Nearly Optimal Competitive Online Replacement Policies

Ran El-Yaniv
Center for Rationality
The Hebrew University
Jerusalem
e-mail: `raniy@math.huji.ac.il`

Richard M. Karp
Department of Computer Science
University of Washington
Seattle
e-mail: `karp@cs.washington.edu`

March, 1996

Abstract

This paper studies the following *online replacement* problem. There is a real function $f(t)$, called the *flow rate*, defined over a finite time horizon $[0, T]$. It is known that $m \leq f(t) \leq M$ for some reals $0 \leq m < M$. At time 0 an *online player* starts to pay money at the rate $f(0)$. At each time $0 < t \leq T$ the player may *changeover* and continue paying money at the rate $f(t)$. The complication is that each such changeover incurs some fixed penalty. The player is called *online* as at each time t the player knows f only over the time interval $[0, t]$. The goal of the player is to minimize the total cost comprised of cumulative payment flow plus changeover costs. This formulation of the replacement problem has various interesting applications among which are: *equipment replacement*, *supplier replacement*, the *menu cost* problem and *mortgage refinancing*.

With respect to the *competitive ratio* performance measure, this paper seeks to determine the best possible competitive ratio achievable by an online replacement policy. Our results include the following: a general lower bound on the performance of any deterministic policy, a policy that is optimal in several special cases and a simple policy that is approximately optimal.

Key words: Replacement, equipment replacement, menu cost, mortgage refinancing, online replacement, online algorithms, competitive analysis

1 Introduction

The replacement problem. A real function $f(t)$, called the *flow rate*, is defined over some finite time horizon $[0,T]$ where T is a positive real. It is given that for all t, $m \leq f(t) \leq M$ where $m, M \in \Re$ and $0 \leq m < M$. A *player* is required to pay money throughout the time interval $[0,T]$ where the payment flow is determined in the following manner. At time 0 the player starts to pay money at the rate $f(0)$. At each time $0 < t \leq T$ the player can *changeover* and continue paying money at the rate $f(t)$. To perform any such changeover the player must pay some fixed positive amount C called the *changeover cost*. The player may choose any number k of *changeover times*, $0 < t_1 < t_2 < \cdots < t_k < T$. Thus, for each such changeover time t_i, the player pays a changeover cost C and throughout the interval $[t_i, t_{i+1})$, $i = 0, 1, \ldots, k$, his payment flow is at the rate $f(t_i)$. (By convention we define $t_0 = 0$ and $t_{k+1} = T$.) For each particular choice of changeover times the *total cost* incurred by the player, comprised of payment flows and changeover costs, is

$$kC + \sum_{i=0}^{k}(t_{i+1} - t_i) f(t_i).$$

Any choice of k and (k) changeover times is called a *replacement policy.* Of course we are interested in replacement policies that minimize the total cost. Given f, it is straightforward to compute an optimal replacement policy via (continuous) dynamic programming [1]. Throughout this paper OPT will denote this optimal policy and for each flow function f, OPT(f) is the (optimum) total cost obtained by OPT with respect to f.

Online replacement. In this paper we are concerned with *online* replacement; that is, the player must determine his changeover times online without knowledge of future values of the flow rate function. Specifically, we consider an *online player* that at each time t knows f only over the interval $[0,t]$.

Measuring the performance of online replacement. Following Sleator and Tarjan [13] we measure the performance of an online replacement policy by its *competitive ratio* defined as follows. Let S be any online replacement policy and denote its total cost with respect to the flow f by $S(f)$. If there exist constants α and r such that for all f

$$S(f) - r \cdot \textsc{opt}(f) \leq \alpha,$$

then we say that S is *r-competitive* or that S *attains a competitive ratio of* r.[1] The least r such that S is r-competitive is called *the competitive ratio* of S (or S's *competitiveness*). If S is r-competitive and r is obtained with $\alpha = 0$ then S is said to be *strictly* r-competitive. Thus, if S is r-competitive it will never pay more than r times the absolute optimum obtained by OPT (up to the additive constant α) and the smaller r is the better S performs compared to OPT (clearly $r > 1$ if S is online).

[1]The "constants" r and α can be functions of the problem parameters (m, M, C and T) but must be independent of the flow rate function f.

Using this performance measure we are interested in determining the optimal competitive ratio for the online replacement problem in terms of the problem parameters m, M, C and T. Optimality here is of course defined in a straightforward manner with respect to the competitive ratio.

Viewing the problem as a two-player game. Using the above competitive ratio objective it will be convenient to view the online replacement problem as the following two-player game. The first player is the online player defined above. The second player is called the *adversary* or the *offline player*. The online player chooses a strategy (or replacement policy) S and makes it known to the adversary. Then, based on S, the adversary chooses a flow function f so as to maximize the competitive ratio. The online player's objective is to minimize the competitive ratio (which means that this game is a zero-sum game).

1.1 Applications

The above replacement problem has various interesting applications. In all of these applications the basic question is when to switch from one activity, investment, or facility, to another more rewarding one, when there is a cost associated with making the switch. More specifically, it is assumed that the (online) player participates in exactly one activity at a time. From time to time a new activity is offered as a possible replacement for the original one. Associated with this new activity is a *changeover cost* and an *operating cost* per unit time. This operating cost corresponds to our *flow rate* defined above.[2] Some striking examples of particular applications are the following.

