OPTIMAL ONLINE ALGORITHMS FOR MINIMAX RESOURCE SCHEDULING[†]

BRADY HUNSAKER[‡], ANTON J. KLEYWEGT[‡ §], MARTIN W. P. SAVELSBERGH[‡], AND CRAIG A. TOVEY[‡]

Abstract. We consider a very general online scheduling problem with an objective to minimize the maximum level of resource allocated. We find a simple characterization of an optimal deterministic online algorithm. We develop further results for two more specific problems, single resource scheduling and hierarchical line balancing. We determine how to compute optimal online algorithms for both problems using linear programming and integer programming, respectively. We show that randomized algorithms can outperform deterministic algorithms, but only if the amount of work done is a non-concave function of resource allocation.

Keywords: Online algorithms, competitive analysis, worst-case analysis, single-machine scheduling, multiprocessor scheduling, line balancing

1. Introduction. Consider the following problem: work with different deadlines arrives over time and has to be performed using a resource. The quantities of work that arrive as well as their deadlines only become known at the times of arrival. At a given set of time points, the decision maker decides how much resource to allocate and which of the available work to perform at that time. The objective is to minimize the maximum amount of resource allocated at any time during the planning period. This problem is called the *online resource minimization problem* (ORMP). It occurs in production scheduling settings in which the major cost component is energy consumption (Kleywegt et al. [11]). The decision maker would like to spread the workload out as evenly as possible over time, but faces the dilemma of uncertainty about future work. That is, the decision maker must make a trade-off between allocating too much resource early, and postponing too much work to be completed with work that arrives later.

Consider another problem: work with different requirements arrives over time and has to be assigned to a collection of machines with different capabilities. The machines form a linear hierarchy based on their capabilities, i.e., machine j has at least the same capabilities as machine $j-1$. The amount of work that arrives as well as the required machine capabilities only become known at the time of arrival. At a given set of decision points, the decision maker decides how to assign the work that has arrived since the previous decision point to the machines. The objective is to minimize the maximum amount of work assigned to any machine. This problem is called the *hierarchical line balancing problem* (HLBP).

In both the ORMP and the HLBP, the quality of an algorithm is evaluated by its competitive ratio, i.e., the worst-case ratio over all possible instances of the value of the solution produced by the algorithm and the value of the optimal solution with perfect information.

In this paper, we introduce a simple parameterized deterministic algorithm, called the α-policy, with parameter α and competitive ratio α, provided it produces a feasible solution. We show that with an appropriate choice of parameter α, the α-policy

[†]Some of these results were reported in Kleywegt, A.J., Nori, V.S., Savelsbergh, M.W.P., and Tovey, C.A., "Online Resource Minimization," *Proceedings of the Tenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA '99)*, Baltimore, MD, January 1999.

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

[§]Supported by the National Science Foundation under grant DMI-9875400.

has as good a competitive ratio as any other deterministic algorithm. Under a convexity assumption, which holds for both the ORMP and the HLBP, the α-policy is also optimal among all randomized algorithms. However, we show that randomized algorithms can outperform deterministic algorithms in other cases.

We also investigate how the optimal parameter can be computed for the ORMP and the HLBP. For the ORMP, we construct linear programs to compute the optimal parameters, which depend on the number of time periods. For the HLBP, the optimal parameters are computed with linear programming if the number of time periods is greater than or equal to the number of machines, and with integer programming for the remaining cases. Interestingly, the resulting parameter values (and hence the optimal competitive ratios) for realistic finite numbers of machines and time periods are substantially lower than the asymptotic values.

The ORMP and the HLBP are special cases of a more general *online min-max problem* (OMMP) and several of our results can be extended to the OMMP.

The ORMP appears to be a new problem, as we have not been able to find any existing literature discussing the problem. On the other hand, the HLBP is a known problem and has been studied, for example, by Bar-Noy et al. [6]. Bar-Noy et al. [6] distinguish a fractional variant and an integral variant. In the integral variant work must be assigned in its entirety to a machine; in the fractional variant work may be split among eligible machines. Our results are for the fractional variant. Bar-Noy et al. [6] give an algorithm with asymptotically optimal competitive ratio e, in the limiting case where the number of machines goes to infinity. When all machines have the same capabilities, i.e., there is no linear hierarchy among machines, the fractional variant of the HLBP is trivial: divide work that arrives equally over the machines. On the other hand, when all machines have the same capabilities, the integral variant of the HLBP remains difficult as it is equivalent to minimizing the makespan on a set of parallel identical machines. Already in 1966, Graham [8] presented a greedy algorithm with a competitive ratio of $2-1/m$ for online makespan minimization on m identical parallel machines. This problem has continued to attract researchers, see for example the recent papers by Albers [1], Bartal et al. [17], Fleischer and Wahl [15], and Seiden [16]. Fleischer and Wahl [15] present the current best deterministic online algorithm and Albers [1] presents the current best randomized algorithm. In Aspnes et al. [2] an 8-competitive algorithm is given for online makespan minimization on related parallel machines, i.e., where the processing requirement of a job is not only determined by the length of the job, but also by the speed of the machine. This was improved by Berman, Charikar, and Karpinski [7]. Azar, Naor, and Rom [5] consider a more general online load balancing problem in which each job can be handled only by a subset of machines and requires a different level of service. In load balancing problems a distinction is made between permanent and temporary jobs. Permanent jobs continue forever after they arrive and load the machine indefinitely, while temporary jobs load the machine only during the interval they are active. Azar, Broder and Karlin [3] and Azar et al. [4] extend the work mentioned above to handle temporary jobs. Finally, we want to mention the work by Hoogeveen and Vestjens [9]. They consider the problem of minimizing the maximum delivery time on a single machine and present an optimal deterministic online algorithm. This is one of the few cases that we are aware of in which an optimal online algorithm is presented. Note that our α-policy, with an appropriate choice of α is also an optimal policy.

The paper is organized as follows. In Section 2, we define the online min-max problem, the online resource minimization problem, and the hierarchical line balancing problem. In Section 3, we introduce an optimal online algorithm called the α-policy. In Section 4, we establish optimality results for the online resource minimization problem, the hierarchical line balancing problem, and we generalize these results for the online min-max problem. In Section 5, we show that randomized algorithms can outperform deterministic algorithms, but only if the amount of work done is a non-concave function of resource allocation. Finally, in Section 6, we point out future research directions.

2. Problem Definition. In this section we define the general problem called the *Online Min-Max Problem* (OMMP), as well as two special cases of the OMMP, called the *Online Resource Minimization Problem* (ORMP), and the *Hierarchical Line Balancing Problem* (HLBP), respectively.

2.1. Online Min-Max Problem. DEFINITION 2.1. (OMMP): An *instance* ω of the OMMP is a finite sequence $(a(1), \ldots, a(T))$ of length T. Each $a(t)$ could be a number, vector, function, set, problem instance, or any other object, and it could be different for different t. The nature of each $a(t)$ is determined by the particular type of OMMP, as illustrated in the examples of Sections 2.2 and 2.3. The set of all instances is denoted by Ω. Often it is of interest to explore the characteristics of the OMMP as a function of problem parameters. For that purpose, the set of all instances with parameter β is denoted by Ω_β. For example, the set of all instances of length T is denoted by Ω_T. For any instance $\omega = (a(1), \ldots, a(T)) \in \Omega_\beta$, and any $t \in \{1, \ldots, T\}$, let $\omega^t \equiv (a(1), \ldots, a(t))$ denote the first t elements of instance ω. Let $\Omega_\beta^t \equiv \{\omega^t : \omega \in \Omega_\beta\}$ denote the set of all such partial instances of the first t elements. For any instance of length T, a *solution* r of the OMMP is a sequence $(r(1), \ldots, r(T)) \in \mathbb{R}_+^T$ of T nonnegative real numbers. Let $\mathbb{R}_+^\infty \equiv \bigcup_{T=1}^\infty \mathbb{R}_+^T$ denote the set of all solutions. For any instance ω, let $\mathcal{R}(\omega)$ denote the set of feasible solutions. A *deterministic algorithm* π for the OMMP is a function $\pi : \Omega \mapsto \mathbb{R}_+^\infty$, such that for any $\omega \in \Omega_T$, $\pi(\omega) \in \mathbb{R}_+^T$, i.e., if ω is of length T, then $\pi(\omega)$ is also of length T. A deterministic algorithm π is called *feasible* if, for every ω, $\pi(\omega) \in \mathcal{R}(\omega)$. Let $\mathcal{B}_+^T$ denote the Borel sets on $\mathbb{R}_+^T$, let $\mathcal{P}^T$ denote the set of probability measures on $\mathcal{B}_+^T$, and let $\mathcal{P} \equiv \bigcup_{T=1}^\infty \mathcal{P}^T$. A *randomized algorithm* π for the OMMP is a function $\pi : \Omega \mapsto \mathcal{P}$, such that for any $\omega \in \Omega_T$, $\pi(\omega) \in \mathcal{P}^T$. The probability that the solution is in a set $B \in \mathcal{B}_+^T$ is denoted by $\pi(\omega)[B]$. We assume that $\mathcal{R}(\omega) \in \mathcal{B}_+^T$ for $\omega \in \Omega_T$, and a randomized algorithm π is called *feasible* if, for every ω, $\pi(\omega)[\mathcal{R}(\omega)] = 1$. Also, for any $t \in \{1, \ldots, T\}$, we will use $\pi(\omega)(t)$ to denote the decision at time t for instance ω under algorithm π. For a deterministic algorithm π, $\pi(\omega)(t)$ is deterministic, and for a randomized algorithm π, $\pi(\omega)(t)$ is a random variable, where the tuple $(\pi(\omega)(1), \ldots, \pi(\omega)(T))$ is distributed according to the probability measure $\pi(\omega)$. When the instance ω has been fixed, we also use $r^\pi(t)$ to denote the decision $\pi(\omega)(t)$ under algorithm π at time t. A *deterministic (randomized) online algorithm* π for the OMMP is a deterministic (randomized) algorithm such that, for each ω and each t, (the probability distribution of) $\pi(\omega)(t)$ depends on ω^t only, i.e., it depends only on the history of instance ω up to time t, and not on the whole instance ω. Let Π^{DO} denote the set of all deterministic online algorithms, and let $\Pi^{RO} \supseteq \Pi^{DO}$ denote the set of all randomized online algorithms for the OMMP.

For any instance $\omega \in \Omega_T$, and any deterministic algorithm π, the *value* $v^\pi(\omega)$ is the maximum norm of $\pi(\omega)$, i.e.,

$$v^\pi(\omega) \quad \equiv \quad \max\left\{r^\pi(1), \ldots, r^\pi(T)\right\}$$

if $\pi(\omega) \in \mathcal{R}(\omega)$, and $v^\pi(\omega) = \infty$ otherwise. Similarly, for any instance $\omega \in \Omega_T$, and any randomized algorithm π, the value $v^\pi(\omega)$ is the expected maximum norm under $\pi(\omega)$, i.e.,

$$v^\pi(\omega) \quad \equiv \quad E^{\pi(\omega)}\left[\max\left\{r^\pi(1),\ldots,r^\pi(T)\right\}\right]$$

if $\pi(\omega)[\mathcal{R}(\omega)] = 1$, and $v^\pi(\omega) = \infty$ otherwise.

For any instance $\omega \in \Omega_T$, the *optimal value with perfect information*, $v^*(\omega)$, is defined by

$$v^*(\omega) \quad \equiv \quad \inf_{r\in\mathcal{R}(\omega)}\left\{\max\{r(1),\ldots,r(T)\}\right\}$$

We assume that $v^*(\omega) < \infty$ for all $\omega \in \Omega$, so $\mathcal{R}(\omega) \neq \varnothing$.

In this paper, the quality of an algorithm π for an instance ω is evaluated by the ratio of the value of the algorithm and the optimal value with perfect information, i.e., $v^\pi(\omega)/v^*(\omega)$. For any class Ω_β of instances, and any algorithm π, the competitive ratio or worst-case ratio ρ^π_β denotes the largest ratio of the value of algorithm π and the optimal value with perfect information over all instances in Ω_β, i.e.,

$$\rho^\pi_\beta \quad \equiv \quad \inf\left\{\rho \geq 1 \,:\, v^\pi(\omega) \leq \rho v^*(\omega) \ \ \forall\, \omega \in \Omega_\beta\right\} \tag{2.1}$$

The convention is that $\inf\varnothing = \infty$. Note that, if $v^*(\omega) > 0$ for all $\omega \in \Omega_\beta$, then

$$\rho^\pi_\beta \quad = \quad \sup_{\omega\in\Omega_\beta}\left\{\frac{v^\pi(\omega)}{v^*(\omega)}\right\}$$

Also note that, if $\rho^\pi_\beta < \infty$, then the infimum in (2.1) is attained, in the sense that $v^\pi(\omega) \leq \rho^\pi_\beta v^*(\omega)$ for all $\omega \in \Omega_\beta$. This criterion is standard in the literature for online algorithms, see for example McGeoch and Sleator [13] and Irani and Karlin [10]. The competitive ratio ρ^π of algorithm π over all instances is given by

$$\rho^\pi \quad \equiv \quad \inf\left\{\rho \geq 1 \,:\, v^\pi(\omega) \leq \rho v^*(\omega) \ \ \forall\, \omega \in \Omega\right\}$$

If $\mathbf{B}$ is the set of parameters β, i.e., $\Omega = \bigcup_{\beta\in\mathbf{B}}\Omega_\beta$, then $\rho^\pi = \sup_{\beta\in\mathbf{B}}\rho^\pi_\beta$.

The optimal competitive ratio ρ^*_β over all deterministic online algorithms, and over all instances in Ω_β, is given by

$$\rho^*_\beta \quad \equiv \quad \inf_{\pi\in\Pi^{DO}}\ \rho^\pi_\beta$$

The optimal competitive ratio ρ^* over all deterministic online algorithms, and over all instances, is given by

$$\rho^* \quad \equiv \quad \inf_{\pi\in\Pi^{DO}}\ \rho^\pi$$

The optimal competitive ratios over all randomized online algorithms are defined similarly.

Alternatively, one may want to define $\rho^* \equiv \sup_{\beta\in\mathbf{B}}\rho^*_\beta$. The question is whether the two definitions of ρ^* are equal, that is, whether $\sup_{\beta\in\mathbf{B}}\inf_{\pi\in\Pi^{RO}}\rho^\pi_\beta = \inf_{\pi\in\Pi^{RO}}\sup_{\beta\in\mathbf{B}}\rho^\pi_\beta$. In general, for any real valued function $f(x,y)$, it holds that $\sup_x\inf_y f(x,y) \leq \inf_y\sup_x f(x,y)$, and the inequality may be strict. Thus $\sup_{\beta\in\mathbf{B}}\inf_{\pi\in\Pi^{RO}}\rho^\pi_\beta \leq$

$\inf_{\pi \in \Pi^{RO}} \sup_{\beta \in \mathbf{B}} \rho^{\pi}_{\beta}$. Lemma 2.2 establishes that if the parameter β is known beforehand by the decision maker, then the two definitions of ρ^* are in fact equal. If the parameter β is not known beforehand, then the inequality may be strict, because for each β, there may be an algorithm π_β that performs particularly well for instances in Ω_β, but there may not exist a single algorithm π, that does not depend on β (because β is not known beforehand), that performs well for all β. For the special cases of the OMMP and the choices of parameter β considered in this paper, β is known beforehand.

LEMMA 2.2. *If β is known beforehand by the decision maker, then*

$$\rho^* \quad = \quad \sup_{\beta \in \mathbf{B}} \rho^*_\beta$$

The proof of Lemma 2.2 is given in the Appendix.

An algorithm $\pi^* \in \Pi^{DO}$ is called *optimal* over deterministic online algorithms if $\rho^{\pi^*}_\beta = \rho^*_\beta$ for all $\beta \in \mathbf{B}$.

For any $r_1, r_2 \in \mathbb{R}^T$, we denote $r_1 \leq r_2$ if $r_1(t) \leq r_2(t)$ for all t. Since the objective is to minimize $\max\{r(1), \ldots, r(T)\}$, it is natural to assume that increasing r does not adversely impact the feasibility of the solution, so that there is a trade-off between smaller values of r for improving the objective value, and larger values of r for improving the feasibility. It is therefore assumed that $\mathcal{R}$ has the *feasibility monotonicity property*, that is, for any $\omega \in \Omega$ and any $r_1 \in \mathcal{R}(\omega)$, it holds that $r_2 \in \mathcal{R}(\omega)$ for all $r_2 \geq r_1$.

2.2. Online Resource Minimization Problem. In this section we define the *Online Resource Minimization Problem* (ORMP), and show that the ORMP is a special case of the OMMP.

DEFINITION 2.3. (ORMP): Work with different deadlines arrives over time and has to be performed using a costly resource. The amount of work that arrives at each point in time as well as their deadlines only become known at the time of arrival. At a given set of decision points, indexed with $t = 1, \ldots, T$, the decision maker decides how much resource to allocate and which of the available work to perform at that time. The objective is to minimize the maximum amount of resource allocated at any time during the planning period. Let $a_u(t) \in \mathbb{R}_+$ denote the amount of work that arrives at time t with deadline u, with $t, u \in \{1, \ldots, T\}$, and let $a(t) \equiv (a_1(t), \ldots, a_T(t))$. Assume that $a_u(t) = 0$ for $u < t$, i.e., work does not arrive after its deadline. Let $r(t) \in \mathbb{R}_+$ denote the total amount of resource allocated at decision point t, and let $r \equiv (r(1), \ldots, r(T))$. Let $q_u(t) \in \mathbb{R}_+$ denote the amount of work with deadline u that is performed at time t, and let $q(t) \equiv (q_1(t), \ldots, q_T(t))$, and $q = (q(1), \ldots, q(T))$. Thus $(r(t), q(t))$ denotes the decision made at time t. For (r, q) to be feasible, $(r(t), q(t))$ must satisfy the following for all t.

$$\sum_{u=1}^{T} q_u(t) \leq r(t) \tag{2.2}$$

$$q_u(t) = 0 \qquad \text{for all } u < t \tag{2.3}$$

$$\sum_{t'=1}^{t} q_t(t') = \sum_{t'=1}^{t} a_t(t') \tag{2.4}$$

$$\sum_{t'=1}^{t} q_u(t') \leq \sum_{t'=1}^{t} a_u(t') \qquad \text{for all } u \tag{2.5}$$

Constraint (2.2) states that the total amount of work performed at time t cannot exceed the amount of work that can be accomplished with $r(t)$ amount of resource. Constraint (2.3) states that no work can be performed after its deadline. Constraint (2.4) states that all work must be performed by the respective deadlines. Constraint (2.5) states that work cannot be performed before it has arrived.

