0) = C > 0$, and the derivative is $F'_\kappa(\alpha) = \frac{1}{\ln 2(1+\alpha)} - \kappa$. Hence, $F_\kappa(\alpha)$ is increasing up to some point α_{\max} and then decreasing, and because $0.332 \leq \kappa \leq 0.918$ there is a single positive root α_K . Moreover, since $\log(1 + \alpha) + C$ is increasing, α_K is a decreasing function in κ . Because κ as a function of K is increasing (the derivative is positive), α_K is decreasing in K , and by numerically solving for α_1 we get $\alpha_K \leq 20.741$.

Since $\lim_{K \rightarrow \infty} \kappa = \frac{3 \log 3 - 2}{3}$, we have $\lim_{K \rightarrow \infty} \alpha_K = \alpha_\infty$, where α_∞ is the positive root of

$$\log(1 + \alpha) + 1 + \frac{1}{\ln 2} - \alpha \frac{3 \log 3 - 2}{3} = 0$$

Numerically solving, we get

$$\alpha_\infty \approx 5.632423693$$

□

There is one issue connected with the previous proof, namely that the number of logical pages in the constructed sequence was unbounded (since every frame used new values, the number of values depended on n). We argue that this feature can be avoided by reusing the values after a constant number of frames, since every optimal algorithm has to replace all pages from a given frame within the next three frames. Assume the contrary, and consider an optimal algorithm that leaves a page p from a frame i in the buffer during frames $i + 1$ and $i + 2$. Obviously, there must be at least $K + 1$ faults in both of them. However, having p in the buffer for frame $i + 3$ can save at most one page fault.

Let us proceed now with the analysis of the answerer mode. First, we give an upper bound by refining Lemma 3:

Theorem 3 *For each $\varepsilon > 0$, $B_A(\text{PAGING}(K)) \leq \frac{1}{2} + \varepsilon$*

Proof. Let $k = \frac{2}{2\varepsilon+1}$. Without loss of generality, let $n = 2ks$. Following Lemma 3, there is an optimal helper algorithm which receives 1 bit of advice each request. Obviously, the same algorithm works also in the answerer mode. We show how to supply this information using $2s$ bits.

In the first step, the algorithm asks, and receives $O(\log n)$ bits of information containing s and n . Amortized, these $O(\log n)$ bits will contribute $O(\frac{\log n}{n})$ bits per request, and thus can be neglected.

During each of the first s steps, the algorithm asks, and receives a non-empty string of advice. Let $\langle \mathbf{a}_1, \dots, \mathbf{a}_s \rangle$ be the s -tuple of advices, such that $\sum_{i=1}^s |\mathbf{a}_i| = 2s$. Because each \mathbf{a}_i is non-empty, the concatenation of \mathbf{a}_i 's provides one bit every step for the first $2s$ steps. The remaining $n - 2s$ bits are encoded in the lengths of \mathbf{a}_i 's as follows: consider all s -tuples of integers $\langle i_1, \dots, i_s \rangle$ such that $\sum_{j=1}^s i_j = s$, ordered lexicographically. Let z be the number of the tuple $\langle |\mathbf{a}_1| - 1, \dots, |\mathbf{a}_s| - 1 \rangle$ in this ordering, then the representation of z as a $(n - 2s)$ -bit-long binary string gives the remaining $n - 2s$ bits.

What remains to be shown is that there are enough possible combinations of lengths to encode all 2^{n-2s} possible binary strings for the last $n - 2s = 2s(k - 1)$ steps. There are

$$X = \binom{2s-1}{s} = \frac{1}{2} \binom{2s}{s}$$

possible s -tuples $\langle i_1, \dots, i_s \rangle$ such that $\sum_{j=1}^s i_j = s$. Using Claim 3 it follows

$$X \geq 2^{2s} \frac{1}{2 \cdot 1.05^2 \cdot \sqrt{\pi s}}$$

for s large enough, it holds

$$X \geq 2^{2s} \frac{1}{2^{s \frac{8\varepsilon}{2\varepsilon+1}}} = 2^{2s(1 - \frac{4\varepsilon}{2\varepsilon+1})} = 2^{2s(k-1)}$$

The bit complexity of the algorithm is

$$B_{\mathcal{A}, \mathcal{O}} = \limsup_{n \rightarrow \infty} \max_{|\mathbf{x}|=n} \frac{B_{\mathcal{A}}(\mathbf{x}, \mathcal{O}(\mathbf{x}))}{n} = \lim_{s \rightarrow \infty} \frac{2s}{2ks} = \frac{1}{2} + \varepsilon$$