Equipment replacement. Here the player needs to use some piece of equipment throughout the time horizon and from time to time, due to *a priori* unknown economic events and/or technological improvements (or equipment deterioration) the player can and may wish to switch to a different or newer equipment that incurs a lower operating cost (or higher payoff). "Typical" examples of equipment for which this application is relevant are cars, computers, industrial machinery, etc. The same formulation applies of course to more abstract types of "equipment" such as jobs, etc. In all these examples the operating cost can be approximated by a fixed rate payment flow (e.g. gasoline consumption rate, salary, etc)

Supplier replacement. A firm is purchasing goods at a constant rate from one supplier. The cost of purchasing the same goods from other suppliers varies with time. The firm can switch to another supplier but at a certain cost. The cost of this switchover can be approximated by some constant that accounts for the paperwork, the wasted time and possibly the costs involved in breaking the contract that are associated with the switch.

The menu cost problem. Many firms are constantly faced with the problem of when to adjust prices of the goods or services they offer. Due to inflationary markets (and/or other economic events) the firm may wish to update its price menu to reflect their "real" values in order to increase its overall payoff. Each of these price adjustments which correspond to our (flow) changeovers, incurs some fixed cost (to physically update the "menu", advertise,

[2] In many applications, this flow rate or operating cost should be viewed as a the rate of *positive* payoff.

etc.).

Mortgage refinancing. The zero-point fixed rate mortgage common in North America is almost exactly modeled by our replacement problem. Here the flow rate corresponds to the payment rate which is based on a fixed interest rate (and the principal).

Our formulation applies to all these situations with various degrees of accuracy. Nevertheless, it is our feeling that our model captures the essential problem in all of them.

1.2 Contextual background and relation to other work

The literature related to online replacement problems (or various of the above applications) is quite extensive and it is beyond the scope of this paper to survey it all. Nevertheless, we note that the common denominator of all previous theoretical work on the subject is that it is based on the conventional "average case analysis"; that is, analyses are typically made under the assumption that the flow rate function follows a particular (usually simple) stochastic process that may or may not be known to the online player. Let us describe two examples.

Derman [4] studies a simplified discrete time replacement problem where his analog of our flow rate function is a piecewise constant function in which the next value is determined via a simple one-stage Markov process.

Sheshinski and Weiss [12] study policies for price adjustments (described above as the menu cost problem) under the assumption that "real" prices are determined by the following two-state process. During each state the price level is changed at a fixed rate (in particular, they assume that at one state the price is fixed and at the second state, the price increases at a fixed rate). The duration of each state is an independent exponentially distributed random variable.

Other examples of average case analyses of replacement problems can be found in [2, 14, 5, 9]. It is not surprising that in all these examples (and in most other analyses of this kind) the optimal policy is heavily dependent on the stochastic assumptions.

Competitive analysis. The use of the competitive ratio performance measure for online problems is called *competitive analysis*. The competitive ratio was first used by computer scientists in the 70's in connection with approximation algorithms for NP-complete problems (e.g. *bin-packing* [8, 11, 10, 15]), and then was explicitly formulated in the 80's in the seminal work of Sleator and Tarjan [13] on *list accessing* and *paging* algorithms. Since then, competitive analysis has been extensively used to analyze and design online algorithms for many online problems naturally appearing in computer science and has gained much recognition as a useful approach for the analysis of online problems.

This paper is the first to study the replacement problem based on the competitive ratio performance measure. To the best of our knowledge our analyses are the first to study replacement problems without relying on particular stochastic models of the flow rate function. Thus, in some sense the replacement policies derived in this paper trade performance for robustness. It is clear that in applications where the nature of the flow rate function is known our replacement policies will be inferior to those that are optimized with respect to the known flow rate function. However, in many instances not much is known of the flow rate function and a large investment of resources is required to gain adequate knowledge (if this is at all possible) in order to devise a realistic stochastic analysis. In these situations our approach to the replacement problem could be the model of choice. Further, a policy that is good on average is not always the most desired one and one can envision situations in which a smaller but almost certain payoff is preferable to a higher uncertain one.

1.3 Problem reduction.

By suitably scaling the time and cost axes, we may assume, that $C = T = 1$. Specifically, given an initial problem set-up with parameters m', M', T' and C' we map each flow rate $x' \in [m', M']$ to $x = x'T'/C'$ and each time $t' \in [0, T']$ to t'/T'. Thus, the new problem set-up is given by $M = M'T'/C'$, $m = m'T'/C'$ and $T = C = 1$. It is not hard to see that any choice of changeover times, $t'_1, t'_2, \ldots, t'_k \in [0, T']$, has total cost A in the initial set-up, if and only if, under the above mappings, the changeover times $t_1, t_2, \ldots, t_k \in [0, 1]$ have total cost A/C' in the new set-up. Therefore, the above scaling preserves the competitive ratio. After scaling, we further assume that $m + 1 < M$. For $M \leq m + 1$ the problem is trivial in the sense that the online player can always achieve a "perfect" competitive ratio of one.

For the rest of this paper we consider the reduced replacement problem with $C = T = 1$ so that the only relevant parameters are m and M.

1.4 Classes of replacement policies

Recall that the job of the (online) player is to choose a finite sequence $t_1, t_2, \ldots, t_k$ of changeover times. At each time, $t \in [0, 1]$, the information available to the online player for this purpose (aside from the parameters m and M) is t, and the "history" of the game until time t. Notice that it does not make sense to changeover too many times. For instance, it is easy to see that changing over more than $\lceil M \rceil$ times, does not make any sense. Motivated by these facts we now define online deterministic replacement policies. First we define *offline* policies and then refine the definition to capture *online* policies.

