Online Resource Minimization

Anton J. Kleywegt Vijay S. Nori
Martin W. P. Savelsbergh Craig A. Tovey*

1 Introduction

Many real-world problems can be modeled as online optimization problems. In these problems, information arrives over time, while decisions must be made at time points based only on the information then available. Surveys of online algorithms are given in Sgall (1997) and Coffman Jr., Garey and Johnson (1996), while Albers (1997) and Ascheuer et al. (1998) describe practical situations where online algorithms can be used.

We investigate a new basic online problem in which work with different deadlines arrives over time, a level of resources must be chosen at each decision point so as to meet all deadlines, and the objective is to minimize the maximum resource usage. The problem is motivated by cases in vehicle fleet planning, warehouse allocation, workspace procurement, and electricity consumption, where the controllable portion of total costs varies with the maximum amount of resource procured. In Section 2, we define our online resource minimization problem, and introduce the notation that will be used throughout.

Online algorithms are typically evaluated by means of their competitive ratio, which is the worst-case ratio of the performance of the online algorithm to the performance of an optimal algorithm with perfect information. We end section 2 by deriving a closed-form expression for the optimal value with perfect information.

In Section 3 we introduce the α-policy, a simple parameterized policy with parameter α and worst-case ratio α, provided it is feasible. The main result of this section is that, with appropriate parameter choice, the α-policy has as good a worst-case ratio as any other policy. This result applies to any online minimax problem satisfying certain natural conditions. For our problem, we also show that an optimal parameter value α^* exists. Hence α^* also equals the optimal competitive ratio. We also show that to find α^* it is sufficient to study the more restricted version of the problem in which all deadlines coincide with the planning horizon.

We introduce two other classes of policies, called the ϕ-policies and the ψ-policies, in Section 4. Analysis of these policies provides an upper bound on the optimal competitive ratio of $\alpha^* \leq 3.45$. The ϕ-policy is shown to achieve competitive ratio at best 4.

We tackle lower bounds on α^* in section 5, by finding a closed-form integral expression, and a coupled differential equation system for continuous approximations of the problem. Analytic and numeric solutions of these continuous models lead to a 10,000 period instance that proves $\alpha^* > 2.51$.

We also conduct a computational study of the α-policy, the ϕ-policy, and the ψ-policy, and find that the ϕ-policy usually performs best, and the α-policy typically performs worst. Thus the average-case performance rankings are the reverse of the worst-case performance rankings. These results are described in Section 6.

2 Problem Definition

Let $T \in \mathbb{Z}_+$ denote the known time horizon, and let the decision points be $\{1, 2, \ldots, T\}$. Let a_{ij} denote the amount of work that arrives at time i with deadline $j \in \{i, i+1, \ldots, T\}$. Without loss of generality, assume that the initial amount of work waiting in the system is zero. Thus an instance I_T with time horizon T of the online resource minimization problem is given by $I_T = (a_{11}, a_{12}, \ldots, a_{TT})$. Let r_i denote the amount of resource made available at decision point i. Thus

*School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, GA 30332-0205.

the maximum amount of work that can be performed at time i is r_i. Let q_{ij} denote the amount of work with deadline j that is performed at time $i \leq j$. Thus $(r_i, q_{ii}, \ldots, q_{iT})$ denotes the decision made at time i. To be feasible, $(r_i, q_{ii}, \ldots, q_{iT})$ must satisfy the following for all i.

$$\sum_{j=i}^{T} q_{ij} \leq r_i \tag{1}$$

$$\sum_{k=1}^{i} q_{ki} = \sum_{k=1}^{i} a_{ki} \tag{2}$$

$$\sum_{k=1}^{i} q_{kj} \leq \sum_{k=1}^{i} a_{kj} \qquad \text{for all } j \tag{3}$$

Constraint (1) states that the total amount of work performed at time i cannot exceed the amount of work that can be accomplished with r_i amount of resource. Constraint (2) states that all work must be performed by the respective deadlines. Constraint (3) states that work cannot be performed before it has arrived. Because we have no information about future arrivals, decision $(r_i, q_{ii}, \ldots, q_{iT})$ can depend on past arrivals only, and not on any arrivals after time i.

Let $w_{ij} \equiv a_{1j} - q_{1j} + a_{2j} - q_{2j} + \cdots + a_{ij}$. Thus, w_{ij} denotes the amount of work with deadline j waiting at time i to be performed, after the arrivals at time i have taken place, but before any work has been performed at time i. Note that $w_{i+1,j} = w_{ij} - q_{ij} + a_{i+1,j}$. Constraint (3) implies that $w_{ij} \geq 0$ for all i and j. Constraint (2) implies that $q_{ii} = w_{ii}$, i.e, all remaining work with deadline i must be completed at time i. Thus $w_{ij} = 0$ for $j < i$, i.e., no work waits to be performed after its deadline. Let $m_i \equiv \max\{0, r_1, \ldots, r_{i-1}\}$ denote the maximum amount of resource made available up to time i, which is the amount of resource procured up to time i. Thus $m_{i+1} = \max\{m_i, r_i\}$.

Let π denote any policy for the online resource minimization problem, and let Π denote the set of all online policies. Let $r_i^\pi(I_T)$ denote the amount of resource procured in period i under policy π for instance I_T. If policy π assigns resource amounts $r_i = r_i^\pi(I_T)$ under instance I_T, then the procurement of policy π is

$$v^\pi(I_T) \equiv \max\{r_1, \ldots, r_T\}$$

if constraints (1), (2), and (3) are satisfied under policy π at all times $i \in \{1, \ldots, T\}$, and $v^\pi(I_T) \equiv \infty$ otherwise. The objective is to minimize the procurement. Thus the optimal value with perfect information is given by

$$v^*(I_T) \equiv \min_{r_1,\ldots,r_T} \max\{r_1, \ldots, r_T\}$$

subject to constraints (1), (2), and (3). For ease we assume $v^*(I_T) > 0$.

