# Weighted Online Problems with Advice * 

Joan Boyar, Lene M. Favrholdt, Christian Kudahl, and Jesper W. Mikkelsen<br>Department of Mathematics and Computer Science, University of Southern Denmark


#### Abstract

Recently, the first online complexity class, AOC, was introduced. The class consists of many online problems where each request must be either accepted or rejected, and the aim is to either minimize or maximize the number of accepted requests, while maintaining a feasible solution. All AOC-complete problems (including Independent Set, Vertex Cover, Dominating Set, and Set Cover) have essentially the same advice complexity. In this paper, we study weighted versions of problems in AOC, i.e., each request comes with a weight and the aim is to either minimize or maximize the total weight of the accepted requests. In contrast to the unweighted versions, we show that there is a significant difference in the advice complexity of complete minimization and maximization problems. We also show that our algorithmic techniques for dealing with weighted requests can be extended to work for non-complete AOC problems such as Matching in the edge arrival model (giving better results than what follow from the general AOC results) and even non-AOC problems such as scheduling.


## 1 Introduction

An online problem is an optimization problem for which the input is divided into small pieces, usually called requests, arriving sequentially. An online algorithm must serve each request, irrevocably, without any knowledge of possible future requests. The quality of online algorithms is traditionally measured using the competitive ratio [12, 16], which is essentially the worst case ratio of the online performance to the performance of an optimal offline algorithm, i.e., an algorithm that knows the whole input sequence from the beginning and has unlimited computational power.

For some online problems such as Independent Set or Vertex Cover, the best possible competitive ratio is linear in the sequence length. This gives rise to the question of what would happen, if the algorithm knew something about future requests. Semi-online settings, where it is assumed that the algorithm has some specific knowledge such as the value of an optimal solution, have been studied (see [6] for many relevant references). The extra knowledge may also be more problem specific such as an access graph for paging [4, 8].

[^0]In contrast to problem specific approaches, advice complexity [3, 9, 11] is a quantitative and standardized way of relaxing the online constraint. The main idea of advice complexity is to provide an online algorithm, AlG, with some partial knowledge of the future in the form of advice bits provided by a trusted oracle which has unlimited computational power and knows the entire request sequence. Informally, the advice complexity of an algorithm is a function of input sequence length, and for a given $n$, it is the maximum number of advice bits read for input sequences of length $n$. The advice complexity of a problem is a function of input sequence length and competitive ratio, and for a given competitive ratio $c$, it is the best possible advice complexity of any $c$-competitive algorithm for the problem. Advice complexity is formally defined in Section 2 ,

Upper bounds on the advice complexity for a problem can sometimes lead to (or come from) semi-online algorithms, and lower bounds can show that such algorithms do not exist. Since its introduction, advice complexity has been a very active area of research. Lower and upper bounds on the advice complexity have been obtained for a large number of online problems; a recent list can be found in [17]. For a survey on advice complexity, see [6].

Recently in [7], the first complexity class for online problems, AOC, was introduced. The class consists of online problems that can be described in the following way: The input is a sequence of requests and each request must either be accepted or rejected. The set of accepted requests is called the solution. For each request sequence, there is at least one feasible solution. The class contains minimization as well as maximization problems. For a minimization problem, the goal is to accept as few requests as possible, while maintaining a feasible solution, and for maximization problems, the aim is to accept as many requests as possible. For minimization problems, any super set of a feasible solution is also a solution, and for maximization problems, any subset of a feasible solution is also a feasible solution. The AOC-complete problems are the hardest problems in the class in terms of their advice complexity. The class AOC is formally defined in Section 3

In this paper, we consider a generalization of the problems in the class AOC in which each request comes with a weight. The goal is now to either minimize or maximize the total weight of the accepted requests. We separately consider the classes of maximization and minimization problems. For AOC-complete maximization problems, we get advice complexity results quite similar to those for the unweighted versions of the problems. On the other hand, for AOC-complete minimization problems, the results are a lot more negative: using less than one advice bit per request leads to unbounded competitive ratios, so this gives a complexity class containing harder problems than AOC. This is in contrast to unweighted AOC-complete problems, where minimization and maximization problems are equally hard in terms of advice complexity. Recently, differences between (unweighted) AOC minimization and maximization problems were found with respect to online bounded analysis [5] and min- and max-induced subgraph problems [13].

Our upper bound techniques are also useful for non-complete AOC problems such as Matching in the edge arrival model, as well as non-AOC problems such as Scheduling.

Previous results. For any AOC-complete problem, $\Theta(n / c)$ advice bits are necessary and sufficient to obtain a competitive ratio of $c$. More specifically, for competitive ratio $c$, the advice complexity is $B(n, c) \pm O(\log n)$, where

$$
\begin{equation*}
B(n, c)=\log \left(1+\frac{(c-1)^{c-1}}{c^{c}}\right) n \tag{1}
\end{equation*}
$$

and $a n / c \leq B(n, c) \leq n / c, a=1 /(e \ln (2)) \approx 0.53$. This is an upper bound on the advice complexity of all problems in AOC. In [7], a list of problems including Independent Set, Vertex Cover, Dominating Set, and Set Cover were proven AOC-complete.

The paper [1 studies a semi-online version of scheduling where it is allowed to keep several parallel schedules and choose the best schedule in the end. The scheduling problem considered is makespan minimization on $m$ identical machines. Using $(1 / \varepsilon)^{O(\log (1 / \varepsilon))}$ parallel schedules, a $(4 / 3+\varepsilon)$-competitive algorithm is obtained. Moreover, a $(1+\varepsilon)$-competitive algorithm which uses $(m / \varepsilon)^{O(\log (1 / \varepsilon) / \varepsilon)}$ parallel schedules is given along with an almost matching lower bound. Note that keeping $s$ different schedules until the end corresponds to working with $s$ different online algorithms. Thus, this particular semi-online model easily translates to the advice model, the advice being which of the $s$ algorithms to run. In this way, the results of [1] correspond to a $(4 / 3+\varepsilon)$-competitive algorithm using $O\left(\log ^{2}(1 / \varepsilon)\right)$ advice bits and a $(1+\varepsilon)$-competitive algorithm using $O(\log (m / \varepsilon) \cdot \log (1 / \varepsilon) / \varepsilon)$ advice bits. In particular, note that this algorithm uses constant advice in the size of the input and only logarithmic advice in the number of machines.

In [15], scheduling on identical machines with a more general type of objective function (including makespan, minimizing the $\ell_{p}$-norm, and machine covering) was studied. The paper considers the advice-with-request model where a fixed number of advice bits are provided along with each request. The main result is a $(1+\varepsilon)$-competitive algorithm that uses $O((1 / \varepsilon) \cdot \log (1 / \varepsilon))$ advice bits per request, totaling $O((n / \varepsilon) \cdot \log (1 / \varepsilon))$ bits of advice for the entire sequence.

Our results. We prove that adding arbitrary weights, AOC-complete minimization problems become a lot harder than AOC-complete maximization problems:

- For AOC-complete maximization problems, the weighted version is not significantly harder than the unweighted version: For any maximization problem in AOC (this includes, e.g., Independent Set), the $c$-competitive algorithm given in $[7$ for the unweighted version of the problem can be converted into a $(1+\varepsilon) c$-competitive algorithm for the weighted version using only $O\left(\left(\log ^{2} n\right) / \varepsilon\right)$ additional advice bits. Thus, a $(1+\varepsilon) c$-competitive algorithm using at most $B(n, c)+O\left(\left(\log ^{2} n\right) / \varepsilon\right)$ bits of advice is obtained. For the
weighted version of non-complete AOC maximization problems, a better advice complexity than $B(n, c)$ may be obtained: For any $c$-competitive algorithm for an AOC maximization problem, P , using $b$ advice bits can be converted into a $O(c \cdot \log n)$-competitive algorithm for the weighted version of P using $b+O(\log n)$ advice bits. For Weighted Matching in the edge arrival model, this implies a $O(\log n)$-competitive algorithm reading $O(\log n)$ bits of advice. We show that this is best possible in the following sense: For a set of weighted AOC problems including Matching, Independent Set and Clique, no algorithm reading $o(\log n)$ bits of advice can have a competitive ratio bounded by any function of $n$. Furthermore, any $O(1)$-competitive algorithm for Matching must read $\Omega(n)$ advice bits.
- For all minimization problems known to be AOC-complete (this includes, e.g., Vertex Cover, Dominating Set, and Set Cover), $n-O(\log n)$ bits of advice are required to obtain a competitive ratio bounded by a function of $n$. This should be contrasted with the fact that $n$ bits of advice trivially yields a strictly 1 -competitive algorithm.
If the largest weight $w_{\max }$ cannot be arbitrarily larger than the smallest weight $w_{\min }$, the $c$-competitive algorithm given in [7] for the unweighted version can be converted into a $c(1+\varepsilon)$-competitive algorithm for the weighted versions using $B(n, c)+O\left(\log ^{2} n+\log \left(\log \left(w_{\max } / w_{\min }\right) / \varepsilon\right)\right)$ advice bits in total.

Our main upper bound technique is a simple exponential classification scheme that can be used to sparsify the set of possible weights. This technique can also be used for problems outside of AOC. For example, for scheduling on related machines, we show that for many important objective functions (including makespan minimization and minimizing the $\ell_{p}$-norm), there exist $(1+\varepsilon)$ competitive algorithms reading $O\left(\left(\log ^{2} n\right) / \varepsilon\right)$ bits of advice. For scheduling on $m$ unrelated machines where $m$ is constant, we get a similar result, but with $O\left((\log n)^{m+1} / \varepsilon^{m}\right)$ advice bits. Finally, for unrelated machines, where the goal is to maximize an objective function, we show that under some mild assumptions on the objective function (satisfied, for example, for machine covering), there is a $(1+\varepsilon)$-competitive algorithm reading $O\left((\log n)^{m+1} / \varepsilon^{m}\right)$ bits of advice.

For scheduling on related and unrelated machines, our results are the first non-trivial upper bounds on the advice complexity. For the case of makespan minimization on identical machines, the algorithm of [1] is strictly better than ours. However, for minimizing the $\ell_{p}$-norm or maximizing the minimum load on identical machines, we exponentially improve the previous best upper bound 15 (which was linear in $n$ ).

## 2 Preliminaries

Throughout the paper, we let $n$ denote the number of requests in the input. We let $\mathbb{R}_{+}$denote the set containing 0 and all positive real numbers. We let log denote the binary $\operatorname{logarithm} \log _{2}$. For $k \geq 1,[k]=\{1,2, \ldots, k\}$. For any bit
string $y$, let $|y|_{0}$ and $|y|_{1}$ denote the number of zeros and the number of ones, respectively, in $y$. We write $x \sqsubseteq y$ if for all indices, $i, x_{i}=1 \Rightarrow y_{i}=1$.

### 2.1 Advice complexity and competitive analysis

In this paper, we use the "advice-on-tape" model [3]. Before the first request arrives, the oracle, which knows the entire request sequence, prepares an advice tape, an infinite binary string. The algorithm Alg may, at any point, read some bits from the advice tape. The advice complexity of ALG is the maximum number of bits read by Alg for any input sequence of at most a given length. Opt is an optimal offline algorithm.

Advice complexity is combined with competitive analysis to determine how many bits of advice are necessary and sufficient to achieve a given competitive ratio.

Definition 1 (Competitive analysis [12, 16] and advice complexity [3]). The input to an online problem, P , is a request sequence $\sigma=\left\langle r_{1}, \ldots, r_{n}\right\rangle$. An online algorithm with advice, ALG, computes the output $y=\left\langle y_{1}, \ldots, y_{n}\right\rangle$, where $y_{i}$ is computed from $\varphi, r_{1}, \ldots, r_{i}$, where $\varphi$ is the content of the advice tape. Each possible output for P is associated with a cost/profit. For a request sequence $\sigma$, $\operatorname{AlG}(\sigma)(\operatorname{Opt}(\sigma))$ denotes the cost/profit of the output computed by Alg (Opt) when serving $\sigma$.

If P is a minimization (maximization) problem, then ALG is $c(n)$-competitive if there exists a constant, $\alpha$, such that, for all $n \in \mathbb{N}, \operatorname{AlG}(\sigma) \leq c(n) \cdot \operatorname{Opt}(\sigma)+\alpha$, ( $\operatorname{Opt}(\sigma) \leq c(n) \cdot \operatorname{AlG}(\sigma)+\alpha)$, for all request sequences, $\sigma$, of length at most $n$. If the relevant inequality holds with $\alpha=0$, we say that AlG is strictly $c(n)$ competitive.

The advice complexity, $b(n)$, of an algorithm, Alg, is the largest number of bits of $\varphi$ read by ALG over all possible request sequences of length at most $n$. The advice complexity of a problem, P , is a function, $f(n, c), c \geq 1$, such that the smallest possible advice complexity of a strictly c-competitive online algorithm for P is $f(n, c)$.

We only consider deterministic online algorithms (with advice). Note that both $b(n)$ and $c(n)$ in the above definition may depend on $n$, but, for ease of notation, we often write $b$ and $c$ instead of $b(n)$ and $c(n)$. Also, with this definition, $c \geq 1$, for both minimization and maximization problems.