Offline replacement policies. Fix m and M. Let $\mathcal{F}$ be the set of all flow rate functions. A *replacement policy* is a sequence $\{M_i\}_{i=1}^{k}$ of *changeover thresholds* where k is any positive integer, and for each i, the changeover threshold M_i is a functional

$$M_i \colon [0, 1] \times \mathcal{F} \to [m, M] \cup \{-1\}.$$

For each choice of $f \in \mathcal{F}$, each of the functionals, $M_i(t, f)$, is a non-increasing function of t, and for all i and t, $M_i(t, f) \geq M_{i+1}(t, f)$. The interpretation is as follows: for a given flow rate function, $f(t)$, the policy will change over at the sequence of times t_i, $i = 1, 2, \ldots, k'$, $k' \leq k$, where t_i is the least t greater than t_{i-1} such that $f(t) \leq M_i(t, f)$. (By convention, we take $t_0 = 0$.) Notice that the range of M_i includes the number -1. This provides the policy with the possibility to 'disable' any threshold at any given time (or in other words, to 'refuse' to change over). Intuitively, for a fixed f, $M_i(t, f)$ should be non-increasing, for the player should become more reluctant to change over as time passes since the later the time of the changeover, the less will he benefit from making the change.

Online replacement policies. We now define an *online* replacement policy. For each $f \in \mathcal{F}$ and $t \in [0, 1]$, let $f_{[0,t]}$ denote the function f restricted to the subinterval $[0, t]$. Let $\mathcal{F}_{[0,t]}$ denote the set of all $f_{[0,t]}$ with $f \in \mathcal{F}$. For each $t \in [0, 1]$ denote by m_t any mapping, $m_t: \mathcal{F}_{[0,t]} \to [m, M] \cup \{-1\}$. An *online replacement policy* is a replacement policy with the following additional requirement. For each i, each $f \in \mathcal{F}$, and each $t \in [0, 1]$, $M_i(t, f) \equiv m_t(f_{[0,t]})$ for some m_t. Intuitively, this simply means that the online policy decides about the ith changeover based only on past and present flow rates (and the fact that it already made $i - 1$ changeovers).

Time independent policies. Perhaps the simplest class of online replacement policies that is still interesting is the following class of *time-independent* policies. A policy in this class is a sequence of constant changeover thresholds (fixed over time and flow rate functions). Thus, a time-independent policy is a decreasing sequence of real numbers,

$$M \geq M_1 > M_2 > \ldots > M_k \geq m.$$

The interpretation is that the online player changes over for the ith time when the flow rate decreases to the level of (or below) M_i.

Time-dependent (refusal) policies. A more sophisticated class of policies is the following class of *time-dependent* or *refusal* policies (we may use either name). A refusal policy is defined as a sequence $\{M_i(t)\}_{i=1}^k$ of functions with domain $[0, 1]$. By definition, each of the functions is non increasing and for all i and t, $M_i(t) \geq M_{i+1}(t)$. Thus, the online policy is willing to make the ith changeover at time t if $f(t) \leq M_i(t)$. Otherwise, it *refuses* to change over. Clearly, a time-independent policy is a rudimentary form of a refusal policy. However, if a refusal policy is also a time-independent policy, it will not usually be called a refusal policy.

In this paper we shall focus on time-independent and refusal policies. Nevertheless, for completeness we mention the following class of replacement policies.

History-dependent policies. The most general class of policies is the class of *history-dependent* policies, where the online player make use of the history of flow rates for making his decisions. Here again, if a history-dependent policy can be represented as a time-independent or a refusal policy it will not be called history-dependent.

As we shall see (and somewhat surprisingly), in all instances of the replacement problem it is possible to obtain optimal or approximately optimal online performance with time-independent and refusal policies (and without resorting to history-dependent policies).

1.5 Outline of results

For the statement of our results (and in the analyses that follow) we use the following standard notational convention for asymptotic relations:

$$f(n) = O(g(n)) \iff |f(n)| \leq c|g(n)| \quad \text{for some } c > 0$$

$$f(n) = \Omega(g(n)) \iff |f(n)| \geq c|g(n)| \quad \text{for some } c > 0$$
$$f(n) = \Theta(g(n)) \iff f(n) = O(g(n)) \text{ and } f(n) = \Omega(g(n))$$

Let us now overview the results of this paper.

A lower bound. In Section 2 we prove a general lower bound on the competitive ratio of any deterministic policy for the replacement problem; This lower bound is obtained by considering a restricted class of adversaries (flow rate functions) against which we can identify an optimal online policy and characterize its competitive ratio. In particular, the restricted class of flow rate functions we consider includes functions that start at time zero at the rate M and then drop "instantaneously" and continuously to some rate chosen by the adversary. It turns out that against such functions the optimal online performance can be attained by a time-independent policy. After establishing this fact we determine the optimal time-independent policy (against such restricted functions) and characterize the asymptotic behavior of its competitive ratio. For a fixed m we show that this competitive ratio is $r^* = \Theta\left(\frac{\ln M}{\ln \ln M}\right)$. The optimal policy is called S^*. Since the competitive ratio of S^* is optimal only against a subclass of all flow rate functions it must be a lower bound on the competitive ratio of any replacement policy against an unrestricted flow rate function.

A characterization theorem. In Section 3 we consider a simple but rich subclass of refusal (time-dependent) policies of the following form. Each policy in this subclass is a time-independent policy equipped with "refusal times"; that is, a (decreasing) sequence of *constant* changeover thresholds $M \geq M_1 > M_2 > \ldots > M_k \geq m$ and for each $1 \leq i \leq k$, the threshold M_i is coupled with a refusal time b_i. The interpretation is that the policy refuses to change over to the threshold M_i if the flow rate function intersects M_i after time b_i. We consider only policies in which the the refusal time sequence $\{b_i\}$ is decreasing. This means that once the policy refuses to change over for the first time, it will refuse to change over for the rest of the game.