□

To conclude this section, the same technique as used in Theorem 2 can be employed to deliver the corresponding lower bound:

Theorem 4 $B_A(\text{PAGING}(2K)) \geq 0.4591 - O\left(\frac{\log K}{K}\right)$

Proof. Consider the same input sequences as in the proof of Theorem 2. For each S_j , there are $\binom{3K-1}{2K-1}$ different possible S_{j+1} , hence there are

$$Y = \binom{3K-1}{2K-1}^i = \left[\frac{2}{3} \binom{3K}{K}\right]^i$$

different inputs of length $2K + 3Ki$.

Let the answerer use s bits overall. Using Lemma 2, there are

$$X = \frac{1}{3} (2^{2s+1} + 1)$$

possible communication patterns, and obviously it must hold $X \geq Y$. Using Claim 3 we get that

$$\log Y = i \log \left[\frac{2}{3} \binom{3K}{K}\right] \geq i \log \left[\frac{3^{3K}}{2^{2K} \sqrt{K}} \frac{1}{1.05^2 \cdot \sqrt{3\pi}}\right] = i \left[K(3 \log 3 - 2) - \frac{\log K}{2} - c_1\right]$$

for some constant c_1 . Because

$$\log X \leq 2s + c_2 \text{ for some constant } c_2$$

from the fact that $\log X \geq \log Y$ it follows that

$$s \geq i \left[K \frac{3 \log 3 - 2}{2} - \frac{\log K}{4} - \frac{c_1}{2}\right] - \frac{c_2}{2}$$

The bit complexity of any optimal algorithm is then

$$\begin{aligned} B_{\mathcal{A}, \mathcal{O}} &= \limsup_{n \rightarrow \infty} \max_{|\mathbf{x}|=n} \frac{B_{\mathcal{A}}(\mathbf{x}, \mathcal{O}(\mathbf{x}))}{n} \geq \lim_{i \rightarrow \infty} \frac{i \left[K \frac{3 \log 3 - 2}{2} - \frac{\log K}{4} - \frac{c_1}{2}\right] - \frac{c_2}{2}}{2K + 3Ki} = \\ &= \frac{3 \log 3 - 2}{6} - \frac{1}{12} \frac{\log K}{K} - \frac{c_1}{6K} \end{aligned}$$

□

4 Diff-Serv

DIFFSERV is another problem widely studied using competitive analysis (see [11],[23] and references therein). The setting involves a server processing an incoming stream of packets of various values. If the processing speed of the server is slower than the arrival rate, some packets must be dropped, ideally those least valuable. For our purposes, following [23], the packets arrive in discrete time steps. In each step a number of packets can arrive, one packet can be processed, and at most K packets can be stored in a buffer. Moreover, it is required that the packets are processed in FIFO manner. The formal definition is as follows:

Definition 10 (DIFF-SERV problem) Consider a sequence of items $\langle p_1, \dots, p_m \rangle$, partitioned into a series of subsequences, called requests. The input is the sequence of requests $\mathbf{x} = \langle \mathbf{x}_1, \dots, \mathbf{x}_n \rangle$, where each $\mathbf{x}_i = \langle p_{j_{i-1}+1}, \dots, p_{j_i} \rangle$ is a (possibly empty) request.

Each item p_i has a value $v(p_i)$. In each step i , the algorithm maintains an ordered buffer $B_i = \langle b_1, \dots, b_K \rangle$ of K items. Upon receiving a request sequence \mathbf{x}_i , the algorithm discards some elements from the sequence $B_i \cdot \mathbf{x}_i$, keeping some subsequence $B'_i \preceq B_i \cdot \mathbf{x}_i$ of length at most $K + 1$. The first item (if B'_i is nonempty) of B'_i is submitted, and the remainder of the sequence forms the new buffer, i.e. $B'_i = y_i \cdot B_{i+1}$. The process ends if there are no more requests¹¹ and the buffer is empty.

The cost of the solution is the sum of the values of all submitted elements, i.e. $\text{COST}(\mathbf{y}) = \sum_{i>0} v(y_i)$.

For the remainder of this section we shall consider only the case of two distinct item values; we shall refer to them as heavy and light items. Without loss of generality we may assume that each request contains at most $K + 1$ heavy items.