### 2.2 Complexity classes

In this paper, we consider the complexity class AOC from [7].
Definition 2 (AOC [7]). A problem, P, is in AOC (Asymmetric Online Covering) if it can be defined as follows: The input to an instance of P consists of a sequence of $n$ requests, $\sigma=\left\langle r_{1}, \ldots, r_{n}\right\rangle$, and possibly one final dummy request. An algorithm for P computes a binary output string, $y=y_{1} \ldots y_{n} \in\{0,1\}^{n}$, where $y_{i}=f\left(r_{1}, \ldots, r_{i}\right)$ for some function $f$.

For minimization (maximization) problems, the score function, s, maps a pair, $(\sigma, y)$, of input and output to a cost (profit) in $\mathbb{N} \cup\{\infty\}(\mathbb{N} \cup\{-\infty\})$. For an input, $\sigma$, and an output, $y, y$ is feasible if $s(\sigma, y) \in \mathbb{N}$. Otherwise, $y$ is infeasible. There must exist at least one feasible output. Let $S_{\min }(\sigma)\left(S_{\max }(\sigma)\right)$ be the set of those outputs that minimize (maximize) sfor a given input $\sigma$.

If P is a minimization problem, then for every input, $\sigma$, the following must hold:

1. For a feasible output, $y, s(\sigma, y)=|y|_{1}$.
2. An output, $y$, is feasible if there exists a $y^{\prime} \in S_{\min }(\sigma)$ such that $y^{\prime} \sqsubseteq y$. If there is no such $y^{\prime}$, the output may or may not be feasible.

If P is a maximization problem, then for every input, $\sigma$, the following must hold:

1. For a feasible output, $y, s(\sigma, y)=|y|_{0}$.
2. An output, $y$, is feasible if there exists a $y^{\prime} \in S_{\max }(\sigma)$ such that $y^{\prime} \sqsubseteq y$. If there is no such $y^{\prime}$, the output may or may not be feasible.

Recall that no problem in AOC requires more than $B(n, c)+O(\log n)$ bits of advice (see Eq. (11) for the definition of $B(n, c)$ ). This result is based on a covering design technique, where the advice indicates a superset of the output bits that are 1 in an optimal solution.

The problems in AOC requiring the most advice are AOC-complete [7]:
Definition 3 (AOC-complete [7]). A problem $\mathrm{P} \in \mathrm{AOC}$ is AOC-complete if for all $c>1$, any c-competitive algorithm for P must read at least $B(n, c)-$ $O(\log n)$ bits of advice.

In [7, an abstract guessing game, minASGk (Minimum Asymmetric String Guessing with Known History), was introduced and shown to be AOC-complete. The minASGK-problem itself is very artificial, but it is well-suited as the starting point of reductions. All minimization problems known to be AOC-complete have been shown to be so via reductions from minASGk.

The input for minASGK is a secret string $x=x_{1} x_{2} \ldots x_{n} \in\{0,1\}^{n}$ given in $n$ rounds. In round $i \in[n]$, the online algorithm must answer $y_{i} \in\{0,1\}$. Immediately after answering, the correct answer $x_{i}$ for round $i$ is revealed to the algorithm. If the algorithm answers $y_{i}=1$, it incurs a cost of 1 . If the algorithm answers $y_{i}=0$, then it incurs no cost if $x_{i}=0$, but if $x_{i}=1$, then the output of the algorithm is declared to be infeasible (and the algorithm incurs a cost of $\infty)$. The objective is to minimize the total cost incurred. Note that the optimal solution has cost $|x|_{1}$. See the appendix for a formal definition of minASGK and for definitions of other AOC-complete problems.

The problem minASGK is based on the binary string guessing problem [2,11]. Binary string guessing is similar to asymmetric string guessing, except that any wrong guess ( 0 instead of 1 or 1 instead of 0 ) gives a cost of 1 .

In Theorem 1, we show a very strong lower bound for a weighted version of minASGk. In Theorem 2 via reductions, we show that this lower bound implies
similar strong lower bounds for the weighted version of other AOC-complete minimization problems.

Definition 4 (Weighted AOC). Let P be a problem in AOC. We define the weighted version of P , denoted $\mathrm{P}_{w}$, as follows: $A \mathrm{P}_{w}$-input $\sigma=\left\langle\left\{r_{1}, w_{1}\right\}\right.$, $\left.\left\{r_{2}, w_{2}\right\}, \ldots,\left\{r_{n}, w_{n}\right\}\right\rangle$ consists of $n$ P-requests, $r_{1}, \ldots, r_{n}$, each of which has a weight $w_{i} \in \mathbb{R}_{+}$. The P -request $r_{i}$ and its weight $w_{i}$ are revealed simultaneously. An output $y=y_{1} \ldots y_{n} \in\{0,1\}^{n}$ is feasible for the input $\sigma$ if and only if $y$ is feasible for the P -input $\left\langle r_{1}, \ldots, r_{n}\right\rangle$. The cost (profit) of an infeasible solution is $\infty(-\infty)$.

If P is a minimization problem, then the cost of a feasible $\mathrm{P}_{w}$-output $y$ for an input $\sigma$ is

$$
s(\sigma, y)=\sum_{i=1}^{n} w_{i} y_{i}
$$

If P is a maximization problem, then the profit of a feasible $\mathrm{P}_{w}$-output $y$ for an input $\sigma$ is

$$
s(\sigma, y)=\sum_{i=1}^{n} w_{i}\left(1-y_{i}\right)
$$

## 3 Weighted Versions of AOC-Complete Minimization Problems

In the weighted version of minASGK, minASGK ${ }_{\mathrm{w}}$, each request is a weight for the current request and the value 0 or 1 of the previous request. Producing a feasible solution requires accepting (answering 1 to) all requests with value 1 , and the cost of a feasible solution is the sum of all weights for requests which are accepted.

We start with a negative result for MINASGK ${ }_{W}$ and then use it to obtain similar results for the weighted online version of Vertex Cover, Set Cover, Dominating Set, and Cycle Finding.

Theorem 1. For minASGK ${ }_{\mathrm{W}}$, no algorithm using less than $n$ bits of advice is $f(n)$-competitive, for any function $f$.

Proof. Let AlG be any algorithm for minASGK ${ }_{\mathrm{W}}$ reading at most $n-1$ bits of advice. We show how an adversary can construct input sequences where the cost of ALG is arbitrarily larger than that of Оpt. We only consider sequences with at least one 1 . It is easy to see that for the unweighted version of the binary string guessing problem, $n$ bits of advice are necessary in order to guess correctly each time: If there are fewer than $n$ bits, there are only $2^{n-1}$ possible advice strings, so, even if we only consider the $2^{n}-1$ possible inputs with at least one 1 , there are at least two different request strings, $x$ and $y$, which get the same advice string. Alg will make an error on one of the strings when guessing the first bit where $x$ and $y$ differ, since up until that point AlG has the same information about both strings.

We describe a way to assign weights to the requests in minASGK ${ }_{w}$ such that if Alg makes a single mistake (either guessing 0 when the correct answer is 1 or vice versa), its performance ratio is unbounded. We use a large number $a>1$, which we allow to depend on $n$. All weights are from the interval $[1, a]$ (note that they are not necessarily integers). We let $x=x_{1}, \ldots, x_{n}$ be the input string and set $w_{1}=a^{1 / 2}$. For $i>1, w_{i}$ is given by:

$$
w_{i}=\left\{\begin{array}{cl}
w_{i-1} \cdot a^{\left(-2^{-i}\right)}, & \text { if } x_{i-1}=0 \\
w_{i-1} \cdot a^{\left(2^{-i}\right)}, & \text { if } x_{i-1}=1
\end{array}\right.
$$

Since the weights are only a function of previous requests, they do not reveal any information to Alg about future requests.

Observation 1 For each $i$, the following hold:
(a) If $x_{i}=0$, then $w_{j} \leq w_{i} \cdot a^{\left(-2^{-n}\right)}$ for all $j>i$.
(b) If $x_{i}=1$, then $w_{j} \geq w_{i} \cdot a^{\left(2^{-n}\right)}$ for all $j>i$.

We argue for each set of inequalities in the observation:
(四): If $x_{i}=0$, for each $j>i, w_{j}=w_{i} \cdot a^{\left(-2^{-(i+1)}\right)} \cdot a^{\sum_{k=i+2}^{j}\left( \pm 2^{-(i+1)}\right)}$, where the plus or minus depends on whether $x_{k}=0$ or $x_{k}=1$. The value $w_{j}$ is largest if all of the $x_{k}$ values are 1 , in which case $w_{j}=w_{i} \cdot a^{\left(-2^{-j}\right)} \leq w_{i} \cdot a^{\left(-2^{-n}\right)}$.
(b): The argument of $x_{i}=1$ is similar, changing minus to plus and vice versa.

We claim that if Alg makes a single mistake, its performance ratio is not bounded by any function of $n$. Indeed, if Alg guesses 0 for a request, but the correct answer is 1 , the solution is infeasible and Alg gets a cost of $\infty$.

We now consider the case where Alg guesses 1 for a request $j$, but the correct answer is 0 . This request gives a contribution of $w_{j}=a^{b}$, for some $0<b<1$, to the cost of the solution produced by Alg. Define $j^{\prime}$ such that $w_{j^{\prime}}=\max \left\{w_{i} \mid x_{i}=1\right\}$. Since Opt only answers 1 if $x_{i}=1$, this is the largest contribution to the cost of Opt from a single request.

If $j^{\prime}>j$, Observation (1) (a) gives that $w_{j^{\prime}} \leq w_{j} \cdot a^{\left(-2^{-n}\right)}=a^{b} \cdot a^{\left(-2^{-n}\right)}=$ $a^{b-2^{-n}}$. The cost of OPT is at most $n \cdot w_{j^{\prime}} \leq n \cdot a^{b-2^{-n}}$. Thus,

$$
\frac{\operatorname{ALG}(x)}{\operatorname{OPT}(x)} \geq \frac{a^{b}}{n \cdot a^{b-2^{-n}}}=\frac{a^{2^{-n}}}{n}
$$

Since $a$ can be arbitrarily large (recall that it can be a function of $n$ ), we see that no algorithm can be $f(n)$-competitive for any specific function $f$.

If $j^{\prime}<j$, Observation (1) gives us that $w_{j} \geq w_{j^{\prime}} \cdot a^{\left(2^{-n}\right)}$. Using $w_{j}=a^{b}$, we get $a^{b-2^{-n}} \geq w_{j^{\prime}}$. We can repeat the argument from the case where $j^{\prime}>j$ to see that no algorithm can be $f(n)$-competitive for any specific function $f$.

In order to show that similar lower bounds apply to all minimization problems known to be complete for AOC, we define a simple type of advice preserving reduction for online problems. These are much less general than those defined
by Sprock in his PhD dissertation [18], mainly because we do not allow the amount of advice needed to change by a multiplicative factor.

Let $\operatorname{OPT}_{\mathrm{P}}(\sigma)$ denote the value of the optimal solution for request sequence $\sigma$ for problem P , and let $|\sigma|$ denote the number of requests in $\sigma$.

Definition 5. Let $\mathrm{P}_{1}$ and $\mathrm{P}_{2}$ be two online minimization problems, and let $\mathcal{I}_{1}$ be the set of request sequences for $\mathrm{P}_{1}$ and $\mathcal{I}_{2}$ be the set of request sequences for $\mathrm{P}_{2}$. For a given function $g: \mathbb{N} \rightarrow \mathbb{R}_{+}$, we say that there is a length preserving $g$-reduction from $\mathrm{P}_{1}$ to $\mathrm{P}_{2}$, if there is a transformation function $f: \mathcal{I}_{1} \rightarrow \mathcal{I}_{2}$ such that

- for all $\sigma \in \mathcal{I}_{1},|\sigma|=|f(\sigma)|$, and
- for every algorithm $\mathrm{AlG}_{2}$ for $\mathrm{P}_{2}$, there is an algorithm $\mathrm{ALG}_{1}$ for $\mathrm{P}_{1}$ such that for all $\sigma_{1} \in \mathcal{I}_{1}$, the following holds:
If $\mathrm{ALG}_{2}$ produces a feasible solution for $\sigma_{2}=f\left(\sigma_{1}\right)$ with advice $\phi\left(\sigma_{2}\right)$, then $\mathrm{ALG}_{1}$, using at most $\left|\phi\left(\sigma_{2}\right)\right|+g\left(\left|\sigma_{2}\right|\right)$ advice bits, produces a feasible solution for $\sigma_{1}$ such that
- $\operatorname{ALG}_{1}\left(\sigma_{1}\right) \leq \operatorname{ALG}_{2}\left(\sigma_{2}\right)+\operatorname{OPT}_{\mathrm{P}_{1}}\left(\sigma_{1}\right)$ and $\operatorname{OPT}_{\mathrm{P}_{1}}\left(\sigma_{1}\right) \geq \operatorname{OPT}_{\mathrm{P}_{2}}\left(\sigma_{2}\right)$, or
- $\operatorname{ALG}_{1}\left(\sigma_{1}\right)=\operatorname{OPT}_{\mathrm{P}_{1}}\left(\sigma_{1}\right)$

Note that the transformation function $f$ is length-preserving in that the lengths of the request sequences for the two problems are identical. This avoids the potential problem that the advice for the two problems could be functions of two different sequence lengths. The amount of advice for the problem being reduced to is allowed to be an additive function, $g(n)$, longer than for the original problem, because this seems to be necessary for some of the reductions showing that problems are AOC-complete. Since the reductions are only used here to show that no algorithm is $F(n)$-competitive for any function $F$, the increase in the performance ratio that occurs with these reductions is insignificant.