A more general version of the ORMP stated above incorporates a productivity function $\eta_t(r)$ that represents the amount of work that can be performed at time t with r amount of resource. Thus the ORMP stated above has productivity function $\eta_t(r) = r$ for all t and r. The ORMP with productivity function $\eta_t(r)$ is the same as the ORMP stated above, except that constraint (2.2) is replaced by

$$\sum_{u=1}^{T} q_u(t) \quad \leq \quad \eta_t(r(t)) \tag{2.6}$$

It is clear that as long as η_t is nondecreasing for all t, $\mathcal{R}$ has the feasibility monotonicity property. It will be stated clearly which results hold for the ORMP with productivity function $\eta_t(r) = r$, and which results hold for more general productivity functions.

The decision maker has to make a sequence of decisions $(r(t), q(t))$ over time, using information as it becomes available. Because the decision maker has no information about future arrivals, decision $(r(t), q_t(t), \ldots, q_T(t))$ can depend on past arrivals only, and not on any arrivals after time t. Thus algorithms are required to be online. Because the objective is to minimize the maximum amount of resource allocated at any time during the planning period, the algorithm evaluation criteria of the ORMP are the same as those of the OMMP. Thus it seems that the ORMP fits into the framework of the OMMP, except that the OMMP includes only a single decision $r(t) \in \mathbb{R}_+$ at each time t, whereas the ORMP includes both resource quantity decision $r(t) \in \mathbb{R}_+$ as well as resource allocation decision $q(t) \in \mathbb{R}_+^T$ at each time t. However, it is clear that one should give preference to available work with earlier deadlines above work with later deadlines in the allocation q of the chosen amounts of resource r. This decision rule for the allocation of the chosen amounts of resource r is called the *earliest deadline first* rule (EDF). Specifically, the EDF rule works as follows. For any time t and any chosen amount of resource $r(t)$, $q_u(t)$ is determined inductively by $q_u(t) = 0$ for all $u < t$, and

$$q_u(t) \quad = \quad \min\left\{\sum_{\tau=1}^{t} a_u(\tau) - \sum_{\tau=1}^{t-1} q_u(\tau),\ \eta_t(r(t)) - \sum_{v=t}^{u-1} q_v(t)\right\} \tag{2.7}$$

for all $u \geq t$. It is easy to see that for any instance ω and any feasible solution (r, q), the solution (r, q'), where q' denotes the resource allocation decisions according to the EDF rule, is both feasible and has the same objective value as solution (r, q). Because EDF performs at least as well as any other allocation rule, attention is restricted to algorithms that use the EDF rule. Thus a solution is specified by r only, and the ORMP is a special case of the OMMP.

It follows from the definition of the EDF rule that constraints (2.6) (or (2.2)), (2.3), and (2.5) cannot be violated by the EDF rule. Thus the only constraint that can be violated by the EDF rule is (2.4), that is, the algorithm can fail to perform all work by the deadlines, in which case the algorithm is infeasible.

The problem parameter β of interest for the ORMP is the length T of the time horizon. Note that ρ^*_T is nondecreasing in T, because for any instance $\omega_T = (a(1), \ldots, a(T))$ of length T there is an instance $\omega_{T+1} = (0, a(1), \ldots, a(T))$ of length $T + 1$, such that the optimal value with perfect information is the same for both instances, $v^*(\omega_T) = v^*(\omega_{T+1})$, and for any feasible solution $r = (r(1), \ldots, r(T+1))$ for ω_{T+1}, the solution $r' = (r'(1), \ldots, r'(T))$ for ω_T with $r'(t) = r(t+1)$ for all $t \in \{1, \ldots, T\}$ is feasible, and $\max\{r'(1), \ldots, r'(T)\} \leq \max\{r(1), \ldots, r(T+1)\}$.

The version of the ORMP in which all deadlines are equal to the planning horizon T is called the single deadline ORMP, and the version with different deadlines is called the multiple deadline ORMP.

2.3. Hierarchical Line Balancing Problem. In this section we define the *Hierarchical Line Balancing Problem* (HLBP), which is another interesting problem that is a special case of the OMMP.

Definition 2.4. (HLBP): Work arrives over time and has to be performed using a set of m machines. The machines can be ordered in a linear hierarchy, that is, the machines can be indexed with the integers 1 through m, so that machine $j \in \{1, \ldots, m\}$ is at least as versatile as any machine $i \in \{1, \ldots, j\}$. Each quantity of work has a specification of the least versatile machine on which the work can be performed. That is, if a quantity of work requires at least machine i for its completion, then the work can be assigned to any one or more than one of machines $i, \ldots, m$. The quantities of work that arrive as well as the specifications of their least versatile machines only become known at the time of arrival. At a given set of decision points, indexed with $t = 1, \ldots, T$, the decision maker decides how to assign the work that has arrived since the previous decision point to the eligible machines. The objective is to minimize the maximum amount of work assigned to any machine. Let $a_i(t) \in \mathbb{R}_+$ denote the amount of work that arrives at time t that requires at least machine i, with $i \in \{1, \ldots, m\}$ and $t \in \{1, \ldots, T\}$, and let $a(t) \equiv (a_1(t), \ldots, a_m(t))$. Let $q_i(t) \in \mathbb{R}_+$ denote the amount of work assigned to machine i at time t, and let $q(t) \equiv (q_1(t), \ldots, q_m(t))$, and $q = (q(1), \ldots, q(T))$. Thus $q(t)$ denotes the decision made at time t. For q to be feasible, $q(t)$ must satisfy the following for all t.

$$\sum_{j=i}^{m} q_j(t) \quad \geq \quad \sum_{j=i}^{m} a_j(t) \qquad \text{for all } i \tag{2.8}$$

Let

$$r_i(t) \quad \equiv \quad \sum_{\tau=1}^{t} q_i(\tau) \tag{2.9}$$

denote the total amount of work assigned to machine i up to time t. Let

$$r(t) \quad \equiv \quad \max\{r_1(t), \ldots, r_m(t)\} \tag{2.10}$$

denote the maximum amount of work assigned to any machine up to time t, and let $r \equiv (r(1), \ldots, r(T))$. Clearly, $\max\{r(1), \ldots, r(T)\} = r(T)$.

As in the OMMP, algorithms are required to be online, because the decision maker has no information about future arrivals. Because the objective is to minimize the maximum amount of work assigned to any machine, $\max\{r_1(T), \ldots, r_m(T)\} \equiv r(T) = \max\{r(1), \ldots, r(T)\}$, the algorithm evaluation criteria of the HLBP are the same as those of the OMMP. Similar to the ORMP, it seems that the HLBP fits into the framework of the OMMP, except that the OMMP describes a decision $r(t) \in \mathbb{R}_+$ at each time t, whereas the HLBP describes a work assignment decision $q(t) \in \mathbb{R}_+^m$ at each time t. However, without loss of optimality one can obtain a decision q from r as follows. Intuitively it is clear that for any chosen value of $r(t)$, the maximum amount of work assigned to any machine up to time t, one should assign work to the least versatile machine that qualifies for that work and which has less than amount $r(t)$ of work assigned to it, until an amount $r(t)$ of work has been assigned to that machine or all the work has been assigned, and if any work remains, one continues to assign it in this fashion. This decision rule for the assignment of work is called the *least versatile first* rule (LVF). Specifically, the LVF rule works as follows. For any time t and any chosen value of $r(t)$, $q_i(t)$ is determined inductively by

$$q_i(t) \quad \equiv \quad \min\left\{\sum_{j=1}^{i} a_j(t) - \sum_{j=1}^{i-1} q_j(t),\ r(t) - \sum_{\tau=1}^{t-1} q_i(\tau)\right\} \tag{2.11}$$

for all $i = 1, \ldots, m$. It is easy to see that for any instance ω and any feasible solution q, the solution q' obtained by choosing $r_i(t) \equiv \sum_{\tau=1}^{t} q_i(\tau)$, and then determining q' according to the LVF rule, is both feasible and has as good an objective value as solution q. Thus, for any given r, the LVF rule performs at least as well as any other work assignment rule. Therefore, attention is restricted to algorithms that use the LVF rule, so that a solution is specified by r only. With the LVF rule, it is easy to see that $\mathcal{R}$ has the feasibility monotonicity property. Thus we have established that the HLBP is a special case of the OMMP.

Although $r(t)$ is treated as a decision from here on, it is clear that one can assume without loss of optimality that $r(t)$ satisfies (2.10). It follows from the definition of the LVF rule that no work is assigned to a machine for which the machine is not eligible, and that $\sum_{\tau=1}^{t} q_i(\tau) \leq r(t)$ for all i and all t. Thus a solution r, with q determined by the LVF rule, is feasible if and only if

$$\sum_{j=1}^{m} q_j(t) \quad = \quad \sum_{j=1}^{m} a_j(t) \tag{2.12}$$

for all t, that is, the LVF rule assigns all the work. Thus the system (2.8) of Tm constraints can be replaced by the system (2.12) of T constraints.

The problem parameter β of interest for the HLBP is the length T of the time horizon, as well as the number m of machines, thus $\beta = (T, m)$. Note that $\rho^*_{T,m}$ is nondecreasing in T and in m, because for any instance $\omega_{T,m} = (a(1), \ldots, a(T))$ of length T with m machines, there is an instance $\omega_{T+1,m} = (0, a(1), \ldots, a(T))$ of length $T+1$ with m machines, and an instance $\omega_{T,m+1} = (a'(1), \ldots, a'(T))$ of length T with $m+1$ machines, with $a'_1(t) = 0$ and $a'_{i+1}(t) = a_i(t)$ for all $i \in \{1, \ldots, m\}$ and all $t \in \{1, \ldots, T\}$, such that $v^*(\omega_{T,m}) = v^*(\omega_{T+1,m}) = v^*(\omega_{T,m+1})$, and for any feasible solution for $\omega_{T+1,m}$ or $\omega_{T,m+1}$, there is a feasible solution for $\omega_{T,m}$ with at least as good an objective value.

3. An Optimal Algorithm. In this section we introduce a simple parameterized algorithm, called the α-policy, with parameter α_β and competitive ratio α_β, provided that it is feasible. The intuition behind the α-policy is as follows. Suppose that at time t, the minimum value over all instances $\omega \in \Omega_\beta$ that start with the part ω^t of the instance observed so far, of the optimal value with perfect information, is $v_\beta(\omega^t)$. (It is easy to compute $v_\beta(\omega^t)$ for both the ORMP and the HLBP, as is shown in Sections 4.1 and 4.2.) If one wants to choose $r(t)$ in such a way that it is guaranteed that the eventual objective value will not exceed the optimal value with perfect information by more than a factor of α, then one must choose $r(t) \leq \alpha v_\beta(\omega^t)$. However, if one chooses α or $r(t)$ too small, the resulting solution may not be feasible.

We show that with appropriate choice of parameters α_β, the α-policy has as good a competitive ratio as any other deterministic algorithm, and that under mild conditions an optimal parameter value exists. Hence, the α-policy is optimal.

For any partial instance $\omega^t \in \Omega_\beta^t$, let $\Omega_\beta(\omega^t) \equiv \{\tilde{\omega} \in \Omega_\beta : \tilde{\omega}^t = \omega^t\}$ denote the set of all instances in Ω_β with first t elements equal to ω^t. Let

$$v_\beta(\omega^t) \quad \equiv \quad \inf_{\tilde{\omega} \in \Omega_\beta(\omega^t)} v^*(\tilde{\omega})$$

denote the best value with perfect information, over all instances in Ω_β that start with ω^t.

Definition 3.1. (α-policy): An α-policy is an algorithm $\pi_\alpha \in \Pi^{DO}$, with parameters $\alpha_\beta \in [1, \infty)$ for each $\beta \in \mathbf{B}$, such that for any instance $\omega \in \Omega_\beta$ and any t,

$$\pi_\alpha(\omega)(t) \quad \equiv \quad \alpha_\beta v_\beta(\omega^t)$$

Recall that, for any instance $\omega \in \Omega_\beta$ and any t, $v_\beta(\omega^t) \leq v^*(\omega)$. Thus, for any instance $\omega \in \Omega_\beta$ and any t, $\pi_\alpha(\omega)(t) \leq \alpha_\beta v^*(\omega)$, and hence

$$v^{\pi_\alpha}(\omega) \quad \equiv \quad \max\Big\{\pi_\alpha(\omega)(1), \ldots, \pi_\alpha(\omega)(T)\Big\} \quad \leq \quad \alpha_\beta v^*(\omega)$$

Therefore, if $\pi_\alpha(\omega)$ is feasible for all $\omega \in \Omega_\beta$, then $\rho_\beta^{\pi_\alpha} \leq \alpha_\beta$. It follows from the definition of $v_\beta(\omega^t)$ that if $v^*(\omega) > 0$ and $\pi_\alpha(\omega)$ is feasible for all $\omega \in \Omega_\beta$, then this bound is tight, and thus $\rho_\beta^{\pi_\alpha} = \alpha_\beta$.

Note that the feasibility monotonicity property implies that if $\pi_{\alpha_1}(\omega) \in \mathcal{R}(\omega)$ for some $\omega \in \Omega$ and some $\alpha_1 \in [1, \infty)$, then $\pi_{\alpha_2}(\omega) \in \mathcal{R}(\omega)$ for all $\alpha_2 \geq \alpha_1$.

Next we show that for any algorithm $\pi \in \Pi^{DO}$, there is an α-policy, with appropriate parameters α_β, that performs as well as π.

Theorem 3.2. *For any algorithm $\pi \in \Pi^{DO}$, if $\rho_\beta^\pi < \infty$, then the α-policy π_α with parameter $\alpha_\beta = \rho_\beta^\pi$ achieves the same competitive ratio, $\rho_\beta^{\pi_\alpha} = \rho_\beta^\pi$.*

Proof. Choose parameter $\alpha_\beta = \rho_\beta^\pi$. We show that the α-policy leads to feasible solutions for all instances $\omega \in \Omega_\beta$, by showing that $\pi_\alpha(\omega)(t) \geq \pi(\omega)(t)$ for all $\omega \in \Omega_\beta$ and all t. This is shown by contradiction; so suppose that $\pi_\alpha(\omega)(t) < \pi(\omega)(t)$ for some $\omega \in \Omega_\beta$ and some t. Choose any instance $\omega' \in \Omega_\beta(\omega^t)$ such that

$$\begin{aligned} & v^*(\omega') < v_\beta(\omega^t) + \frac{\pi(\omega)(t) - \pi_\alpha(\omega)(t)}{\alpha_\beta} \\ \Rightarrow \quad & \alpha_\beta v^*(\omega') < \alpha_\beta v_\beta(\omega^t) + \pi(\omega)(t) - \pi_\alpha(\omega)(t) \\ & = \pi(\omega)(t) \quad = \quad \pi(\omega')(t) \\ & \leq v^\pi(\omega') \end{aligned}$$

The last equality follows because instances ω and ω' have the same history up to time t, and π is an online algorithm. It follows that $v^\pi(\omega') > \rho_\beta^\pi v^*(\omega')$, which contradicts algorithm π having competitive ratio $\rho_\beta^\pi < \infty$. Hence $\pi_\alpha(\omega)(t) \geq \pi(\omega)(t)$ for all

$\omega \in \Omega_\beta$ and all t, and thus it follows from feasibility monotonicity and $\rho_\beta^\pi < \infty$ that the α-policy with $\alpha_\beta = \rho_\beta^\pi$ leads to feasible solutions for all $\omega \in \Omega_\beta$. Therefore $\rho_\beta^{\pi_\alpha} \leq \alpha_\beta = \rho_\beta^\pi$. Also, $\pi_\alpha(\omega)(t) \geq \pi(\omega)(t)$ for all $\omega \in \Omega_\beta$ and all t, implies that $\rho_\beta^{\pi_\alpha} \geq \rho_\beta^\pi$. Thus $\rho_\beta^{\pi_\alpha} = \rho_\beta^\pi$. □

COROLLARY 3.3. *To determine ρ_β^* (and ρ^*), it is sufficient to consider only the α-policy. That is,*

$$\rho_\beta^* \quad = \quad \inf_{\alpha \geq 1} \; \rho_\beta^{\pi_\alpha}$$

Also,

$$\rho_\beta^* \quad = \quad \inf \left\{ \alpha \geq 1 : \rho_\beta^{\pi_\alpha} < \infty \right\}$$

where $\inf \varnothing = \infty$.

However, the α-policy with parameter $\alpha_\beta = \rho_\beta^*$ may not be feasible for all $\omega \in \Omega_\beta$, in which case there is no optimal algorithm, as stated in Corollary 3.4.

COROLLARY 3.4. *If there exists an algorithm that is optimal for instances in Ω_β, and $\rho_\beta^* < \infty$, then*

1. *ρ_β^* is the least parameter for which the α-policy is feasible for instances in Ω_β, and ρ_β^* is therefore the optimal parameter for the α-policy, and*
2. *the α-policy with parameter $\alpha_\beta = \rho_\beta^*$ is optimal for instances in Ω_β among all deterministic online algorithms.*

Next it is natural to ask under which conditions an optimal algorithm exists, that is, under which conditions the α-policy with $\alpha_\beta = \rho_\beta^*$ is feasible. Proposition 3.5 shows that if $\mathcal{R}(\omega)$ is closed for all $\omega \in \Omega_\beta$, then the set $\left\{ \alpha \geq 1 : \rho_\beta^{\pi_\alpha} < \infty \right\}$ of feasible α-values is closed, and thus the α-policy with $\alpha_\beta = \rho_\beta^*$ is feasible.