The worst-case ratio ρ_T^π of policy π over all instances with time horizon T is given by

$$\rho_T^\pi \equiv \sup_{I_T} \frac{v^\pi(I_T)}{v^*(I_T)}$$

and the worst-case ratio ρ^π of policy π over all instances is given by

$$\rho^\pi \equiv \sup_{T \in \mathbb{Z}_+} \rho_T^\pi$$

The optimal worst-case ratio ρ_T^* over all online policies for instances with time horizon T is given by

$$\rho_T^* \equiv \inf_{\pi \in \Pi} \rho_T^\pi$$

and the optimal worst-case ratio ρ^* over all online policies and instances is given by

$$\rho^* \equiv \sup_{T \in \mathbb{Z}_+} \rho_T^*$$

Lemma 1

$$\rho^* = \inf_{\pi \in \Pi} \rho^\pi$$

An online policy $\pi^* \in \Pi$ is called *optimal* if $\rho_T^{\pi^*} = \rho_T^*$ for all $T \in \mathbb{Z}_+$.

Lemma 2 *The optimal value $v^*(I_T)$ with perfect information is given by*

$$v^*(I_T) = \max_{\{s,t \in \{1,\ldots,T\}: s \leq t\}} \frac{1}{t-s+1} \sum_{i=s}^{t} \sum_{j=i}^{t} a_{ij}$$

Proof: Let

$$\gamma^*(I_T) = \max_{s \leq t} \frac{1}{t-s+1} \sum_{i=s}^{t} \sum_{j=i}^{t} a_{ij}$$

Clearly $v^*(I_T) \geq \gamma^*(I_T)$. It remains to show that $v^*(I_T) \leq \gamma^*(I_T)$. Consider the solution (with perfect information) inductively defined as follows.

Let $W_i \equiv w_{ii} + \cdots + w_{iT}$. Thus W_i denotes the total amount of work waiting to be processed after the arrivals at time i have taken place, but before any processing at time i. If $W_i \leq \gamma^*(I_T)$, then let $r_i = W_i$, and process all W_i work. Else, if $W_i > \gamma^*(I_T)$, then let $r_i = \gamma^*(I_T)$, and process $\gamma^*(I_T)$ work in earliest due date (EDD) order. Thus this solution never allocates more than $\gamma^*(I_T)$ amount of resource at a time. The solution is feasible and has objective value $\gamma^*(I_T)$ if and only if all work is completed by the deadlines, which is established next.

For any time t, let $\ell(t)$ denote the last time $\tau \in \{1, \ldots, t\}$ that there is no work with deadline less than or equal to t waiting to be processed after the processing at time τ has taken place. If there is no such time $\tau \in \{1, \ldots, t\}$, then let $\ell(t) = 0$ (which will turn out never to be the case). Thus the objective is to show that $\ell(t) = t$, which implies that all work with deadline less than or equal to t has been processed at the end of time t. Suppose $\ell(t) < t$. Then the amount of work with deadline less than or equal to t that has to be processed in $\{\ell(t)+1, \ldots, t\}$ is $\sum_{i=\ell(t)+1}^{t} \sum_{j=i}^{t} a_{ij}$, which is less than or equal to $\gamma^*(I_T)(t - \ell(t))$ from the definition of $\gamma^*(I_T)$. However, from the definition of $\ell(t)$ and the solution, there is work with deadline less than or equal to t remaining at the end of each of the times in $\{\ell(t)+1, \ldots, t\}$, and thus $\gamma^*(t - \ell(t))$ work with deadline less than or equal to t is processed in $\{\ell(t)+1, \ldots, t\}$. Thus the amount of work with deadline less than or equal to t that is processed in $\{\ell(t)+1, \ldots, t\}$ is at least as much as the amount that needs to be processed to finish all such work by time t. Thus all work is finished by the deadlines, and $v^*(I_T) \leq \gamma^*(I_T)$. □

3 Optimal Online Algorithm

Let $I_T^i = (a_{11}, a_{12}, \ldots, a_{iT}, 0, \ldots, 0)$ denote the truncated instance with the same arrivals up to time i as instance I_T, and 0 arrivals thereafter.

Definition 1 *The α-policy is an online policy with parameter $\alpha \geq 1$, and defined by the procurement*

$$r_i^\alpha(I_T) = \alpha v^*(I_T^i) \quad \forall i.$$

At period i the α-policy assigns α times the optimal value with perfect information of the truncated instance I_T^i. This policy can miss deadlines if α is too small.

Given T, if there is a value for the parameter α that provides the best competitive ratio for the α-policy, it is denoted α_T^*. We show in this section that these optimal parameters exist, and we let $\alpha^* \equiv \sup_T \alpha_T^*$.

The α-policy turns out to be optimal, though we do not know the optimal value α_T^* of the parameter. This optimality property holds for a general class of problems where the maximum quantity of resource procured is to be minimized. An instance I_T of a problem consists of work arriving each time period, in a form depending on the problem; resources procured are denoted $R = \{r_1, \ldots, r_T\}$. A binary-valued function $F(I_T, R)$ defines feasibility of procurements for the problem. The function $F(I_T, R) =$ true iff all the jobs in I_T can be completed on time with resource allocation R. The definitions of $m_i, v^\pi(I_T), \rho_T^*$, etc. remain the same.

We only require that F be monotonic in resources and work. That is,

resource monotonicity Suppose $F(I, R) =$ true. If $\hat{R} \geq R$, then $F(I, \hat{R}) =$ true.

work monotonicity Suppose $F(I, R) =$ true. If $\hat{I} \preceq I$, then $F(\hat{I}, R) =$ true.

Theorem 1 *Suppose an online minimax resource procurement problem satisfies resource and work monotonicity. For any online policy π, if $\rho_T^\pi < \infty$, then the α-policy with parameter $\alpha = \rho_T^\pi$ achieves the same competitive ratio ρ_T^π.*

Corollary 1 *If there exists an optimal online policy, then for any T, ρ_T^* is the least parameter α for which the α-policy is always feasible.*

Proof: Choose parameter $\alpha = \rho_T^\pi$. First we claim that $\forall I_T \forall t : 1 \leq t \leq T, \;\; r_t^\alpha(I_T) \geq r_t^\pi(I_T)$. Proof of claim: suppose not. Then $\exists I_T, 1 \leq t \leq T$ s.t.

$$r_t^\pi(I_T) > r_t^\alpha(I_T).$$

Consider the behavior of π on the instance I_T^t. Since π does not look ahead,

$$r_t^\pi(I_T) = r_t^\pi(I_T^t).$$

Since the cost function is the maximum of r_i,

$$r_t^\pi(I_T^t) \leq v^\pi(I_T^t).$$

Therefore $r_t^\alpha(I_T) < v^\pi(I_T^t)$. On the other hand, by definition, $r_t^\alpha(I_T) = \alpha v^*(I_T^t) = \rho_T^\pi v^*(I_T^t)$. So

$$\frac{v^\pi(I_T^t)}{v^*(I_T^t)} > \rho_T^\pi.$$

We've exhibited an instance I_T^t for which π has performance ratio $> \rho_T^\pi$, contradicting the assumption about π. This proves the claim.