The following lemma shows how length-preserving reductions can be used.
Lemma 1. Let $\mathrm{P}_{1}$ and $\mathrm{P}_{2}$ be online minimization problems. Suppose that at least $b_{1}(n, c)$ advice bits are required to be $(c+1)$-competitive for $\mathrm{P}_{1}$ and suppose there is a length preserving $g(n)$-reduction from $\mathrm{P}_{1}$ to $\mathrm{P}_{2}$. Then, at least $b_{1}(n, c)-$ $g(n)$ advice bits are needed for an algorithm for $\mathrm{P}_{2}$ to be c-competitive.

Proof. Let $f$ be the transformation function associated with $g$. Suppose for the sake of contradiction that there is a (strictly) c-competitive algorithm $\mathrm{ALG}_{2}$ for $\mathrm{P}_{2}$ with advice complexity $b_{2}(n, c)<b_{1}(n, c)-g(n)$. Then there exists a constant $\alpha$ such that for any request sequence $\sigma_{1} \in \mathcal{I}_{1}$, either $\operatorname{ALG}_{1}\left(\sigma_{1}\right)=\operatorname{OPT}_{\mathrm{P}_{1}}\left(\sigma_{1}\right)$ or

$$
\begin{aligned}
\operatorname{ALG}_{1}\left(\sigma_{1}\right) & \leq \operatorname{ALG}_{2}\left(\sigma_{2}\right)+\operatorname{OPT}_{\mathrm{P}_{1}}\left(\sigma_{1}\right) \\
& \leq c \cdot \operatorname{OPT}_{\mathrm{P}_{2}}\left(\sigma_{2}\right)+\alpha+\operatorname{OPT}_{\mathrm{P}_{1}}\left(\sigma_{1}\right) \\
& \leq(c+1) \cdot \operatorname{OPT}_{\mathrm{P}_{1}}\left(\sigma_{1}\right)+\alpha,
\end{aligned}
$$

where $\sigma_{2}=f\left(\sigma_{1}\right)$. Thus, ALG ${ }_{1}$ is (strictly) $(c+1)$-competitive, with less than $b_{1}(n, c)$ bits of advice, a contradiction.

All known AOC-complete problems were proven complete using length-preserving reductions from MINASGK, so the following holds for the weighted versions of all such problems:

Theorem 2. For the weighted online versions of Vertex Cover, Cycle Finding, Dominating Set, Set Cover, an algorithm reading less than $n-O(\log n)$ bits of advice cannot be $f(n)$-competitive for any function $f$.

Proof. The reductions in [7] showing that these problems are AOC-complete are length preserving $O(\log n)$-reductions from minASGk, and hence, the theorem follows from Lemma 1 For Vertex Cover, the following $O(\log n)$-reduction can be used (the other three reductions are given in the Appendix (B):

Each input $\sigma=\left\langle x_{1}, x_{2}, \ldots, x_{n}\right\rangle$ to the problem MINASGK, is transformed to $f(\sigma)=\left\langle v_{1}, v_{2}, \ldots, v_{n}\right\rangle$, where $V=\left\{v_{1}, v_{2}, \ldots, v_{n}\right\}$ is the vertex set of a graph with edge set $E=\left\{\left(v_{i}, v_{j}\right): x_{i}=1\right.$ and $\left.i<j\right\}$. Let $V_{1}=\left\{v_{i} \in V: x_{1}=1\right\}$. Note that $V_{1} \backslash\left\{v_{n}\right\}$ is a minimum vertex cover of the graph and that no algorithm can reject more than one vertex from $V_{1}$, since $V_{1}$ induces a clique.

The advice used by the minASGk algorithm ALG $_{1}$ consists of the advice used by the Vertex Cover algorithm $\mathrm{ALG}_{2}$ and $O(\log n)$ bits that are either all 0 or give (an encoding of) an index to a position $i$ in the input sequence, such that $v_{i} \in V_{1}$ and $\mathrm{AlG}_{2}$ rejects $v_{i}$.

Let $V_{\mathrm{ALG}_{2}} \subseteq V$ be the vertex cover constructed by $\mathrm{ALG}_{2}$ and let $X_{\mathrm{ALG}_{1}}$ be the set of requests on which $\mathrm{ALG}_{1}$ returns a 1. Then either $X_{\mathrm{ALG}_{1}}=V_{\mathrm{ALG}_{1}}$ or $X_{\mathrm{ALG}_{1}}=$ $V_{\mathrm{ALG}_{2}} \cup\left\{v_{i}\right\}$, where $\left\{v_{i}\right\}=V_{1} \backslash V_{\mathrm{ALG}_{1}}$. Thus, $\operatorname{ALG}_{1}(\sigma) \leq \operatorname{ALG}_{2}(f(\sigma))+\operatorname{Opt}(\sigma)$, since $w_{i} \leq \operatorname{Opt}(\sigma)$.

## 4 Exponential Sparsification

Assume that we are faced with an online problem for which we know how to obtain a reasonable competitive ratio, possibly using advice, in the unweighted version (or when there are only few possible different weights). We use exponential sparsification, a simple technique which can be of help when designing algorithms with advice for weighted online problems by reducing the number of different possible weights the algorithm has to handle. The first step is to partition the set of possible weights into intervals of exponentially increasing length, i.e., for some small $\varepsilon, 0<\varepsilon<1$,

$$
\mathbb{R}_{+}=\bigcup_{k=-\infty}^{\infty}\left[(1+\varepsilon)^{k},(1+\varepsilon)^{k+1}\right)
$$

How to proceed depends on the problem at hand. We now informally explain the meta-algorithm that we repeatedly use in this paper. Note that if $w_{1}, w_{2} \in$ $\left[(1+\varepsilon)^{k},(1+\varepsilon)^{k+1}\right)$ and $w_{1} \leq w_{2}$, then $w_{1} \leq w_{2} \leq(1+\varepsilon) w_{1}$. For many online problems, this means that an algorithm can treat all requests whose weights belong to this interval as if they all had weight $(1+\varepsilon)^{k+1}$ with only a small loss in competitiveness.

Consider now a set of weights and let $w_{\max }$ denote the largest weight in the set. Let $k_{\max }$ be the integer for which $w_{\max } \in\left[(1+\varepsilon)^{k_{\max }},(1+\varepsilon)^{k_{\max }+1}\right)$. We say that a request with weight $w \in\left[(1+\varepsilon)^{k},(1+\varepsilon)^{k+1}\right)$ is unimportant if $k<k_{\max }-\left\lceil\log _{1+\varepsilon}\left(n^{2}\right)\right\rceil$. Furthermore, we will often categorize the request as important if $k_{\max }-\left\lceil\log _{1+\varepsilon}\left(n^{2}\right)\right\rceil \leq k<k_{\max }+1$ and as huge if $k \geq k_{\max }+1$. Each unimportant request has weight $w \leq(1+\varepsilon)^{k+1} \leq(1+\varepsilon)^{k_{\max }-\left\lceil\log _{1+\varepsilon}\left(n^{2}\right)\right\rceil-1+1} \leq$ $w_{\max } / n^{2}$, so the total sum of the unimportant weights is $O\left(w_{\max } / n\right)$. For many weighted online problems, this means that an algorithm can easily serve the requests with unimportant weights, as follows. In maximization problems, this is done by rejecting them. In minimization problems, it is done by accepting them. Thus, exponential sparsification (when applicable) essentially reduces the problem of computing a good approximate solution for a problem with $n$ distinct weights to that of computing a good approximate solution with only $O\left(\log _{1+\varepsilon} n\right)$ distinct weights.

For a concrete problem, several modifications of this meta-algorithm might be necessary. Often, the most tricky part is how the algorithm can learn $k_{\text {max }}$ without using too much advice. One approach that we often use is the following: The oracle encodes the index $i$ of the first request whose weight is close enough to $(1+\varepsilon)^{k_{\max }}$ that the algorithm only needs a little bit of advice to deduce $k_{\max }$ from the weight of this request. If it is somehow possible for the algorithm to serve all requests prior to $i$ reasonably well, then this approach works well.

Our main application of exponential sparsification is to weighted AOC problems. We begin by considering maximization problems. Note that no assumptions are made about the weights of $\mathrm{P}_{\mathrm{w}}$ in Theorem 3,

Theorem 3. If $\mathrm{P} \in \mathrm{AOC}$ is a maximization problem, then for any $c>1$ and $0<\varepsilon \leq 1, \mathrm{P}_{w}$ has a strictly $(1+\varepsilon) c$-competitive algorithm using $B(n, c)+$ $O\left(\varepsilon^{-1} \log ^{2} n\right)$ advice bits.

Proof. Fix $\varepsilon>0$. Let $\sigma=\left\langle\left\{r_{1}, w_{1}\right\}, \ldots,\left\{r_{n}, w_{n}\right\}\right\rangle$ be the input and let $x=$ $x_{1} \ldots x_{n} \in\{0,1\}^{n}$ specify an optimal solution for $\sigma$, with zeros indicating membership in the optimal solution. Throughout most of this proof, we assume that $n$ is sufficiently large. The necessary conditions are discussed at the end of the proof, along with how to handle small $n$.

Define $s=1+\varepsilon / 2$. Let $V_{\mathrm{OPT}}=\left\{i: x_{i}=0\right\}$. Note that $V_{\mathrm{OPt}}$ contains exactly those rounds in which Opt answers 0 and thus accepts. Furthermore, for $k \in \mathbb{Z}$, let $V^{k}=\left\{i: s^{k} \leq w(i)<s^{k+1}\right\}$ and let $V_{\mathrm{OPT}}^{k}=V_{\mathrm{OPT}} \cap V^{k}$. Finally, let $i_{\max } \in V_{\mathrm{OPT}}$ be such that $w\left(i_{\max }\right) \geq w(i)$ for every $i \in V_{\text {OPT }}$.

The oracle computes the unique $m \in \mathbb{Z}$ such that $i_{\max } \in V_{\mathrm{OPT}}^{m}$. We say that a request $r_{i}$ is unimportant if $w(i)<s^{m-\left\lceil\log _{s}\left(n^{2}\right)\right\rceil \text {, important if } s^{m-\left\lceil\log _{s}\left(n^{2}\right)\right\rceil} \leq, ~}$ $w(i)<s^{m+1}$, and huge if $w(i) \geq s^{m+1}$. The oracle computes the index $i^{\prime}$ of the first important request in the input sequence. Assume that $i^{\prime} \in V^{m^{\prime}}$. The oracle writes the length $n$ of the input onto the advice tape using a self-delimiting encoding ${ }^{11}$, and then writes the index $i^{\prime}$ and the integer $m-m^{\prime}$ (which is at most

[^1]$\left.\left\lceil\log _{s}\left(n^{2}\right)\right\rceil\right)$ onto the tape, using a total of $O(\log n)$ bits. This advice allows the algorithm to learn $m$ as soon as the first important request arrives. From there on, the algorithm will know if a request is important, unimportant, or huge. Whenever an unimportant or a huge request arrives, the algorithm answers 1 (rejects the request). We now describe how the algorithm and oracle work for the important requests.

For each $0 \leq j \leq\left\lceil\log _{s}\left(n^{2}\right)\right\rceil$, let $n_{m-j}=\left|V^{m-j}\right|$. For the requests (whose indices are) in $V^{m-j}$, we use the covering design based $c$-competitive algorithm for unweighted AOC-problems. This requires $B\left(n_{m-j}, c\right)+O\left(\log n_{m-j}\right)$ bits of advice. Since $B(n, c)$ is linear in $n$, this means that we use a total of

$$
b=\sum_{j=0}^{\left\lceil\log _{s}\left(n^{2}\right)\right\rceil}\left(B\left(n_{m-j}, c\right)+O\left(\log n_{m-j}\right)\right) \leq B(n, c)+O\left(\log _{s} n \cdot \log n\right)
$$

bits of advice. Note that $\log _{s}(n) \leq 2 \varepsilon^{-1} \log n$ for $\varepsilon / 2 \leq 1$, giving the bound on the advice in the statement of the theorem.