PROPOSITION 3.5. *If $\rho_\beta^* < \infty$ and for some $\omega \in \Omega_\beta$, $\mathcal{R}(\omega)$ is closed, then $\pi_{\rho_\beta^*}(\omega) \in \mathcal{R}(\omega)$. Thus, if $\mathcal{R}(\omega)$ is closed for all $\omega \in \Omega_\beta$, then the α-policy with parameter $\alpha_\beta = \rho_\beta^*$ is feasible for all $\omega \in \Omega_\beta$.*

Proof. Consider $\omega \in \Omega_\beta$. Choose any sequence $\{\alpha_n\}$ such that $\alpha_n > \rho_\beta^*$ for all n and $\alpha_n \to \rho_\beta^*$ as $n \to \infty$. Thus $\pi_{\alpha_n}(\omega) = (\alpha_n v_\beta(\omega^1), \ldots, \alpha_n v_\beta(\omega^T)) \to (\rho_\beta^* v_\beta(\omega^1), \ldots, \rho_\beta^* v_\beta(\omega^T)) = \pi_{\rho_\beta^*}(\omega)$ as $n \to \infty$. It follows from Corollary 3.3 and feasibility monotonicity that $\pi_{\alpha_n}(\omega) \in \mathcal{R}(\omega)$ for all n. Then it follows from $\mathcal{R}(\omega)$ being closed that $\pi_{\rho_\beta^*}(\omega) \in \mathcal{R}(\omega)$. □

4. Optimal Competitive Ratios.

The results in Section 3 are all existential in nature. They do not show how to compute ρ_β^*, and therefore the optimal parameters α_β for the α-policy. In this section we show how the optimal parameters for the α-policy can be computed, first for the ORMP in Section 4.1, then for the HLBP in Section 4.2, and then we show how some of these results generalize for the OMMP in Section 4.3.

4.1. Online Resource Minimization Problem.

In this section we investigate the application of the α-policy to the ORMP, including the calculation of the optimal competitive ratios and optimal parameters for the α-policy, ρ_T^* and ρ^*. (The problem parameter β of interest for the ORMP is T.)

Recall that, for a given instance $\omega \in \Omega_T$, the decisions under the α-policy are given by $\pi_\alpha(\omega)(t) \equiv \alpha_T v_T(\omega^t)$. To implement the α-policy, one has to determine the optimal value of α_T, that is, ρ_T^*, as well as $v_T(\omega^t)$. These two issues are addressed next.

4.1.1. Calculation of $v_T(\omega^t)$. It is easy to compute $v_T(\omega^t) \equiv \inf_{\omega\in\Omega_T(\omega^t)} v^*(\omega)$ for the ORMP. The optimal value $v^*(\omega)$ with perfect information can be computed for any ω using Proposition 4.1. Next, for any partial instance ω^t, the best instance $\tilde{\omega} \in \Omega_T$ that starts with ω^t can be determined, and $v_T(\omega^t)$ can be computed, as shown in Proposition 4.4.

PROPOSITION 4.1. *For any instance $\omega = (a(1), \ldots, a(T))$ of the ORMP with nondecreasing productivity function $\eta_t(r)$, the optimal value $v^*(\omega)$ with perfect information is given by*

$$v^*(\omega) \quad = \quad \inf\left\{r \geq 0 \,:\, \sum_{t=i}^{j} \eta_t(r) \geq \sum_{t=i}^{j}\sum_{u=t}^{j} a_u(t) \quad \forall\, i,j \in \{1,\ldots,T\}, i \leq j\right\}$$

where $\inf \varnothing = \infty$.

The sum $\sum_{t=i}^{j} \eta_t(r)$ is the total amount of work that can be done between time periods i and j inclusive, while the sum $\sum_{t=i}^{j}\sum_{u=t}^{j} a_u(t)$ is the total amount of work that arrives at or after time i and is due at or before time j. Clearly the former sum must be at least as great as the latter for all pairs i, j. The proof of Proposition 4.1 consists of a straightforward verification that this is all that is required. It is given in the Appendix.

COROLLARY 4.2. *For any instance $\omega = (a(1), \ldots, a(T))$ of the ORMP with productivity function $\eta_t(r) = r$, the optimal value $v^*(\omega)$ with perfect information is given by*

$$v^*(\omega) \quad = \quad \max_{\{i,j\in\{1,\ldots,T\}\,:\,i\leq j\}} \left\{\frac{1}{j-i+1}\sum_{t=i}^{j}\sum_{u=t}^{j} a_u(t)\right\}$$

For any $\omega = (a(1), \ldots, a(T)), \omega' = (a'(1), \ldots, a'(T))$, we denote $\omega \leq \omega'$ if $a_u(t) \leq a'_u(t)$ for all t and u.

COROLLARY 4.3. *The optimal value $v^*(\omega)$ with perfect information of the ORMP with productivity function $\eta_t(r)$ is nondecreasing, that is, for any $\omega, \omega' \in \Omega_T$ with $\omega \leq \omega'$, it holds that $v^*(\omega) \leq v^*(\omega')$.*

It follows from Corollary 4.3 that for any partial instance $\omega^t = (a(1), \ldots, a(t))$, instance $(a(1), \ldots, a(t), 0, \ldots, 0) \in \Omega_T$ has the best optimal value $v^*(\tilde{\omega})$ with perfect information among all $\tilde{\omega} \in \Omega_T$ that start with ω^t. This result makes it easy to compute $v_T(\omega^t) \equiv \inf_{\omega\in\Omega_T(\omega^t)} v^*(\omega)$.

PROPOSITION 4.4. *For any partial instance $\omega^t = (a(1), \ldots, a(t))$ of the ORMP with productivity function $\eta_t(r)$, $v_T(\omega^t) = v^*(a(1), \ldots, a(t), 0, \ldots, 0)$. Specifically, for the ORMP with nondecreasing productivity function $\eta_t(r)$,*

$$v_T(\omega^t) \quad = \quad \inf\left\{r \geq 0 \,:\, \sum_{\tau=i}^{j} \eta_\tau(r) \geq \sum_{\tau=i}^{\min\{j,t\}}\sum_{u=\tau}^{j} a_u(\tau) \quad \forall\, i \in \{1,\ldots,t\}, j \in \{i,\ldots,T\}\right\}$$

and for the ORMP with productivity function $\eta_t(r) = r$,

$$v_T(\omega^t) \quad = \quad \max_{\{i\in\{1,\ldots,t\},j\in\{i,\ldots,T\}\}} \left\{\frac{1}{j-i+1}\sum_{\tau=i}^{\min\{j,t\}}\sum_{u=\tau}^{j} a_u(\tau)\right\}$$

Corollary 4.3 and Proposition 4.4 lead to the following result.

Corollary 4.5. *For any partial instances $\omega_1^t, \omega_2^t \in \Omega_T^t$ of the ORMP with productivity function $\eta_t(r)$, with $\omega_1^t \leq \omega_2^t$, it holds that $v_T(\omega_1^t) \leq v_T(\omega_2^t)$. Specifically, for any instance $\omega \in \Omega_T$, $v_T(\omega^t)$ is nondecreasing in t. It follows that for any α-policy and any instance $\omega \in \Omega_T$, $\pi_\alpha(\omega)(t) \equiv \alpha_T v_T(\omega^t)$ is nondecreasing in t.*

Thus the α-policy takes full advantage of resource that has already been allocated, since it allocates at least as much at time $t+1$ as at time t.

4.1.2. Optimal Parameter Values. Next we address the determination of the optimal value of α_T. For any $\omega \in \Omega$, the feasible region $\mathcal{R}(\omega)$ is determined by linear constraints (2.2), (2.3), (2.4), and (2.5). Thus $\mathcal{R}(\omega)$ is a polyhedron, and is closed. (Proposition B.1 in the Appendix establishes the more general result that if the productivity function η_t is upper semicontinuous for all t, then $\mathcal{R}(\omega)$ is closed for all $\omega \in \Omega$.) Therefore it follows from Corollary 3.4 and Proposition 3.5 that the optimal value of α_T is ρ_T^*, and that the α-policy with parameter $\alpha_T = \rho_T^*$ is feasible and optimal among all deterministic online algorithms. (It is shown in Section 5.2, Proposition 5.3, that if the productivity function η_t is concave for all t, then the α-policy with parameters $\alpha_T = \rho_T^*$ is optimal among all randomized online algorithms.)

First we show that, to determine ρ_T^* for the multiple deadline ORMP, it is sufficient to consider the single deadline ORMP. This result simplifies the calculation of ρ_T^*.

We introduce the following notation to distinguish between the single deadline ORMP and the multiple deadline ORMP. Consider the function $\theta : \Omega \mapsto \Omega$ that postpones the deadlines of all work to the end of the time horizon. That is, for any instance $\omega = (a_1(1), a_2(1), \ldots, a_T(T)) \in \Omega_T$ of the multiple deadline ORMP, $\theta(\omega) = (a'_1(1), a'_2(1), \ldots, a'_T(T))$, where

$$a'_u(t) \quad = \quad \begin{cases} 0 & \text{if} \quad u < T \\ \sum_{v=t}^{T} a_v(t) & \text{if} \quad u = T \end{cases}$$

When considering the single deadline ORMP, we simplify the notation slightly, by letting $a(t) \in \mathbb{R}_+$ denote the amount of work arriving at time t (with deadline T), and $q(t) \in \mathbb{R}_+$ denote the amount of work performed at time t.

Recall that Ω_T denotes the set of instances of length T for the multiple deadline ORMP. Note that the set of instances of length T for the single deadline ORMP is given by $\theta(\Omega_T) \subset \Omega_T$. Also note that, for any instance $\theta(\omega)$, the set of feasible solutions of the single deadline ORMP is the same as the set $\mathcal{R}(\theta(\omega))$ of feasible solutions for the same instance of the multiple deadline ORMP. In addition, for any instance $\theta(\omega)$, the optimal value with perfect information of the single deadline ORMP is equal to the optimal value $v^*(\theta(\omega))$ with perfect information for the same instance of the multiple deadline ORMP. The following corollary for the single deadline ORMP follows from Proposition 4.1 and Corollary 4.2.

Corollary 4.6. *For any instance $\omega = (a(1), \ldots, a(T))$ of the single deadline ORMP, the optimal value $v^*(\omega)$ with perfect information is given by the following: With nondecreasing productivity function $\eta_t(r)$,*

$$v^*(\omega) \quad = \quad \inf\left\{ r \geq 0 \; : \; \sum_{t=i}^{T} \eta_t(r) \geq \sum_{t=i}^{T} a(t) \quad \forall \; i \in \{1, \ldots, T\} \right\}$$

and with productivity function $\eta_t(r) = r$,

$$v^*(\omega) \quad = \quad \max_{\{i\in\{1,\ldots,T\}\}} \left\{ \frac{1}{T-i+1} \sum_{t=i}^{T} a(t) \right\}$$

It follows from Proposition 4.4 that, for any partial instance $\theta(\omega)^t$, the best value with perfect information, over all instances in $\theta(\Omega_T)$ of the single deadline ORMP that start with $\theta(\omega)^t$, is equal to the best value $v_T(\theta(\omega)^t)$ with perfect information, over all instances in Ω_T of the multiple deadline ORMP that start with $\theta(\omega)^t$.

COROLLARY 4.7. *For any partial instance* $\omega^t = (a(1), \ldots, a(t))$ *of the single deadline ORMP with nondecreasing productivity function* $\eta_t(r)$,

$$v_T(\omega^t) \quad = \quad \inf \left\{ r \geq 0 \, : \, \sum_{\tau=i}^{T} \eta_\tau(r) \geq \sum_{\tau=i}^{t} a(\tau) \ \ \forall \, i \in \{1, \ldots, t\} \right\}$$

and for the single deadline ORMP with productivity function $\eta_t(r) = r$,

$$v_T(\omega^t) \quad = \quad \max_{\{i\in\{1,\ldots,t\}\}} \left\{ \frac{1}{T-i+1} \sum_{\tau=i}^{t} a(\tau) \right\}$$

Thus, for any instance $\theta(\omega) \in \theta(\Omega_T)$ of the single deadline ORMP, the α-policy prescribes exactly the same decisions for the single deadline ORMP as for the multiple deadline ORMP:

$$\pi_\alpha(\theta(\omega))(t) \quad \equiv \quad \alpha_T v_T(\theta(\omega)^t)$$

Also note that, for all $\omega \in \Omega_T$, $v^*(\theta(\omega)) \leq v^*(\omega)$, and $v_T(\theta(\omega)^t) \leq v_T(\omega^t)$ for all $t \in \{1, \ldots, T\}$.

We want to show that ρ^*_T is the same for the multiple deadline ORMP and the single deadline ORMP, and thus that the optimal parameter α_T is the same for the multiple deadline ORMP and the single deadline ORMP. That is, we want to show that

$$\begin{aligned} \rho^*_T &= \inf_{\alpha\geq 1} \inf \left\{ \rho \geq 1 \, : \, v^{\pi_\alpha}(\omega) \leq \rho v^*(\omega) \ \ \forall \, \omega \in \Omega_T \right\} \\ &= \inf_{\alpha\geq 1} \inf \left\{ \rho \geq 1 \, : \, v^{\pi_\alpha}(\omega) \leq \rho v^*(\omega) \ \ \forall \, \omega \in \theta(\Omega_T) \right\} \end{aligned}$$

The first equality follows from Corollary 3.3, and it remains to establish the second equality. We do that by recalling that ρ^*_T is nondecreasing in T, and we show that if $\alpha \in (\rho^*_{T-1}, \rho^*_T)$, so that the α-policy with parameter α is infeasible for some instance $\omega \in \Omega_T$, then the α-policy with parameter α is also infeasible for instance $\theta(\omega) \in \theta(\Omega_T)$. The performance of the α-policy on instances in Ω_{T-1} enter into the evaluation, and we introduce the following notation for that purpose.

Consider the function $\vartheta : \Omega \mapsto \Omega$ that removes all work with deadline equal to T. That is, for any instance $\omega = (a_1(1), a_2(1), \ldots, a_T(T)) \in \Omega_T$, $\vartheta(\omega) = (a''_1(1), a''_2(1), \ldots, a''_{T-1}(T-1))$, where $a''_u(t) = a_u(t)$ for all $t \in \{1, \ldots, T-1\}$ and $u \in \{t, \ldots, T-1\}$. Note that for any $\omega \in \Omega_T$, $\vartheta(\omega) \in \Omega_{T-1}$. Thus, for any $\omega \in \Omega_T$, the α-policy prescribes decisions $\pi_\alpha(\vartheta(\omega))(t) = \alpha_{T-1} v_{T-1}(\vartheta(\omega)^t)$ for $\vartheta(\omega) \in \Omega_{T-1}$. For any instance $\omega = (a_1(1), a_2(1), \ldots, a_T(T)) \in \Omega_T$, let $\tilde{\omega} = (\tilde{a}_1(1), \tilde{a}_2(1), \ldots, \tilde{a}_T(T)) \in \Omega_T$

be given by $\tilde{a}_u(t) = a_u(t)$ for all $t \in \{1, \ldots, T-1\}$ and $u \in \{t, \ldots, T-1\}$, and $\tilde{a}_T(t) = 0$ for all $t \in \{1, \ldots, T\}$. Note that $\tilde{\omega} \leq \omega$. Thus, if $\eta_t(r) \geq 0$ for all $r, t \geq 0$, then $v^*(\vartheta(\omega)) = v^*(\tilde{\omega}) \leq v^*(\omega)$, and $v_{T-1}(\vartheta(\omega)^t) = v_T(\tilde{\omega}^t) \leq v_T(\omega^t)$ for all $t \in \{1, \ldots, T-1\}$.

For any $r \in \mathbb{R}_+^T$, let r^{T-1} denote the first $T-1$ components of r.

LEMMA 4.8. *For the ORMP with productivity function $\eta_t(r)$, any $\omega \in \Omega_T$, and any $r \in \mathbb{R}_+^T$, $r \in \mathcal{R}(\theta(\omega))$ and $r^{T-1} \in \mathcal{R}(\vartheta(\omega))$ imply that $r \in \mathcal{R}(\omega)$.*

Proof. Consider any instance $\omega = (a_1(1), a_2(1), \ldots, a_T(T))$, $\theta(\omega) = (a'_1(1), a'_2(1), \ldots, a'_T(T))$, and $\vartheta(\omega) = (a''_1(1), a''_2(1), \ldots, a''_{T-1}(T-1))$. Consider any r such that $r \in \mathcal{R}(\theta(\omega))$ and $r^{T-1} \in \mathcal{R}(\vartheta(\omega))$. As usual, available work is performed in EDF order. Let $w_u(t) \equiv a_u(1) - q_u(1) + a_u(2) - q_u(2) + \cdots + a_u(t)$, $w'_T(t) \equiv a'_T(1) - q'_T(1) + a'_T(2) - q'_T(2) + \cdots + a'_T(t)$, and $w''_u(t) \equiv a''_u(1) - q''_u(1) + a''_u(2) - q''_u(2) + \cdots + a''_u(t)$ denote the remaining amount of work at time t with deadline u for instances ω, $\theta(\omega)$, and $\vartheta(\omega)$ respectively.

It is shown by induction on t that $w_u(t) = w''_u(t)$ and $q_u(t) = q''_u(t)$ for all $t = 1, \ldots, T-1$, $u = t, \ldots, T-1$, and $w'_T(t) = \sum_{u=t}^T w_u(t)$ and $q'_T(t) = \sum_{u=t}^T q_u(t)$ for all $t = 1, \ldots, T$. Clearly the hypothesis holds for $t = 1$. Suppose the hypothesis holds for t. Note that $q_t(t) = q''_t(t) = w''_t(t) = w_t(t)$ from the assumption that $r^{T-1} \in \mathcal{R}(\vartheta(\omega))$. Then $w_u(t+1) = w_u(t) - q_u(t) + a_u(t+1) = w''_u(t) - q''_u(t) + a''_u(t+1) = w''_u(t+1)$ for all $u = t+1, \ldots, T-1$. Because available work is performed in EDF order, $q_u(t+1) = q''_u(t+1)$ for all $u = t+1, \ldots, T-1$. Also, $w'_T(t+1) = w'_T(t) - q'_T(t) + a'_T(t+1) = \sum_{u=t}^T w_u(t) - \sum_{u=t}^T q_u(t) + \sum_{u=t+1}^T a_u(t+1) = \sum_{u=t+1}^T (w_u(t) - q_u(t) + a_u(t+1)) = \sum_{u=t+1}^T w_u(t+1)$, and the hypothesis has been established.

Recall that the EDF rule ensures that constraints (2.6), (2.3), and (2.5) are satisfied. Thus, to show that $r \in \mathcal{R}(\omega)$, it remains to verify that solution r satisfies (2.4) for instance ω. ¿From the assumption that $r^{T-1} \in \mathcal{R}(\vartheta(\omega))$, it follows that r^{T-1} satisfies (2.4) for $\vartheta(\omega)$, and thus it follows from the hypothesis established above that r satisfies (2.4) for ω, for all $t = 1, \ldots, T-1$. It remains to be shown that r satisfies (2.4) for ω at $t = T$. ¿From the assumption that $r \in \mathcal{R}(\theta(\omega))$ and the hypothesis, it follows that $w_T(T) = w'_T(T) \leq \eta_T(r(T))$, and thus r satisfies (2.4) for ω at $t = T$. □

Next it is shown that, to determine ρ_T^*, it is sufficient to consider only the single deadline ORMP.