Second, from the claim, the α-policy at every period allocates at least as much resource as π does. By assumption $\rho_T^\pi < \infty$, so by resource monotonicity of F, the α-policy always generates feasible resource allocations. Then by definition,

$$v^\alpha(I_T) = \max_t r_t^\alpha(I_T) \quad \forall I_T$$

By work monotonicity of F,

$$v^*(I_T^t) \le v^*(I_T)$$

because any allocation feasible for I_T is also feasible for I_T^t. Therefore, $r_t^\alpha(I_T) = \alpha v^*(I_T^t) \le \alpha v^*(I_T)$. This is true for all t, and thus $v^\alpha(I_T) \le \alpha v^*(I_T)$. Hence

$$\frac{v^\alpha(I_T)}{v^*(I_T)} \le \alpha = \rho_T^\pi.$$

This proves that the α-policy achieves the same worst-case performance as π. Thus if there is a least α value which always gives feasible allocations, it provides an optimal online policy. □

The monotonicity properties are obviously satisfied by our basic problem, defined in Section 2, so Theorem 1 applies. We now turn to the question of existence of optimal online policies, so that we may apply Corollary 1.

Proposition 1 *If $\rho_T^* < \infty$ and the feasibility function F is upper semicontinuous in R, then there exists an optimal α-policy. That is, the α-policy with $\alpha = \rho_T^*$ is feasible.*

Proof: Consider any work stream I_T. For any α, let $R^\alpha \equiv (\alpha v^*(I_T^1), \ldots, \alpha v^*(I_T^T))$. Choose any $\varepsilon \in (0,1)$. Because F is upper semicontinuous in R, there exists $\delta > 0$ such that $F(I_T, R^{\rho_T^*}) > F(I_T, R) - \varepsilon$ for all R with $\|R^{\rho_T^*} - R\| < \delta$. Let $\overline{v}(I_T) \equiv \max\{1, v^*(I_T)\}$. Because $\rho_T^* < \infty$, there exists α such that $F(I_T, R^\alpha) = 1$ and $\rho_T^* \le \alpha < \rho_T^* + \delta/(T\overline{v}(I_T))$. Then $\rho_T^* v^*(I_T^i) \le \alpha v^*(I_T^i) < \rho_T^* v^*(I_T^i) + \delta/T$ for all $i = 1, \ldots, T$. Thus $\|R^{\rho_T^*} - R^\alpha\| < \delta$, and hence $F(I_T, R^{\rho_T^*}) > F(I_T, R^\alpha) - \varepsilon = 1 - \varepsilon$, which implies that $F(I_T, R^{\rho_T^*}) = 1$. □

For our problem, we have the following result. The proof is omitted.

Proposition 2 *For the problem defined in Section 2, assuming q is chosen according to EDD assignment order, the feasibility function F is upper semicontinuous in R. This implies that an optimal parameter α_T^* exists, and $\alpha_T^* = \rho_T^*$. Also, $\alpha^* = \rho^*$.*

Theorem 2 *If for some α and some instance I_T, the α-policy gives an infeasible procurement ($v^\alpha(I_T) = \infty$), then there exists an instance $\hat{I}_T$, in which all work is due the last period T, for which the α-policy with the same parameter α does not give a feasible procurement. Therefore, ρ^* is the same for the multiple deadline problem and for the single deadline problem in which all work is due in the last time period.*

Proof: Let I_T be an instance with minimum T satisfying the hypothesis. Let R^α denote the procurements generated by the α-policy on I_T. By hypothesis they are not feasible. Let I_{T_t} denote the instance I_T with all work due after period t removed. Let R_{T-1}^α denote the procurements of the α-policy on $I_{T_{T-1}}$. By the minimality of T, and by a padding argument, R_{T-1}^α is feasible for the instance $I_{T_{T-1}}^{T-1} = I_{T_{T-1}}$.

For every t, any R feasible for I_T^t is also feasible for $I_{T_{T-1}}^t$ by work monotonicity. Therefore $v^*(I_{T_{T-1}}^t) \le v^*(I_T^t)$ for all t. By definition of the α-policy, this implies $R_{T-1}^\alpha \le R^\alpha$. So by resource monotonicity, R^α is a feasible procurement for $I_{T_{T-1}}$.

Let $\hat{I}_T$ denote the instance I_T, altered so all work has deadline T, and let $\hat{R}^\alpha$ denote the alpha policy procurements for instance $\hat{I}_T$. If work is scheduled by EDD, it follows from the hypothesis that R^α is not feasible for $\hat{I}_T$. Now a similar argument to the above shows $v^*(\hat{I}_T^t) \le v^*(I_T^t)$ for all t. Hence $\hat{R}^\alpha \le R^\alpha$. But since R^α is not feasible for $\hat{I}_T$, neither is $\hat{R}^\alpha$, by resource monotonicity. Thus we have exhibited a single deadline instance, namely $\hat{I}_T$, for which the α-policy fails. Therefore, α^* for the single deadline problem is not less than α^* for the multiple deadline problem. But the reverse is obvious because a single deadline instance is a special case of the multiple deadline problem. Thus $\rho^* = \alpha^*$ is the same for the single and multiple deadline problems. □

4 Other Online Policies

From numerical experiments, it was found that the α-policies do not have good average case performance. In this section we study other classes of policies, called the ϕ-policies and the ψ-policies. It seems that these policies have much better average case performance than the α-policies. The numerical experiments are discussed in Section 6. These policies also provide upper bounds on ρ^*_T and ρ^*. From Theorem 2, ρ^*_T is the same for the single and multiple deadline versions of the online resource minimization problem. Therefore most of the analysis of these policies will be for the single deadline version.