We provide a characterization theorem that gives necessary and sufficient conditions for establishing the competitive ratio of any policy in this subclass. Given any refusal policy S (of the above subclass) with k changeover thresholds, and a number $r > 1$ the theorem specifies a set of $O(k^2)$ inequalities such that they are all satisfied if and only if S is r-competitive. This characterization theorem proves to be a strong tool for proving upper bounds and it also provides a computational tool for determining the competitive ratio of any policy numerically.

Upper bounds. In Section 4 we apply the characterization theorem of Section 3 to establish four upper bounds. First we construct in Section 4.1 a refusal policy based on the time-independent policy S^* by coupling it with an appropriate sequence of refusal times. We call the resulting policy S^{**}. We prove the following results:

- In Section 4.2 we prove that S^{**} is optimal whenever $\sqrt{M/(m+1)} \leq (m+2)/(m+1)$.
- In section 4.3 we state (without a proof) a theorem that establishes the optimality of S^{**} whenever $m = 0$ (for all values of M). Both these results of optimality for S^{**}

are established by showing that S^{**} attains the ratio r^* (i.e. the lower bound for the restricted problem).

- In Section 4.4 we find that S^{**} cannot always attain the ratio r^*. Nevertheless, based on numerical evidence we conjecture that S^{**} is always optimal within a very small constant factor.
- In Section 4.5 we prove that S^{**} is $(r^*)^2$-competitive for almost all values of M when $m > 0$. Of course, this bound appears to be weak based on our conjecture of Section 4.4.
- In Section 4.6 we take a different direction and construct a new time-independent policy. This policy, called S^{***} is proven to be approximately optimal for almost all values of M when $m > 0$; that is, the competitive ratio attained by S^{***} is within a constant factor of r^*.

Finally, in Section 5 we conclude and indicate some directions for future work.

2 A Lower bound

In this section we derive a lower bound on possible competitive ratios for the replacement problem. To this end, we confine ourselves to a restricted class of adversaries, corresponding to a restricted class of flow rate functions. The *optimal* competitive ratio for this restricted version of the problem is a lower bound on the competitive ratio for the original problem. Most of this section will then be devoted to identifying the optimal online policy for the restricted problem and to characterizing its competitive ratio.

Fix m and M. For each $0 < \delta < 1/2$, set $\varepsilon_0 = \frac{\delta}{M-m}$. For each $0 < \varepsilon < \varepsilon_0$ and $\mu \in [m, M]$, consider a flow rate function, f, satisfying: (i) $f(0) = M$; (ii) f decreases continuously over the time interval $[0, \varepsilon]$ to the value μ; and (iii) $f(t) = M$ for all $t \in (\varepsilon, 1]$. In other words, the flow rate swoops down continuously to μ at the initial time interval, and then jumps to its maximum possible value and remains there. Let $F_{\mu,\varepsilon}$ be the set of all such flow rate functions. Let $\mathcal{F}_\delta$ be the union over all $\mu \in [m, M]$ and $\varepsilon \in (0, \varepsilon_0)$ of the sets $F_{\mu,\varepsilon}$. An adversary is called a δ-*adversary* if he is restricted to choose only one of the flow rate functions in $\mathcal{F}_\delta$.

Notice that no sensible online policy will ever change over if $\mu \in (M-1, M]$; the savings resulting from such a changeover never exceeds 1, which is the penalty incurred to perform the replacement. Hence, without loss of generality, we shall further restrict our attention to $\mu \in [m, M-1]$, and update the above definitions accordingly.

Assume that the online player knows that he plays against a δ-adversary. We now seek an optimal policy for the online player. We shall first argue that in the limit $\delta \to 0$, every deterministic online policy (and in particular, the optimal one) is captured by a time-independent policy. More formally, it will be shown that there exists a constant c such that for every δ, if the online player confines himself to one of the best *time-independent* policies, he can guarantee a competitive ratio smaller then $r + c \cdot \delta$ where r is the (general) optimal competitive ratio against the δ-adversary.

Lemma 2.1 *Let S be any online policy against the δ-adversary and assume that r is its competitive ratio. Then, there exists a time-independent policy, $\hat{S}$, that is $\hat{r}$-competitive against the δ-adversary and $\hat{r} < r + c \cdot \delta$ where c is a constant independent of δ.*

Proof. Let $S = \{M_i(\cdot)\}_{i=1}^{k}$, an arbitrary replacement policy, be given. Choose $g \in \mathcal{F}_\delta$ such that $\frac{S(g)}{\mathrm{OPT}(g)} = r$. Hence, $g \in F_{\mu',\varepsilon'}$ for some $0 < \varepsilon' < \varepsilon_0$ and $m \leq \mu' \leq M-1$. Choose any $\hat{\varepsilon}$ with $\varepsilon' < \hat{\varepsilon} < \varepsilon_0$ and choose a flow rate function, $\hat{f} \in F_{m,\hat{\varepsilon}}$, such that $\hat{f}_{[0,\varepsilon']} = g_{[0,\varepsilon']}$. That is, $\hat{f}$ is an "extension" of g that decreases to the minimum possible rate. Clearly $\hat{f} \in \mathcal{F}_\delta$. We shall define $\hat{S}$, a time-independent policy, as follows. Let $\varepsilon_1, \varepsilon_2, \ldots, \varepsilon_{k'}$, $(k' \leq k)$ be the times S changes over when played against $\hat{f}$. For $1 \leq i \leq k'$, set $\hat{M}_i = \hat{f}(\varepsilon_i)$, and define $\hat{S} \stackrel{\mathrm{def}}{=} \{\hat{M}_i\}_{i=1}^{k'}$.