THEOREM 4.9. *Suppose that the productivity function $\eta_t(r)$ is nondecreasing. If the α-policy with parameter α gives an infeasible solution for some instance ω and a feasible solution for instance $\vartheta(\omega)$, then the α-policy with parameter α gives an infeasible solution for instance $\theta(\omega)$.*

Proof. Consider any $\omega \in \Omega_T$. Recall that $v_{T-1}(\vartheta(\omega)^t) \leq v_T(\omega^t)$, and thus $\pi_\alpha(\vartheta(\omega))(t) = \alpha v_{T-1}(\vartheta(\omega)^t) \leq \alpha v_T(\omega^t) = \pi_\alpha(\omega)(t)$, for all $t \in \{1, \ldots, T-1\}$. It follows from the assumption that $\pi_\alpha(\vartheta(\omega)) \in \mathcal{R}(\vartheta(\omega))$ and from feasibility monotonicity that $\pi_\alpha(\omega)^{T-1} \in \mathcal{R}(\vartheta(\omega))$. ¿From the assumption that $\pi_\alpha(\omega) \notin \mathcal{R}(\omega)$ and the contrapositive of the result in Lemma 4.8, it follows that $\pi_\alpha(\omega) \notin \mathcal{R}(\theta(\omega))$. Recall that $v_T(\theta(\omega)^t) \leq v_T(\omega^t)$, and thus $\pi_\alpha(\theta(\omega))(t) = \alpha v_T(\theta(\omega)^t) \leq \alpha v_T(\omega^t) = \pi_\alpha(\omega)(t)$, for all $t \in \{1, \ldots, T\}$. It follows from feasibility monotonicity and $\pi_\alpha(\omega) \notin \mathcal{R}(\theta(\omega))$ that $\pi_\alpha(\theta(\omega)) \notin \mathcal{R}(\theta(\omega))$. Thus the α-policy with parameter α gives an infeasible solution for instance $\theta(\omega)$. □

THEOREM 4.10. *Suppose that the productivity function $\eta_t(r)$ is nondecreasing.*

Then

$$\rho_T^* \;=\; \inf_{\alpha\geq 1} \inf\left\{\rho \geq 1 \,:\, v^{\pi_\alpha}(\omega) \leq \rho v^*(\omega) \;\;\forall\, \omega \in \theta(\Omega_T)\right\}$$

That is, to determine ρ_T^* (and ρ^*), it is sufficient to consider only the α-policy, and only the instances in $\theta(\Omega_T)$.

Proof. Corollary 3.3 established that

$$\rho_T^* \;=\; \inf_{\alpha\geq 1} \inf\left\{\rho \geq 1 \,:\, v^{\pi_\alpha}(\omega) \leq \rho v^*(\omega) \;\;\forall\, \omega \in \Omega_T\right\}$$

Let

$$\rho_T^\theta \;\equiv\; \inf_{\alpha\geq 1} \inf\left\{\rho \geq 1 \,:\, v^{\pi_\alpha}(\omega) \leq \rho v^*(\omega) \;\;\forall\, \omega \in \theta(\Omega_T)\right\}$$

Note that, because $\theta(\Omega_T) \subset \Omega_T$, it follows that $\rho_T^\theta \leq \rho_T^*$. Also note that, similar to ρ_T^*, ρ_T^θ is nondecreasing in T. We show by induction on T that $\rho_T^\theta = \rho_T^*$. For $T = 1$, $\theta(\Omega_1) = \Omega_1$, and thus $\rho_1^\theta = \rho_1^*$. Suppose that $\rho_{T-1}^\theta = \rho_{T-1}^*$. Then $\rho_{T-1}^* = \rho_{T-1}^\theta \leq \rho_T^\theta \leq \rho_T^*$. Thus, if $\rho_{T-1}^* = \rho_T^*$, then $\rho_T^\theta = \rho_T^*$.
Otherwise, if $\rho_{T-1}^* < \rho_T^*$, then consider any $\alpha \in (\rho_{T-1}^*, \rho_T^*)$. Then there exists an $\omega \in \Omega_T$ such that the α-policy with parameter α gives an infeasible solution for instance ω, i.e., $\pi_\alpha(\omega) \notin \mathcal{R}(\omega)$. Also, the α-policy with parameter α gives a feasible solution for instance $\vartheta(\omega)$, i.e., $\pi_\alpha(\vartheta(\omega)) \in \mathcal{R}(\vartheta(\omega))$, because $\alpha > \rho_{T-1}^*$ and $\vartheta(\omega) \in \Omega_{T-1}$. Then it follows from Theorem 4.9 that $\pi_\alpha(\theta(\omega)) \notin \mathcal{R}(\theta(\omega))$, and thus $v^{\pi_\alpha}(\theta(\omega)) = \infty$. Hence $\{\rho \geq 1 : v^{\pi_\alpha}(\omega) \leq \rho v^*(\omega)\; \forall \omega \in \theta(\Omega_T)\} = \varnothing$. Thus by feasibility monotonicity $\inf_{\alpha<\rho_T^*} \inf\{\rho \geq 1 : v^{\pi_\alpha}(\omega) \leq \rho v^*(\omega)\; \forall \omega \in \theta(\Omega_T)\} = \infty$. Next consider any $\alpha > \rho_T^*$. It follows from Theorem 3.2 and feasibility monotonicity that $v^{\pi_\alpha}(\omega) = \alpha v^*(\omega)$ for all $\omega \in \Omega_T$, and thus for all $\omega \in \theta(\Omega_T) \subset \Omega_T$. Hence, noting that $v^*(\omega) > 0$ for some $\omega \in \theta(\Omega_T)$, it follows that $\inf\{\rho \geq 1 : v^{\pi_\alpha}(\omega) \leq \rho v^*(\omega)\; \forall \omega \in \theta(\Omega_T)\} = \alpha$. Thus $\rho_T^\theta = \inf_{\alpha>\rho_T^*} \inf\{\rho \geq 1 : v^{\pi_\alpha}(\omega) \leq \rho v^*(\omega)\; \forall \omega \in \theta(\Omega_T)\} = \inf_{\alpha>\rho_T^*} \alpha = \rho_T^*$. ☐

Next we show how ρ_T^* is given by the optimal value of a linear program for the ORMP with a linear productivity function.

Theorem 4.10 simplifies the calculation of ρ_T^* by establishing that it is sufficient to consider the single deadline ORMP. Theorem 4.12 further simplifies the calculation of ρ_T^* by establishing that the parameter α_T is too small if and only if there exists an instance $\omega = (a(1), \ldots, a(T))$ such that the total amount of resource allocated under the α-policy, $\alpha_T \sum_{t=1}^T v_T(\omega^t)$, is less than the total amount of work to be performed, $\sum_{t=1}^T a(t)$. Such instances can be identified with a linear program, as shown later. Theorem 4.12 follows directly from Lemma 4.11.

Lemma 4.11. *Suppose that the α-policy with parameter α_T gives an infeasible solution with instance $\omega' = (a'(1), \ldots, a'(T)) \in \Omega_T$. Then there exists an instance $\omega = (a(1), \ldots, a(T))$ such that*

$$\alpha_T \sum_{t=1}^{T} v_T(\omega^t) \;\;<\;\; \sum_{t=1}^{T} a(t) \tag{4.1}$$

Proof. Let $q'(t)$ denote the amount of work performed at time t on instance ω'. Let $\tau \equiv \max\{t \,:\, \sum_{l\leq t} q'(l) = \sum_{l\leq t} a'(l)\}$. Since ω' is infeasible, $\tau < T$. For $t > \tau$, $q'(t) = \alpha_T v_T(\omega'_t)$, so if $\tau = 0$, we have proven the desired inequality with $\omega = \omega'$.

If $\tau > 0$, then the part of instance ω' before τ does not contribute to the solution being infeasible. Define $\omega = (a(1), \ldots, a(T))$ by

$$a(t) = \begin{cases} 0 & \text{if} \quad t \leq \tau \\ a'(t) & \text{if} \quad t > \tau \end{cases}$$

Note that $\omega \leq \omega'$, so $v_T(\omega^t) \leq v(\omega'^t)$ for all t.

Since $\sum_{t=1}^T q'(t) < \sum_{t=1}^T a'(t)$ and $\sum_{t \leq \tau} q'(t) = \sum_{t \leq \tau} a'(t)$, we have $\sum_{t > \tau} q'(t) < \sum_{t > \tau} a'(t)$. This gives

$$\alpha_T \sum_{t>\tau} v_T(\omega^t) \leq \alpha_T \sum_{t>\tau} v_T(\omega'^t) = \sum_{t>\tau} q'(t) < \sum_{t>\tau} a'(t).$$

For $t \leq \tau$, we know $v_T(\omega^t) = a(t) = 0$, and for $t > \tau$, $a(t) = a'(t)$. Thus, $\alpha_T \sum_{t=1}^T v_T(\omega^t) < \sum_{t=1}^T a(t)$, which is (4.1). ◻

THEOREM 4.12. *Parameter α_T for the α-policy is too small ($\alpha_T < \rho_T^*$) if and only if there exists an instance $\omega = (a(1), \ldots, a(T))$ such that $\alpha_T \sum_{t=1}^T v_T(\omega^t) < \sum_{t=1}^T a(t)$. Also,*

$$\rho_T^* = \inf \left\{ \rho \in [1, \infty) \,:\, \rho \sum_{t=1}^T \max_{\{i \in \{1,\ldots,t\}\}} \left\{ \frac{1}{T-i+1} \sum_{j=i}^t a(j) \right\} \geq \sum_{t=1}^T a(t) \quad \forall \ (a(1), \ldots, a(T)) \in \mathbb{R}_+^T \right\} \tag{4.2}$$

The α-policy provides the optimal solution with perfect information for the zero instance $\omega = (0, \ldots, 0)$. Thus, to determine ρ_T^*, one can restrict attention to instances $(a(1), \ldots, a(T)) \in \mathbb{R}_+^T$ for which $\sum_{t=1}^T a(t) > 0$. From (4.2) it follows that ρ_T^* is determined by instances $(a(1), \ldots, a(T)) \in \mathbb{R}_+^T$ that minimize

$$\sum_{t=1}^T \max_{\{i \in \{1,\ldots,t\}\}} \left\{ \frac{1}{T-i+1} \sum_{j=i}^t \frac{a(j)}{\sum_{k=1}^T a(k)} \right\}$$

It follows that one can restrict attention to instances $(a(1), \ldots, a(T)) \in \mathbb{R}_+^T$ for which $\sum_{t=1}^T a(t) = 1$.

COROLLARY 4.13. *The optimal competitive ratio ρ_T^* for the ORMP can be calculated by solving the following linear program with decision variables $a(t)$ and $x(t)$, $t \in \{1, \ldots, T\}$:*

$$\begin{aligned}
\text{(LP)} \quad \text{minimize} \ & \sum_{t=1}^T x(t) \\
\text{subject to} \ & \sum_{t=1}^T a(t) = 1 \\
& x(t) \geq \frac{1}{T-i+1} \sum_{j=i}^t a(j) \quad \forall\, t \in \{1, \ldots, T\}, \ \forall\, i \in \{1, \ldots, t\} \qquad (4.3) \\
& a(t) \geq 0 \qquad \forall\, t \in \{1, \ldots, T\}
\end{aligned}$$

Proof. In an optimal solution (a^*, x^*) of the LP, $a^* = (a^*(1), \ldots, a^*(T))$ represents a worst-case instance of length T for the α-policy, and each $x^*(t)$ represents the corresponding value of $v_T(a^{*t})$. Also, it follows from (4.2) that

$$\rho_T^* \quad = \quad \inf\left\{\rho \in [1,\infty) \,:\, \rho \sum_{t=1}^{T} x^*(t) \geq 1\right\}$$

That is, $\rho_T^* = \max\left\{1, 1/\sum_{t=1}^{T} x^*(t)\right\}$. Actually, $\rho_T^* = 1/\sum_{t=1}^{T} x^*(t)$, as shown next. It is easily checked that (a, x) with $a(t) = 1/T$ and $x(t) = t/T^2$ is a feasible solution for the LP. Thus $\sum_{t=1}^{T} x^*(t) \leq \sum_{t=1}^{T} t/T^2 = (1 + 1/T)/2$. Hence $\rho_T^* = 1/\sum_{t=1}^{T} x^*(t) \geq 2/(1 + 1/T) \geq 1$. □

There are $2T$ decision variables, but quadratically many constraints, because of the $T(T+1)/2$ allocation constraints (4.3). As the number of periods T increases, working with the full LP becomes prohibitive, in terms of memory requirements as well as computation times. Fortunately, most of the allocation constraints are inactive at an optimal solution, which suggests that a cutting plane method may provide an effective solution approach. To determine a good initial set of constraints, we analyzed the active allocation constraints for the instance with $T = 50$. In Figure 4.1, each grid cell (t, i) corresponds to an allocation constraint from (4.3:t, i), and each of the black dots shows an active allocation constraint at optimality.

Let $\gamma(t) \equiv \arg\max_{\{i \in \{1,\ldots,t\}\}} \left\{\sum_{j=i}^{t} a(j)/(T - i + 1)\right\}$. Then the black dots in Figure 4.1 can also be viewed as a map of $\gamma(t)$ versus t. (Note that $\gamma(t)$ is a set valued function.) Although the precise pattern of $\gamma(t)$ versus t is quite complicated, it does follow a crude pattern. For $1 \leq t \leq \delta T \approx T/3$, $\gamma(t) = 1$. This means that for the first approximately $T/3$ time periods t, the optimum value $v_T(\omega^t)$ of the truncated instance ω^t is the amount of resource required to do all the work that has arrived since the first time period. For larger t, the slightly smaller amount of work that has arrived since a later time period $\gamma(t)$, divided by the smaller value $T - \gamma(t) + 1$, provides a larger bound on the required resource. As $t \to T$, the value of $\gamma(t) \to T$ also. This means that in a worst-case instance, so much work arrives in the last time periods, that the amount of resource required is determined by the few most recent work arrivals only.

By estimating δ and lower and upper bounds on $\gamma(t)$, we extrapolated the pattern in Figure 4.1 to larger values of T, and thereby identified relatively small subsets of allocation constraints which we hoped would contain all the active allocation constraints for an optimal solution. We solved the relaxed LP. Then we checked the remaining allocation constraints to see if our guess was correct. If not, we added all the violated allocation constraints that were initially left out and repeated. For $T \leq 200$, no cut generation was necessary, but for $T > 200$, several rounds of cut generation were needed prior to finding the optimal solution.

Table 4.1 gives the values of ρ_T^* for some values of T, as determined by the LP.

4.1.3. Asymptotic Behavior. Sometimes it is interesting to determine $\rho^* = \sup_{\beta \in \mathbf{B}} \rho_\beta^*$. Thus, for the ORMP, we are interested in $\rho^* = \sup_{T \in \mathbb{Z}_+} \rho_T^*$. As pointed out in Section 2.2, ρ_T^* is nondecreasing in T, and thus $\rho^* = \lim_{T\to\infty} \rho_T^*$. It is convenient to study ρ^* with a continuous time model for the ORMP, as in Kleywegt et al. [12]. Also, by Theorem 4.10, to determine ρ^*, it is sufficient to consider the single deadline version of the ORMP.

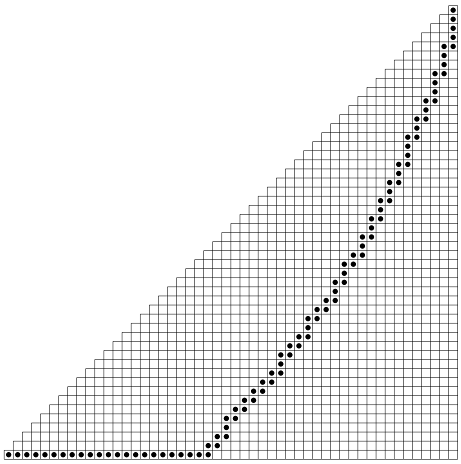

FIG. 4.1. *Incidence of active constraints at optimality for $T = 50$*

T	ρ_T^*
1	1
2	4/3
3	3/2
4	44/27
5	1.71329
6	1.77778
7	1.82765
8	1.86880
9	1.90547
10	1.93576
25	2.14951
50	2.26470
75	2.31800
100	2.35061
200	2.41585
300	2.44663
400	2.46592
500	2.47956
750	2.501833

TABLE 4.1
Values of ρ_T^* for the ORMP.

In the continuous time model, $t \in [0, 1]$. An instance is given by an integrable function $a : [0, 1] \mapsto [0, \infty)$, with $a(t)$ denoting the rate at which work (with deadline 1) arrives at time t. Let Ω denote the set of all such instances. Thus $A(t) \equiv \int_0^t a(\tau)\, d\tau$ is the total amount of work that has arrived by time t. A solution is given by an integrable function $r : [0, 1] \mapsto [0, \infty)$, with $r(t)$ denoting the rate at which resources are allocated at time t. Let $W(t)$ denote the amount of unfinished work at time t.

Thus W satisfies the following differential equation.

$$\frac{dW}{dt}(t) \quad = \quad \begin{cases} a(t) - r(t) & \text{if} \quad W(t) > 0 \\ \max\{a(t) - r(t),\, 0\} & \text{if} \quad W(t) = 0 \end{cases}$$

with boundary condition $W(0) = 0$. A solution r is feasible for an instance a if $W(1) = 0$, for which it is necessary that $\int_0^1 a(t)\,dt \le \int_0^1 r(t)\,dt$. The value of a solution r for an instance a is $\sup_{t\in[0,1]} r(t)$ if r is feasible for a, and the value is ∞ if r is not feasible for a. Algorithms are defined similar to those for the discrete time ORMP. It is easy to verify that, for any instance a, the set $\mathcal{R}(a)$ of feasible solutions is convex, and it will follow from Theorem 5.2 that it is sufficient to restrict attention to deterministic algorithms.