It was shown by Lotker and Patt-Shamir [23] that Algorithm ?? is optimal for the DIFFSERV problem:

Algorithm 1 OPT

- 1: sort all items by decreasing value, within each value class by increasing time
 - 2: $S \leftarrow$ empty schedule
 - 3: **while** the list is non empty **do**
 - 4: $p \leftarrow$ head of the sorted list
 - 5: **if** $S \cup \{p\}$ is a feasible schedule **then** $S \leftarrow S \cup \{p\}$
 - 6: remove head of the list
 - 7: output S
-

We first present another optimal offline algorithm, and then show how to transform it to an online algorithm with a helper.

Let us start with a simple greedy algorithm that never discards more items than necessary (Algorithm 1 without line 4). This algorithm is not optimal in situations where it is favorable to discard leading light items even if the buffer would not be filled¹². These situations, however, can be easily recognized:

Definition 11 Consider a buffer B at time t_0 and the remainder $\{\mathbf{x}_{t_0+i}\}_{i=1}^n$ of the input sequence. Let a_0 be the number of heavy elements in B (before \mathbf{x}_{t_0+1} has arrived), and $a_i \leq K + 1$ be the number of heavy elements in x_{t_0+i} . The remainder of sequence \mathbf{x} is called critical (w.r.t. B), if there exists $t > 0$ such that $\sum_{i=0}^t a_i \geq K + t$, and for each t' such that $0 < t' \leq t$ it holds $\sum_{i=0}^{t'} a_i \geq t'$.

Informally, an input sequence is critical w.r.t. an initial buffer if the buffer gradually fills with heavy items even if the algorithm submits a heavy item in each step.

Our algorithm processes requests sequentially. Each request is processed as shown in Algorithm 1, and it can be proven that this algorithm is optimal.

Lemma 5 Algorithm 1 is optimal.

¹¹ in this case some number of virtual empty requests is added until the buffer is emptied

¹² Consider a situation with a buffer of size 3 containing one light and two heavy items. If there are no more requests, the best solution is to submit all three of them in the next three steps. However, if there is another request coming, containing two heavy items, the best solution is to discard the light one and submit heavy items in the next four steps.

Algorithm 2 Processing of a request x_i with a buffer B

- 1: $B' \leftarrow B \cdot x_i$
 - 2: starting from left, discard light items from B' until $|B'| = K + 1$ or there are no light items left.
 - 3: **if** $|B'| > K + 1$ **then** discard last $|B'| - K - 1$ (heavy) items
 - 4: **if** *the remainder of the input sequence is critical and there are some heavy items in B'* **then** discard leading light items from B'
 - 5: submit the first item of B' (if exists)
 - 6: $B \leftarrow$ remainder of B'
-

Proof. First, we prove that if Algorithm 1 does not submit any item in step i , no optimal algorithm can submit anything in step i . Consider a step i such that Algorithm 1 does not submit (i.e. has empty buffer). Let j be the last time before i when the algorithm had full buffer (or, $j = 0$ if the buffer has never been full). We argue that during the steps $j, j + 1, \dots, i$, no items have been discarded by the algorithm. The algorithm only discards items in such a way that the length of the resulting buffer is less than K only if the input is critical. Moreover, if the buffer is shorter than K in the case of critical input, only light items are discarded. However, starting from a critical input, the buffer eventually (after t steps) fills with heavy items, and is never emptied in-between (if the buffer is shorter than K , no heavy items are discarded). Since j was the last time before i when the buffer was full, clearly there was no critical input between j and i , and hence no item was discarded (since the length of the buffer has always been strictly less than K).

Now consider an arbitrary optimal algorithm A . It is easy to see that during steps $j, j + 1, \dots, i$, the length of the buffer of A is not longer than that of Algorithm 1. Hence, A has empty buffer in step i , too.

Second, we prove that the number of submitted heavy items is optimal. Consider the ordering of the heavy items as they appear in the input sequence. Let an item p be discarded. Let S_p be the set of submitted items ordered before p . We show that it is not possible to submit all items from $S_p \cup \{p\}$. Let i be the arrival time of the item p . Because p was discarded, the buffer was full of heavy items after step i , and the algorithm submitted a heavy item. Consider the longest sequence of steps $j, j + 1, \dots, i$ such that in each step a heavy item was submitted. At the beginning of the sequence, the buffer contained no heavy elements: the input was critical, so a heavy element would have been submitted. Because in all consecutive steps, heavy items were submitted, it is not possible to submit all $S_p \cup \{p\}$ items. Hence, Algorithm 1 maintains for the heavy items the same optimality property as Algorithm ??, and the number of submitted heavy items is optimal.