We now prove that the algorithm achieves the desired competitiveness. We can ignore the huge requests, since neither Alg nor Opt accepts any of them. Let $V_{\mathrm{ALG}}$ be those rounds in which ALG answers 0 and let $V_{\mathrm{ALG}}^{k}=V_{\mathrm{ALG}} \cap V^{k}$. We consider the important requests first. Fix $0 \leq j \leq\left\lceil\log _{s}\left(n^{2}\right)\right\rceil$. Let $n_{m-j}^{\mathrm{Opt}}=$ $\left|V_{\text {Opt }}^{m-j}\right|$, i.e., $n_{m-j}^{\mathrm{Opt}}$ is the number of requests in $V^{m-j}$ which are also in the optimal solution $V_{\mathrm{OPT}}$. By construction, we have $n_{m-j}^{\mathrm{OPT}} \leq c\left|V_{\mathrm{ALG}}^{m-j}\right|$. Since the largest possible weight of a request in $V^{m-j}$ is at most $s$ times larger than the smallest possible weight of a request in $V^{m-j}$, this implies that $w\left(V_{\text {OPT }}^{m-j}\right) \leq$ $s \cdot c \cdot w\left(V_{\mathrm{ALG}}^{m-j}\right)$. Thus, we get that

$$
\begin{equation*}
\sum_{j=0}^{\left\lceil\log _{s}\left(n^{2}\right)\right\rceil} w\left(V_{\mathrm{OPT}}^{m-j}\right) \leq \sum_{j=0}^{\left\lceil\log _{s}\left(n^{2}\right)\right\rceil} s \cdot c \cdot w\left(V_{\mathrm{ALG}}^{m-j}\right)=s \cdot c \cdot \operatorname{ALG}(\sigma) . \tag{2}
\end{equation*}
$$

We now consider the unimportant requests. If $r_{i}$ is unimportant, then $w(i) \leq$ $s^{m^{\prime}-\left\lceil\log _{s}\left(n^{2}\right)\right\rceil} \leq s^{m^{\prime}} / n^{2} \leq w\left(x_{i_{\max }}\right) / n^{2} \leq \operatorname{OpT}(\sigma) / n^{2}$. This implies that

$$
\begin{equation*}
\sum_{j=\left\lceil\log _{s}\left(n^{2}\right)\right\rceil+1}^{\infty} w\left(V_{\mathrm{OPT}}^{m-j}\right) \leq n \frac{\operatorname{OPT}(\sigma)}{n^{2}}=\frac{\operatorname{OPT}(\sigma)}{n} \tag{3}
\end{equation*}
$$

We conclude that

$$
\operatorname{OPT}(\sigma)=w\left(V_{\mathrm{OPT}}\right)=\sum_{j=0}^{\left\lceil\log _{s}\left(n^{2}\right)\right\rceil} w\left(V_{\mathrm{OPT}}^{m-j}\right)+\sum_{j=\left\lceil\log _{s}\left(n^{2}\right)\right\rceil+1}^{\infty} w\left(V_{\mathrm{OPT}}^{m-j}\right)
$$

By Eq. (3),

$$
\left(1-\frac{1}{n}\right) \operatorname{OPT}(\sigma) \leq \sum_{j=0}^{\left\lceil\log _{s}\left(n^{2}\right)\right\rceil} w\left(V_{\mathrm{OPT}}^{m-j}\right)
$$

so by Eq. (2), $\operatorname{Opt}(\sigma) \leq \frac{n}{n-1} \cdot s \cdot c \cdot \operatorname{AlG}(\sigma)$.
Note that for $n \geq n_{0}=\frac{2+2 \varepsilon}{\varepsilon},\left(\frac{n}{n-1}\right)(1+\varepsilon / 2) \leq(1+\varepsilon)$. For inputs of length less than $n_{0}$, the oracle writes an optimal solution onto the advice tape, using at most $n_{0}$ bits. Since $n_{0} \leq \frac{4}{\varepsilon}, b \in O\left(\varepsilon^{-1} \log ^{2} n\right)$ as required. For inputs of length at least $n_{0}$, we use the algorithm described above. Thus, for every input $\sigma$, it holds that $\operatorname{OPT}(\sigma) \leq(1+\varepsilon) c \operatorname{AlG}(\sigma)$. Since $\varepsilon$ was arbitrary, this proves the theorem.

It may be surprising that adding weights to AOC-complete maximization problems has almost no effect, while adding weights to AOC-complete minimization problems drastically changes the advice complexity. In particular, one might wonder why the technique used in Theorem 3 does not work for minimization problems. The key difference lies in the beginning of the sequence. Let $w_{\max }$ be the largest weight of a request accepted by Opt.

For maximization problems, the algorithm can safely reject all requests before the first important one. For minimization problems, this approach does not work, since the algorithm must accept a superset of what OPT accepts in order to ensure that its output is feasible. Thus, rejecting an unimportant request that Opt accepts may result in an infeasible solution. This essentially means that the algorithm is forced into accepting all requests before the first important request arrives. Accepting all unimportant requests is no problem, since they will not contribute significantly to the total cost. However, accepting even a single huge request can give an unbounded contribution to the algorithm's cost. As shown in Theorem it is not possible in general for the algorithm to tell if a request in the beginning of the sequence is unimportant or huge without using a lot of advice.

However, if the ratio of the largest to the smallest weight is not too large, exponential sparsification is also useful for minimization problems in AOC. Essentially, when this ratio is bounded, it is possible for the algorithm to learn a good approximation of $w_{\max }$ when the first request arrives. This is formalized in Theorem 4, the proof of which is very similar to the proof of Theorem 3.

Theorem 4. If $\mathrm{P} \in \mathrm{AOC}$ is a minimization problem and $0<\varepsilon \leq 1$, then $\mathrm{P}_{w}$ with all weights in $\left[w_{\min }, w_{\max }\right]$ has a $(1+\varepsilon) c$-competitive algorithm with advice complexity at most

$$
B(n, c)+O\left(\varepsilon^{-1} \log ^{2} n+\log \left(\varepsilon^{-1} \log \frac{w_{\max }}{w_{\min }}\right)\right)
$$

Proof. Fix $\varepsilon>0$. Let $\sigma=\left\langle\left\{r_{1}, w_{1}\right\}, \ldots,\left\{r_{n}, w_{n}\right\}\right\rangle$ be the input and let $x=$ $x_{1} \ldots x_{n} \in\{0,1\}^{n}$ specify an optimal solution for $\sigma$, with ones indicating membership in the optimal solution. Define $s=1+\varepsilon / 2$. Let $V_{\mathrm{Opt}}=\left\{i: x_{i}=1\right\}$. Note that $V_{\text {Opt }}$ contains exactly those rounds in which Opt answers 1 and thus accepts. Furthermore, for $k \in \mathbb{Z}$, let $V^{k}=\left\{i: s^{k} \leq w(i)<s^{k+1}\right\}$ and let $V_{\text {OPT }}^{k}=V_{\text {OPT }} \cap V^{k}$. Finally, let $i_{\max } \in V_{\text {OPT }}$ be such that $w\left(i_{\max }\right) \geq w(i)$ for every $i \in V_{\text {Opt }}$.

The oracle computes the unique $m \in \mathbb{Z}$ such that $i_{\max } \in V_{\text {OPT }}^{m}$. We say that a request $r_{i}$ is unimportant if $w(i)<s^{m-\left\lceil\log _{s}\left(n^{2}\right)\right\rceil \text {, important if } s^{m-\left\lceil\log _{s}\left(n^{2}\right)\right\rceil} \leq, ~ \text { in }}$ $w(i)<s^{m+1}$, and huge if $w(i) \geq s^{m+1}$. The oracle also computes the unique $m \in \mathbb{Z}$ such that $s^{m^{\prime}} \leq w_{1}<s^{m^{\prime}+1}$ and writes the values $n$ and $m-m^{\prime}$ on the tape in a self-delimiting encoding.

The number of advice bits needed to write $m-m^{\prime}$ is $O\left(\log \left(m-m^{\prime}\right)\right)$.

$$
\begin{aligned}
\log \left(m-m^{\prime}\right) & \leq \log \left(\log _{s} w\left(i_{\max }\right)-\log _{s} w_{1}\right)+1 \\
& \leq \log \log _{s} \frac{w_{\max }}{w_{\min }}+1 \\
& \leq \log \left(2 \varepsilon^{-1} \log \frac{w_{\max }}{w_{\min }}\right)+1, \text { since } \log _{s} n \leq 2 \varepsilon^{-1} \log n, \text { for } \varepsilon / 2 \leq 1
\end{aligned}
$$

Note that since the length of $m-m^{\prime}$ is not known, we need to use a self-delimiting encoding, which means that we use $O\left(\log n+\log \left(\varepsilon^{-1} \log \frac{w_{\max }}{w_{\min }}\right)\right)$ advice bits at the beginning.

This advice allows the algorithm to learn $m$ as soon as the first request arrives. From there on, the algorithm will know if a request is important, unimportant, or huge. Whenever a huge request arrives, the algorithm answers 0 (rejects the request). When an unimportant request arrives, the algorithm answers 1 (accepts the request). We now describe how the algorithm and oracle work for the important requests.

For the important requests (whose indices are) in $V^{m-j}$, we use the covering design based $c$-competitive algorithm for unweighted AOC-problems. This is similar to what we do in the proof of Theorem 3. The same calculations yield an upper bound on this advice of $B(n, c)+O\left(\log _{s} n \cdot \log n\right)$. Note that $\log _{s}(n) \leq 2 \varepsilon^{-1} \log n$ for $\varepsilon / 2 \leq 1$, giving the bound on the advice in the statement of the theorem.

First, we note that the solution produced is valid, since it is a superset of the solution of Opt.

We now argue that the cost of the solution is at most $(1+\varepsilon) c$ times the cost of Opt. Following the proof of Theorem 3 and switching the roles of Opt and Alg, we have by construction that the cost of the important requests for the algorithm is at most $s c$ times larger than the cost for Opt on the important requests. For the huge requests, both this algorithm and Opt incur a cost of zero.

We now consider the unimportant requests. If $r_{i}$ is unimportant, then

$$
w(i)<s^{m-\left\lceil\log _{s}\left(n^{2}\right)\right\rceil} \leq s^{m} / n^{2} \leq w\left(x_{i_{\max }}\right) / n^{2} \leq \operatorname{OPT}(\sigma) / n^{2}
$$

This implies that

$$
\begin{equation*}
\sum_{j=\left\lceil\log _{s}\left(n^{2}\right)\right\rceil+1}^{\infty} w\left(V^{m-j}\right) \leq n \frac{\operatorname{OPT}(\sigma)}{n^{2}}=\frac{\operatorname{Opt}(\sigma)}{n} \tag{4}
\end{equation*}
$$

Thus, even if the algorithm accepts all unimportant requests and Opt accepts none of them, it only accepts an additional $\frac{\operatorname{OPT}(\sigma)}{n}$. In total, the algorithm gets
a cost of at most $\left(1+\frac{1}{n}\right)(1+\varepsilon / 2) c \operatorname{Opt}(\sigma)$. For $n \geq n_{0}=\frac{2+\varepsilon}{\varepsilon}$, this is at most $(1+\varepsilon) c \operatorname{Opt}(\sigma)$. For inputs of length less than $n_{0}$, the oracle will write an optimal solution onto the advice tape, using at most $n_{0}$ bits. Since $n_{0} \leq \frac{3}{\varepsilon}, b \in$ $O\left(\varepsilon^{-1} \log ^{2} n\right)$ as required. For inputs of length at least $n_{0}$, we use the algorithm described above. Thus, for every input $\sigma$, it holds that $\operatorname{OPT}(\sigma) \leq(1+\varepsilon) c \operatorname{AlG}(\sigma)$.

## 5 Matching and Other Non-Complete AOC Problems

We first provide a general theorem that works for all maximization problems in AOC, giving better results in some cases than that in Theorem 3 ,

Theorem 5. Let $\mathrm{P} \in \mathrm{AOC}$ be a maximization problem. If there exists a $c$ competitive P -algorithm reading $b$ bits of advice, then there exists a $O(c \cdot \log n)$ competitive $\mathrm{P}_{w}$-algorithm reading $O(b+\log n)$ bits of advice.

Proof. Use exponential sparsification on the weights with an arbitrary $\varepsilon$, say $\varepsilon=1 / 2$, and let $s=1+\varepsilon$. For a given request sequence, $\sigma$, let $w_{\max }$ be the maximum weight that $\mathrm{OPT}_{\mathrm{P}_{\mathrm{w}}}$ accepts. The oracle computes the unique $m \in \mathbb{Z}$ such that $w_{\max } \in\left[s^{m}, s^{m+1}\right)$. The important requests are those with weight $w$, where $s^{m-\left\lceil\log _{s}\left(n^{2}\right)\right\rceil} \leq w<s^{m+1}$.

We consider only the $\left\lceil\log _{s}\left(n^{2}\right)\right\rceil+1$ important intervals, i.e., the intervals $\left[s^{i}, s^{i+1}\right), m-\left\lceil\log _{s}\left(n^{2}\right)\right\rceil \leq i \leq m$, and index them by $i$. Let $k$ be the index of the interval of weights contributing the most weight to $\mathrm{OPT}_{\mathrm{P}_{\mathrm{w}}}(\sigma)$. The advice is a self-delimiting encoding of the index, $j$, of the first request with weight $w \in\left[s^{k}, s^{k+1}\right)$, plus the advice used by the given $c$-competitive P -algorithm. This requires at most $b+O(\log (n))$ bits of advice.

The algorithm rejects all requests before the $j$ th. From the $j$ th request, the algorithm calculates the index $k$. The algorithm accepts those requests which would be accepted by the P-algorithm when presented with the subsequence of $\sigma$ consisting of the requests with weights in $\left[s^{k}, s^{k+1}\right)$. Since, by exponential sparsification, $\mathrm{OPT}_{\mathrm{P}_{\mathrm{w}}}$ accepts total weight at most $\frac{1}{n} \mathrm{OPT}_{\mathrm{P}_{\mathrm{w}}}(\sigma)$ from requests with unimportant weights, and it accepts at least as much from interval $k$ as from any of the other $\left\lceil\log _{s}\left(n^{2}\right)\right\rceil+1$ intervals considered, $\mathrm{OPT}_{\mathrm{P}_{\mathrm{w}}}$ accepts weight at least $\left(1-\frac{1}{n}\right) \frac{\operatorname{OPT}_{\mathrm{P}_{\mathrm{w}}}(\sigma)}{\left[\log _{s}\left(n^{2}\right)\right]+1}$ from interval $k$. The algorithm, ALG, described here accepts at least $\frac{1}{c}$ as many requests as $\mathrm{OPT}_{\mathrm{P}_{\mathrm{w}}}$ does in this interval, and each of the requests it accepts is at least a fraction $\frac{1}{s}$ as large as the largest weight in this interval. Thus, $c(1+\varepsilon) \operatorname{ALG}(\sigma) \geq\left(\frac{1-\frac{1}{n}}{\left\lceil\log _{s}\left(n^{2}\right)\right\rceil+1}\right) \operatorname{OPT}_{\mathrm{P}_{\mathrm{w}}}(\sigma)$, so ALG is $O(c \log n)$-competitive.