For any instance a, the optimal value $v^*(a)$ with perfect information is given by

$$v^*(a) \quad = \quad \sup_{\gamma\in[0,1)} \frac{1}{1-\gamma} \int_\gamma^1 a(t)\; dt$$

As before, a^t denotes the instance a truncated at time t, and the optimal value $v_\infty(a^t)$ with perfect information among all completions of truncated instance a^t is given by

$$v_\infty(a^t) \quad = \quad \sup_{\gamma\in[0,t)} \frac{1}{1-\gamma} \int_\gamma^t a(x)\; dx \quad = \quad \sup_{\gamma\in[0,t)} \frac{A(t) - A(\gamma)}{1-\gamma} \quad \equiv \quad \sup_{\gamma\in[0,t)} f_t(\gamma)$$

The α-policy allocates resource at rate

$$r^\alpha(t) \quad = \quad \alpha v_\infty(a^t) \tag{4.4}$$

For the α-policy with parameter α to be feasible, it is necessary that

$$\int_0^1 a(t)\,dt \quad \le \quad \int_0^1 r^\alpha(t)\,dt \quad = \quad \alpha \int_0^1 v_\infty(a^t)\,dt$$

for all instances a. Suppose the α-policy with parameter α_∞ gives an infeasible solution with instance a', that is, $W'(1) > 0$. Let $\tau \equiv \sup\{t \in [0,1] \,:\, W'(t) = 0\}$. Let

$$a(t) \quad = \quad \begin{cases} 0 & \text{if} \quad t \in [0,\tau) \\ a'(t) & \text{if} \quad t \in [\tau, 1] \end{cases}$$

Then, similar to Lemma 4.11, it follows that

$$\alpha_\infty \int_0^1 v_\infty(a^t)\,dt \quad < \quad \int_0^1 a(t)\,dt$$

Thus, for the α-policy with parameter α to be feasible, it is necessary and sufficient that

$$\int_0^1 a(t)\,dt \quad \le \quad \alpha \int_0^1 v_\infty(a^t)\,dt \tag{4.5}$$

for all instances a. Therefore

$$\rho^* \quad = \quad \inf\left\{\rho \in [1,\infty) \,:\, \rho \int_0^1 \sup_{\gamma\in[0,t)} \frac{1}{1-\gamma} \int_\gamma^t a(y)\,dy\,dt \ge \int_0^1 a(t)\,dt \;\; \forall\, a \in \Omega\right\} \tag{4.6}$$

Thus ρ^* can be calculated by solving the following optimal control problem with integrable controls $a(t)$ and $x(t)$, $t \in [0,1]$:

$$
\begin{aligned}
\text{minimize} \quad & \int_0^1 x(t)\,dt \\
\text{subject to} \quad & \int_0^1 a(t)\,dt = 1 \\
& x(t) \geq \frac{1}{1-\gamma}\int_\gamma^t a(y)\,dy \qquad \forall\, t \in [0,1], \ \ \forall\, \gamma \in [0,t) \\
& a(t) \geq 0 \qquad \forall\, t \in [0,1]
\end{aligned} \tag{4.7}
$$

Rewriting (4.5) gives

$$
\rho^* \quad \geq \quad \frac{\int_0^1 a(t)\,dt}{\int_0^1 v_\infty(a^t)\,dt} \tag{4.8}
$$

for all instances a such that $\int_0^1 a(t)\,dt > 0$. To obtain a lower bound on ρ^*, one can substitute particular instances a into (4.8). Consider instances $a : [0,1] \mapsto [0,\infty)$ that are continuous and nondecreasing. Consider any continuous nondecreasing extension $a : (-\infty,1] \mapsto [0,\infty)$, so that $\lim_{t\to-\infty} a(t) = 0$. Then, for any $t \in (0,1)$, it follows from a being continuous that f_t is differentiable, and it follows from a being nondecreasing that f_t is at first nondecreasing and then nonincreasing on $(-\infty,t)$, and f_t attains a maximum at a point $\gamma(t) \in (-\infty,t)$ where $f_t'(\gamma) = 0$. That is, for any $t \in (0,1)$, $\gamma(t)$ is a solution of

$$
A(t) \quad = \quad A(\gamma) + (1-\gamma)a(\gamma) \quad \equiv \quad g(\gamma) \tag{4.9}
$$

It follows from a being nondecreasing that g is nondecreasing. Also, A is nondecreasing, and thus, although the solution of (4.9) is not necessarily unique, γ can be chosen to be nondecreasing. Note that $\gamma(t) > 0$ for t sufficiently close to 1. If $\gamma(t) < 0$ for some $t \in (0,1)$, let t' be a crossing point of $\gamma(t) = 0$, that is, $\gamma(t) \leq 0$ for all $t < t'$ and $\gamma(t) \geq 0$ for all $t > t'$; otherwise, let $t' \equiv 0$. Then, for all $t < t'$, $v_\infty(a^t) = \sup_{\gamma\in[0,t)} f_t(\gamma) = f_t(0) = A(t)$. Also, for all $t > t'$, $v_\infty(a^t) = f_t(\gamma(t)) = [A(t) - A(\gamma(t))]/[1-\gamma(t)] = a(\gamma(t))$. Hence,

$$
\int_0^1 v_\infty(a^t)\,dt \quad = \quad \int_0^{t'} A(t)\,dt + \int_{t'}^1 a(\gamma(t))\,dt
$$

Therefore,

$$
\rho^* \quad \geq \quad \frac{\int_0^1 a(t)\,dt}{\int_0^{t'} A(t)\,dt + \int_{t'}^1 a(\gamma(t))\,dt}
$$

For example, consider a linear arrival rate $a(t) = t$. Then $A(t) = t^2/2$, $\gamma(t) = 1 - \sqrt{1-t^2}$, $t' = 0$, and thus $\rho^* \geq \int_0^1 t\,dt / \int_0^1 \left(1 - \sqrt{1-t^2}\right)\,dt = 0.5/(1-\pi/4) \approx 2.3298$ (Kleywegt et al. [12]). As another example, one can consider $a(t;k) = e^{k(1-t)^{1/k}}$. In the context of the HLBP, Bar-Noy et al. [6] arrive at the same optimal control problem to determine ρ^*_{HLBP}. Using a function similar to $a(t;k)$ they prove a lower bound $\rho^*_{HLBP} \geq e$. They also showed that this bound is tight by demonstrating an algorithm with a competitive ratio of e in the continuous model. Thus, $\rho^* = \lim_{T\to\infty} \rho^*_T = e$ for the single deadline version of the ORMP. Applying Theorem (4.10), we conclude that the same ρ^* value applies to the ORMP with multiple deadlines.

Convergence to the limit e is quite slow, as shown earlier in Table 4.1. For example, values such as $\rho^*_{50} \approx 2.26470$ and $\rho^*_{100} \approx 2.35061$ offer considerably better performance than e.

4.2. Hierarchical Line Balancing Problem. In this section we present optimality results for the HLBP; we investigate the application of the α-policy to the HLBP, including the calculation of the optimal competitive ratios and optimal parameters for the α-policy, $\rho^*_{T,m}$ and ρ^*. (The problem parameter β of interest for the HLBP is (T, m).) We show that if $T \geq m$, then $\rho^*_{T,m}$ for the HLBP is equal to ρ^*_m for the ORMP, but if $T < m$ then $\rho^*_{T,m}$ for the HLBP can be strictly between ρ^*_T and ρ^*_m for the ORMP.

As before, for a given instance $\omega \in \Omega_{T,m}$, the decisions under the α-policy are given by $\pi_\alpha(\omega)(t) \equiv \alpha_{T,m} v_{T,m}(\omega^t)$. To implement the α-policy, one has to determine the optimal value of $\alpha_{T,m}$, that is, $\rho^*_{T,m}$, as well as $v_{T,m}(\omega^t)$. These two issues are addressed next.

4.2.1. Computing $v_{T,m}(\omega^t)$. Similar to the ORMP, it is easy to compute $v_{T,m}(\omega^t) \equiv \inf_{\omega\in\Omega_{T,m}(\omega^t)} v^*(\omega)$ for the HLBP. The optimal value $v^*(\omega)$ with perfect information can be computed for any ω using Proposition 4.14. Next, for any partial instance ω^t, the best instance $\tilde{\omega} \in \Omega_{T,m}$ that starts with ω^t can be determined, and $v_{T,m}(\omega^t)$ can be computed, as shown in Proposition 4.16.

PROPOSITION 4.14. *For any instance $\omega = (a(1), \ldots, a(T))$ of the HLBP, the optimal value $v^*(\omega)$ with perfect information is given by*

$$v^*(\omega) \quad = \quad \max_{\{i\in\{1,\ldots,m\}\}} \left\{ \frac{1}{m-i+1} \sum_{t=1}^{T} \sum_{j=i}^{m} a_j(t) \right\}$$

The proof of Proposition 4.14 is similar to that of Proposition 4.1, and is omitted.

COROLLARY 4.15. *The optimal value $v^*(\omega)$ with perfect information for the HLBP is nondecreasing, that is, for any $\omega, \omega' \in \Omega_{T,m}$ with $\omega \leq \omega'$, it holds that $v^*(\omega) \leq v^*(\omega')$.*

It follows from Corollary 4.15 that for any partial instance $\omega^t = (a(1), \ldots, a(t))$, instance $(a(1), \ldots, a(t), 0, \ldots, 0)$ has the best optimal value $v^*(\tilde{\omega})$ with perfect information among all $\tilde{\omega} \in \Omega_{T,m}$ that start with ω^t. This leads to the following result, which lets us compute $v_{T,m}(\omega^t) \equiv \inf_{\tilde{\omega}\in\Omega_{T,m}(\omega^t)} v^*(\tilde{\omega})$.

PROPOSITION 4.16. *For any partial instance $\omega^t = (a(1), \ldots, a(t)) \in \Omega^t_{T,m}$ of the HLBP, $v_{T,m}(\omega^t) = v^*(a(1), \ldots, a(t), 0, \ldots, 0)$. That is,*

$$v_{T,m}(\omega^t) \quad = \quad \max_{\{i\in\{1,\ldots,m\}\}} \left\{ \frac{1}{m-i+1} \sum_{\tau=1}^{t} \sum_{j=i}^{m} a_j(\tau) \right\}$$

Note that $v_{T_1,m}(\omega^t) = v_{T_2,m}(\omega^t)$ for any $T_1, T_2 \geq t$ and any ω^t. Corollary 4.15 and Proposition 4.16 lead to the following result, which is used to compute $\rho^*_{T,m}$.

COROLLARY 4.17. *For any partial instances $\omega^t_1, \omega^t_2 \in \Omega^t_{T,m}$ with $\omega^t_1 \leq \omega^t_2$, it holds that $v_{T,m}(\omega^t_1) \leq v_{T,m}(\omega^t_2)$. Specifically, for any instance $\omega \in \Omega_{T,m}$, $v_{T,m}(\omega^t)$*

is nondecreasing in t. It follows that, for any α-policy and any instance $\omega \in \Omega_{T,m}$, $\pi_\alpha(\omega)(t) \equiv \alpha_{T,m} v_{T,m}(\omega^t)$ *is nondecreasing in t.*

4.2.2. Optimal Parameter Values. Next we address the determination of the optimal value of $\alpha_{T,m}$. For any $\omega \in \Omega$, the feasible region $\mathcal{R}(\omega)$ is determined by linear constraints (2.8), (2.9), and (2.10). Thus $\mathcal{R}(\omega)$ is a polyhedron, and is closed. Therefore it follows from Corollary 3.4 and Proposition 3.5 that the optimal value of $\alpha_{T,m}$ is $\rho^*_{T,m}$, and that the α-policy with parameter $\alpha_{T,m} = \rho^*_{T,m}$ is feasible and optimal among all deterministic online algorithms. (It is shown in Section 5.2, Proposition 5.4, that the α-policy with parameters $\alpha_{T,m} = \rho^*_{T,m}$ is optimal among all randomized online algorithms.)

Next we show how $\rho^*_{T,m}$ is given by the optimal value of an integer program, and by a linear program in the case with $T \geq m$. For any instance $\omega = (a(1), \ldots, a(T)) \in \Omega_{T,m}$ and any machine i, let

$$\tau(i) \quad \equiv \quad \max\left\{t \,:\, \sum_{j=1}^{i} a_j(t) > 0\right\}$$

denote the last time that work arrives for which machine i is eligible. $\tau(i) = 0$ if there is no such time. Then the maximum amount of work that can be assigned to machine i under the α-policy is $\alpha_{T,m} v_{T,m}(\omega^{\tau(i)})$, and thus the maximum amount of work that can be assigned under the α-policy is $\alpha_{T,m} \sum_{i=1}^m v_{T,m}(\omega^{\tau(i)})$. In Lemma 4.18 we show that, if the parameter $\alpha_{T,m}$ is too small, then there exists an instance $\omega = (a(1), \ldots, a(T))$ such that the maximum amount of work that can be assigned under the α-policy, $\alpha_{T,m} \sum_{i=1}^m v_{T,m}(\omega^{\tau(i)})$, is less than the total amount of work to be performed,
$\sum_{t=1}^T \sum_{i=1}^m a_i(t)$.

LEMMA 4.18. *Suppose that the α-policy with parameter* $\alpha_{T,m}$ *gives an infeasible solution with instance* $\omega' \in \Omega_{T,m}$. *Then there exists an instance* $\omega = (a(1), \ldots, a(T)) \in \Omega_{T,m}$ *such that*

$$\alpha_{T,m} \sum_{i=1}^{m} v_{T,m}(\omega^{\tau(i)}) \quad < \quad \sum_{t=1}^{T}\sum_{i=1}^{m} a_i(t) \tag{4.10}$$

where $\tau(i) \equiv \max\left\{t \,:\, \sum_{j=1}^{i} a_j(t) > 0\right\}$.

Proof. Instance ω is constructed from ω' by discarding work that does not contribute to infeasibility. This is done by inductively constructing a sequence of instances $\omega^m, \ldots, \omega^1$. Let $a^i_j(t)$ denote the amount of work in instance ω^i that arrives at time t that requires at least machine j. As before, $\omega^{i,t}$ denotes the first t components of instance ω^i, and $v_{T,m}(\omega^{i,0}) \equiv 0$. Let $q^i_j(t)$ denote the amount of work assigned to machine j at time t under the α-policy with parameter $\alpha_{T,m}$ and for instance ω^i. Let $\tau^i(j) \equiv \max\left\{t \,:\, \sum_{k=1}^{j} a^i_k(t) > 0\right\}$ denote the last time in instance ω^i that work arrives for which machine j is eligible, with $\tau^i(j) \equiv 0$ if $\sum_{t=1}^T \sum_{k=1}^j a^i_k(t) = 0$. We know that $\sum_{t=1}^T \sum_{j=1}^m q'_j(t) < \sum_{t=1}^T \sum_{j=1}^m a'_j(t)$, so we will create a sequence of instances where $\sum_{t=1}^T q_j(t) = \alpha_{T,m} v_{T,m}(\omega^{\tau(j)})$ for each $j = m, \ldots, 1$. Specifically, it is shown that

$$\sum_{t=1}^{T}\sum_{j=1}^{i-1} q^i_j(t) + \alpha_{T,m} \sum_{j=i}^{m} v_{T,m}(\omega^{i,\tau^i(j)}) \quad < \quad \sum_{t=1}^{T}\sum_{j=1}^{m} a^i_j(t) \tag{4.11}$$

for all $i = m, \ldots, 1$, which implies (4.10) with $\omega = \omega^1$.

Let

$$t(m) \quad \equiv \quad \min \left\{ t' \,:\, \sum_{j=1}^{m} q'_j(t') < \sum_{j=1}^{m} a'_j(t') \right\}$$

Such a t' must exist because the α-policy with parameter $\alpha_{T,m}$ gives an infeasible solution with instance ω'.

Since at time $t(m)$, the allocation $r(t(m))$ is insufficient to complete all work available at $t(m)$, it must be that machine m is being 'fully used', i.e. $q'_m(t(m)) = \alpha_{T,m} v_{T,m}(\omega'^{t(m)}) - \sum_{t<t(m)} q'_m(t)$.

Thus,

$$\sum_{t=1}^{t(m)} \sum_{j=1}^{m-1} q'_j(t) + \alpha_{T,m} v_{T,m}(\omega'^{t(m)}) \quad < \quad \sum_{t=1}^{t(m)} \sum_{j=1}^{m} a'_j(t)$$

Instance ω^m is obtained from ω' by discarding work arriving after time $t(m)$, that is, $a^m_j(t) = 0$ for all j and all $t > t(m)$, and $a^m_j(t) = a'_j(t)$ for all j and all $t \le t(m)$. Hence $\tau^m(m) = t(m)$, and

$$\sum_{t=1}^{T} \sum_{j=1}^{m-1} q^m_j(t) + \alpha_{T,m} v_{T,m}(\omega^{m,\tau^m(m)}) \quad < \quad \sum_{t=1}^{T} \sum_{j=1}^{m} a^m_j(t)$$

which shows (4.11) for $i = m$.

As induction hypothesis, assume that (4.11) holds for some $i \in \{2, \ldots, m\}$. Let

$$t(i-1) \quad \equiv \quad \max \left\{ t \,:\, \sum_{\tau=1}^{t} q^i_{i-1}(\tau) = \alpha_{T,m} v_{T,m}(\omega^{i,t}) \right\}$$

if such a t exists; otherwise $t(i-1) \equiv 0$. Note that $t(i-1)$ is the latest time that machine $i-1$ did the maximum amount of work it could, so later work that requests machines $i-1$ or lower does not contribute to the infeasibility of the instance. Instance ω^{i-1} is obtained from ω^i by discarding work arriving after time $t(i-1)$ that requests machines $j \le i-1$, that is, $a^{i-1}_j(t) = 0$ for all $j \le i-1$ and $t > t(i-1)$, and $a^{i-1}_j(t) = a^i_j(t)$ for all other j and t.