For each $\mu \in [m, M-1]$, denote by $\varepsilon(\mu)$ the (unique) ε such that $\hat{f}(\varepsilon) = \mu$, and let $\hat{f}_\mu \in F_{\mu,\varepsilon(\mu)}$ be a function such that $(\hat{f}_\mu)_{[0,\varepsilon(\mu)]} = \hat{f}_{[0,\varepsilon(\mu)]}$. It is evident, by the choice of $\hat{S}$, and since S is an online policy, that for each μ, the performance of S and $\hat{S}$ against $\hat{f}_\mu$ is identical; that is, $S(\hat{f}_\mu) = \hat{S}(\hat{f}_\mu)$.

By the choice of r and g, it is clear that for each μ,

$$\frac{S(\hat{f}_\mu)}{\mathrm{OPT}(\hat{f}_\mu)} \leq r. \tag{1}$$

Fix μ. As ε varies, the total cost of $\hat{S}$ with respect to a function in $F_{\mu,\varepsilon}$ does not vary by much. Since $\hat{S}$ is time-independent, the number of changeovers it performs is independent of the choice of ε. Clearly, for $\varepsilon \in (0, \varepsilon_0)$,

$$0 < \sup_{\substack{f,\varepsilon,\\ f \in F_{\mu,\varepsilon}}} \hat{S}(f) - \inf_{\substack{f,\varepsilon,\\ f \in F_{\mu,\varepsilon}}} \hat{S}(f) < \varepsilon_0(M-\mu). \tag{2}$$

From (2) and by the definition of ε_0 we obtain that for each choice of $f \in \mathcal{F}_\delta$ with $f \in F_{\mu,\varepsilon}$,

$$\hat{S}(f) \leq \hat{S}(\hat{f}_\mu) + \varepsilon_0(M-\mu) < \hat{S}(\hat{f}_\mu) + \delta. \tag{3}$$

The bounds (2) are also true for OPT since $\sup_{f,\varepsilon} \mathrm{OPT}(f)$ $(f \in F_{\mu,\varepsilon})$ is certainly not greater than $(M-\mu)\varepsilon_0 + \mu + 1$ and $\inf_{f,\varepsilon} \mathrm{OPT}(f)$ $(f \in F_{\mu,\varepsilon})$ is not smaller than $\min\{M, \mu+1\} \geq \mu+1$. Thus,

$$0 < \sup_{\substack{f,\varepsilon,\\ f \in F_{\mu,\varepsilon}}} \mathrm{OPT}(f) - \inf_{\substack{f,\varepsilon,\\ f \in F_{\mu,\varepsilon}}} \mathrm{OPT}(f) < \varepsilon_0(M-\mu).$$

Therefore,

$$\mathrm{OPT}(f) \geq \mathrm{OPT}(\hat{f}_\mu) - \varepsilon_0(M-\mu) > \mathrm{OPT}(\hat{f}_\mu) - \delta. \tag{4}$$

Combining (1), (3) and (4), plus the facts that $r \leq M$ always (a degenerate policy that never changes over is clearly M-competitive), that $\mathrm{OPT}(\hat{f}_\mu) > 1$, and $\delta < 1/2$, we obtain

$$\frac{\hat{S}(f)}{\mathrm{OPT}(f)} < \frac{\hat{S}(\hat{f}_\mu) + \delta}{\mathrm{OPT}(\hat{f}_\mu) - \delta}$$

$$
\begin{aligned}
&\leq \frac{r \cdot \text{OPT}(\hat{f}_\mu) + \delta}{\text{OPT}(\hat{f}_\mu) - \delta} \\
&= r + \delta \cdot \frac{r+1}{\text{OPT}(\hat{f}_\mu) - \delta} \\
&< r + \delta \cdot 2(M+1).
\end{aligned}
$$

∎

Remark 2.1 Ultimately, we would like to consider a δ-adversary where δ is "infinitesimally small". Such a δ is, of course, non-existent. Hence, we should investigate this problem as δ approaches zero. Formally, one should obtain all results for a fixed δ and then calculate the appropriate limits. However, to simplify the analysis, we shall avoid this formality whenever possible. We shall pretend that the flow rate function decreases continuously in "zero-time" to its final value, μ, and measure all the quantities that depend on δ by their limits, as $\delta \to 0$. For instance, we shall pretend that the optimal competitive ratio (for the restricted replacement problem) is attained by a time-independent policy. This is not proven to be true, of course, but as $\delta \to 0$, by Lemma 2.1, the competitive ratio of the *optimal* time-independent policy approaches the (general) optimal ratio against the limit behavior of the δ-adversary. It is to be understood that by taking this approach the only sacrifice we make is that our final result, the lower bound, will not be a real number, but a limit.

In continuation to Remark 2.1, we shall consider for the rest of this section only time-independent policies (that may be referred to simply as 'policies'). For each $\mu \in [m, M-1]$, let f_μ denote the fictitious flow rate function that drops continuously in "zero-time" to the rate μ and then jumps to M and remains constant. Formally, when we measure costs of policies against such f_μ, we actually refer to the limit, $\varepsilon \to 0$, of these costs against functions in $F_{\mu,\varepsilon}$.

For each μ, the cost of OPT against f_μ is characterized by the following lemma.