For the single deadline version, let a_i denote the amount of work that arrives at time i. An instance I_T of the problem is then given by $I_T = (a_1, a_2, \ldots, a_T)$. A decision is fully specified by the amount of resource r_t allocated at time t. Let $w_t \equiv a_1 - r_1 + a_2 - r_2 + \cdots + a_t$ denote the amount of work waiting to be performed at time t, after the arrival has taken place but before any work has been performed at time t.

The optimal value with perfect information for the single deadline version is given by

$$v^*(I_T) \;=\; \max_{t\in\{1,\ldots,T\}} \frac{\sum_{i=t}^{T} a_i}{T-t+1} \tag{4}$$

The following lemma, which is used later, establishes characteristics of an instance based on its optimal value with perfect information.

Lemma 3 *For any instance $I_T = (a_1, a_2, \ldots, a_T)$, let $a = \sum_{t=1}^{T} a_t/T$. If the optimal value with perfect information $v^*(I_T) = a+\Delta$ $(\Delta \geq 0)$, then*

$$\sum_{i=t}^{T} a_i \;\leq\; (T+1-t)(a+\Delta), \quad t = 1,\ldots,T \tag{5}$$

$$\sum_{i=1}^{t} a_i \;\geq\; ta - (T-t)\Delta, \quad t = 1,\ldots,T \tag{6}$$

4.1 ϕ-policies

In this section we study another parameterized class of policies for the online resource minimization problem, called the ϕ-policies, which had the best average case performance in numerical experiments.

Policy ϕ_1 *At each decision point divide the amount of waiting work by the number of remaining periods and allocate that amount of resource, i.e., $r_t^{\phi_1} \equiv w_t/(T-t+1)$.*

It can be verified that

$$r_t^{\phi_1} = \sum_{i=1}^{t} \frac{a_i}{T-i+1}, \quad t = 1,\ldots,T. \tag{7}$$

Since the values of $r_t^{\phi_1}$ are nondecreasing in t, it follows that $v^{p_1}(I_T) \equiv \max_{t\in\{1,\ldots,T\}} r_t^{\phi_1} = r_T^{\phi_1}$.

Theorem 3 *For policy ϕ_1, $\rho_T^{\phi_1} = \sum_{t=1}^{T} 1/t$.*

Proof: We give an intuitive sketch of the proof. Consider the class of instances I_T with $\sum_{t=1}^{T} a_t/T = a$ and $v^*(I_T) = a + \Delta$. We want to determine values a_i for $k = 1,\ldots,T$, that maximize r_T over this class of instances. Since the coefficients of a_i in (7), i.e., $1/T, 1/(T-1), \ldots, 1/2, 1$, form an increasing sequence, we want to successively make $a_i, i = 1,\ldots,T$, as small as possible to make a_T and r_T as large as possible From (6) we have $a_1 \geq a - (T-1)\Delta$. We set $a_1 = a - (T-1)\Delta$. Consequently, $a_2 \geq a + \Delta$, since $a_1 + a_2 \geq 2a - (T-2)\Delta$ (by (6)). We set $a_2 = a + \Delta$. Proceeding similarly, we find $a_i = a + \Delta, i = 2,\ldots,T$. Hence, the worst-case ratio over all instances in the class is

$$\frac{1}{a+\Delta}\left[\frac{a-(T-1)\Delta}{T} + (a+\Delta)\sum_{t=1}^{T-1}\frac{1}{t}\right]$$
$$= \sum_{t=1}^{T}\frac{1}{t} - \frac{\Delta}{a+\Delta}$$

Since $\Delta \geq 0$, it follows that $\rho_T^{\phi_1} = \sum_{t=1}^{T} 1/t$. □

Because the supremum over $\Delta \geq 0$ is attained at $\Delta = 0$, and the resulting expression is independent of a, it follows that any instance of the form $I_T = (a,\ldots,a)$ for any positive a is a worst-case instance with time horizon T for policy ϕ_1. Also, the policy does not have a bounded worst-case ratio because $\lim_{T\to\infty}\sum_{t=1}^{T} 1/t = \infty$, and thus $\rho^{\phi_1} = \infty$.

At every decision point, policy ϕ_1 acts as though there will be no work arriving in the future. The next policy also determines the amount of resource to allocate based on the amount of waiting work and the number of remaining periods, but allocates *more* than ϕ_1 to compensate for work that may arrive in the remaining periods.

Policy ϕ_p *At the first $T-p$ decision points, divide the amount of waiting work by the number of remaining periods and allocate p ($p \in \{2,\ldots,T\}$) times that amount of resource. At the last p decision points, allocate as much resource as there is waiting work. That is,*

$$r_t^{\phi_p} \equiv \min\left\{\frac{pw_t}{T-t+1}, w_t\right\}.$$

Note that $r_t^{\phi_p} = w_t$ for all $t \in \{T-p+1,\ldots,T\}$, i.e., at each time $t \in \{T-p+1,\ldots,T\}$, all waiting work is completed, and thus $r_t^{\phi_p} = a_t$ for all $t \in \{T-p+2,\ldots,T\}$. The number of resources allocated at time t can be written as

$$r_t^{\phi_p} = \begin{cases} p\left(\prod_{i=t+1}^{t+p-1}(T-i+1)\right) \times \\ \quad \left(\sum_{i=1}^{t} a_i \prod_{j=0}^{p-1}\left(\frac{1}{T+1-(i+j)}\right)\right), \\ \qquad t = 1,\ldots,T-p+1 \\ a_t, \quad t = T-p+2,\ldots T \end{cases} \tag{8}$$

Next we characterize worst-case instances under policy ϕ_p.