We have proved that Algorithm 1 submits the maximal possible amount of heavy items, and moreover, for an optimal algorithm A , if A submits in some step, Algorithm 1 submits in this step, too. So it follows that Algorithm 1 gives the optimal cost. \square

Now we turn this offline algorithm into an online algorithm with helper. We are going to simulate Algorithm 1 with an algorithm and a helper. The only place where the algorithm needs information about the future is on line 4, where the algorithm tests the criticality of the input. Clearly, one bit per request (indicating whether the input is critical or not) is sufficient to achieve optimality. However, we show that situation in which a bit must be sent can occur at most once in every $K + 1$ steps.

Theorem 5 $B_H(\text{DIFFSERV}(K)) \leq \frac{1}{K+1}$

Proof. The presented algorithm will operate in two modes: *standard* and *critical*. It starts in standard mode and behaves the same way as Algorithm 1 (leaving out line 4). If it receives a bit from the helper (indicating that the current input forms a critical sequence), it switches into critical mode where it always discards leading light items and submits only heavy items. Once the buffer is full of heavy items, it switches back into standard mode. Obviously, this algorithm is optimal.

Moreover, the helper has to send a bit only when the algorithm is in standard mode, the input is critical and there are leading light items. How often a bit can be sent? Clearly, after the algorithm switches into critical mode, the buffer eventually fills with heavy items. Then there are at least K requests where there are no leading light items so no bit is sent. Hence, the bit is sent at works every $K + 1$ steps. \square

Using a technique similar to the proof of Theorem 2, we can show the following:

Theorem 6 *For $K \geq 4$ it holds $B_H(\text{DIFFSERV}(K)) \geq \frac{1}{\gamma_K \cdot K}$, where $\gamma_K \leq 6.13$ and $\lim_{K \rightarrow \infty} \gamma_K = 1$*

Proof. Consider a sequence of frames of two types. Both types start with a request with one light item followed by one heavy item. The frame of type A continues with a request with $K + 1$ heavy items and K empty requests. The type B frame continues with one empty request. The input consists of β frames of type A and $(c - 1)\beta$ frames of type B for some $c > 1$, so the length of the input is $\beta(K + 2c)$. First note that for two different inputs the communication patterns have to be different: if this is not the case consider two different inputs with the same communication pattern, and focus on the first frame where the inputs differ. Clearly, the algorithm is in the same state in both executions and since it is deterministic, it must make the same decision. However, the optimal decision in frame A is to discard the initial light item, and submit $K + 2$ heavy items in the whole frame, whereas the only optimal decision in frame B is to submit the light and heavy item.

Hence, we have

$$Y = \frac{(c\beta)!}{\beta!((c-1)\beta)!}$$

different inputs of length $n + 1 = \beta(K + 2c)$, each of them requiring a different communication pattern. Using Claim 3 we get

$$\log Y \geq \beta(c \log c - (c - 1) \log(c - 1)) + \frac{1}{2}(\log c - \log(c - 1)) - \frac{1}{2} \log \beta + r$$

for some constant r . Let us denote

$$\psi_c = c \log c - (c - 1) \log(c - 1) \tag{8}$$

Since $\log c - \log(c - 1) < 1$ we get $\log Y \geq \beta\psi_c - \frac{1}{2} \log \beta + r'$ for some constant r' . Replacing the parameter β by the length of the input we get

$$\log Y \geq \frac{n + 1}{K + 2c} \psi_c - \frac{1}{2} \log \frac{n + 1}{K + 2c} + r' \tag{9}$$

However, following Lemma 1 there are at most

$$X = \sum_{a=0}^s 2^a \binom{a + n}{a} \leq s 2^s \binom{s + n}{s}$$

different communication patterns that use at most s bits, and

$$\log X \leq \frac{n}{\alpha} \left(\log(1 + \alpha) + 1 + \frac{1}{\ln 2} \right) + \frac{1}{2} \left[\log \left(1 + \frac{1}{\alpha} \right) + \log \frac{n}{\alpha} \right] + r'' \quad (10)$$

where $\alpha = \frac{n}{s}$ and r'' is a constant. Since $X \geq Y$, by combining (??) and (??) we get

$$0 \leq \frac{n}{\alpha} \left[\log(1 + \alpha) + 1 + \frac{1}{\ln 2} - \frac{\alpha \psi_c}{K + 2c} \right] + \frac{1}{2} [\log n + \log(n + 1) + \log(\alpha + 1) - 2 \log \alpha] + Q_{K,c} \quad (11)$$

where $Q_{K,c}$ does not depend on n . Since n is the leading term in this expression and α is positive, it must hold

$$0 \leq F(\alpha) := \log(1 + \alpha) + 1 + \frac{1}{\ln 2} - \alpha \frac{\psi_c}{K + 2c} \quad (12)$$