In the online matching problem, edges arrive one by one. Each request contains the names of the edge's two endpoints (the set of endpoints is not known from the beginning, but revealed gradually as the edges arrive). The algorithm must irrevocably accept or reject them as they arrive, and the goal is to maximize the number of edges accepted. The natural greedy algorithm for this problem is
well known to be 2-competitive. In terms of advice, the problem is known to be in AOC, but is not AOC-complete [7]. We remark that a version of unweighted online matching with vertex arrivals (incomparable to our weighted matching with edge arrivals) has been studied with advice in [10].

Corollary 1. There exists a $O(\log n)$-competitive algorithm for Weighted Matching reading $O(\log n)$ bits of advice.

Proof. The result follows from Theorem [5 since there exists a 2 -competitive algorithm without advice for (unweighted) Matching.

### 5.1 Lower bounds

First, we present a result which holds for the weighted versions of many maximization problems in AOC. It also holds for the weighted versions of AOC-complete minimization problems, but Theorem 2 gives a much stronger result.

Theorem 6. For the weighted online versions of Independent Set, Clique, Disjoint Path Allocation, and Matching, an algorithm reading o $(\log n)$ bits of advice cannot be $f(n)$-competitive for any function $f$.

To prove Theorem 6, we start by proving the following lemma from which the theorem easily follows.

Lemma 2. Let $\mathrm{P} \in \mathrm{AOC}$ and suppose there exists a family $\left(\sigma_{n}\right)_{n \in \mathbb{N}}$ of P -inputs with the following properties:

1. $\sigma_{n}=\left\langle r_{1}, r_{2}, \ldots, r_{n}\right\rangle$ consists of $n$ requests.
2. $\sigma_{n+1}$ is obtained by adding a single request to the end of $\sigma_{n}$.
3. If P is a maximization problem, the feasible solutions are those in which at most one request is accepted.
If P is a minimization problem, the feasible solutions are those in which at least one request is accepted.

Then, no algorithm for the weighted problem $\mathrm{P}_{w}$ reading o $(\log n)$ bits of advice can be $f(n)$-competitive for any function $f$.

Proof. Let Alg be a $\mathrm{P}_{\mathrm{w}}$-algorithm reading at most $b=o(\log n)$ bits of advice. Let $f(n)>0$ be an arbitrary non-decreasing function of $n$. We will show that for all sufficiently large $n$, there exists an input of length $n$ such that the profit obtained by Opt is at least $f(n)$ times as large as the profit obtained by Alg. Since $f(n)$ was arbitrary, it follows that ALG is not $f(n)$-competitive for any function $f$.

Since $b=o(\log n)$, there exists an $N \in \mathbb{Z}$ such that for any $n \geq N$, ALG reads less than $\log (n)-1$ bits of advice on inputs of length at most $n$. Fix an $n \geq N$. For $1 \leq i \leq n$, define the $\mathrm{P}_{\mathrm{w}}$-input $\widehat{\sigma}_{i}=\left\langle\left\{r_{1}, f(n)\right\},\left\{r_{2}, f(n)^{2}\right\}, \ldots\right.$, $\left.\left\{r_{i}, f(n)^{i}\right\}\right\rangle$. Consider the set of inputs $\left\{\widehat{\sigma}_{1}, \ldots, \widehat{\sigma}_{n}\right\}$. For every $1 \leq i \leq n$, the number of advice bits read by Alg on the input $\widehat{\sigma}_{i}$ is at most $\log (n)-1$ (since
the length of the input $\widehat{\sigma}_{i}$ is $\left.i \leq n\right)$. Thus, by the pigeonhole principle, there must exist two integers $n_{1}, n_{2}$ with $n_{1}<n_{2}$ such that ALG reads the same advice on $\widehat{\sigma}_{n_{1}}$ and $\widehat{\sigma}_{n_{2}}$. If ALG rejects all requests in $\widehat{\sigma}_{n_{1}}$, then it achieves a profit of 0 while Opt obtains a profit of $f(n)^{n_{1}}$. If AlG accepts a request in $\widehat{\sigma}_{n_{1}}$, then it obtains a profit of at most $f(n)^{n_{1}}$. Since Alg reads the same advice on $\widehat{\sigma}_{n_{1}}$ and $\widehat{\sigma}_{n_{2}}$ and since the two inputs are indistinguishable for the first $n_{1}$ requests, this means that AlG also obtains a profit of at most $f(n)^{n_{1}}$ on the input $\widehat{\sigma}_{n_{2}}$. But $\operatorname{Opt}\left(\widehat{\sigma}_{n_{2}}\right)=f(n)^{n_{2}}$, and hence $\operatorname{Opt}\left(\widehat{\sigma}_{n_{2}}\right) / \operatorname{ALG}\left(\widehat{\sigma}_{n_{2}}\right) \geq f(n)^{n_{2}-n_{1}} \geq f(n)$.

For minimization problems, we can use the same arguments and the input sequence $\widehat{\sigma}_{i}=\left\langle\left\{r_{1}, f(n)^{-1}\right\},\left\{r_{2}, f(n)^{-2}\right\}, \ldots,\left\{r_{i}, f(n)^{-n}\right\}\right\rangle$.

Proof (Proof of Theorem (6). For Independent Set, we can use the above lemma with a family of cliques $\left(K_{n}\right)_{n \in \mathbb{N}}$, and for Clique, we can use a family of independent sets. For Matching, we can use a family of stars $\left(K_{1, n}\right)_{n \in \mathbb{N}}$. For Disjoint Path Allocation, we use a path $P_{2 n}=\left\langle v_{1}, v_{2}, \ldots, v_{2 n}\right\rangle$ and $r_{i}=\left\langle v_{i}, v_{i+1}, \ldots v_{i+n}\right\rangle$.

Returning to the example of Weighted Matching, we now know that $O(\log n)$ bits suffice to be $O(\log n)$-competitive, and that no algorithm can be $f(n)$ competitive for any function $f$ with $o(\log n)$ bits of advice. In order to prove that a linear number of advice bits is necessary to achieve constant competitiveness for Weighted Matching, we use a direct product theorem from [14]. This uses the concept defined in 14 of a problem being $\Sigma$-repeatable. Informally, this means that it is always possible to combine $r$ (sufficiently profitable) input sequences $I_{1}, I_{2}, \ldots, I_{r}$ into a single input $g\left(I_{1}, I_{2}, \ldots, I_{r}\right)$ such that serving this single input gives profit close to that of serving each of the $I_{i}$ independently and adding the profits.

Definition 6. Let $P$ be an online maximization problem and $I$ be the set of possible input sequences. Assume that for every input in $I$, there are only a finite number of valid outputs. Let $I^{*}$ be the set of concatenations of sequences (rounds) from I. P is $\Sigma$-repeatable with parameters $\left(k_{1}, k_{2}, k_{3}\right)$ if there exists a function $g: I^{*} \rightarrow I$ satisfying the following:

- For every $\sigma^{*} \in I^{*}$ with r rounds, $\left|g\left(\sigma^{*}\right)\right| \leq\left|\sigma^{*}\right|+k_{1} r$
- For every deterministic algorithm Alg for P, there is a deterministic algorithm ALG* for sequences from $I^{*}$, such that for every $\sigma^{*} \in I^{*}$ with $r$ rounds, $\operatorname{ALG}^{*}\left(\sigma^{*}\right) \geq \operatorname{AlG}\left(g\left(\sigma^{*}\right)-k_{2} r\right.$,
- Let Opt* denote an optimal algorithm for sequences from $I^{*}$. For every $\sigma^{*} \in I^{*}$ with $r$ rounds, $\operatorname{OPT}^{*}\left(\sigma^{*}\right) \leq \operatorname{Opt}\left(g\left(\sigma^{*}\right)\right)+k_{3} r$.

Theorem 7. An $O(1)$-competitive algorithm for Weighted Matching must read $\Omega(n)$ bits of advice.

Proof. We prove the lower bound using a direct product theorem [14]. According to [14, it suffices to show that: (i) Weighted Matching is $\Sigma$-repeatable, and (ii) for every $c$, there exists a probability distribution $p_{c}$ with finite support such that for every deterministic algorithm DET without advice, it holds that
$\mathbb{E}_{p_{c}}[\operatorname{Opt}(\sigma)] \geq c \cdot \mathbb{E}_{p_{c}}[\operatorname{DET}(\sigma)]$. Also, there must be a finite upper bound on the profit an algorithm can obtain on an input in the support of $p_{c}$.

It is trivial to see that Weighted Matching is $\Sigma$-repeatable. Fix $c \geq 1$ and let $k=2 c-1$. We define the probability distribution $p_{c}$ by specifying a probabilistic adversary: The input graph will be a star $K_{1, m}$ consisting of $m$ edges for some $1 \leq m \leq k$. In round $i$, the adversary reveals the edge $e_{i}=\left(v, v_{i}\right)$ where $v_{i}$ is a new vertex and $v$ is the center vertex of the star. The edge $e_{i}$ has weight $2^{i}$. If $i<k$, then with probability $1 / 2$ the adversary will proceed to round $i+1$, and with probability $1 / 2$ the input sequence will end. If the adversary reaches round $k$, it will always stop after revealing the edge $e_{k}$ of round $k$. Note that the support of $p_{c}$ and the largest profit an algorithm can obtain on any input in the support of $p_{c}$ are both finite.

Let $X$ be the random variable which denotes the number of edges revealed by the adversary. Note that $\operatorname{Pr}(X=j)=2^{-j}$ if $1 \leq j<k$. Consequently,

$$
\begin{equation*}
\operatorname{Pr}(X=k)=1-\operatorname{Pr}(X<k)=1-\sum_{i=1}^{k-1} 2^{-i}=2^{-(k-1)} \tag{5}
\end{equation*}
$$

Let Det be a deterministic algorithm without advice. We may assume that Det decides in advance on some $1 \leq j \leq k$ and accepts the edge $e_{j}$ (the only other possible deterministic strategy it to never accept an edge, but this is always strictly worse than following any of the $k$ strategies that accepts an edge). If $X<j$, then the profit obtained by Det is zero. If $X \geq j$, then Det obtains a profit of $2^{j}$. It follows that

$$
\mathbb{E}[\operatorname{DET}(\sigma)]=\operatorname{Pr}(X \geq j) 2^{j}=(1-\operatorname{Pr}(X<j)) 2^{j}=2^{-(j-1)} 2^{j}=2
$$

The optimal algorithm Opt always accepts the last edge of the input. Thus, if $X=j$, then the profit of Opt is $2^{j}$. It follows that

$$
\mathbb{E}[\operatorname{OPT}(\sigma)]=\sum_{j=1}^{k} \operatorname{Pr}(X=j) 2^{j}=\sum_{j=1}^{k-1}\left(2^{-j} 2^{j}\right)+2^{-(k-1)} 2^{k}=k+1
$$

Thus, we conclude that $\mathbb{E}[\operatorname{Opt}(\sigma)] \geq \frac{k+1}{2} \mathbb{E}[\operatorname{DET}(\sigma)]=c \mathbb{E}[\operatorname{DET}(\sigma)]$.
In particular, we cannot achieve constant competitiveness using $O(\log n)$ bits of advice for Weighted Matching. We leave it as an open problem to close the gap between $\omega(1)$ and $O(\log n)$ on the competitiveness of Weighted Matching algorithms with advice complexity $O(\log n)$.

## 6 Scheduling with Sublinear Advice

For the scheduling problems studied, the requests are jobs, each characterized by its size. Each job must be assigned to one of $m$ available machines. If the machines are identical, the load of a job on any machine is simply its size. If
the machines are related, each machine has a speed, and the load of a job, $J$, assigned to a machine with speed $s$ is the size of $J$ divided by $s$. If the machines are unrelated, each job arrives with a vector specifying its load on each machine.