Lemma 2.2 *For sufficiently small $0 < \varepsilon < \varepsilon_0$, for each $\mu \in [m, M-1]$, and each $f \in F_{\mu,\varepsilon}$,*

$$\text{OPT}(f) = \min\{M, \mu + 1 + \varepsilon(M-\mu)\}.$$

Proof. We argue that the two options OPT has are either to avoid change overs and keep paying at the rate M, or to change over once, at time ε, and to keep paying at the rate μ. First we show that OPT will change over at most once. Suppose OPT changes over exactly once at time $0 < t \leq \varepsilon$, to the rate $f(t)$. An additional changeover at time t', $0 < t' < t$, will cut the flow rate component of the total cost by $(t-t')(M - f(t')) < \varepsilon_0(M-m) < \delta < 1$. However, the penalty to perform this changeover is 1. It follows that OPT will change over at most once.

Next, let us compare the cost of changing over once at time ε to the rate μ, with the cost of changing over once to some flow rate $a > \mu$, say, at time $\varepsilon' < \varepsilon$. By changing over once to the rate a instead of changing over once to the rate μ, OPT will gain $G = (\varepsilon - \varepsilon')(M-a)$ but will lose $L = (a - \mu)(1 - \varepsilon)$. It is not hard to see that when $\varepsilon \to 0$, either $L > G$ or $a \to \mu$. Hence, we may assume that between these two alternatives, OPT will choose to change over to μ. ∎

Since $\lim_{\varepsilon \to 0} \mu + 1 + \varepsilon(M - \mu) = \mu + 1 \le M$, a corollary of Lemma 2.2 is that for each $\mu \in [m, M - 1]$, $\text{OPT}(f_\mu) = \mu + 1$.

Let $S = \{M_i\}_{i=1}^{k}$ be any (time-independent) policy that achieves a competitive ratio r. As noted in a previous discussion, we may assume that $M_1 \le M - 1$. Also, we may assume that the last threshold, M_k, is strictly greater than m. The reason is that the adversary can always choose μ to be greater than m, but arbitrarily close to m. By convention, we set $M_0 = M$ and if the smallest threshold of S is M_k, we set $M_{k+1} = m$. Thus,

$$m = M_{k+1} < M_k < \cdots < M_1 \le M_0 - 1 = M - 1.$$

If μ lies in the interval $(M_{i+1}, M_i]$, the online player's cost with respect to f_μ is $M_i + i$ and the optimal offline cost is $\mu + 1$. A simple observation is that an optimal choice of $\mu \in (M_{i+1}, M_i]$ by the adversary is larger than but arbitrarily close to M_{i+1}. Hence, we assume that the optimal offline cost for such μ is arbitrarily close to $M_{i+1} + 1$. Putting all this together we have

Corollary 2.1 *S is r-competitive if and only if*

$$M_i + i \le r \cdot (M_{i+1} + 1), \qquad i = 0, 1, \ldots, k. \tag{5}$$

Lemma 2.3 *If the competitive ratio r is attainable by some policy, then it can be attained by a policy $\{M_i\}_{i=1}^{k}$ such that $M_i + i = r(M_{i+1} + 1)$, $i = 0, 1, \ldots, k - 1$, and $M_k + k \le r(m + 1)$.*

Proof. Let $S = \{M_i\}_{i=1}^{k}$ be any r-competitive policy. By Corollary 2.1, $M_i + i \le r(M_{i+1} + 1)$ for all $0 \le i \le k$. Let j be the minimal index such that $M_j + j < r(M_{j+1} + 1)$. In other words, $M_{j+1} > \frac{M_j + j}{r} - 1$. Set $M'_{j+1} \stackrel{\text{def}}{=} \frac{M_j + j}{r} - 1$. The modified policy, $\{M_1, \ldots, M_j, M'_{j+1}, M_{j+2}, \ldots, M_k\}$, must also attain a competitive ratio r since $M'_{j+1} + j + 1 < M_{j+1} + j + 1 \le r(M_{j+2} + 1)$. Similarly, we can continue to modify the policy, step by step, replacing the thresholds $M_{j+2}, M_{j+3}, \ldots$, etc. Since we replace each of these thresholds by a *lower* threshold, we may end up with some thresholds smaller than m. As the final policy we shall take all remaining thresholds that are still in $(m, M - 1]$. Let $S' = \{M'_i\}_{i=1}^{k'}$ be this final policy, where $k' \le k$ is the maximum index i such that M_i was still replaced by a threshold $M'_i > m$. By construction, $M'_i + i = r(M'_{i+1} + 1)$, $i = 0, 1, \ldots, k' - 1$. Although $M'_{k'+1}$ ($\le m$) was not included in S', by construction, $M'_{k'} + k' = r(M'_{k'+1} + 1)$, so $M_{k'} + k' \le r(m + 1)$. Hence, by Corollary 2.1 S' is r-competitive and the proof is completed. ∎

2.1 The sequence $S(r, a)$ and time-independent policies

For any reals $a > 1$ and $r > 1$, consider the following sequence:

$$\begin{cases} m_0(r) &= a; \\ m_{i+1}(r) &= \frac{m_i(r) + i}{r} - 1, \qquad i \ge 0. \end{cases} \tag{6}$$

Denote this infinite sequence, $\{m_i(r)\}_i$, by $S(r,a)$. The nature of the sequence $S(r,a)$ is tied with the nature of r-competitive time-independent policies. For instance, if there exists an r-competitive refinancing policy, then, by Lemma 2.3, there exists an r-competitive policy $S = \{M_i\}_{i=1}^k$ with $M_i = m_i(r)$, $0 \le i \le k-1$ and $m_{k+1}(r) \le m$, where the $m_i(r)$'s are elements of $S(r,M)$.