First, we try to find a value of c such that the term $\frac{\psi_c}{K+2c}$ is maximized. Let

$$h(c) := \frac{\psi_c}{K + 2c} = \frac{c \log c - (c - 1) \log(c - 1)}{K + 2c} > \frac{\log(c - 1) + \frac{1}{\ln 2}}{K + 2c} =: g(c)$$

The last inequality is due to the fact that $f(c) - f(c - 1) > f'(c - 1)$ where $f(x) = x \log x$ and $f'(x) = \log x + \frac{1}{\ln 2}$. Since

$$g'(c) = \frac{K + 2 - 2(c - 1) \ln(c - 1)}{\ln 2 (K + 2c)^2 (c - 1)}$$

the maximum is attained for

$$c - 1 = \exp \left[\mathcal{W} \left(\frac{K + 2}{2} \right) \right]$$

where \mathcal{W} is the Lambert function.¹³ Taking into account first few terms in the series, we set

$$c := \frac{K + 2}{2 \ln \frac{K+2}{2}} + 1$$

Summarizing, we get

$$h(c) > \frac{\log(c - 1) + \frac{1}{\ln 2}}{K + 2c} = \frac{\ln(K + 2) - \ln \left(2 \ln \frac{K+2}{2} \right) + 1}{\ln 2 (K + 2) \left[1 + \frac{1}{2 \ln \frac{K+2}{2}} \right]}$$

In the sequel, we would like the value $h(c) > \frac{\log K}{\varepsilon K}$ for some ε . By solving

$$\frac{\ln(K + 2) - \ln \left(2 \ln \frac{K+2}{2} \right) + 1}{\ln 2 (K + 2) \left[1 + \frac{1}{2 \ln \frac{K+2}{2}} \right]} \geq \frac{\log K}{\varepsilon K}$$

¹³ The Lambert function \mathcal{W} is the inverse function of $y \mapsto y \cdot e^y$, i.e. $\mathcal{W}(x) \cdot e^{\mathcal{W}(x)} = x$. The asymptotics is given by

$$\mathcal{W}(x) = \ln x - \ln \ln x + \frac{\ln \ln x}{\ln x} - \frac{\ln \ln x - \frac{(\ln \ln x)^2}{2}}{(\ln x)^2} + O \left(\ln \left(\frac{1}{x} \right)^{-3} \right)$$

we get

$$\varepsilon \geq \frac{\ln K(K+2) \left(1 + \frac{1}{2 \ln \frac{K+2}{2}}\right)}{K \left[\ln(K+2) - \ln \left(2 \ln \frac{K+2}{2}\right) + 1\right]}$$

The right hand side converges to 1; it is increasing for $K < 11$ and then decreasing (the actual position of the maximum is around 9.929), and is always smaller than 1.6003.

Now, since $h(c) > \frac{\log K}{\varepsilon K}$, we can write

$$0 \leq F(\alpha) < \log(1 + \alpha) + 1 + \frac{1}{\ln 2} - \frac{\alpha \log K}{\varepsilon K}$$

Let us substitute $\alpha = t \cdot K$ for some $t > 1$ and consider a function

$$G(t) := \log(1 + tK) + 1 + \frac{1}{\ln 2} - \frac{t}{\varepsilon} \log K$$

Since the derivative

$$G'(t) = \frac{K}{\ln 2(tK + 1)} - \frac{\log K}{\varepsilon}$$

and $G(0) > 0$, we have that $G(t)$ is increasing up to some point and then monotonically decreasing into $-\infty$. Hence, if we find any t_0 such that $G(t_0) < 0$, it must follow that $\alpha < t_0 K$.