Consider a sequence $\sigma=\left\langle r_{1}, \ldots, r_{n}\right\rangle$ of $n$ jobs that arrive online. Each job $r_{i} \in \sigma$ has an associated weight-function $w_{i}:[m] \rightarrow \mathbb{R}_{+}$. Upon arrival, a job must irrevocably be assigned to one of the $m$ machines. The load $L_{j}$ of a machine $j \in[m]$ is defined as $L_{j}=\sum_{i \in M_{j}} w_{i}(j)$ where $M_{j}$ is the set of (indices of) jobs scheduled on machine $j$. The total load of a schedule for $\sigma$ is the vector $\mathbf{L}=\left(L_{1}, \ldots, L_{m}\right)$. We say that $\left(L_{1}, \ldots, L_{m}\right) \leq\left(L_{1}^{\prime}, \ldots, L_{m}^{\prime}\right)$ if and only if $L_{i} \leq L_{i}^{\prime}$ for $1 \leq i \leq m$. A scheduling problem of the above type is specified by an objective function $f: \mathbb{R}_{+}^{m} \rightarrow \mathbb{R}_{+}$and by specifying if the goal is to minimize or maximize $f(\mathbf{L})=f\left(L_{1}, \ldots, L_{m}\right) \in \mathbb{R}_{+}$. We assume that $f$ is non-decreasing, i.e., $f(\mathbf{L}) \leq f\left(\mathbf{L}^{\prime}\right)$ for all $\mathbf{L} \leq \mathbf{L}^{\prime}$. Some of the classical choices of objective function include:

- Minimizing the $\ell_{p}$-norm $f_{p}(\mathbf{L})=f_{p}\left(L_{1}, \ldots, L_{m}\right)=\left\|\left(L_{1}, \ldots, L_{m}\right)\right\|_{p}$ for some $1 \leq p \leq \infty$. That is, for $1 \leq p<\infty$, the goal is to minimize $\left(\sum_{j \in[m]} L_{j}^{p}\right)^{1 / p}$ and for $p=\infty$, the goal is to minimize the makespan $\max _{j \in[m]} L_{j}$.
- Maximizing the minimum load $f(\mathbf{L})=\min _{j \in[m]} L_{j}$. This is also known as machine covering. Note that this objective function is not a norm ${ }^{2}$, but it does satisfy that $f(\alpha \mathbf{L})=\alpha f(\mathbf{L})$ for every $\alpha \geq 0$ and $\mathbf{L} \in \mathbb{R}_{+}^{m}$.

We begin with a result for unrelated machines.
Theorem 8. Let P be a scheduling problem on $m$ unrelated machines where the goal is to minimize an objective function $f$. Assume that $f$ is a non-decreasing norm. Then, for $0<\varepsilon \leq 1$, there exists a $(1+\varepsilon)$-competitive P -algorithm reading $O\left(\left(\frac{4}{\varepsilon} \log (n)+2\right)^{m} \log (n)\right)$ bits of advice. In particular, if $m=O(1)$ and $\varepsilon=\Omega(1)$, then there exists a $(1+\varepsilon)$-competitive algorithm reading $O(\operatorname{poly} \log (n))$ bits of advice.

Proof. Since the objective function $f$ is a norm on $\mathbb{R}^{m}$, we will denote it by $\|\cdot\|$. Let $\mathbf{1}_{j}$ be the $j$ th unit vector (the vector with a 1 in the $j$ th coordinate and 0 elsewhere).

Fix an input sequence $\sigma$. The oracle starts by computing an arbitrary optimal schedule for $\sigma$. Throughout most of this proof, we assume that $n$ is sufficiently large. The necessary conditions are discussed at the end of the proof, along with how to handle small $n$.

Let $\mathbf{L}_{\text {Opt }}$ be the load-vector of this schedule. Thus, $\operatorname{Opt}(\sigma)=\left\|\mathbf{L}_{\text {Opt }}\right\|$. Let $s=1+\varepsilon / 2$ and let $k$ be the unique integer such that $s^{k} \leq\left\|\mathbf{L}_{\text {Opt }}\right\|<s^{k+1}$. A job $r_{i} \in \sigma$ is said to be unimportant if there exists a machine $j \in[m]$ such that $\left\|w_{i}(j) \mathbf{1}_{j}\right\|<s^{k-\left\lceil\log _{s}\left(n^{2}\right)\right\rceil}$. A job which is not unimportant is important.

The oracle uses $O(\log n)$ bits to encode $n$ using a self-delimiting encoding. It then writes the index $i^{\prime}$ of the first important job $r_{i^{\prime}}$ onto the advice tape (or indicates that $\sigma$ contains no important jobs) using $\lceil\log (n+1)\rceil$ bits. Let

[^2]$j^{\prime}$ be the machine minimizing $\left\|w_{i^{\prime}}\left(j^{\prime}\right) \mathbf{1}_{j^{\prime}}\right\|$, where ties are broken arbitrarily. The oracle also writes $\Delta_{i^{\prime}}=k-k_{i^{\prime}}$, where $k_{i^{\prime}}$ is the unique integer such that $s^{k_{i^{\prime}}} \leq\left\|w_{i^{\prime}}\left(j^{\prime}\right) \mathbf{1}_{j^{\prime}}\right\|<s^{k_{i^{\prime}}+1}$ onto the advice tape using $O\left(\log \log _{s} n\right)$ bits.

Scheduling unimportant jobs. If a job $r_{i} \in \sigma$ is unimportant, then the algorithm schedules the job on the machine $j$ minimizing $\left\|w_{i}(j) \mathbf{1}_{j}\right\|$ where ties are broken arbitrarily. We now explain how the algorithm knows if a job is unimportant or not. If $r_{i} \in \sigma$ is a job that arrives before the first important job, i.e., if $i<i^{\prime}$, then $r_{i}$ is unimportant by definition. When job $r_{i^{\prime}}$ arrives, the algorithm can deduce $k$ since it knows $\Delta_{i^{\prime}}$ from the advice and since it can compute $\min _{j}\left\|w_{i^{\prime}}(j) \mathbf{1}_{j}\right\|$ without help. Knowing $k$ (and the number of jobs $n$ ), the algorithm is able to tell if a job is unimportant or not.

Scheduling important jobs. We now describe how the algorithm schedules the important jobs. To this end, we define the type of an important job. For an important job $r_{i}$, let $\Delta_{i}(1), \ldots, \Delta_{i}(m)$ be defined as follows: For $1 \leq j \leq m$, if there exists an integer $k_{i}(j) \leq k$ such that $s^{k_{i}(j)} \leq\left\|w_{i}(j) \mathbf{1}_{j}\right\|<s^{k_{i}(j)+1}$, then $\Delta_{i}(j)=k-k_{i}(j)$ (since $r_{i}$ is important, $\Delta_{i}(j) \leq\left\lceil\log _{s}\left(n^{2}\right)\right\rceil$ ). If no such integer exists, then it must be the case that $\left\|w_{i}(j) \mathbf{1}_{j}\right\| \geq s^{k+1}>\left\|\mathbf{L}_{\text {OPT }}\right\|$. In this case, we let $\Delta_{i}(j)=\perp$ be a dummy symbol. The type of $r_{i}$ is the vector $\boldsymbol{\Delta}_{i}=$ $\left(\Delta_{i}(1), \ldots, \Delta_{i}(m)\right)$. Note that there are only $\left(\left\lceil\log _{s}\left(n^{2}\right)\right\rceil+2\right)^{m}$ different types. For each possible type $\boldsymbol{\Delta}=(\Delta(1), \ldots, \Delta(m))$, the oracle writes the number, $a_{\boldsymbol{\Delta}}$, of jobs of type $\boldsymbol{\Delta}$ onto the advice tape. This requires at most $\left(\left\lceil\log _{s}\left(n^{2}\right)\right\rceil+\right.$ 2) ${ }^{m}\lceil\log (n+1)\rceil$ bits of advice.

Note that since $\|\cdot\|$ is a norm, if $r_{i} \in \sigma$ is of type $\boldsymbol{\Delta}_{i}=\left(\Delta_{i}(1), \ldots, \Delta_{i}(m)\right)$, then $s^{k-\Delta_{i}(j)}\left\|\mathbf{1}_{j}\right\|^{-1} \leq w_{i}(j) \leq s^{k-\Delta_{i}(j)+1}\left\|\mathbf{1}_{j}\right\|^{-1}$ if $\Delta_{i}(j) \neq \perp$ and $w_{i}(j)>$ $\left\|\mathbf{L}_{\text {OPT }}\right\|\left\|\mathbf{1}_{j}\right\|^{-1}$ if $\Delta_{i}(j)=\perp$. The algorithm computes an optimal schedule $\widehat{S}_{\text {imp }}$ for the input $\widehat{\sigma}$ which for each possible type $\boldsymbol{\Delta}$ contains $a_{\Delta}$ jobs with weightfunction $\widehat{w}_{\boldsymbol{\Delta}}$ where $\widehat{w}_{\boldsymbol{\Delta}}(j)=s^{k-\Delta(j)+1}\left\|\mathbf{1}_{j}\right\|^{-1}$ if $k(j) \neq \perp$ and $\widehat{w}_{\boldsymbol{\Delta}}(j)=\infty$ otherwise. This choice of weight-function ensures that if $r_{i} \in \sigma$ is a job of type $\Delta_{i}$, then for each $j$ with $\Delta_{i}(j) \neq \perp$,

$$
\begin{equation*}
w_{i}(j)<\widehat{w}_{\Delta_{i}}(j) \leq s \cdot w_{i}(j) \tag{6}
\end{equation*}
$$

When an important job of $\sigma$ arrives, the algorithm computes the type of the job. Based solely on this type, the algorithm schedules the important jobs in $\sigma$ by following the schedule $\widehat{S}_{\text {imp }}$ for $\widehat{\sigma}$. Let $\mathbf{L}_{\text {imp }}$ be the load-vector of the important jobs of $\sigma$ scheduled by Alg. Note that by Eq. (6), the weight-function of an important job of $\sigma$ is strictly smaller (for all machines) than the weight-function of the corresponding job of $\widehat{\sigma}$. Thus, since $f$ is non-decreasing $\left\|\mathbf{L}_{\text {imp }}\right\|$ is bounded from above by the cost of the schedule $\widehat{S}_{\text {imp }}$ for $\widehat{\sigma}$.

Putting it all together. The optimal schedule for $\sigma$ computed by the oracle induces a schedule of $\widehat{\sigma}$. Let $\widehat{\mathbf{L}}$ be the load-vector of this schedule. By Eq. (6), we get that $\|\widehat{\mathbf{L}}\| \leq s\left\|\mathbf{L}_{\text {OPt }}\right\|$. Thus, the cost of $\widehat{S}_{\text {imp }}$ (which was an optimal scheduling of $\widehat{\sigma})$ is at most $\|\widehat{\mathbf{L}}\| \leq s\left\|\mathbf{L}_{\text {OPT }}\right\|$.

Let $\mathbf{L}_{\text {unimp }}$ be the load-vector of the unimportant jobs scheduled by Alg. Furthermore, let $M_{j}$ be the set of indices of the unimportant jobs scheduled by

Alg on machine $j$. By subadditivity,

$$
\begin{aligned}
\left\|\mathbf{L}_{\text {unimp }}\right\| & =\left\|\sum_{j=1}^{m} \sum_{i \in M_{j}} w_{i}(j) \mathbf{1}_{j}\right\| \leq \sum_{j=1}^{m} \sum_{i \in M_{j}}\left\|w_{i}(j) \mathbf{1}_{j}\right\|<\sum_{j=1}^{m} \sum_{i \in M_{j}} s^{k-\left\lceil\log _{s}\left(n^{2}\right)\right\rceil} \\
& \leq n \frac{\left\|\mathbf{L}_{\mathrm{OPT}}\right\|}{n^{2}} \leq \frac{\left\|\mathbf{L}_{\mathrm{OPT}}\right\|}{n}
\end{aligned}
$$

We are finally able to bound the cost of the entire schedule created by ALG:

$$
\operatorname{ALG}(\sigma)=\left\|\mathbf{L}_{\text {imp }}+\mathbf{L}_{\text {unimp }}\right\| \leq\left\|\mathbf{L}_{\text {imp }}\right\|+\left\|\mathbf{L}_{\text {unimp }}\right\| \leq(s+1 / n)\left\|\mathbf{L}_{\text {OPT }}\right\|
$$

Recall that $s=1+\varepsilon / 2$. Thus, if $n \geq 2 / \varepsilon$, then $\operatorname{AlG}(\sigma) \leq(1+\varepsilon) \operatorname{Opt}(\sigma)$. For inputs of length less than $2 / \varepsilon$, the oracle can simply encode the optimal solution using at most $\frac{2}{\varepsilon}\lceil\log m\rceil$ bits of advice. The total amount of advice used by our algorithm is at most

$$
\left(\left\lceil\log _{s}\left(n^{2}\right)\right\rceil+2\right)^{m}\lceil\log (n+1)\rceil+O\left(\log n+\log ^{\log } n\right)=O\left(\left(4 \varepsilon^{-1} \log (n)+2\right)^{m} \log (n)\right) .
$$

For the following discussion, assume that $\varepsilon=\Theta(1)$. We remark that the $(1+\varepsilon)$-competitive algorithm in Theorem $\mathbb{\square}$ is only of interest if the number of machines $m$ is small compared to the number of jobs $n$. As already noted, the most interesting aspect of Theorem 8 is that our algorithm uses only $\operatorname{polylog}(n)$ bits of advice if $m$ is a constant. More generally, if $m=o(\log n / \log \log n)$, then our algorithm will use $o(n)$ bits of advice. On the other hand, if $m=$ $\Theta(\log n)$, then our algorithm uses $\Omega\left(\log (n)^{\log (n)}\right)$ bits of advice, which is worse than the trivial 1-competitive algorithm which uses $n\lceil\log m\rceil=O(n \log \log n)$ bits of advice when $m=\Theta(\log n)$.