The following is a closed form formula for $m_i(r)$.

$$m_i(r) = \left(a + \frac{r^2}{(r-1)^2}\right) r^{-i} + \frac{i}{r-1} - \frac{r^2}{(r-1)^2}. \tag{7}$$

This formula can be verified by induction on i.

The following two lemmata establish a few basic properties of the sequence $S(r,a)$ that will be used later.

Lemma 2.4 *For any $a > 1$ and $r > 1$, the sequence $S(r,a) = \{m_i(r)\}$ is unimodal (that is, it has a unique minimum).*

Proof. Fix any $a > 1$ and $r > 1$. Set $m_i = m_i(r)$, $i = 0, 1, \ldots$. We shall prove the following three claims. The lemma readily follows from the latter two. (i) for all i, $m_i > -1$; (ii) $\{m_i\}$ is initially (strictly) decreasing and then stops decreasing; and (iii) once the sequence $\{m_i\}$ stop decreasing it strictly increases forever.

We prove (i) by induction on i. By assumption, $m_0 = a > 1$. Assume for the induction hypothesis that, $m_i > -1$ for some $i > 0$. Then, $m_{i+1} = \frac{m_i + i}{r} - 1 > \frac{-1+1}{r} - 1 = -1$.

We now prove that (ii) holds. It is not hard to see that, initially, the sequence is strictly decreasing. For example, $m_1 = m_0/r - 1 < m_0$. Assume by way of contradiction that the sequence is always decreasing. Hence, by (i), for all i, $-1 < m_i < m_0$. Take large enough i so that $i/r \ge m_i(1 - 1/r) + 1$. Such an i must exist since the m_i's are bounded. For this i we obtain $m_{i+1} = \frac{m_i}{r} - 1 + \frac{i}{r} \ge \frac{m_i}{r} - 1 + m_i(1 - \frac{1}{r}) + 1 = m_i$. This proves (ii). Finally, to prove (iii) we show that for every i, if $m_i \le m_{i+1}$, then $m_{i+1} < m_{i+2}$. But if $m_i \le m_{i+1}$, then $m_{i+1} = \frac{m_i + i}{r} - 1 < \frac{m_{i+1} + i + 1}{r} - 1 = m_{i+2}$. ∎

The next lemma describes the dependency of each element of $S(r,a)$ in r.

Lemma 2.5 *Fix $a > 1$. Then for all $r > 1$, $S(r,a)$ has the property that for each $i \ge 1$, (i) $m_i(r)$ is decreasing if and only if r is increasing; and (ii) $m_i(r)$ is increasing if and only if r is decreasing.*

Proof. We prove by induction on i that (i) holds. For the base case, consider the equation $m_1(r) = \frac{m_0}{r} - 1$. Since $m_0 = a$ is fixed, (i) trivially holds. To complete the induction step, consider the equation $m_{i+1}(r) = \frac{m_i(r) + i}{r} - 1$. We start with the 'only if' direction. If $m_{i+1}(r)$ decreases, it must be that r increases and/or $m_i(r)$ decreases. In the latter case, by the induction hypothesis, r increases, so in either case the induction step is completed for this case. To complete the induction case for the 'if' case, notice that if r increases then by the induction hypothesis, $m_i(r)$ decreases so it must be that $m_{i+1}(r)$ decreases. The statement (ii) is proven similarly. ∎

Set $a = M$. For each $r > 1$, consider $S(r, M)$ and set $m(r) \stackrel{\text{def}}{=} \min_i\{m_i(r)\}$. By Lemma 2.4, $m(r)$ is well defined. For each $r > 1$ the initial decreasing segment of $S(r, M)$, $m_0, m_1, \ldots, m(r)$, induces a time-independent replacement policy. Specifically, we define the *policy induced by* $S(r, M)$ *with respect to* m to be the decreasing part of the sequence from the second element, up to, but not including, the first element that is less than or equal to m. If $m(r) \geq m$, the induced policy consists of all elements from the second, up to (and including) the minimum element, $m(r)$.

Lemma 2.6 *For each $r > 1$, let $S = \{M_i\}_{i=1}^k$ be the policy induced by $S(r, M)$. Then S is r-competitive if and only if $m \geq m(r)$.*

Proof. If $m \geq m(r)$, then by the construction of $S(r, M)$, for any choice of $0 \leq i < k$, and any choice of μ in the interval $(M_{i+1}, M_i]$, S will pay $M_i + i$ against f_μ, and OPT will pay at least $M_{i+1} + 1$. But $M_i + i = r(M_{i+1} + 1)$, so S attains the ratio r. For a choice $i = k$, for each $\mu \in (m, M_k]$, S will pay $M_k + k$ and OPT, at least $m + 1$. But since $m_{k+1}(r) = m(r) \leq m$, $m_k(r) + k = r(m_{k+1}(r) + 1) \leq r(m + 1)$. So here again, S attains the ratio r.