Using the fact that $\log(1 + x) \leq \log x + \frac{1}{x \ln 2}$ we compute

$$G(t) \leq \log t + \log K \left(1 - \frac{t}{\varepsilon}\right) + \frac{1}{\ln 2} \left(\frac{1}{tK} + 1\right) + 1$$

and, in order to ensure $G(t_0) < 0$, we want to find t_0 such that

$$t_0 \geq \frac{\varepsilon}{\ln K} \ln t_0 + \varepsilon \left[\frac{1}{\ln K} \left(1 + \frac{1}{t_0 K}\right) + \frac{1}{\log K} + 1 \right]$$

Since $t_0 > 1$, we can write

$$t_0 \geq P \ln t_0 + Q$$

where

$$P = \frac{\varepsilon}{\ln K} \quad \text{and} \quad Q = \varepsilon \left[\frac{1}{\ln K} \left(1 + \frac{1}{K}\right) + \frac{1}{\log K} + 1 \right]$$

It holds

$$2P + Q = \varepsilon \left[\frac{1}{\ln K} \left(3 + \frac{1}{K}\right) + \frac{1}{\log K} + 1 \right]$$

Since $K \geq 4$ and $\varepsilon < 1.6003$ it holds

$$2P + Q < 6.13 < e^2$$

Let $t_0 = 2P + Q$; we verify that

$$2P + Q \geq P \ln(2P + Q) + Q$$

From this we know that $\alpha < (2P + Q)K < 6.13$, and that the bit complexity is at least $1/\alpha$. Moreover the value of t_0 is a decreasing function in K which converges to 1. □

In a similar fashion, the following results can be shown for the answerer mode:

Theorem 7 $B_A(\text{DIFFSERV}(K)) \leq \frac{1+\log(K+1)}{K+1}$

Proof. Consider the helper algorithm from Theorem 5, in which the helper sends one bit at certain time intervals that are at least $K + 1$ steps apart. Modify the algorithm in such a way that, initially, the algorithm asks in which step t_1 the helper would send a bit for the first time. In step t_1 the algorithm asks for the next step t_2 in which the helper would send the next bit etc. Let $t_0 = 0$ and $\delta_i = t_i - t_{i-1}$. Then clearly the number of communicated bits is

$$B_n \leq \sum_{i=1}^r \lceil \log \delta_i \rceil \leq r + \sum_{i=1}^r \log \delta_i$$

where r is the number of "critical decisions".¹⁴ Clearly, it holds

$$\sum_{i=1}^r \delta_i = n$$

and

$$r \leq \frac{n}{K+1} + 1$$

Given these constraints, the function $\sum_{i=1}^r \log \delta_i$ is maximized when all δ_i are equal, and r is biggest possible. The maximum value is

$$B_n \leq r + r \log \frac{n}{r} \leq \left(1 + \frac{n}{K+1}\right) \left(1 + \log \frac{n(K+1)}{n+K+1}\right)$$

Taking the limit $\lim_{n \rightarrow \infty} B_n/n$, the proof is concluded □

Theorem 8 For each fixed $K \geq 4$ there exists a $\gamma_K \leq 3.822$ such that $B_A(\text{DIFFSERV}(K)) \geq \frac{\log(K+2)}{\gamma_K(K+2)}$ Moreover $\lim_{K \rightarrow \infty} \gamma_K = 2$.

Proof. By copying the proof of Theorem 6, and using Lemma 2 we get

$$0 \leq \log X - \log Y \leq \frac{n}{\alpha} \left[2 - \frac{\alpha \psi_c}{K+2c} \right] + O(\log n) + Q_{K,c}$$

So it must hold

$$\alpha \geq \frac{2(K+2c)}{\psi_c}$$

Minimizing the same way as in Theorem 6 we set

$$c := \frac{K+2}{2 \ln \frac{K+2}{2}} + 1$$

from which follows

$$\alpha \leq \frac{2 \left(K + \frac{K+2}{\ln \left(\frac{K+2}{2} \right)} + 2 \right)}{\log(K+2) - \log \left(\ln \left(\frac{K+2}{2} \right) \right) + \frac{1}{\ln 2} - 1} \leq \frac{\gamma_K(K+2)}{\log(K+2)}$$

¹⁴ In order to simplify things let us consider also the end of input as a critical decision point.

where

$$\gamma_k = \frac{2 \left(1 + \frac{1}{\ln(K+2) - \ln 2} \right)}{1 - \frac{\ln \ln(K+2)}{\ln(K+2)} - \frac{\alpha}{\log(K+2)}}$$

for

$$\alpha = 1 - \log \ln 2 - \frac{1}{\ln 2} \approx 0.086$$

Since $K \geq 4$, it holds $\gamma_K \leq 2(1 + 0.911) = 3.822$.