The advice complexity of the algorithm in Theorem 8 depends on the number of machines $m$ because we want the result to hold even when the machines are unrelated. We now show that when restricting to related machines, we can obtain a $(1+\varepsilon)$-competitive algorithm using $O\left(\varepsilon^{-1} \log ^{2} n\right)$ bits of advice, independent of the number of machines. The proof resembles that of Theorem 8 . The main difference is that we are able to reduce the number of types to $O\left(\log ^{2} n\right)$.

Theorem 9. Let P be a scheduling problem on $m$ related machines where the goal is to minimize an objective function $f$. Assume that $f$ is a non-decreasing norm. Then, for $0<\varepsilon \leq 1$, there exists a $(1+\varepsilon)$-competitive P -algorithm with advice complexity $O\left(\varepsilon^{-1} \log ^{2} n\right)$.

Proof. Since the objective function $f$ is a norm on $\mathbb{R}^{m}$, we will denote it by $\|\cdot\|$. Fix an input sequence $\sigma$. The oracle starts by computing an arbitrary optimal schedule for $\sigma$. Let $s=1+\varepsilon / 2$. The oracle uses $O(\log n)$ bits to encode $n$ using a self-delimiting encoding.

Let $C_{1}, \ldots, C_{m}$ be the speeds of the $m$ machines. Assume without loss of generality that $\left\|\mathbf{1}_{j}\right\| / C_{j}$ attains its minimum value when $j=1$. Define $B=\left\|\mathbf{1}_{1}\right\| / C_{1}$.

Let $\mathbf{L}_{\text {Opt }}$ be the load-vector of the fixed optimal schedule. Thus, $\operatorname{Opt}(\sigma)=$ $\left\|\mathbf{L}_{\text {Opt }}\right\|$. Let $k$ be the unique integer such that $s^{k} \leq\left\|\mathbf{L}_{\text {Opt }}\right\|<s^{k+1}$. A job $r_{i} \in \sigma$ is said to be unimportant if its weight, $w_{i}$, satisfies $w_{i} B<s^{k-\left\lceil\log _{s}\left(n^{2}\right)\right\rceil}$. A job which is not unimportant is important. Note that $w_{i} B$ is always bounded from above by $\left\|\mathbf{L}_{\text {Opt }}\right\|$ since $r_{i}$ must be placed on some machine. The oracle writes the index $i^{\prime}$ of the first important job $r_{i^{\prime}}$ onto the advice tape (or indicates that $\sigma$ contains no important jobs) using $\lceil\log n\rceil+1$ bits. The oracle also writes the unique integer $k^{\prime}$ such that $s^{k-k^{\prime}} \leq w_{i^{\prime}} B<s^{k-k^{\prime}+1}$ onto the advice tape, using $O\left(\log \log _{s}(n)\right)$ bits.

We now explain how the algorithm knows if a job is unimportant or not. If $r_{i} \in \sigma$ is a job that arrives before the first important job, i.e., if $i<i^{\prime}$, then $r_{i}$ is unimportant by definition. When job $r_{i^{\prime}}$ arrives, the algorithm can deduce $k$ since it knows $k^{\prime}$ from the advice and since it can compute $w_{i^{\prime}} B$ without help. Knowing $k$ (and the number of jobs $n$ ), the algorithm is able to tell if a job is unimportant or not.

If a job $r_{i} \in \sigma$ is unimportant, then the algorithm schedules the job on machine 1.

Scheduling important jobs. We now describe how the algorithm schedules the important jobs. To this end, we define the type of an important job. The type of an important job $r_{i}$ is the non-negative integer $t_{i}$ such that $s^{k-t_{i}} \leq w_{i} B<$ $s^{k-t_{i}+1}$. Note that there are only $\left\lceil\log _{s}\left(n^{2}\right)\right\rceil+1$ different types. For each possible type $0 \leq t \leq\left\lceil\log _{s}\left(n^{2}\right)\right\rceil$, the oracle writes the number of jobs $a_{t}$ of that type onto the advice tape. This requires at most $O\left(\log _{s}\left(n^{2}\right) \log (n)\right)$ bits of advice.

Note that since $\|\cdot\|$ is a norm, if $r_{i} \in \sigma$ is of type $t_{i}$, then $s^{k-t_{i}} B^{-1} \leq w_{i} \leq$ $s^{k-t_{i}+1} B^{-1}$. The algorithm computes an optimal schedule $\widehat{S}_{\mathrm{imp}}$ for the input $\overline{\widehat{\sigma}}$ which for each possible type $0 \leq t \leq\left\lceil\log _{s}\left(n^{2}\right)\right\rceil$ contains $a_{t}$ jobs with weight $\widehat{w}_{t}=s^{k-k_{i}+1} B^{-1}$. This choice of weight ensures that if $r_{i} \in \sigma$ is a job of type $t_{i}$, then,

$$
\begin{equation*}
w_{i}<\widehat{w}_{t_{i}} \leq s \cdot w_{i} \tag{7}
\end{equation*}
$$

When an important job of $\sigma$ arrives, the algorithm computes the type of the job. Based solely on this type, the algorithm schedules the important jobs in $\sigma$ by following the schedule $\widehat{S}_{\text {imp }}$ for $\widehat{\sigma}$. Let $\mathbf{L}_{\text {imp }}$ be the load-vector of the important jobs of $\sigma$ scheduled by Alg. Note that by Eq. (7), the weight of an important job of $\sigma$ is strictly smaller than the weight of the corresponding job of $\widehat{\sigma}$. Thus, $\left\|\mathbf{L}_{\text {imp }}\right\|$ is bounded from above by the cost of the schedule $\widehat{S}_{\text {imp }}$ for $\widehat{\sigma}$.

Putting it all together. The fixed optimal schedule for $\sigma$ induces a scheduling of $\widehat{\sigma}$. Let $\widehat{\mathbf{L}}$ be the load-vector of this schedule. By Eq. (7), we get that $\widehat{\mathbf{L}} \leq$ $s \mathbf{L}_{\text {Opt }}$. Thus, the cost of $\widehat{S}_{\text {imp }}$ (which was an optimal scheduling of $\widehat{\sigma}$ ) is at most $\|\widehat{\mathbf{L}}\| \leq s\left\|\mathbf{L}_{\text {OPT }}\right\|$.

Let $W_{u}$ be the total weight of unimportant jobs scheduled on machine 1 by Alg. We have that

$$
\left\|\left(W_{u} / C_{1}\right) \mathbf{1}_{1}\right\|=W_{u} B \leq n s^{k-\left\lceil\log _{s}\left(n^{2}\right)\right\rceil} \leq n \frac{\left\|\mathbf{L}_{\mathrm{OPT}}\right\|}{n^{2}} \leq \frac{\left\|\mathbf{L}_{\mathrm{OPT}}\right\|}{n}
$$

We are finally able to bound the cost of the entire schedule $S_{\mathrm{imp}} \cup S_{\text {unimp }}$ created by Alg:

$$
\operatorname{ALG}(\sigma)=\left\|\mathbf{L}_{\mathrm{imp}}+\left(W_{u} / C_{1}\right) \mathbf{1}_{1}\right\| \leq\left\|\mathbf{L}_{\mathrm{imp}}\right\|+\left\|\left(W_{u} / C_{1}\right) \mathbf{1}_{1}\right\| \leq(s+1 / n)\left\|\mathbf{L}_{\mathrm{OPT}}\right\|
$$

Recall that $s=1+\varepsilon / 2$. Thus, if $n>2 / \varepsilon$, then $\operatorname{AlG}(\sigma) \leq(1+\varepsilon) \operatorname{Opt}(\sigma)$. For inputs of length less than $2 / \varepsilon$, the oracle can simply encode the optimal solution using at most $\frac{2}{\varepsilon}\lceil\log m\rceil$ bits of advice. The total amount of advice used by our algorithm is $O\left(\varepsilon^{-1} \log ^{2} n\right)$.

We now consider scheduling problems where the goal is to maximize an objective function $f$. Recall that we assume that the objective function is nondecreasing. The most notable example is when $f$ is the minimum load. In the following theorem, we show how to schedule almost optimally on unrelated machines with only a rather weak constraint on $f$ (weaker than $f$ being a norm).

Theorem 10. Let P be a scheduling problem on $m$ unrelated machines where the goal is to maximize an objective function $f$. Assume that $f$ is non-decreasing, that $f(\alpha \mathbf{L}) \leq \alpha f(\mathbf{L})$ for every $\alpha \geq 0$, and $\mathbf{L} \in \mathbb{R}_{+}^{m}$. Then, for every $0<$ $\varepsilon \leq 1$, there exists a $(1+\varepsilon)$-competitive P -algorithm with advice complexity $O\left(\left(\frac{4}{\varepsilon} \log (n)+2\right)^{m} m^{2} \log n\right)$. In particular, if $m=O(1)$ and $\varepsilon=\Omega(1)$, the advice complexity is $O(\operatorname{polylog}(n))$.

Proof. Fix an input sequence $\sigma$ and an arbitrary optimal schedule. Let $s=$ $1+\varepsilon / 2$. The oracle uses $O(\log n)$ bits to encode $n$ using a self-delimiting encoding.

For $1 \leq j \leq m$, let $L_{j}$ be the load on machine $j$ in the optimal schedule. Furthermore, let $k_{j}$ be the unique integer such that $s^{k_{j}} \leq L_{j}<s^{k_{j}+1}$. We say that a job $r_{i}$ is unimportant to machine $j$ if $w_{i}(j)<s^{k_{j}-\left\lceil\log _{s}\left(n^{2}\right)\right\rceil \text {, important to }}$ machine $j$ if $s^{k_{j}-\left\lceil\log _{s}\left(n^{2}\right)\right\rceil} \leq w_{i}(j)<s^{k_{j}+1}$ and huge to machine $j$ if $w_{i}(j) \geq$ $s^{k_{j}+1}$. Note that if $r_{i}$ is huge to machine $j$, Opt does not schedule $r_{i}$ on machine $j$. A job which is important to at least one machine is called important. All other jobs are called unimportant. Note that, by definition, any unimportant job is unimportant (and not huge) to the machine where it is scheduled by Opt. We number the machines such that the first job which is important to machine $j$ arrives no later than the first job which is important to machine $j^{\prime}$ for every $j<j^{\prime}$. This numbering is written to the advice tape, using $O(m \log m)$ advice bits.

The algorithm works in $m+1$ phases (some of which might be empty). Phase 0 begins when the first request arrives. For $1 \leq j \leq m$, phase $j-1$ ends and phase $j$ begins when the first important job for machine $j$ arrives. Note that the same job could be the first important job for more than one machine. Phase $m$ ends with the last request of $\sigma$. For each phase, $j$, the oracle writes the index, $i$, of the request starting the phase and the unique integer $\Delta_{i}(j)$ such that $s^{k_{j}-\Delta_{i}(j)} \leq w_{i}(j)<s^{k_{j}-\Delta_{i}(j)+1}$.

The unimportant jobs are scheduled arbitrarily by our algorithm (it will become clear from the analysis of the algorithm that any choice will do).

We now describe how the algorithm schedules the important jobs in phase $j$ for $1 \leq j \leq m$. By definition, at any point in phase $j$, we have received an important job for machines $1,2, \ldots, j$ and no important job for machine $j+1$ has yet arrived.

The type of a job $r_{i}$ in phase $j$ is a vector $\boldsymbol{\Delta}_{i}=\left(\Delta_{i}(1), \ldots, \Delta_{i}(j)\right)$ where $\Delta_{i}\left(j^{\prime}\right)$ is the interval of $r_{i}$ on machine $j^{\prime}$ (so $s^{k_{j}-\Delta_{i}\left(j^{\prime}\right)} \leq w_{i}\left(j^{\prime}\right)<s^{k_{j}-\Delta_{i}\left(j^{\prime}\right)+1}$ ) or $\perp$ if the job is not important to machine $j^{\prime}$. Note that there are $\lceil(2+$ $\left.\left.\log _{s}\left(n^{2}\right)\right)^{j}\right\rceil$ possible job types in phase $j$. The oracle considers how the jobs in phase $j$ are scheduled in the fixed optimal schedule. For each job type $\boldsymbol{\Delta}$ and each machine $1 \leq j^{\prime} \leq j$, the oracle encodes the number of jobs of that type which are scheduled on machine $j^{\prime}$ during phase $j$. This can be done using $O\left(\left\lceil 2+\log _{s}\left(n^{2}\right)\right\rceil^{m} m \log n\right)$ bits of advice for a single phase, and $O(\lceil 2+$ $\left.\left.\log _{s}\left(n^{2}\right)\right\rceil^{m} m^{2} \log n\right)$ bits of advice for all $m$ phases.

Equipped with the advice described above, the algorithm simply schedules the important jobs in the current phase based on their types. This ensures that, for each machine $j$, the total load of important jobs that Opt schedules on machine $j$ is at most $s$ times as large as the total load of important jobs scheduled by AlG on machine $j$ (since if $r_{i}$ and $r_{i^{\prime}}$ are important to machine $j$ and of the same type, then $\left.w_{i}(j)<s \cdot w_{i^{\prime}}(j)\right)$.