Lemma 4 *Let $I_T = (a_1, a_2, \ldots, a_T)$ and $I'_T = (a_1, a_2, \ldots, a_{T-p}, a_{T-p+1} + \ldots + a_T, 0, \ldots, 0)$. Then $v^{\phi_p}(I_T)/v^*(I_T) \le v^{\phi_p}(I'_T)/v^*(I'_T)$.*

Proof: From (4) it can be verified that $v^*(I'_T) \le v^*(I_T)$. Up to time $T-p$, I_T and I'_T are the same. Thus

$$\begin{aligned} v^{\phi_p}(I_T) &= \max\Big\{r_1^{\phi_p},\ldots,r_{T-p}^{\phi_p}, \\ &\qquad w(h_{T-p}) - r_{T-p} + a_{T-p+1}, a_{T-p+2},\ldots,a_T\Big\} \\ v^{\phi_p}(I'_T) &= \max\Bigg\{r_1^{\phi_p},\ldots,r_{T-p}^{\phi_p}, \\ &\qquad w(h_{T-p}) - r_{T-p} + a_{T-p+1} + \sum_{t=T-p+2}^{T} a_t, 0,\ldots,0\Bigg\} \end{aligned}$$

Hence, $v^{\phi_p}(I'_T) \ge v^{\phi_p}(I_T)$, and $v^{\phi_p}(I_T)/v^*(I_T) \le v^{\phi_p}(I'_T)/v^*(I'_T)$. □

Let $\mathcal{C}_T$ denote the class of instances I_T with $a_{T-p+2} = \ldots = a_T = 0$. Lemma 4 shows that to determine $\rho_T^{\phi_p}$, it is sufficient to consider only instances in class $\mathcal{C}_T$, i.e.,

$$\rho_T^{\phi_p} = \sup_{I_T \in \mathcal{C}_T} \frac{v^{\phi_p}(I_T)}{v^*(I_T)}.$$

Note that for any instance $I_T \in \mathcal{C}_T$, $r_{T-p+2}^{\phi_p} = \ldots = r_T^{\phi_p} = 0$.

To determine a bound on the worst-case ratio ρ^{ϕ_p} of policy ϕ_p, we initially allow instances with negative a_t. It is shown later that instances with positive a_t achieve the same worst-case ratio as instances which allow negative a_t. Let $\mathcal{C}_{T,a,\Delta}$ denote the class of instances $I_T \in \mathcal{C}_T$ such that $\sum_{t=1}^{T} a_t/T = a$ and $v^*(I_T) = a + \Delta$.

Lemma 5 *Let $I_T^{(k)} = (a_1^{(k)},\ldots,a_k^{(k)},0,\ldots,0) \in \mathcal{C}_{T,a,\Delta}$ be defined by $a_1^{(1)} = aT$ and $a_t^{(1)} = 0$ for all $t \in \{2,\ldots,T\}$ for $k = 1$ (and $\Delta = 0$), and*

$$a_t^{(k)} = \begin{cases} a-(T-1)\Delta & t = 1 \\ a+\Delta & t = 2,\ldots,k-1 \\ (T-k+1)(a+\Delta) & t = k \\ 0 & t = k+1,\ldots,T \end{cases} \tag{9}$$

for $k \in \{2,\ldots,T-p+1\}$ for any $\Delta \ge 0$. Consider any instance $I'_T \in \mathcal{C}_{T,a,\Delta}$. Let $r_t^{(k)}$ (r'_t) denote the amount of resources allocated by policy ϕ_p at time t, under instance $I_T^{(k)}$ (I'_T). Then

$$r'_k \;\le\; r_k^{(k)}$$

Proof: We give an intuitive sketch of the proof. To get the largest possible value for $r_k^{(k)}$, we want to have as much work as possible arriving before or at decision point k. Therefore, we set $a_t^{(k)} = 0$ for $t = k+1,\ldots,T$. Since the coefficients of a_t in (8) form an increasing sequence, we want to successively make a_t for $t = 1,\ldots,k-1$ as small as possible to make $a_k^{(k)}$ and $r_k^{(k)}$ as large as possible. From (6), $a_1 \ge a-(T-1)\Delta$. Thus we set $a_1^{(k)} = a-(T-1)\Delta$. Consequently, $a_2 \ge a+\Delta$, since $a_1 + a_2 \ge 2a-(T-2)\Delta$. Thus we set $a_2^{(k)} = a+\Delta$. Proceeding similarly, $a_t^{(k)} = a+\Delta$, for $t = 2,\ldots,k-1$, and we set $a_k^{(k)} = Ta - ((k-2)(a+\Delta) + a - (T-1)\Delta) = (T-k+1)(a+\Delta)$. When $k = 1$, all work arrives in the first period, i.e., $a_1^{(k)} = Ta$. □

Theorem 4 *For policy ϕ_p*

$$\rho_T^{\phi_p} = \max\left\{p, \max_{k\in\{2,\ldots,T-p+1\}}\left\{p\prod_{i=1}^{p-1}(T-k-i+1)\times \sum_{i=1}^{k}\prod_{j=0}^{p-1}\left[\frac{1}{T+1-(i+j)}\right]+p\frac{T-k}{T-k+1}\right\}\right\}$$

Proof:

$$\begin{aligned}\rho_T^{\phi_p} &= \sup_{I_T\in\mathcal{C}_T}\frac{v^{\phi_p}(I_T)}{v^*(I_T)}\\ &= \sup_{I_T\in\mathcal{C}_T}\max_{k\in\{1,\ldots,T-p+1\}}\frac{r_k^{\phi_p}}{v^*(I_T)}\\ &= \max_{k\in\{1,\ldots,T-p+1\}}\sup_a\sup_{\Delta\geq 0}\sup_{I_T\in\mathcal{C}_{T,a,\Delta}}\frac{r_k^{\phi_p}}{v^*(I_T)}\\ &= \max_{k\in\{1,\ldots,T-p+1\}}\sup_a\sup_{\Delta\geq 0}\frac{r_k^{(k)}}{a+\Delta}\end{aligned}$$

from which the result follows by substituting in (8) and (9). Note that the worst-case ratio $\rho_T^{\phi_p}$ is attained at $\Delta=0$ and any $a>0$. □

Theorem 5 *For $p\in\{2,\ldots,T\}$, $\rho^{\phi_p}=p^2/(p-1)$.*

Proof: Define f_1, f_2 by

$$\begin{aligned}f_1(k) &\equiv p\prod_{i=1}^{p-1}(T-k-i+1)\times \sum_{i=1}^{k}\prod_{j=0}^{p-1}\left[\frac{1}{T+1-(i+j)}\right]\\ f_2(k) &\equiv p\frac{T-k}{T-k+1}\end{aligned}$$