□

5 Conclusion

We have proposed a new way to evaluate online problems, based on the communication complexity. While the competitive analysis is an algorithmic measure evaluating the output quality degradation incurred by the requirements to produce the output online, our measure is a structural one quantifying the amount of additional information about the input needed to produce optimal output in an online fashion. The study of the relation between those two measures can lead to a deeper understanding of the nature of online problems. We have shown that there are problems like PAGING and DIFFSERV where the advice complexity (in the helper mode) is proportional to the competitive ratio. On the other hand, there are problems with simple structure like SKIRENTAL [21], which has competitive ratio $2 - \varepsilon$, but a single bit of information is sufficient to solve the problem optimally (i.e. it has zero advice complexity).

Studying advice complexity of a problem can lead to exposure of the critical decisions to be made (like in Algorithm 1 for DIFFSERV) and subsequently to better understanding of the problem and possibly more efficient algorithms. Moreover, we expect that in certain situations involving cooperating devices of uneven computational power communicating over a costly medium (as e.g. in sensor networks), the advice complexity might be of practical interest.

The proposed topic presents a number of intriguing open questions. Is it, for example, possible to characterize a class of problems where the competitive ratio is proportional to the advice complexity? Another whole research area is to study the tradeoff between the amount of communicated information and the achieved competitive ratio.

There is also a number of variations of the model that could be investigated. One potential modification would be to limit the size of advice given in one step. In our model this size is unbounded, and this fact is heavily relied upon (sending the length of the input in one step). However, for modelling potentially infinite inputs it would be more appealing to limit the size of advice to be independent of the input size.

References

1. D. Achlioptas, M. Chrobak, J. Noga: Competitive Analysis of Randomized Paging Algorithms. *Theor. Comp. Sci.* 234 , pp. 203 - 218, 2000
2. S. Albers: On the influence of lookahead in competitive paging algorithms. *Algorithmica* 18, pp. 283-305, 1997
3. S. Albers: Online algorithms: A survey. *Mathematical Programming* 97, pp. 3-26, 2003
4. L.A. Belady: A study of replacement algorithms for virtual storage computers. *IBM Systems Journal* 5, pp. 78-101, 1966