In order to finish the proof, we need to show that the contribution of unimportant jobs to $\operatorname{Opt}(\sigma)$ is negligible (recall that all jobs are either important or unimportant). To this end, let $L_{j}^{\text {unimp }}$ (resp. $L_{j}^{\text {imp }}$ ) be the load on machine $j$ of the unimportant (resp. important) jobs scheduled on that machine in the optimal schedule. Note that $L_{j}=L_{j}^{\text {unimp }}+L_{j}^{\text {imp }}$. By the definition of an unimportant job (and since there trivially can be no more than $n$ unimportant jobs), we find that for every $1 \leq j \leq m$,

$$
L_{j}^{\text {unimp }}<n \cdot s^{-\left\lceil\log _{s}\left(n^{2}\right)\right\rceil} \cdot L_{j} \leq \frac{L_{j}}{n} .
$$

Thus, $L_{j}=L_{j}^{\mathrm{unimp}}+L_{j}^{\mathrm{imp}} \leq L_{j} / n+L_{j}^{\mathrm{imp}}$ from which $L_{j} \leq \frac{n}{n-1} \cdot L_{j}^{\mathrm{imp}}$ follows, assuming that $n>1$. Since this holds for all machines, and since as previously argued $\mathbf{L}_{\mathrm{OPT}}^{\mathrm{imp}} \leq s \cdot \mathbf{L}_{\mathrm{ALG}}^{\mathrm{imp}} \leq s \cdot \mathbf{L}_{\mathrm{ALG}}$, we get that

$$
\mathbf{L}_{\mathrm{OPT}} \leq \frac{n}{n-1} \mathbf{L}_{\mathrm{OPT}}^{\mathrm{imp}} \leq s \frac{n}{n-1} \mathbf{L}_{\mathrm{ALG}} .
$$

By assumption, the objective function $f$ satisfies $f(\alpha \mathbf{L}) \leq \alpha f(\mathbf{L})$ and is nondecreasing. Thus, we conclude that

$$
\operatorname{OPT}(\sigma)=f\left(\mathbf{L}_{\mathrm{OPT}}\right) \leq f\left(s \frac{n}{n-1} \mathbf{L}_{\mathrm{ALG}}\right) \leq s \frac{n}{n-1} f\left(\mathbf{L}_{\mathrm{ALG}}\right)=s \frac{n}{n-1} \operatorname{ALG}(\sigma) .
$$

For $n \geq 2+\frac{2}{\varepsilon}$, this gives a ratio of at most $1+\varepsilon$.

## References

1. Susanne Albers and Matthias Hellwig. Online makespan minimization with parallel schedules. In SWAT, volume 8503 of $L N C S$, pages 13-25, 2014.
2. Hans-Joachim Böckenhauer, Juraj Hromkovič, Dennis Komm, Sacha Krug, Jasmin Smula, and Andreas Sprock. The string guessing problem as a method to prove lower bounds on the advice complexity. Theor. Comput. Sci., 554:95-108, 2014.
3. Hans-Joachim Böckenhauer, Dennis Komm, Rastislav Královič, Richard Královič, and Tobias Mömke. On the advice complexity of online problems. In ISAAC, volume 5878 of LNCS, pages 331-340, 2009.
4. Allan Borodin, Sandy Irani, Prabhakar Raghavan, and Baruc h Schieber. Competitive paging with locality of reference. Journal of Computer and System Sciences, 50(2):244-258, 1995.
5. Joan Boyar, Leah Epstein, Lene M. Favrholdt, Kim S. Larsen, and Asaf Levin. Online bounded analysis. In CSR, volume 9691 of $L N C S$, pages 131-145, 2016.
6. Joan Boyar, Lene M Favrholdt, Christian Kudahl, Kim S Larsen, and Jesper W Mikkelsen. Online algorithms with advice: A survey. ACM Computing Surveys (CSUR), 50(2):19, 2017.
7. Joan Boyar, Lene M. Favrholdt, Christian Kudahl, and Jesper W. Mikkelsen. Advice complexity for a class of online problems. In STACS, volume 30 of LIPIcs, pages 116-129, 2015. Full paper to appear in Theory of Computing Systems.
8. Marek Chrobak and John Noga. LRU is better than FIFO. Algorithmica, 23(2):180-185, 1999.
9. Stefan Dobrev, Rastislav Královič, and Dana Pardubská. Measuring the problemrelevant information in input. RAIRO - Theor. Inf. Appl., 43(3):585-613, 2009.
10. Christoph Dürr, Christian Konrad, and Marc P. Renault. On the power of advice and randomization for online bipartite matching. In $E S A$, pages 37:1-37:16, 2016.
11. Yuval Emek, Pierre Fraigniaud, Amos Korman, and Adi Rosén. Online computation with advice. Theor. Comput. Sci., 412(24):2642-2656, 2011.
12. Anna R. Karlin, Mark S. Manasse, Larry Rudolph, and Daniel D. Sleator. Competitive snoopy caching. Algorithmica, 3:77-119, 1988.
13. Dennis Komm, Rastislav Královič, Richard Královič, and Christian Kudahl. Advice complexity of the online induced subgraph problem. In MFCS, volume 58 of LIPIcs, pages 59:1-59:13, 2016.
14. Jesper W. Mikkelsen. Randomization can be as helpful as a glimpse of the future in online computation. In ICALP, volume 55 of LIPIcs, pages 39:1-39:14, 2016.
15. Marc P. Renault, Adi Rosén, and Rob van Stee. Online algorithms with advice for bin packing and scheduling problems. Theor. Comput. Sci., 600:155-170, 2015.
16. Daniel D. Sleator and Robert E. Tarjan. Amortized efficiency of list update and paging rules. Commun. ACM, 28(2):202-208, 1985.
17. Jasmin Smula. Information Content of Online Problems: Advice versus Determinism and Randomization. PhD thesis, ETH, Zürich, 2015.
18. Andreas Sprock. Analysis of hard problems in reoptimization and online computation. PhD thesis, ETH, Zürich, 2013.

## Appendix

## A AOC-Complete Problems

For completeness, we state the full definition of minASGk from [7]:
Definition 7 ( [7]). The minimum asymmetric string guessing problem with known history, minASGK, has input $\left\langle ?, x_{1}, \ldots, x_{n}\right\rangle$, where $x=x_{1} \ldots x_{n} \in$ $\{0,1\}^{n}$, for some $n \in \mathbb{N}$. For $1 \leq i \leq n$, round $i$ proceeds as follows:

1. If $i>1$, the algorithm learns the correct answer, $x_{i-1}$, to the request in the previous round.
2. The algorithm answers $y_{i}=f\left(x_{1}, \ldots, x_{i-1}\right) \in\{0,1\}$, where $f$ is a function defined by the algorithm.

The output $y=y_{1} \ldots y_{n}$ computed by the algorithm is feasible, if $x \sqsubseteq y$. Otherwise, $y$ is infeasible. The cost of a feasible output is $|y|_{1}$, and the cost of an infeasible output is $\infty$.

In addition to minASGk, the class of AOC-complete problems also contains many graph problems. The following four graph problems are studied in the vertex-arrival model, so the requests are vertices, each presented together with its edges to previous vertices. The first three problems are minimization problems and the last one is a maximization problem. In Vertex Cover, an algorithm must accept a set of vertices which constitute a vertex cover, so for every edge in the requested graph, at least one of its endpoints is accepted. For Dominating Set, the accepted vertices must constitute a dominating set, so every vertex in the requested graph must be accepted, or one its neighbors must be accepted. In Cycle Finding, an algorithm must accept a set of vertices inducing a cyclic graph. For Independent Set, the accepted vertices must form an independent set, i.e., no two accepted vertices share an edge.

For Disjoint Path Allocation a path $P$ is given, and the requests are subpaths of $P$. The aim is to accept as many edge disjoint paths as possible.

For Set Cover, the requests are finite subsets from a known universe, and the union of the accepted subsets must be the entire universe. The aim is to accept as few subsets as possible.

## B Reductions for Theorem 2

In the proof of Theorem 2, a reduction sketch was given for the weighted online version of Vertex Cover. Here we include sketches for the reductions for the weighted versions of Cycle Finding, Dominating Set and Set Cover.

Cycle Finding Each input $\sigma=\left\langle x_{1}, x_{2}, \ldots, x_{n}\right\rangle$ to the problem minASGk, is transformed to $f(\sigma)=\left\langle v_{1}, v_{2}, \ldots, v_{n}\right\rangle$, where $V=\left\{v_{1}, v_{2}, \ldots, v_{n}\right\}$ is the vertex set of a graph with edge set $E=\left\{\left(v_{j}, v_{i}\right): f^{\prime}\left(x_{i}\right)=j\right\} \cup\left\{\left(v_{\text {MIN }}, v_{\text {MAX }}\right)\right\}$, where $f^{\prime}\left(x_{i}\right)$ is the largest $j<i$ such that $x_{j}=1$, MAX is the largest $i$ such that $x_{i}=1$,
and MIN is the smallest $i$ such that $x_{i}=1$. If $|\sigma|_{1}>2$, the vertices corresponding to 1 s form the only cycle in the graph.

The advice used by the minASGk algorithm ALG $_{1}$ consists of the advice used by the Cycle Finding algorithm $\mathrm{ALG}_{2}$ in combination with 1 bit indicating whether or not $|\sigma|_{1} \leq 2$ and in this case (an encoding of) one or two indices of 1 s in the input sequence. If $|\sigma|_{1}>2$, then $A_{L G}$ accepts some vertices, and $\mathrm{ALG}_{1}$ returns a 1 for the $x_{i}$ corresponding to each of those vertices.

If $\mathrm{AlG}_{1}$ returns a non-optimal feasible set, $\mathrm{AlG}_{2}$ does too, and the sets have the same weights, so $\mathrm{ALG}_{1}(\sigma) \leq \mathrm{ALG}_{2}(f(\sigma))+\operatorname{Opt}(\sigma)$. In this case, the weights of the optimal solutions for $\sigma$ and $f(\sigma)$ are both the sum of the weights of the elements corresponding to 1 s in $\sigma$, so $f$ is a length preserving $O(f(n))$-reduction.

Dominating Set Each input $\sigma=\left\langle x_{1}, x_{2}, \ldots, x_{n}\right\rangle$ to the problem minASGk, is transformed to $f(\sigma)=\left\langle v_{1}, v_{2}, \ldots, v_{n}\right\rangle$, where $V=\left\{v_{1}, v_{2}, \ldots, v_{n}\right\}$ is the vertex set of a graph with edge set $E=\left\{\left(v_{i}, v_{\text {MAX }}\right)\right\}$., where MAX is the largest $i$ such that $x_{i}=1$.

The advice used by the minASGK algorithm ALG $_{1}$ consists of the advice used by the Dominating Set algorithm $\mathrm{ALG}_{2}$ in combination with 1 bit indicating whether or not $|\sigma|_{1}=0$. If $|\sigma|_{1} \geq 1$, then there is another bit of advice indicating whether or not $\mathrm{ALG}_{2}$ accepted $v_{\mathrm{MAX}}$. If $\mathrm{ALG}_{2}$ did not accept $v_{\mathrm{MAX}}$, the advice also contains an index of a vertex corresponding to a 0 in $\sigma$ which was accepted, plus the index of $v_{\text {max }}$.

If $A_{1 G}$ 's solution is feasible, but not optimal, then $|\sigma|_{1}>0$ and $A_{1} G_{2}$ accepts some vertices, and $\mathrm{ALG}_{1}$ returns a 1 for the $x_{i}$ corresponding to each of those vertices (though, in the case where $v_{\text {max }}$ was rejected, it answers 1 for $x_{\text {MAX }}$ and answers 0 for the earlier request indicated by the advice).

If $|\sigma|_{1}>0$ a minimum weight dominating set for $f(\sigma)$ consists of exactly those vertices corresponding to 1 s in $\sigma$, so the weights of the optimal solutions for $\sigma$ and $f(\sigma)$ are both the sum of the weights of the elements corresponding to 1s in $\sigma$, unless $\mathrm{ALG}_{2}$ did not accept $v_{\mathrm{MAX}}$. However, the weight of $x_{\mathrm{MAX}} \leq \operatorname{Opt}(\sigma)$, so $\operatorname{Alq}_{1}(\sigma) \leq \operatorname{ALG}_{2}(f(\sigma))+\operatorname{Opt}(\sigma)$. Thus, $f$ is a length preserving $O(\log (n)$ reduction.

Set Cover This reduction is very similar to that for Dominating Set. Each input $\sigma=\left\langle x_{1}, x_{2}, \ldots, x_{n}\right\rangle$ to the problem minASGk, MAX is the largest $i$ such that $x_{i}=1$. In the set cover instance, the universe is $\{1, \ldots, n\}$, and $f(\sigma)$ is a set of $n$ requests, where request $i$ is $\{i\}$, unless $i=$ mAx, in which case, the set consists of MAX and all of the $j$ where $x_{j}=0$.

As with the reduction to Dominating Set, this is a length preserving $O(\log (n)$ reduction.


[^0]:    * This work was partially supported by the Villum Foundation, grant VKR023219, and the Danish Council for Independent Research, Natural Sciences, grant DFF-1323-00247.

[^1]:    ${ }^{1}$ For example, $\lceil\log n\rceil$ could be written in unary ( $\lceil\log n\rceil$ ones, followed by a zero) before writing $n$ itself in binary.

[^2]:    ${ }^{2} f$ is a norm if $f(\alpha \mathbf{v})=|\alpha| f(\mathbf{v}), f(\mathbf{u}+\mathbf{v}) \leq f(\mathbf{u})+f(\mathbf{v})$, and $f(\mathbf{v})=0 \Rightarrow \mathbf{v}=\mathbf{0}$.