In the new instance, machine $i-1$ is 'fully used', so we have $\sum_{t=1}^{T} q^{i-1}_{i-1}(t) = \alpha_{T,m} v_{T,m}(\omega^{i,t(i-1)})$. Note that the total amount of work done by machines 1 through $i-1$ was reduced by the amount actually done after time $t(i-1)$, so we have

$$\sum_{t=1}^{T} \sum_{j=1}^{i-2} q^{i-1}_j(t) + \alpha_{T,m} v_{T,m}(\omega^{i,t(i-1)}) = \sum_{t=1}^{T} \sum_{j=1}^{i-1} q^i_j(t) - \sum_{t=t(i-1)+1}^{T} \sum_{j=1}^{i-1} a^i_j(t)$$

The final ingredients in the proof are inequalities relating $v_{T,m}$ for the various instances and truncated instances under consideration. Note that $\omega^{i-1} \le \omega^i$ implies $\omega^{i-1,t} \le \omega^{i,t}$ as well as $\tau^{i-1}(j) \le \tau^i(j)$. In addition, we know that $\tau^{i-1}(i-1) \le t(i-1)$. Putting all of this together with Corollary 4.17 gives

$$v_{T,m}(\omega^{i-1,\tau^{i-1}(j)}) \le v_{T,m}(\omega^{i-1,\tau^i(j)}) \le v_{T,m}(\omega^{i,\tau^i(j)})$$

and

$$v_{T,m}(\omega^{i-1,\tau^{i-1}(i-1)}) \leq v_{T,m}(\omega^{i-1,t(i-1)}) \leq v_{T,m}(\omega^{i,t(i-1)}).$$

Therefore, we can establish the induction hypothesis as follows:

$$\begin{aligned}
&\sum_{t=1}^{T}\sum_{j=1}^{i-2} q_j^{i-1}(t) + \alpha_{T,m}\sum_{j=i-1}^{m} v_{T,m}(\omega^{i-1,\tau^{i-1}(j)}) \\
&\leq \sum_{t=1}^{T}\sum_{j=1}^{i-2} q_j^{i-1}(t) + \alpha_{T,m} v_{T,m}(\omega^{i,t(i-1)}) + \alpha_{T,m}\sum_{j=i}^{m} v_{T,m}(\omega^{i,\tau^{i}(j)}) \\
&= \sum_{t=1}^{T}\sum_{j=1}^{i-1} q_j^{i}(t) - \sum_{t=t(i-1)+1}^{T}\sum_{j=1}^{i-1} a_j^{i}(t) + \alpha_{T,m}\sum_{j=i}^{m} v_{T,m}(\omega^{i,\tau^{i}(j)}) \\
&< \sum_{t=1}^{T}\sum_{j=1}^{m} a_j^{i}(t) - \sum_{t=t(i-1)+1}^{T}\sum_{j=1}^{i-1} a_j^{i}(t) \\
&= \sum_{t=1}^{T}\sum_{j=1}^{m} a_j^{i-1}(t)
\end{aligned}$$

□

Theorem 4.19. *Parameter $\alpha_{T,m}$ for the α-policy is too small ($\alpha_{T,m} < \rho^*_{T,m}$) if and only if there exists an instance $\omega = (a(1), \ldots, a(T)) \in \Omega_{T,m}$ such that $\alpha_{T,m}\sum_{i=1}^{m} v_{T,m}(\omega^{\tau(i)}) < \sum_{t=1}^{T}\sum_{i=1}^{m} a_i(t)$. Also,*

$$\rho^*_{T,m} = \inf\left\{\rho \in [1,\infty) : \rho\sum_{i=1}^{m}\max_{\{j\in\{1,\ldots,m\}\}}\left\{\frac{1}{m-j+1}\sum_{t=1}^{\tau(i)}\sum_{k=j}^{m} a_k(t)\right\} \geq \sum_{t=1}^{T}\sum_{i=1}^{m} a_i(t) \quad \forall\ (a(1),\ldots,a(T)) \in \mathbb{R}_+^{Tm}\right\} \quad (4.12)$$

The α-policy provides the optimal solution with perfect information for the zero instance $\omega = (0, \ldots, 0)$. Thus, to determine $\rho^*_{T,m}$, one can restrict attention to instances $(a(1), \ldots, a(T)) \in \mathbb{R}_+^{Tm}$ for which $\sum_{t=1}^{T}\sum_{i=1}^{m} a_i(t) > 0$. From (4.12) it follows that $\rho^*_{T,m}$ is determined by instances $(a(1), \ldots, a(T)) \in \mathbb{R}_+^{Tm}$ that minimize

$$\sum_{i=1}^{m}\max_{\{j\in\{1,\ldots,m\}\}}\left\{\frac{1}{m-j+1}\sum_{t=1}^{\tau(i)}\sum_{k=j}^{m}\frac{a_k(t)}{\sum_{t'=1}^{T}\sum_{i'=1}^{m} a_{i'}(t')}\right\}$$

It is also clear that one can restrict attention to instances $(a(1), \ldots, a(T)) \in \mathbb{R}_+^{Tm}$ for which $\sum_{t=1}^{T}\sum_{i=1}^{m} a_i(t) = 1$.

Corollary 4.20. *The optimal competitive ratio $\rho^*_{T,m}$ for HLBP can be calculated by solving the following integer linear program with decision variables $a_i(t)$, $x(t)$, y_i,* *and* $z_i(t)$, $i \in \{1, \dots, m\}$, $t \in \{1, \dots, T\}$*:*

$$
\begin{array}{llll}
\textit{(IP)} & \textit{minimize} & \displaystyle\sum_{i=1}^{m} y_i & \\
& \textit{subject to} & \displaystyle\sum_{t=1}^{T}\sum_{i=1}^{m} a_i(t) = 1 & \\
& & \displaystyle x(t) \geq \frac{1}{m-i+1}\sum_{\tau=1}^{t}\sum_{j=i}^{m} a_j(\tau) & \forall\ i, \ \ \forall\ t \\
& & y_i \geq x(t) + z_i(t) - 1 & \forall\ i, \ \ \forall\ t \\
& & \displaystyle z_i(t) \geq \sum_{j=1}^{i} a_j(t) & \forall\ i, \ \ \forall\ t \\
& & a_i(t) \geq 0 & \forall\ i, \ \ \forall\ t \\
& & y_i \geq 0 & \forall\ i \\
& & z_i(t) \in \{0, 1\} & \forall\ i, \ \ \forall\ t
\end{array}
$$

Proof. In an optimal solution (a^*, x^*, y^*, z^*) of the IP, $a^* = (a^*(1), \dots, a^*(T))$ represents a worst-case instance in $\Omega_{T,m}$ for the α-policy. Without loss of generality we can assume that $x^*(t) = v_{T,m}(a^*(t))$, which is at most 1 for all t. It follows from Corollary 4.17 that $x^*(t)$ is nondecreasing in t, and thus each y_i^* is equal to $x^*(t)$ for the largest value of t for which $z_i^*(t) = 1$, that is, y_i^* represents the corresponding value of $v_{T,m}(a^*(\tau(i)))$. Also, it follows from (4.12) that

$$
\rho^*_{T,m} \quad = \quad \inf\left\{\rho \in [1, \infty) \ : \ \rho \sum_{i=1}^{m} y_i^* \geq 1\right\}
$$

That is, $\rho^*_{T,m} = \max\{1, 1/\sum_{i=1}^m y_i^*\}$. Actually, $\rho^*_{T,m} = 1/\sum_{i=1}^m y_i^*$, as shown next. If $T \geq m$, then it is easily checked that (a, x, y, z) with

$$
\begin{aligned}
a_i(i) &= 1/m \text{ for } i \in \{1, \dots, m\} \\
a_i(t) &= 0 \text{ otherwise} \\
x(t) &= \begin{cases} t/m^2 & \text{for } t \leq m \\ 1/m & \text{for } t \geq m \end{cases} \\
y_i &= i/m^2 \text{ for } i \in \{1, \dots, m\} \\
z_i(t) &= \begin{cases} 1 & \text{for } t \leq i \\ 0 & \text{for } t > i \end{cases}
\end{aligned}
$$

is a feasible solution for the IP. Thus $\sum_{i=1}^m y_i^* \leq \sum_{i=1}^m i/m^2 = (1+1/m)/2$. Similarly,

T	m	$\rho^*_{T,m}$
2	4	4/3
2	5	4/3
2	6	4/3
3	4	3/2
3	5	1.511629
3	6	1.511629
3	9	1.520549
3	14	1.522063
3	19	1.522063
4	5	1.629631
4	6	1.629631
4	7	1.630138
4	8	1.630138

TABLE 4.2
Values of $\rho^*_{T,m}$ for the HLBP.

if $T < m$, then it is easily checked that (a, x, y, z) with

$$\begin{aligned} a_{m-T+t}(t) &= 1/T \text{ for } t \in \{1, \ldots, T\} \\ a_i(t) &= 0 \text{ otherwise} \\ x(t) &= t/T^2 \text{ for } t \in \{1, \ldots, T\} \\ y_i &= \begin{cases} 0 & \text{for } i \in \{1, \ldots, m-T\} \\ (i-m+T)/T^2 & \text{for } i \in \{m-T+1, \ldots, m\} \end{cases} \\ z_i(t) &= \begin{cases} 1 & \text{for } t \leq i-m+T \\ 0 & \text{for } t > i-m+T \end{cases} \end{aligned}$$

is a feasible solution for the IP. Thus $\sum_{i=1}^m y_i^* \leq \sum_{i=m-T+1}^m (i-m+T)/T^2 = \sum_{t=1}^T t/T^2 = (1+1/T)/2$. Therefore $\rho^*_{T,m} = 1/\sum_{i=1}^m y_i^* \geq 2/(1+1/\min\{T,m\}) \geq 1$. □

Table 4.2 gives the values of $\rho^*_{T,m}$ for some values of T and m, as determined by the IP.

There are a total of $2Tm + T + m$ decision variables and $3Tm + 1$ constraints, so the IP grows fairly rapidly. Next it is shown that if $T \geq m$, then $\rho^*_{T,m}$ for the HLBP is equal to ρ^*_m for the ORMP, which can be computed with the linear program given in Section 4.1.2. First Lemma 4.21 shows that, to determine $\rho^*_{T,m}$, one can restrict attention to instances in which at most one machine is requested in each time period and in which machines are requested from least versatile to most versatile.

LEMMA 4.21. *Suppose that the α-policy with parameter $\alpha_{T,m}$ gives an infeasible solution with instance $\omega' = (a'(1), \ldots, a'(T)) \in \Omega_{T,m}$. Then there exists an instance $\hat{\omega} = (\hat{a}(1), \ldots, \hat{a}(\lambda)) \in \Omega_{\lambda,m}$ such that the following hold.*

1. *For each t, let $m'(t) \equiv \min\{i : a'_i(t) > 0\}$ (ignoring all t s.t. $\sum_i a'_i(t) = 0$). Then $\lambda \equiv |\{m'(1), \ldots, m'(T)\}| \leq \min\{T, m\}$.*
2. *For all $t \in \{1, \ldots, \lambda\}$, $\hat{a}_i(t) > 0$ for exactly one i.*
3. *Let $\hat{m}(t) \equiv \min\{i : \hat{a}_i(t) > 0\}$, that is, $\hat{m}(t)$ is the unique i such that $\hat{a}_i(t) > 0$. Then $\hat{m}$ is strictly increasing.*

4.

$$\alpha_{T,m}\sum_{i=1}^{m} v_{\lambda,m}(\hat{\omega}^{\hat{\tau}(i)}) \quad < \quad \sum_{t=1}^{\lambda}\hat{a}_{\hat{m}(t)}(t)$$

where $\hat{\tau}(i) \equiv \max\left\{t \,:\, \sum_{j=1}^{i}\hat{a}_j(t) > 0\right\}$.

Proof. It was shown in Lemma 4.18 that if the α-policy with parameter $\alpha_{T,m}$ gives an infeasible solution with instance $\omega' \in \Omega_{T,m}$, then there exists an instance $\omega = (a(1), \ldots, a(T)) \in \Omega_{T,m}$ such that

$$\alpha_{T,m}\sum_{i=1}^{m} v_{T,m}(\omega^{\tau(i)}) \quad < \quad \sum_{t=1}^{T}\sum_{i=1}^{m} a_i(t)$$

For each t, let $m(t) \equiv \min\{i \,:\, a_i(t) > 0\}$ denote the least versatile machine requested at time t (simply ignore all t such that $\sum_{i=1}^{m} a_i(t) = 0$). For each t, let $\tilde{a}_{m(t)}(t) \equiv \sum_{i=1}^{m} a_i(t)$, and $\tilde{a}_i(t) = 0$ for all $i \neq m(t)$, that is, at each time t, all the work requests machine $m(t)$. Note that, for each i, $\tilde{\tau}(i) = \tau(i)$, and for all i and all t, $\sum_{j=i}^{m}\tilde{a}_j(t) \leq \sum_{j=i}^{m} a_j(t)$. Thus it follows that $v_{T,m}(\tilde{\omega}^{\tilde{\tau}(i)}) \leq v_{T,m}(\omega^{\tau(i)})$ for all i. Also, $\sum_{t=1}^{T}\sum_{i=1}^{m}\tilde{a}_i(t) = \sum_{t=1}^{T}\sum_{i=1}^{m} a_i(t)$, from which we have

$$\alpha_{T,m}\sum_{i=1}^{m} v_{T,m}(\tilde{\omega}^{\tilde{\tau}(i)}) \quad < \quad \sum_{t=1}^{T}\sum_{i=1}^{m}\tilde{a}_i(t)$$

¿From the construction of $\tilde{\omega}$ it follows that $\tilde{m}(t) \equiv \min\{i \,:\, \tilde{a}_i(t) > 0\} = m'(t)$, that is, $\tilde{m}(t) = m'(t)$ is the unique i such that $\tilde{a}_i(t) > 0$. For each $t \in \{1, \ldots, \lambda\}$, let $i(t)$ be the tth smallest element in $\{\tilde{m}(1), \ldots, \tilde{m}(T)\}$. Let $\hat{a}_{i(t)}(t) \equiv \sum_{\tau=1}^{T}\tilde{a}_{i(t)}(\tau)$, and $\hat{a}_i(t) = 0$ for all $i \in \{1, \ldots, m\}$ and all $t \in \{1, \ldots, \lambda\}$ with $i \neq i(t)$. Thus $\hat{m}(t) = i(t)$, which is by definition strictly increasing in t. For any $i, j \in \{1, \ldots, m\}$,

$$\sum_{t=1}^{\hat{\tau}(i)}\sum_{k=j}^{m}\hat{a}_k(t) \quad = \quad \sum_{t=1}^{\lambda}\sum_{k=j}^{i}\hat{a}_k(t) \quad = \quad \sum_{t=1}^{T}\sum_{k=j}^{i}\tilde{a}_k(t) \quad = \quad \sum_{t=1}^{\tilde{\tau}(i)}\sum_{k=j}^{i}\tilde{a}_k(t) \quad \leq \quad \sum_{t=1}^{\tilde{\tau}(i)}\sum_{k=j}^{m}\tilde{a}_k(t)$$

and thus

$$\begin{aligned}
v_{\lambda,m}(\hat{\omega}^{\hat{\tau}(i)}) &= \max_{\{j\in\{1,\ldots,m\}\}}\left\{\frac{1}{m-j+1}\sum_{t=1}^{\hat{\tau}(i)}\sum_{k=j}^{m}\hat{a}_k(t)\right\} \\
&\leq \max_{\{j\in\{1,\ldots,m\}\}}\left\{\frac{1}{m-j+1}\sum_{t=1}^{\tilde{\tau}(i)}\sum_{k=j}^{m}\tilde{a}_k(t)\right\} \\
&= v_{T,m}(\tilde{\omega}^{\tilde{\tau}(i)})
\end{aligned}$$

Therefore

$$\alpha_{T,m}\sum_{i=1}^{m} v_{\lambda,m}(\hat{\omega}^{\hat{\tau}(i)}) \quad \leq \quad \alpha_{T,m}\sum_{i=1}^{m} v_{T,m}(\tilde{\omega}^{\tilde{\tau}(i)}) \quad < \quad \sum_{t=1}^{T}\sum_{i=1}^{m}\tilde{a}_i(t) \quad = \quad \sum_{t=1}^{\lambda}\hat{a}_{\hat{m}(t)}(t)$$

□

It follows from Lemma 4.21 that if the α-policy with parameter $\alpha_{T,m}$ gives an infeasible solution with instance $\omega' \in \Omega_{T,m}$, then there exists an instance $\omega = (a(1), \ldots, a(m)) \in \Omega_{m,m}$ such that $a_i(t) = 0$ for all $i \neq t$, and

$$\alpha_{T,m} \sum_{i=1}^{m} v_{m,m}(\omega^i) \quad < \quad \sum_{i=1}^{m} a_i(i)$$

This observation leads to Theorem 4.22.

THEOREM 4.22. *Parameter $\alpha_{T,m}$, with $T \geq m$, for the α-policy is too small ($\alpha_{T,m} < \rho^*_{T,m}$) if and only if there exists an instance $\omega = (a(1), \ldots, a(m)) \in \Omega_{m,m}$ such that $a_i(t) = 0$ for all $i \neq t$, and $\alpha_{T,m} \sum_{i=1}^{m} v_{m,m}(\omega^i) < \sum_{i=1}^{m} a_i(i)$. Also, for $T \geq m$,*

$$\rho^*_{T,m} \quad = \quad \rho^*_{m,m} = \inf \left\{ \rho \in [1, \infty) \, : \, \rho \sum_{i=1}^{m} \max_{\{j \in \{1,\ldots,i\}\}} \left\{ \frac{1}{m-j+1} \sum_{k=j}^{i} a(k) \right\} \right. \\ \left. \geq \sum_{i=1}^{m} a(i) \quad \forall \ (a(1), \ldots, a(m)) \in \mathbb{R}^m_+ \right\} \quad (4.13)$$

Corollary 4.23 follows by comparing (4.13) and (4.2).

COROLLARY 4.23. *For $T \geq m$, $\rho^*_{T,m}$ for the HLBP is equal to ρ^*_m for the ORMP, which can be computed with the linear program (LP).*

4.2.3. Asymptotic Behavior. For the HLBP, we are interested in $\rho^* = \sup_{T \in \mathbb{Z}_+, m \in \mathbb{Z}_+} \rho^*_{T,m}$, as well as $\rho^*_{\infty,m} \equiv \sup_{T \in \mathbb{Z}_+} \rho^*_{T,m}$, and $\rho^*_{T,\infty} \equiv \sup_{m \in \mathbb{Z}_+} \rho^*_{T,m}$. As pointed out in Section 2.3, $\rho^*_{T,m}$ is nondecreasing in T and m, and thus $\rho^* = \lim_{m \to \infty} \rho^*_{m,m}$. Thus it follows from Corollary 4.23 that ρ^* for the HLBP is equal to ρ^* for the ORMP, which is equal to e. It also follows that $\sup_{T \in \mathbb{Z}_+} \rho^*_{T,m} = \lim_{T \to \infty} \rho^*_{T,m} = \rho^*_{m,m}$, which is equal to ρ^*_m for the ORMP. Also, $\sup_{m \in \mathbb{Z}_+} \rho^*_{T,m} = \lim_{m \to \infty} \rho^*_{T,m}$; however, this asymptotic behavior is not well understood yet, except for the bounds $\lim_{m \to \infty} \rho^*_{T,m} \geq \rho^*_T$ for the ORMP, and $\lim_{m \to \infty} \rho^*_{T,m} \leq \rho^* = e$. By using a discrete time, continuous machine model of the HLBP, it can be shown that $\lim_{m \to \infty} \rho^*_{2,m} = \rho^*_2 = 4/3$. It also follows from a comparison of Tables 4.1 and 4.2 that $\lim_{m \to \infty} \rho^*_{T,m} > \rho^*_T$ for some T. We conjecture that $\lim_{m \to \infty} \rho^*_{T,m} < \lim_{m \to \infty} \rho^*_{T+1,m}$, which would imply that $\lim_{m \to \infty} \rho^*_{T,m} < \rho^* = e$ for all T.