5. S. Ben-David, A. Borodin: A new measure for the study of on-line algorithms. *Algorithmica* 11(1), pp. 73-91, 1994
6. A. Borodin, R. El-Yaniv: *Online Computation and Competitive Analysis*. Cambridge University Press, 1998
7. A. Borodin, S. Irani, P. Raghavan and B. Schieber: Competitive paging with locality of reference. *In Proc. 23rd Annual ACM Symp. on Theory of Computing*, pp. 249-259, 1991
8. J. Boyar, M.R. Ehmsen, and K.S. Larsen: Theoretical Evidence for the Superiority of LRU-2 over LRU for the Paging Problem. *In Fourth Workshop on Approximation on Online Algorithms*, Lecture Notes in Computer Science. Springer-Verlag. Accepted for publication.
9. J. Boyar, K.S. Larsen, and M.N. Nielsen: The Accommodating Function: a generalization of the competitive ratio. *SIAM Journal on Computing* 31(1), pp. 233-258, 2001
10. J. Boyar, L.M. Favrholdt: The Relative Worst Order Ratio for Online Algorithms, *Algorithms and Complexity, 5th Italian Conference, CIAC 2003, Rome, Italy, Lecture Notes in Computer Science 2653*, pp.58-69, 2003
11. M. Englert and M. Westermann: Lower and Upper Bounds on FIFO Buffer Management in QoS Switches. *In Proc. ESA 2006*, LNCS 4168, pp. 352-363, Springer Verlag, 2006
12. A. Fiat, R.M. Karp, M. Luby, L.A. McGeoch, D.D. Sleator, N. E. Young: Competitive Paging Algorithms. *J. Algorithms* 12, pp.685-699, 1991
13. P. Fraigniaud, C. Gavoille, D. Ilcinkas, A. Pelc: Distributed computing with advice: information sensitivity of graph coloring, *In Proc. 34th International Colloquium on Automata, Languages and Programming (ICALP 2007)*, 2007
14. P. Fraigniaud, D. Ilcinkas, A. Pelc: Tree exploration with an oracle, *In Proc. 31st International Symposium on Mathematical Foundations of Computer Science, (MFCS 2006)*, LNCS 4162, pp. 24-37, 2006
15. P. Fraigniaud, D. Ilcinkas, A. Pelc: Oracle size: a new measure of difficulty for communication problems, *In Proc. 25th Ann. ACM Symposium on Principles of Distributed Computing (PODC 2006)*, pp. 179-187, 2006
16. R.L. Graham: Bounds for Certain Multiprocessing Anomalies. *Bell Systems Technical Journal* 45, pp. 1563-1581, 1966
17. S. Irany, A.R. Karlin: Online Computation. *In: D.S. Hochbaum, Ed.: Approximation Algorithms for NP-Hard Problems*, pp. 521-564, PWS Publishing Company, 1997
18. S. Irani, A.R. Karlin and S. Phillips. Strongly competitive algorithms for paging with locality of reference. *In Proc. 3rd Annual ACM-SIAM Symp. on Discrete Algorithms*, pp. 228-236, 1992
19. B. Kalyanasundaram, K. Pruhs: Speed is as Powerful as Clairvoyance. *IEEE Symposium on Foundations of Computer Science*, pp. 214-221, 1995
20. A.R. Karlin, M.S. Manasse, L. Rudolph, D.D. Sleator: Competitive Snoopy Caching. *Algorithmica* 3, pp.79-119, 1988
21. R. Karp: On-line algorithms versus off-line algorithms: how much is it worth to know the future? *Proc. IFIP 12th World Computer Congress*, Vol. 1, pp. 416-429, 1992
22. E. Koutsoupias and C.H. Papadimitriou. Beyond competitive analysis. *In Proc. 34th Annual Symp. on Foundations of Computer Science*, pp. 394-400, 1994.
23. Z. Lotker, B. Patt-Shamir: Nearly Optimal FIFO Buffer Management for DiffServ. *PODC 2002*, pp. 134-143, 2002
24. M.M. Manasse, L.A. McGeoch, D.D. Sleator: Competitive Algorithms for Online Problems. *In Proc. 20th Annual Symposium on the Theory of Computing*, pp. 322-333, 1988
25. C.A. Philips, C. Stein, E. Torng, J. Wein: Optimal Time-Critical Scheduling via Resource Augmentation. *In Proc. 29th Annual ACM Symp on the Theory of Computing*, pp. 140-149, 1997
26. U.M. O'Reilly, N. Santoro: The Expressiveness of Silence: Tight Bounds for Synchronous Communication of Information Using Bits and Silence. *In Proc. of 18th International Workshop on Graph-Theoretic Concepts in Computer Science*, pp. 321-332, 1992
27. P. Raghavan. A statistical adversary for on-line algorithms. *In On-Line Algorithms, DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, pp. 79-83, 1991

28. D.D. Sleator, R.E. Tarjan: Amortized Efficiency of Update and Paging Rules. *Comm. of the ACM* 28(2), pp. 202-208, 1985
29. E. Torng: A Unified Analysis of Paging and Caching. *Algorithmica* 20, pp. 175-200, 1998
30. N. Young: On-line Paging against Adversially Biased Random Inputs. *J. Algorithms* 37, pp.218-235, 2000
31. N. Young: The k -server dual and loose competitiveness for paging, *Algorithmica* 11, pp. 525-541, 1994

Appendix

In this section we show how the requirement of self-delimited advice strings would alter our upper bounds. The following inequality comes from the Stirling formula:

Claim 3 *For each $n \geq 4$ it holds*

$$\sqrt{2\pi n} \left(\frac{n}{e}\right)^n < n! < 1.05 \cdot \sqrt{2\pi n} \left(\frac{n}{e}\right)^n$$

Now consider an algorithm \mathcal{A} with an oracle \mathcal{O} , and an input sequence \mathbf{x} of length n . Let \mathcal{O} send B bits in n_0 different strings a_1, \dots, a_{n_0} . Now consider an algorithm \mathcal{A}' with oracle \mathcal{O}' such that all the advice strings are in a self-delimited form: $\lceil \log |a_i| \rceil$ ones, one zero, then a string of $\lceil \log |a_i| \rceil$ bits encoding the length $|a_i|$ in binary, and, finally, a_i itself. Hence, \mathcal{A}' communicates

$$B' = \sum_{i=1}^{n_0} |a_i| + 2\lceil \log |a_i| \rceil + 1 \leq B + 3n_0 + 2 \sum_{i=1}^{n_0} \log |a_i| \leq B + 3n_0 + 2n_0 \log(B/n_0)$$

bits. Since $n_0 \leq B$, the maximum is attained for $n_0 = B$, so $B' \leq 4B$.