It follows that

$$\begin{aligned}f_1(k) &= \frac{p}{p-1}\left[1-\frac{\Gamma(T-k+1)\Gamma(T-p+2)}{\Gamma(T-k-p+2)\Gamma(T+1)}\right]\\ &\leq \frac{p}{p-1}\end{aligned}$$

for all $T\geq 2$. Since $f_2'(k)=-p/(T-k+1)^2<0$, $f_2(k)$ is strictly decreasing, and $f_2(k)\leq f_2(2)=p(T-2)/(T-1)<p$ for all $k\in[2,T-p+1]$. Hence,

$$\begin{aligned}\rho^{\phi_p} &= \max\{p, \sup_{T\geq 2}\max_{k\in\{2,\ldots,T-p+1\}}\{f_1(k)+f_2(k)\}\}\\ &\leq \max\{p, p/(p-1)+p\} = p^2/(p-1)\end{aligned}$$

By choosing $k=\lceil T-\sqrt{T}\rceil$ and letting $T\to\infty$, it follows that $\rho^{\phi_p}\geq p^2/(p-1)$. □

For $p\in\{2,\ldots,T\}$, $p^2/(p-1)$ attains its minimum value at $p=2$, since it is an increasing function. Hence, policy ϕ_2 has the smallest worst-case ratio among the ϕ-policies.

In the policies considered so far, we have ignored the fact that the resources procured in earlier periods are available in later periods at no extra cost. It is intuitively clear that these resources should be used whenever possible. The next policy is the extension of policy ϕ_p in which we do precisely that.

Policy φ_p *At the first $T-p$ decision points, divide the amount of waiting work by the number of remaining periods and allocate p ($p\in\{2,\ldots,T\}$) times that many resources, unless the amount of available resource is larger. If the amount of available resource is larger, allocate all available resource or as much resource as there is waiting work, whichever is smaller. At the last p decision points, allocate as much resource as there is waiting work. That is,*

$$r_t^{\varphi_p}=\min\left\{\max\left\{m_t, \frac{p}{T-t+1}w_t\right\}, w_t\right\}.$$

Similar to policy ϕ_p, policy φ_p has $r_t^{\varphi_p}=a_t$ for all $t\in\{T-p+2,\ldots,T\}$.

Theorem 6 *For any instance I_T,*

$$v^{\varphi_p}(I_T)\leq v^{\phi_p}(I_T)$$

Theorem 7 *For $p\in\{2,\ldots,T\}$, $\rho^{\varphi_p}=p^2/(p-1)$.*

Proof: Let a worst-case instance under policy ϕ_p be $I_T=(a,\ldots,a,(T-k+1)a,0,\ldots,0)$ for some $a>0$ and some $k\in\{1,\ldots,T-p+1\}$. Then,

$$r_t^{\phi_p}-r_{t-1}^{\phi_p}=\begin{cases}\frac{\Gamma(T-p+2)\Gamma(T-t+1)}{\Gamma(T-t-p+3)\Gamma(T+1)}pa, & t\in\{2,\ldots,k-1\}\\ \frac{\Gamma(T-p+2)\Gamma(T-t+1)}{\Gamma(T-t-p+3)\Gamma(T+1)}ap+\frac{T-t}{T-t+1}pa, & t=k\end{cases}$$

which is positive for $t\in\{2,\ldots,k\}$. Thus $r_1^{\phi_p}<r_2^{\phi_p}<\ldots<r_k^{\phi_p}$, which implies that $r_t^{\varphi_p}=r_t^{\phi_p}$ for all $t=1,\ldots,k$. □

4.2 ψ-policies

The class of ψ-policies divide the time horizon in powers of 2 and schedule work to be completed based on arrival times only. Without loss of generality, let $T = 2^K - 1, K \in \mathbb{Z}_+$. If not, we make the instance longer by letting $K = \lceil \log(T + 1) \rceil$ and setting arrivals in the initial $2^K - 1 - T$ time periods to zero.

Policy ψ_1 *Allocate resources uniformly to complete work arriving in periods* $\{2^K - 2^i + 1, \ldots, 2^K - 2^{i-1}\}$ *in periods* $\{2^K - 2^{i-1}, \ldots, 2^K - 2^{i-2} - 1\}$*,* $i = K, \ldots, 2$*. Hence,*

$$r_k = \begin{cases} 0, & k \in \{1, \ldots, 2^{K-1} - 1\}, \\ \frac{1}{2^{n-2}} \sum_{i=2^K-2^n+1}^{2^K-2^{n-1}} a_i, & \\ \quad k \in \{2^K - 2^{n-1}, \ldots, 2^K - 2^{n-2} - 1\}, & \\ \quad n \in \{K, \ldots, 2\}, & \\ a_k, & k = 2^K - 1. \end{cases} \tag{10}$$

Theorem 8 *For policy* ψ_1*, the competitive ratio* $\rho^{\psi_1} = 4$.

Proof: Consider any instance I_T. A lower bound on the optimal value with perfect information is

$$\begin{aligned} v^*(I_T) &\geq \frac{1}{2^n - 1} \sum_{i=2^K-2^n+1}^{2^K-1} a_i \\ &\geq \frac{1}{2^n - 1} \sum_{i=2^K-2^n+1}^{2^K-2^{n-1}} a_i, \quad \forall\, n \in \{1, \ldots, K\}. \end{aligned} \tag{11}$$

Dividing (10) by (11), a bound on the competitive ratio $\rho_T^{\psi_1}$ is obtained.

$$\rho_T^{\psi_1} \leq \max_{n \in \{2, \ldots, K\}} \frac{2^n - 1}{2^{n-2}} \leq 4.$$

Consider an instance I'_T with $a_1 = A$, $a_i = 0, i = 2, \ldots, 2^K - 1$.

$$\begin{aligned} v^{\psi_1}(I'_T) = \frac{A}{2^{K-2}}, \ v^*(I'_T) &= \frac{A}{2^K - 1} \\ \Rightarrow \frac{v^{\psi_1}(I'_T)}{v^*(I'_T)} &= 4 - \frac{1}{2^{K-2}}. \end{aligned}$$

Thus $v^{\psi_1}(I'_T)/v^*(I'_T) \to 4$ as $K \to \infty$. Hence, $\rho^{\psi_1} = 4$. □

A drawback with policy ψ_1 is that it waits for almost half the time horizon before assigning resources to complete work. So, by starting to process the work earlier and matching the completion times of policy ψ_1, we should get a policy with competitive ratio that is not worse than that of policy ψ_1.