4.3. Generalizations for the OMMP. In this section we briefly show how some of the results in Section 4.1 for the ORMP can be generalized for the OMMP.

Let $\theta : \Omega \mapsto \Omega$ be any function such that for any $\omega \in \Omega_T$, $\theta(\omega) \in \Omega_T$. For any partial instance $\theta(\omega)^t$, let $\Omega^\theta_T(\theta(\omega)^t)$ denote the set of all instances in $\theta(\Omega_T)$ with first t elements equal to $\theta(\omega)^t$. Let

$$v^\theta_T(\theta(\omega)^t) \quad \equiv \quad \inf_{\omega \in \Omega^\theta_T(\theta(\omega)^t)} v^*(\omega)$$

Assume that $v^\theta_T(\theta(\omega)^t) \leq v_T(\omega^t)$ for all $t \in \{1, \ldots, T\}$.

Let $\vartheta : \Omega \mapsto \Omega$ be any function such that for any $\omega \in \Omega_T$, $\vartheta(\omega) \in \Omega_{T-1}$. For any partial instance $\vartheta(\omega)^t$, $t \leq T-1$, let $\Omega^\vartheta_T(\vartheta(\omega)^t)$ denote the set of all instances in $\vartheta(\Omega_T)$ with first t elements equal to $\vartheta(\omega)^t$. Let

$$v^\vartheta_T(\vartheta(\omega)^t) \quad \equiv \quad \inf_{\omega \in \Omega^\vartheta_T(\vartheta(\omega)^t)} v^*(\omega)$$

Assume that $v_T^{\vartheta}(\vartheta(\omega)^t) \leq v_T(\omega^t)$ for all $t \in \{1, \dots, T-1\}$.

Thus both θ and ϑ map an instance ω to instances $\theta(\omega)$ and $\vartheta(\omega)$ that are "smaller" than ω. As before, for any $r \in \mathbb{R}_+^T$, r^{T-1} denotes the first $T-1$ components of r.

DEFINITION 4.24. (extensibility property): We say that $(\mathcal{R}, \theta, \vartheta)$ has the extensibility property if, for any $\omega \in \Omega_T$ and any $r \in \mathbb{R}_+^T$, $r \in \mathcal{R}(\theta(\omega))$ and $r^{T-1} \in \mathcal{R}(\vartheta(\omega))$, imply that $r \in \mathcal{R}(\omega)$.

Intuitively, the extensibility property states that although instances $\theta(\omega)$ and $\vartheta(\omega)$ are smaller than ω, together the feasibility of a solution r for $\theta(\omega)$ and r^{T-1} for $\vartheta(\omega)$ are sufficient to establish the feasibility of r for ω.

Theorem 4.25 follows along the lines of Theorem 4.9.

THEOREM 4.25. *Suppose that $(\mathcal{R}, \theta, \vartheta)$ has the extensibility property. If the α-policy with parameter α gives an infeasible solution for some instance ω and a feasible solution for instance $\vartheta(\omega)$, then the α-policy with parameter α gives an infeasible solution for instance $\theta(\omega)$.*

Next it is shown that if $(\mathcal{R}, \theta, \vartheta)$ has the extensibility property, then to determine ρ_T^*, it is sufficient to consider only the instances in the image $\theta(\Omega_T)$ of θ.

THEOREM 4.26. *Suppose that $(\mathcal{R}, \theta, \vartheta)$ has the extensibility property, and that ρ_T^* is nondecreasing in T. Then*

$$\rho_T^* \;=\; \inf_{\alpha \geq 1} \inf \left\{ \rho \geq 1 \,:\, v^{\pi_\alpha}(\omega) \leq \rho v^*(\omega) \quad \forall\, \omega \in \theta(\Omega_T) \right\}$$

That is, to determine ρ_T^* (and ρ^*), it is sufficient to consider only the α-policy, and only the instances in $\theta(\Omega_T)$.

5. Randomized Algorithms.

So far we have analyzed only deterministic algorithms. However, for some problems, randomized algorithms can have better competitive ratios than deterministic algorithms (Motwani and Raghavan [14], and Hoogeveen and Vestjens [9]), and we investigate that possibility here. First, we show that a randomized algorithm can have better competitive ratios for the OMMP than any deterministic algorithm. Second, we show that if the feasible set $\mathcal{R}(\omega)$ is convex for all ω, then randomized algorithms do not have better competitive ratios than deterministic algorithms.

5.1. Randomized Algorithms May Have Better Competitive Ratios.

We give an example of an OMMP for which a randomized algorithm has a better competitive ratio than any deterministic algorithm. The example is for an ORMP with strictly convex productivity function η.

PROPOSITION 5.1. *Consider the single deadline ORMP with a stationary productivity function of the form $\eta_t(r) = cr^p$, $c > 0$, $p > 1$, for all t. For this problem there exists a randomized algorithm $\pi \in \Pi^{RO}$ such that $v^{\pi}(\omega) < v^{\pi_\alpha}(\omega)$ for every $\omega \in \Omega_2$, $\omega > 0$.*

The idea is to choose a randomized algorithm π that prescribes almost the same decisions as the α-policy, except that it randomly chooses to do $\pm\varepsilon$ extra work in the first time period, and $\mp\varepsilon$ extra work in the second time period, where ε is a sufficiently small value. The solution is still feasible since the total amount of work performed is the same, and the maximum amount of work performed will be equal to the maximum amount under the α-policy, $\mp\epsilon$. When one takes η^{-1} to determine the amount of resource required to perform the work, and averages over the $+\varepsilon$ and $-\varepsilon$ outcomes, one gets a number that is lower than the maximum amount of resource under the α-policy, due to the strict convexity of η. The proof of Proposition 5.1 breaks into several tedious cases, and is given in the Appendix.

5.2. A Sufficient Condition for Optimality of Deterministic Algorithms. In this section we present a sufficient condition for the (deterministic online) α-policy to be optimal among all randomized online algorithms for the OMMP. Thereafter we apply the result to show that the α-policy is optimal among all randomized online algorithms for the ORMP with concave productivity functions η_t, as well as for the HLBP.

THEOREM 5.2. *Suppose that $\mathcal{R}(\omega)$ is convex for all $\omega \in \Omega_\beta$. For any algorithm $\pi \in \Pi^{RO}$, if $\rho^\pi_\beta < \infty$, then the α-policy with parameter $\alpha_\beta = \rho^\pi_\beta$ achieves the same competitive ratio, $\rho^{\pi_\alpha}_\beta = \rho^\pi_\beta$.*

Proof. We show that the α-policy with parameter $\alpha_\beta = \rho^\pi_\beta$ leads to feasible solutions for all instances $\omega \in \Omega_\beta$, by showing that $\pi_\alpha(\omega)(t) \geq E[\pi(\omega)(t)]$ for all $\omega \in \Omega_\beta$ and all t. For this, we mimic the first part of the proof of Theorem 3.2, replacing $\pi(\omega)(t)$ with $E[\pi(\omega)(t)]$.

Because $\rho^\pi_\beta < \infty$, it holds for all $\omega \in \Omega_\beta$ that $\pi(\omega)[\mathcal{R}(\omega)] = 1$, that is, with probability 1, the solution $(\pi(\omega)(1), \ldots, \pi(\omega)(T))$ is in $\mathcal{R}(\omega)$. Then, because $\mathcal{R}(\omega)$ is convex and $E[\pi(\omega)(t)] \leq \rho^\pi_\beta < \infty$ for all $\omega \in \Omega_\beta$ and all t, the solution $(E[\pi(\omega)(1)], \ldots, E[\pi(\omega)(T)])$ is in $\mathcal{R}(\omega)$. Thus it follows from feasibility monotonicity that $(\pi_\alpha(\omega)(1), \ldots, \pi_\alpha(\omega)(T)) \in \mathcal{R}(\omega)$, that is, the α-policy with $\alpha_\beta = \rho^\pi_\beta$ leads to feasible solutions for all $\omega \in \Omega_\beta$. Therefore $\rho^{\pi_\alpha}_\beta \leq \alpha_\beta = \rho^\pi_\beta$. Also, it follows from the definition of the α-policy and $\rho^\pi_\beta < \infty$ that $\rho^{\pi_\alpha}_\beta \geq \rho^\pi_\beta$. Thus $\rho^{\pi_\alpha}_\beta = \rho^\pi_\beta$. □

Next we show that the α-policy is optimal among all randomized online algorithms for the ORMP with quasi-concave productivity functions η_t, as well as for the HLBP. These results follow directly from applying Theorem 5.2 to the ORMP and the HLBP. Recall that for any quasi-concave function $f : \mathbb{R}^n \mapsto \mathbb{R}$ and for any l, the set $\{x \in \mathbb{R}^n : f(x) \geq l\}$ is convex.

PROPOSITION 5.3. *If the ORMP productivity functions η_t are quasi-concave for all t, then*

1. *the set of feasible solutions $\mathcal{R}(\omega)$ is convex for all ω, and*
2. *the α-policy with parameters $\alpha_T = \rho^*_T$ is optimal among all randomized online algorithms.*

PROPOSITION 5.4. *For the HLBP,*

1. *the set of feasible solutions $\mathcal{R}(\omega)$ is convex for all ω, and*
2. *the α-policy with parameters $\alpha_{T,m} = \rho^*_{T,m}$ is optimal among all randomized online algorithms.*

6. Conclusions and Further Research. The α-policy theory developed in this paper is a powerful tool for finding worst-case optimal algorithms for online min-max problems. It makes analysis easier by turning optimization questions into feasibility questions, and by focusing on properties such as feasibility monotonicity and extensibility.

Our work suggests several questions and possible extensions. For the ORMP, the solutions to the LP show that the convergence of ρ^*_T to ρ^* is quite slow. It would be nice to have a theoretical explanation of this. For the HLBP, the asymptotic values of $\rho^*_{T,m}$ as $m \to \infty$ are not known. Given the slow convergence for the ORMP, and the computational limitations of the IP, it may be difficult to make inferences about asymptotic values from known values of $\rho^*_{T,m}$. The HLBP can be generalized to nonlinear server hierarchies [6] such as rooted in-trees and general partial orders. Values of $\rho^*_{T,m}$ for these cases are unknown.

If the future is completely unknown, and the decision maker is risk-averse, then the online criterion optimized in this paper would be quite appropriate. If future arrivals are known in distribution, the problem would become a Markov decision process, and could be attacked with dynamic programming methods. We suspect that the most realistic models would involve an intermediate level of information. It would be interesting to see whether an algorithm which is a blend of the α-policy and other algorithms could perform well on such a model, or whether altogether new algorithms are needed.

REFERENCES

[1] S. Albers, *On randomized online scheduling*, in Proceedings of the 34th Annual ACM Symposium on the Theory of Computing (STOC), 2002, pp. 134–143.

[2] J. Aspnes, Y. Azar, A. Fiat, S. Plotkin, and O. Waarts., *On-line load balancing with applications to machine scheduling and virtual circuit routing*, Journal of the ACM, 44 (1997), pp. 486–504.

[3] Y. Azar, A. Z. Broder, and A. R. Karlin, *On-line load balancing*, Theoretical Computer Science, 130 (1994), pp. 73–84.

[4] Y. Azar, B. Kalyanasundaram, S. Plotkin, K. Pruhs, and O. Waarts, *Online load balancing of temporary tasks*, Journal of Algorithms, 22 (1997), pp. 93–110.

[5] Y. Azar, J. Naor, and R. Rom, *The competitiveness of on-line assignments*, Journal of Algorithms, 18 (1995), pp. 221–237.

[6] A. Bar-Noy, A. Freund, and J. Naor, *On-line load balancing in a hierarchical server topology*, in Algorithms—ESA '99, 7th Annual European Symposium, Prague, Czech Republic, July 16–18, J. Nesetril, ed., vol. 1643, Springer-Verlag, Berlin, Germany, 1999, pp. 77–88.

[7] P. Berman, M. Charikar, and M. Karpinski, *On-line load balancing for related machines*, in Proc. 5th Workshop on Algorithms and Data Structures, vol. 1272 of Lecture Notes in Computer Science, Springer, 1997, pp. 116–125.

[8] R. L. Graham, *Bounds for certain multiprocessing anomalies*, Bell System Technical Journal, 45 (1966), pp. 1563–1581.

[9] J. A. Hoogeveen and A. P. A. Vestjens, *A best possible deterministic on-line algorithm for minimizing maximum delivery time on a single machine*, SIAM Journal on Discrete Mathematics, 13 (2000), pp. 56–63.

[10] S. Irani and A. R. Karlin, *On online computation*, in Approximation Algorithms for NP-Hard Problems, D. S. Hochbaum, ed., PWS Publishing Company, Boston, MA, 1997, ch. 13.

[11] A. J. Kleywegt, V. S. Nori, M. W. P. Savelsbergh, and C. A. Tovey, *Dynamic and stochastic resource minimization*, Tech. Report 97-05, The Logistics Institute, School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, GA 30332-0205, 1997.

[12] ———, *Online resource minimization*, in Proceedings of the Tenth Annual ACM-SIAM Symposium on Discrete Algorithms, Baltimore, MD, January 1999, pp. 576–585.

[13] L. A. McGeoch and D. D. Sleator, eds., *On-line Algorithms, DIMACS: Series in Discrete Mathematics and Theoretical Computer Science*, vol. 7, American Mathematical Society, Association for Computing Machinery, 1992.

[14] R. Motwani and P. Raghavan, *Randomized Algorithms*, Cambridge University Press, New York, NY, 1995.

[15] F. R. and W. M., *Online scheduling revisited*, in Proceedings of the 8th Annual Europeann Symposium on Algorithms (ESA), 2000, pp. 202–210.

[16] S. S., *Online randomized multiprocessor scheduling*, Algorithmica, 28 (2000), pp. 173–216.

[17] B. Y., L. S., M.-S. J., S. J., and S. L., *Multiprocessor scheduling with rejection*, SIAM Journal on Discrete Mathematics, 13 (2000), pp. 64–78.

Appendix for Referees

Appendix A. Proof of Lemma 2.2.
Lemma 2.2 *If β is known beforehand, then*

$$\rho^* = \sup_{\beta \in \mathbf{B}} \rho^*_\beta$$

Proof. It remains to be shown that

$$\rho^* \equiv \inf_{\pi \in \Pi^{RO}} \sup_{\beta \in \mathbf{B}} \rho^\pi_\beta \leq \sup_{\beta \in \mathbf{B}} \inf_{\pi \in \Pi^{RO}} \rho^\pi_\beta \equiv \hat{\rho}$$

If $\hat{\rho} = \infty$, the result follows immediately. Suppose $\hat{\rho} < \infty$. Thus $\inf_{\pi \in \Pi^{RO}} \rho^\pi_\beta \leq \hat{\rho} < \infty$ for all β. Hence, for any $\beta \in \mathbf{B}$ and any $\varepsilon > 0$, there exists an algorithm $\pi^\varepsilon_\beta \in \Pi^{RO}$ such that $\rho^{\pi^\varepsilon_\beta}_\beta < \hat{\rho} + \varepsilon$. Choose algorithm π^ε to be the same as algorithm π^ε_β if $\omega \in \Omega_\beta$. Because β is known in advance, π^ε is an online algorithm. Then $\rho^{\pi^\varepsilon}_\beta < \hat{\rho} + \varepsilon$ for all $\beta \in \mathbf{B}$. Thus $\sup_{\beta \in \mathbf{B}} \rho^{\pi^\varepsilon}_\beta \leq \hat{\rho} + \varepsilon$, and hence $\inf_{\pi \in \Pi^{RO}} \sup_{\beta \in \mathbf{B}} \rho^\pi_\beta \leq \hat{\rho} + \varepsilon$ for any $\varepsilon > 0$. Therefore $\inf_{\pi \in \Pi^{RO}} \sup_{\beta \in \mathbf{B}} \rho^\pi_\beta \leq \hat{\rho}$. □

Appendix B. Proposition B.1.

PROPOSITION B.1. *For the ORMP, if the productivity function η_t is upper semi-continuous for all $t \in \{1, \ldots, T\}$, then the set of feasible solutions $\mathcal{R}(\omega)$ is closed for all instances $\omega \in \Omega_T$.*

Proof. Consider any instance $\omega \in \Omega_T$, and any limit point $\hat{r} = (\hat{r}(1), \ldots, \hat{r}(T))$ of $\mathcal{R}(\omega)$. It is to be shown that $\hat{r} \in \mathcal{R}(\omega)$. Because of the EDF assignment of resources to work, the only constraint that may be violated by solution $\hat{r}$ is (2.4). Consider any $\varepsilon > 0$. Because η_t is upper semicontinuous, there exists $\delta_t > 0$ such that $\eta_t(\hat{r}(t)) > \eta_t(r(t)) - \varepsilon/T$ for all $r(t)$ with $|\hat{r}(t) - r(t)| < \delta_t$. Let $\delta \equiv \min\{\delta_1, \ldots, \delta_T\}$. There exists $r^\delta = (r^\delta(1), \ldots, r^\delta(T)) \in \mathcal{R}(\omega)$ such that $\|\hat{r} - r^\delta\| < \delta$. Thus $|\hat{r}(t) - r^\delta(t)| < \delta_t$ and $\eta_t(\hat{r}(t)) > \eta_t(r^\delta(t)) - \varepsilon/T$ for all t. Let $\hat{w}_u(t)$ ($w^\delta_u(t)$) denote the amount of work with deadline u waiting at time t to be performed under solution $\hat{r}$ (r^δ) and EDF assignment, after the arrivals at time t have taken place, but before any work has been performed at time t. Let $\hat{W}_u(t) \equiv \hat{w}_1(t) + \cdots + \hat{w}_u(t)$, and let $W^\delta_u(t) \equiv w^\delta_1(t) + \cdots + w^\delta_u(t)$. It is shown by induction on t that $\hat{W}_u(t) < W^\delta_u(t) + \varepsilon t/T$ for all $t, u \in \{1, \ldots, T\}$, and that $\hat{w}_u(t) = 0$ for all $u < t$. For $t = 1$, $\hat{W}_u(1) = a_1(1) + \cdots + a_u(1) = W^\delta_u(1) < W^\delta_u(1) + \varepsilon/T$. As induction hypothesis, suppose that $\hat{W}_u(t) < W^\delta_u(t) + \varepsilon t/T$ for all $u \in \{1, \ldots, T\}$. For all $u < t$, $w^\delta_u(t) = 0$ because $r^\delta \in \mathcal{R}(\omega)$, and thus for all $u < t$, $\hat{W}_u(t) < W^\delta_u(t) + \varepsilon t/T = \varepsilon t/T$. Hence, for all $u < t$, $\hat{w}_u(t) < \varepsilon t/T$, for all $\varepsilon > 0$. Thus $\hat{w}_u(t) = 0$ for all $u < t$, $\hat{W}_u(t) = 0$, and $\hat{W}_u(t+1) = 0 < W^\delta_u(t+1) + \varepsilon(t+1)/T$ for all $u < t$. Consider $u \geq t$. If $\hat{W}_u(t) \leq \eta_t(\hat{r}(t))$, then $\hat{W}_u(t+1) = a_{t+1}(t+1) + \cdots + a_u(t+1) \leq W^\delta_u(t+1) < W^\delta_u(t+1) + \varepsilon(t+1)/T$. Otherwise, if $\hat{W}_u(t) > \eta_t(\hat{r}(t))$, then $\hat{W}_u(t+1) = \hat{W}_u(t) - \eta_t(\hat{r}(t)) + a_{t+1}(t+1) + \cdots + a_u(t+1) < W^\delta_u(t) + \varepsilon t/T - \eta_t(r^\delta(t)) + \varepsilon/T + a_{t+1}(t+1) + \cdots + a_u(t+1) \leq W^\delta_u(t+1) + \varepsilon(t+1)/T$, and the induction hypothesis has been established. Hence $\hat{w}_u(t) = 0$ for all $u < t$, and therefore $\hat{r} \in \mathcal{R}(\omega)$. □

Appendix C. Proof of Proposition 4.1.

Proposition 4.1 *For any instance $\omega = (a(1), \ldots, a(T))$ of the ORMP with nondecreasing productivity function $\eta_t(r)$, the optimal value $v^*(\omega)$ with perfect infor-*

mation is given by

$$v^*(\omega) \;=\; \inf\left\{r \geq 0 \,:\, \sum_{t=i}^{j} \eta_t(r) \geq \sum_{t=i}^{j}\sum_{u=t}^{j} a_u(t) \;\;\forall\, i,j \in \{1,\ldots,T\}, i \leq j\right\}$$

where $\inf \varnothing = \infty$.

Proof. Let

$$\gamma^*(\omega) \;\equiv\; \inf\left\{r \geq 0 \,:\, \sum_{t=i}^{j} \eta_t(r) \geq \sum_{t=i}^{j}\sum_{u=t}^{j} a_u(t) \;\;\forall\, i,j \in \{1,\ldots,T\}, i \leq j\right\}$$

It follows from the productivity function $\eta_t(r)$ being nondecreasing for all t that $v^*(\omega) \geq \gamma^*(\omega)$. It remains to show that $v^*(\omega) \leq \gamma^*(\omega)$. If $\gamma^*(\omega) = \infty$, the result follows immediately. Otherwise, fix any $\varepsilon > 0$. ¿From the definition of $\gamma^*(\omega)$, for any $\varepsilon > 0$ there exists an $r^\varepsilon \geq 0$ such that $r^\varepsilon < \gamma^*(\omega) + \varepsilon$ and $\sum_{t=i}^{j} \eta_t(r^\varepsilon) \geq \sum_{t=i}^{j}\sum_{u=t}^{j} a_u(t)$ for all $i, j \in \{1,\ldots,T\}, i \leq j$. Consider the solution that makes r^ε amount of resource available at each time t. Let $w(t)$ denote the total amount of work waiting to be processed after the arrivals at time t have taken place, but before any processing at time t. Thus if $w(t) \leq \eta_t(r^\varepsilon)$, then all $w(t)$ amount of work is performed at time t. Else, if $w(t) > \eta_t(r^\varepsilon)$, then $\eta_t(r^\varepsilon)$ amount of work is performed at time t, according to the EDF rule. Hence this solution never uses more than r^ε amount of resource at a time. The solution is feasible and has objective value $r^\varepsilon < \gamma^*(\omega) + \varepsilon$ 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 deadlines 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 deadlines 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 deadlines less than or equal to t that has to be processed in $\{\ell(t)+1,\ldots,t\}$ is $\sum_{\tau=\ell(t)+1}^{t}\sum_{u=\tau}^{t} a_u(\tau)$, which is less than or equal to $\sum_{\tau=\ell(t)+1}^{t} \eta_\tau(r^\varepsilon)$ from the definition of r^ε. However, from the definition of $\ell(t)$, there is work with deadlines less than or equal to t remaining at the end of each of the times in $\{\ell(t)+1,\ldots,t\}$, and thus from the definition of the solution with the EDF rule, $\sum_{\tau=\ell(t)+1}^{t} \eta_\tau(r^\varepsilon)$ amount of work with deadlines less than or equal to t is processed in $\{\ell(t)+1,\ldots,t\}$. Thus the amount of work with deadlines 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 $v^*(\omega) \leq r^\varepsilon < \gamma^*(\omega) + \varepsilon$ for arbitrary $\varepsilon > 0$. Therefore $v^*(\omega) \leq \gamma^*(\omega)$. □

Appendix D. Proof of Lemma 4.8.

Lemma 4.8 *For the ORMP with productivity function* $\eta_t(r)$, $(\mathcal{R}, \theta, \vartheta)$ *has the extensibility property.*

Proof. Consider any instance $\omega = (a_1(1), a_2(1), \ldots, a_T(T))$, $\theta(\omega) = (a'_1(1), a'_2(1), \ldots, a'_T(T))$, and $\vartheta(\omega) = (a''_1(1), a''_2(1), \ldots, a''_{T-1}(T-1))$. Consider any r such that $r \in \mathcal{R}(\theta(\omega))$ and $r^{T-1} \in \mathcal{R}(\vartheta(\omega))$. As usual, available work is performed in EDF order. Let $w_u(t) \equiv a_u(1) - q_u(1) + a_u(2) - q_u(2) + \cdots + a_u(t)$, $w'_T(t) \equiv a'_T(1) - q'_T(1) + a'_T(2) - q'_T(2) + \cdots + a'_T(t)$, and $w''_u(t) \equiv a''_u(1) - q''_u(1) + a''_u(2) - q''_u(2) + \cdots + a''_u(t)$.

It is shown by induction on t that $w_u(t) = w''_u(t)$ and $q_u(t) = q''_u(t)$ for all $t = 1,\ldots,T-1$, $u = t,\ldots,T-1$, and $w'_T(t) = \sum_{u=t}^{T} w_u(t)$ and $q'_T(t) = \sum_{u=t}^{T} q_u(t)$ for all $t = 1, \ldots, T$. Clearly the hypothesis holds for $t = 1$. Suppose the hypothesis holds for t. Note that $q_t(t) = q''_t(t) = w''_t(t) = w_t(t)$ from the assumption that $r^{T-1} \in \mathcal{R}(\vartheta(\omega))$. Then $w_u(t+1) = a_u(1) - q_u(1) + \cdots + a_u(t+1) = a''_u(1) - q''_u(1) + \cdots + a''_u(t+1) = w''_u(t+1)$ for all $u = t+1, \ldots, T-1$. Because available work is performed in EDF order, $q_u(t+1) = q''_u(t+1)$ for all $u = t+1, \ldots, T-1$. Also, $w'_T(t+1) = w'_T(t) - q'_T(t) + a'_T(t+1) = \sum_{u=t}^{T} w_u(t) - \sum_{u=t}^{T} q_u(t) + \sum_{u=t+1}^{T} a_u(t+1) = \sum_{u=t+1}^{T} (w_u(t) - q_u(t) + a_u(t+1)) = \sum_{u=t+1}^{T} w_u(t+1)$, and the hypothesis has been established.

¿From the assumption that $r^{T-1} \in \mathcal{R}(\vartheta(\omega))$, it follows that constraints (2.6), (2.3), (2.4) and (2.5) are satisfied by $\vartheta(\omega)$ and r^{T-1}, and thus by ω and r for all $t = 1, \ldots, T-1$. It remains to be shown that constraints (2.6), (2.3), (2.4) and (2.5) are satisfied by ω and r at $t = T$. ¿From the hypothesis and the assumption that $r \in \mathcal{R}(\theta(\omega))$, it follows that $w_T(T) = w'_T(T) \leq \eta_T(r(T))$, and thus constraints (2.6), (2.3), (2.4) and (2.5) are satisfied by ω and r at $t = T$. ▯

Appendix E. Proof of Proposition 5.1.

The proof of Proposition 5.1 is based on the use of a strictly convex productivity function with the ORMP. It is important to the proof that our techniques for determining the optimal competitve ratio ρ^*_T can be adapted for a certain class of such functions. Those results are presented in the next section, which is followed by the proof of Proposition 5.1.

E.1. Other productivity functions for the ORMP. The determination of the optimal competitve ratio for the ORMP in Section 4.1.2 applied to the case $\eta_t(r) = r$ for all $t = 1, \ldots, T$. In general, the analysis does not easily extend to other productivity functions, though there are a few easy cases. Assume that η_t is increasing, and therefore invertible.

The general equation for $q(t)$ is

$$q(t) = \eta_t(\alpha_T v_T(\omega^t)).$$

To "solve" for α_T as we did in Section 4.1.2, we would like to factor it out of the sum

$$\sum_{t=1}^{T} \eta_t(\alpha_T v_T(\omega^t)),$$

which is not possible in general, even if $\eta_t = \eta$ is stationary.

It is possible, however, for stationary productivity functions that satisfy $\eta(kr) = f(k)\eta(r)$ for some function $f(k)$ and all $k, r > 0$. This class of functions includes functions of the form $\eta(r) = cr^p$, where $p > 0$ (with $f(k) = k^p$). In these cases, we have

$$\sum_{t=1}^{T} \eta(\alpha_T v_T(\omega^t)) = f(\alpha_T) \sum_{t=1}^{T} \eta(v_T(\omega^t))$$

Note that for these productivity functions, which are stationary and nondecreasing, one can derive from Proposition 4.1 that

$$v_T(\omega^t) = \eta^{-1} \left(\max_{\{i \in \{1,\ldots,t\}\}} \frac{1}{T-i+1} \sum_{t=i}^{t} a(t) \right)$$

	Decision 1		Decision 2	
	Period 1	Period 2	Period 1	Period 2
$a(1) \leq a(2)$	$\eta^{-1}\left(\frac{2}{3}a(1)+\varepsilon\right)$	$\eta^{-1}\left(\frac{4}{3}a(2)-\varepsilon\right)$	$\eta^{-1}\left(\frac{2}{3}a(1)-\varepsilon\right)$	$\eta^{-1}\left(\frac{4}{3}a(2)+\varepsilon\right)$
$3\varepsilon \leq a(2) < a(1)$		$\eta^{-1}\left(\frac{2}{3}a(1)+\frac{2}{3}a(2)-\varepsilon\right)$		$\eta^{-1}\left(\frac{2}{3}a(1)+\frac{2}{3}a(2)+\varepsilon\right)$
$a(2) < 3\varepsilon$		$\eta^{-1}\left(\frac{2}{3}a(1)+\varepsilon\right)$		$\eta^{-1}\left(\frac{2}{3}a(1)-\varepsilon\right)$

TABLE E.1
Algorithm π for Proposition 5.1 ($\varepsilon = a(1)/15$)

Therefore, $\eta(v_T(\omega^t))$ is just the the maximization expression in the parentheses above.

If f is invertible (as it is for this class of functions), then one can write

$$\alpha_T \geq f^{-1}\left(\frac{\sum_{t=1}^{T} a(t)}{\sum_{t=1}^{T} \max_{\{i\in\{1,\ldots,t\}\}} \frac{1}{T-i+1}\sum_{\tau=i}^{t} a(\tau)}\right)$$

Assuming f^{-1} is monotonically increasing (as it is for this class of functions), one can still use the LP formulation (LP) to find the optimal value of ρ_T^*:

$$\rho_T^* = f^{-1}\left(\frac{1}{\text{optimal solution of LP}}\right)$$

In particular, this means that *any* stationary linear productivity function ($\eta(r) = cr$) has the same ρ_T^* values as previously determined. For a stationary productivity function of the form $\eta(r) = cr^p$, ρ_T^* is the pth root of the reciprocal of the optimal LP value. We know from Bar-Noy et al. [6] that the limiting LP value reciprocal is $\rho^* = e$, which means that as $p \to \infty$, the corresponding value of ρ^* goes to 1. Similarly, as $p \to 0$, the corresponding value of ρ^* increases without bound.

E.2. Proof of Proposition 5.1.

Proposition 5.1 *Consider the single deadline ORMP with a stationary productivity function of the form $\eta_t(r) = cr^p$, $c > 0$, $p > 1$, for all t. For this problem there exists a randomized algorithm $\pi \in \Pi^{RO}$ such that $v^{\pi}(\omega) < v^{\pi_\alpha}(\omega)$ for every $\omega \in \Omega_2$, $\omega > 0$.*

Proof. Note that $\eta(r) = cr^p$ with $c > 0$ and $p > 1$ is strictly convex, and strictly increasing on $[0, \infty)$.

¿From Section E.1, we know that for η of this form, $\rho_2^* = f^{-1}(4/3)$, where $f(k) = k^p$, so the optimal α-policy (and thus the optimal deterministic algorithm) will allocate $f^{-1}(4/3)\eta^{-1}(a(1)/2) = \eta^{-1}(2a(1)/3)$ in the first time period. Similarly, the optimal α-policy allocates $\max\left\{\eta^{-1}\left((4/3)(a(1)+a(2))/2\right), \eta^{-1}\left(4a(2)/3\right)\right\}$ in the second time period.

To show that the α-policy is not optimal among all randomized online algorithms for this problem, we construct a randomized algorithm π that achieves a better value than the α-policy on any instance.

Algorithm π randomly chooses one of two decisions, each with probability 1/2. If decision 1 is chosen, more work than under the α-policy is done in the first time period, and if decision 2 is chosen, less work than under the α-policy is done in the first time period. Specifically, algorithm π allocates resources according to Table E.1, with $\varepsilon = a(1)/15$.

Next we prove that π is always feasible and results in a better value than the α-policy.

First note that the *work attempted* under algorithm π in the first time period, by which we mean η of the resource allocated, is at most $2a(1)/3+\varepsilon = 2a(1)/3+a(1)/15 <$

$a(1)$, so π does not attempt to do work before it arrives. We consider the following three cases that come up in the second time period: $a(2) \geq a(1)$, $3\varepsilon \leq a(2) < a(1)$, and $a(2) < 3\varepsilon$.

Consider the first two cases together. It is easy to see that in these two cases, the total work attempted is the same as the total work attempted under the α-policy. Since we already showed that all work attempted in the first time period is actually accomplished, this means that algorithm π is feasible for both of these two cases (since the α-policy is feasible).

We claim that in both these cases, the maximum resource allocation occurs in the second time period. Recall that this statement is true for the α-policy. Therefore, our only concern is decision 1: specifically, the possibility that adding ε work to the first time period gives a greater total than subtracting ε work from the second time period. Thus it will suffice to show that $\eta(v_T(\omega^2))$, whose value depends on whether $a(2) \geq a(1)$, is at least 2ε more than $\eta(v_T(\omega^1))$, which is always $2a(1)/3$. If $a(2) \geq a(1)$, then $\eta(v_T(\omega^2)) = 4a(2)/3$ and we have

$$\frac{4}{3}a(2) - \frac{2}{3}a(1) \geq \frac{4}{3}a(1) - \frac{2}{3}a(1) = \frac{2}{3}a(1) > \frac{2}{15}a(1) = 2\varepsilon$$

If $3\varepsilon \leq a(2) < a(1)$, then $\eta(v_T(\omega^2)) = 2a(1)/3 + 2a(2)/3$ and we have

$$\frac{2}{3}a(1) + \frac{2}{3}a(2) - \frac{2}{3}a(1) = \frac{2}{3}a(2) \geq \frac{2}{3}(3\varepsilon) = 2\varepsilon$$

Therefore, the value of algorithm π on any instance in which $a(1) \leq a(2)$ is the expected value of the second time period's allocation, which is

$$\frac{1}{2}\eta^{-1}\left(\frac{4}{3}\frac{(a(1)+a(2))}{2} + \varepsilon\right) + \frac{1}{2}\eta^{-1}\left(\frac{4}{3}\frac{(a(1)+a(2))}{2} - \varepsilon\right) < \eta^{-1}\left(\frac{4}{3}\frac{(a(1)+a(2))}{2}\right)$$

where we have strict inequality because of the strict concavity of η^{-1} (due to the strict convexity of η). Thus, the expected value of algorithm π is always less than the value of the α-policy.

Similarly, if $3\varepsilon \leq a(2) < a(1)$, the value of algorithm π is

$$\frac{1}{2}\eta^{-1}\left(\frac{4}{3}a(2) + \varepsilon\right) + \frac{1}{2}\eta^{-1}\left(\frac{4}{3}a(2) - \varepsilon\right) < \eta^{-1}\left(\frac{4}{3}a(2)\right),$$

so algorithm π has a better value in this case as well.

It now remains to consider the third case: $a(2) < 3\varepsilon$. For feasibility, note that decision 1 allocates more at both times than decision 2, so it suffices to show that decision 2 allocates enough resources to meet the deadline. The total work attempted is

$$2\left(\frac{2}{3}a(1) - \varepsilon\right) = \frac{4}{3}a(1) - 2\varepsilon = \frac{4}{3}a(1) - \frac{2}{15}a(1)$$
$$= \frac{18}{15}a(1) = a(1) + \frac{3}{15}a(1) \geq a(1) + a(2).$$

This shows that at least as much work is attempted as is available, and the work attempted in the first time period does not exceed $a(1)$, so the algorithm is feasible.

Now note that the value of algorithm π in this case is

$$\frac{1}{2}\eta^{-1}\left(\frac{2}{3}a(1) + \varepsilon\right) + \frac{1}{2}\eta^{-1}\left(\frac{2}{3}a(1) - \varepsilon\right) < \eta^{-1}\left(\frac{2}{3}a(1)\right)$$

So the value of algorithm π is strictly less than the first time period's allocation by the α-policy, which must be no more than the value of the α-policy (since the second time period's allocation is at least as great as the first time period's).

So, in all three cases, algorithm π has a value strictly less than the α-policy, which completes the proof. $\square$