Policy ψ_2 *Allocate resources uniformly to complete work arriving in periods* $\{2^K - 2^i + 1, \ldots, 2^K - 2^{i-1} - 2^{i-2}\}$ *and* $\{2^K - 2^{i-1} - 2^{i-2} + 1, \ldots, 2^K - 2^{i-1}\}$ *in periods* $\{2^K - 2^{i-1} - 2^{i-2}, \ldots, 2^K - 2^{i-2} - 1\}$ *and* $\{2^K - 2^{i-1}, \ldots, 2^K - 2^{i-2} - 1\}$*, respectively, for* $i = K, \ldots, 2$.

Theorem 9 *For policy* ψ_2*, the competitive ratio* $\rho^{\psi_2} = 3.5$.

Proof: Consider any instance I_T. A lower bound on the optimal value with perfect information is

$$v^*(I_T) \geq \max \left\{ \frac{1}{2^{n+1} - 1} \sum_{i=2^K-2^{n+1}+1}^{2^K-2^{n-1}-2^{n-2}} a_i, \ \frac{1}{2^n + 2^{n-1} - 1} \sum_{i=2^K-2^n-2^{n-1}+1}^{2^K-2^{n-1}-2^{n-2}} a_i \right\}, \quad \forall\, n \in \{2, \ldots, K - 1\}. \tag{12}$$

There are two cases, depending on which term in (12) is larger. Algebraic manipulation on the cases results in $\rho_T^{\psi_2} \leq 3.5$. An instance I'_T with $a_1 = A$, $a_{2^{K-2}+1} = 3A$, and remaining $a_i = 0$, has $v^{\psi_2}(I'_T)/v^*(I'_T) = 3.5(1 - 1/(3 \cdot 2^{K-2}))$. Hence, $\rho^{\psi_2} = 3.5$. □

Based on the improvement in competitive ratio achieved with policy ψ_2 over policy ψ_1, we can carry the argument further and begin processing work even earlier than in policy ψ_2.

Policy ψ_3 *Allocate resources uniformly to complete work arriving in periods* $\{2^K - 2^i + 1, \ldots, 2^K - 2^{i-1} - 2^{i-2} - 2^{i-3}\}$*,* $\{2^K - 2^{i-1} - 2^{i-2} - 2^{i-3} + 1, \ldots, 2^K - 2^{i-1} - 2^{i-2}\}$ *and* $\{2^K - 2^{i-1} - 2^{i-2} + 1, \ldots, 2^K - 2^{i-1}\}$ *in periods* $\{2^K - 2^{i-1} - 2^{i-2} - 2^{i-3}, \ldots, 2^K - 2^{i-2} - 1\}$*,* $\{2^K - 2^{i-1} - 2^{i-2}, \ldots, 2^K - 2^{i-2} - 1\}$ *and* $\{2^K - 2^{i-1}, \ldots, 2^K - 2^{i-2} - 1\}$*, respectively, for* $i = K, \ldots, 3$*. Additionally,* a_{2^K-3} *is completed in periods* $\{2^K - 3, 2^K - 2\}$*,* a_{2^K-2} *in period* $2^K - 2$*, and* a_{2^K-1} *in period* $2^K - 1$.

Theorem 10 *For policy ψ_3, the competitive ratio $\rho^{\psi_3} = 3.45$.*

Proof: The proof is similar to the proof of Theorem 9. We now have three cases which need to be considered for $v^*(I_T)$, which results in $\rho_T^{\psi_3} \leq 3.45$. An instance I'_T with $a_1 = A$, $a_{2^{K-3}+1} = A$ and $a_{2^{K-2}+1} = 6A$, and remaining $a_i = 0$, gives the required competitive ratio. □

The competitive ratio for policies which allow resources to be allocated even earlier than policy ψ_3 may be better than ρ^{ψ_3}. However, the analysis for that is tedious.

5 Lower Bounds

In view of Theorem 1, to obtain a lower bound L on ρ^*, it suffices to find an instance for which the α-policy with $\alpha = L$ is infeasible. To find such instances, we develop a continuous time version of the online resource minimization problem.

Let time t range from 0 to 1. Let work arrive at rate $a(t) \geq 0$. So $A(t) \equiv \int_0^t a(\tau)\,d\tau$ is the total amount of work that has arrived by time t. Let $r(t)$ be the rate of resources allocated at time t. Under the α-policy

$$r(t) = \alpha \sup_{0 \leq \tau \leq t} \frac{1}{1-\tau} \int_\tau^t a(x)dx. \qquad (13)$$

Under appropriate smoothness conditions the supremum occurs when the derivative of its argument is zero.

$$\frac{-(1-\tau)a(\tau) + (A(t) - A(\tau))}{(1-\tau)^2} = 0$$

$$\Rightarrow A(t) = A(\tau) + (1-\tau)a(\tau) \qquad (14)$$

Implicit differentiation of (14) gives

$$a(t)dt = (1-\tau)a'(\tau)d\tau \qquad (15)$$

When τ satisfies (13),

$$r(t) = \frac{\alpha}{1-\tau}[A(t) - A(\tau)] = \alpha a(\tau). \qquad (16)$$

Assuming that the resources are not starved for work, $\int_0^1 a(t)dt = \int_0^1 r(t)dt$. Combining with (16)

$$\int_0^1 a(t)dt = \alpha \int_0^1 a(\tau)dt \qquad (17)$$

The variable of integration in the right side of (17) can be changed from t to τ using (14) and (15)

$$\int_0^1 a(t)dt = \alpha \int_0^\gamma \frac{a(\tau)(1-\tau)a'(\tau)}{a(f(\tau))}d\tau$$

where γ is the solution to $A(1) = A(\gamma) + (1-\gamma)a(\gamma)$, and $t = f(\tau)$ is a solution of (14). Hence, the α-policy will fail if

$$\alpha < \frac{\int_0^1 a(t)dt}{\int_0^\gamma \frac{a(\tau)(1-\tau)a'(\tau)}{a(f(\tau))}d\tau}. \qquad (18)$$

For example, for a linear arrival rate, $a(t) = t$ and $A(t) = t^2/2$. Hence, $a'(\tau) = 1$, $f(\tau) = \sqrt{2\tau - \tau^2} = a(f(\tau))$, $\gamma = 1$, and α must not be less than $0.5/(1-\pi/4) \approx 2.3298$. A discrete instance with a linear arrival rate and 7500 periods requires $\alpha > 2.329$. When the arrival rate is quadratic, $\alpha \geq 2.392$, while a discrete instance with 10,000 periods requires $\alpha > 2.391$.

Since it appears that growing arrival rates give rise to worse instances, we tried various other functional forms, such as $a(t) = \exp(5t)$. Unfortunately, such arrival rates starve the resources and (18) becomes incorrect. However, we were able to construct a discrete instance with 10,000 periods that shows that $\rho^* > 2.51$.

6 Computational Experiments

The competitive ratios for φ- and ψ-policies computed in Section 4 may have little or no bearing on how they perform on average. Also, the policies analyzed considered a single deadline. In this section, we present multiple deadline versions of the φ- and ψ-policies and then conduct computational experiments to gain insight in the average performance of these policies, as well as that of the α-policy.

The φ-policy for the multiple deadline version of the problem allocates resources as

$$r_t^{\varphi_p} = \min\left\{ \max\left\{ m_t, \max_{t \leq j \leq T} \frac{p}{j-t+1} \sum_{k=t}^{j} w_{tk} \right\}, \sum_{k=t}^{T} w_{tk} \right\},$$

with work being processed in EDD order. The ψ-policy allocates resources for the multiple deadline problem by applying policy ψ_q on a_{tj} and considering a time horizon from t to j, for $t \leq j \leq T$.

Pol	A			B			C			D		
	(I)	(II)	(III)	(I)	(II)	(III)	(I)	(II)	(III)	(I)	(II)	(III)
$\alpha_{2.75}$	0	2.337	2.330	0	2.338	2.330	0	2.387	2.331	0	2.365	2.330
α_3	0	2.452	2.444	0	2.453	2.444	0	2.490	2.438	0	2.472	2.443
$\alpha_{3.25}$	0	2.541	2.532	0	2.542	2.532	0	2.581	2.527	0	2.562	2.531
$\alpha_{3.5}$	0	2.601	2.590	0	2.600	2.590	0	2.675	2.592	0	2.649	2.590
φ_2	99.99	1.480	1.469	100	1.477	1.469	87.45	1.524	1.473	95.95	1.498	1.471
$\varphi_{2.5}$	0.01	1.491	1.481	0	1.490	1.481	12.05	1.533	1.485	4.03	1.509	1.483
φ_3	0	1.514	1.501	0	1.510	1.501	0.49	1.559	1.506	0.02	1.536	1.503
$\varphi_{3.5}$	0	1.537	1.522	0	1.532	1.522	0	1.577	1.527	0	1.557	1.524
φ_4	0	1.558	1.543	0	1.553	1.542	0.01	1.615	1.548	0	1.581	1.545
ψ_2	0	1.630	1.578	0	1.598	1.576	0	1.719	1.616	0	1.663	1.594
ψ_1	0	1.721	1.659	0	1.683	1.658	0	1.833	1.699	0	1.753	1.677

Table 1: Values of (I), (II) and (III) over 10,000 instances with $T = 75$ generated using characteristics A, B, C and D of arrival distribution.

A truncated and discretized Normal distribution was used to generate the arrivals with different deadlines in each period. To assess the impact of the mean and coefficient of variation, four classes of instances were generated with combinations of low and high means and coefficients of variation. The values of (μ, σ^2, cv) are A $(1.01, 0.01, 0.11)$, B $(10.5, 1.08, 0.1)$, C $(9.58, 26.20, 0.53)$ and D $(1.54, 0.25, 0.32)$. Ten thousand instances were generated for each class of instances, and planning horizon $T \in \{10, 15, 25, 50, 75, 100\}$. The policies considered were α-policy with $\alpha \in \{2.75, 3, 3.25, 3.5\}$, φ-policy with $p \in \{2, 2.5, 3, 3.5, 4\}$, and ψ policy with $q \in \{1, 2\}$. Values of the parameter for the α-policy were chosen based on lower and upper bounds on ρ^* $(2.51 < \rho^* \leq 3.45)$ computed earlier.

The measures for comparison were (I) the percentage of instances where a particular policy performed better than all other policies, (II) the maximum performance ratio v^π/v^* of a policy, and (III) the average performance ratio for a policy. Table 1 shows values of (I), (II) and (III) for 10,000 instances with $T = 75$ generated in classes A, B, C and D. The values obtained are representative of the results obtained for other values of T.

The results indicate that policy φ_2 performs better than the other policies for most of the instances. Based on the values of α that were tried, it does not seem as though the α-policy performs well in practice. The essential difference between the α-, φ-, and ψ-policies is that the α- and ψ-policies do not take the amount of remaining work into account when making decisions, whereas the φ-policies do, and this seems to be important.

References

Albers, S. 1997. Competitive Online Algorithms. *Optima: Mathematical Programming Society Newsletter*, **54**, 1–7.

Ascheuer, N., Grötschel, M., Rambau, J. and Kamin, N. 1998. Combinatorial Online Optimization in Practice. *Optima: Mathematical Programming Society Newsletter*, **57**, 1–6.

Coffman Jr., E. G., Garey, M. R. and Johnson, D. S. 1996. Approximation Algorithms for Bin Packing: A Survey, *Technical report*, AT&T Bell Laboratories, Murray Hill, NJ.

Sgall, J. 1997. Online Scheduling, *Technical report*, Mathematical Institute, AV CR, Zitna, Praha.