On the other hand, if $m < m(r)$ then suppose, by contradiction, that S is r-competitive. Choose any μ greater than but very close to m. Then, against f_μ, S pays $m(r) + k + 1$, and OPT incurs a cost arbitrarily close to $m + 1$. As S is r-competitive, it must be that $m(r) + k + 1 \leq r(m+1)$. Hence, $m \geq \frac{m(r)+k+1}{r} - 1 = m_{k+2}(r)$. This contradicts the fact that $m(r) = m_{k+1}(r)$ is the minimum element of $S(r, M)$. Therefore, S is not r-competitive. ∎

Example 2.1 Suppose $M = 100$, and $r = 2$. The resulting sequence, $S(2, 100)$, is:

$$100, 49, 24, 12, 6.5, 4.25, 3.625, 3.813, 4.406, 5.203, \ldots$$

The minimum element of this sequence is 3.625. For $m = 3.625$ the induced policy is

$$S_1 = \{49, 24, 12, 6.5, 4.25\}.$$

For any $6.5 < m < 12$, the induced policy is

$$S_2 = \{49, 24, 12\}.$$

By Lemma 2.6, the competitive ratio 2 is attainable by the induced policy if and only if $m \geq m(2) = 3.625$. Therefore, both S_1 and S_2 are 2-competitive. However, for $m < 3.625$, the induced policy, which is again S_1, cannot attain the ratio 2.

Lemma 2.7 *Fix any $m \geq 0$. For each $M > m + 1$, let r be the optimal competitive ratio. Then, the minimum element of $S(r, M)$, $m(r)$, equals m, and the policy induced by $S(r, M)$ with respect to m is r-competitive.*

Proof. As there exists an r-competitive policy, Lemma 2.3 implies that the policy induced by $S(r, M)$ with respect to m, is r-competitive and that $m(r) \leq m$. Let l denote the index of the minimum element, $m(r)$. Assume, by way of contradiction, that $m > m(r)$. Then it is possible to increase $m_l(r)$ (together with all elements $m_i(r)$, $1 \leq i < l$) to obtain a new sequence $S(r', M)$ with $m(r) < m(r') \leq m$. By Lemma 2.6, the policy induced by $S(r', M)$, S', is r'-competitive, and by Lemma 2.5, $r' < r$. This contradicts the optimality of r. Hence, $m(r) = m$. ∎

For each $m \geq 0$ and each $M > m + 1$, denote by $r^*(m, M)$ the optimal competitive ratio for the restricted replacement problem. Denote by $S^*(m, M) = \{M_i\}_{i=1}^k$ the policy induced by $S(r^*(m, M), M)$. By Lemma 2.7, $S^*(m, M)$ is $r^*(m, M)$-competitive and $m(r^*(m, M)) = m_{k+1}(r^*(m, M)) = m$. Denote by $k^*(m, M)$ the number, k, of changeover thresholds in $S^*(m, M)$.

Whenever there is no confusion we may interchange the $m_i = m_i(r)$'s with the M_i's, the changeover thresholds of the policy induced by $S(r, M)$. The next lemma characterizes the number $k^*(m, M)$.

Lemma 2.8 *Let $r = r^*(m, M)$ and $k = k^*(m, M)$. Then,*

$$k = \lceil (m + 1)(r - 1) \rceil.$$

Proof. Consider the policy $S^*(m, M) = \{M_i\}_{i=1}^k$. We know that $M_{k+1} = m$ and $m < M_k$. Also, $M_k + k = r(M_{k+1} + 1)$. Hence, $m + k < M_k + k = r(M_{k+1} + 1) = r(m + 1)$. Therefore,

$$k < r(m + 1) - m. \tag{8}$$

By unimodality, $m \leq m_{k+2} = \frac{m+k+1}{r} - 1$. Thus,

$$k \geq rm + r - m - 1, \tag{9}$$

and, together with (8), and the fact the k is an integer, the proof is completed. ∎

We are now interested in characterizing, $r^*(m, M)$, the optimal competitive ratio. It is not hard to obtain an explicit expression for sufficiently small optimal ratios.

Lemma 2.9 *Let $r = r^*(m, M)$, $k = k^*(m, M)$. If $r \leq \frac{m+2}{m+1}$, then $r = \sqrt{\frac{M}{m+1}}$ and $k = 1$.*

Proof. First, if $r \leq \frac{m+2}{m+1}$, then $r(m+1) - m \leq 2$. By inequality (8), $k < 2$. On the other hand, inequality (9) implies that $k \geq 1$. Therefore, $S^*(m, M)$ consists of one changeover threshold, M_1. Set $\rho = \sqrt{\frac{M}{m+1}}$ and consider the policy induced by $S(\rho, M)$. By definition,

$$\begin{aligned} m_1(\rho) &= m_0(\rho)/\rho - 1 = \sqrt{M}\sqrt{m+1} - 1. \\ m_2(\rho) &= (m_1(\rho) + 1)/\rho - 1 = \frac{\sqrt{M}\sqrt{m+1}}{\sqrt{M}/\sqrt{m+1}} - 1 = m. \end{aligned}$$

Therefore, the policy $\{M_1\}$ achieves a competitive ratio ρ. By Lemma 2.5, if we take $\rho' < \rho$ then $m(\rho') > m(\rho) = m$, so, it is implied, by Lemmas 2.3 and 2.6, that no policy can attain the ratio ρ'. Hence, ρ is optimal. Thus, $r = \rho = \sqrt{\frac{M}{m+1}}$. ∎

For small ratios, the following lemma proves some relationships between r, k, m, and M that better characterize the lower bound at its small values.

Lemma 2.10 *Set $r = r^*(m, M)$, $k = k^*(m, M)$. The following conditions are equivalent: (a) $r \leq \frac{m+2}{m+1}$; (b) $r = \sqrt{\frac{M}{m+1}}$; (c) $k = 1$; and (d) $M \leq \frac{(m+2)^2}{m+1}$.*

Proof. The following four statements easily derive the lemma:

- (a) $\Rightarrow$ (b), (c) and (d): this is established by Lemma 2.9. Notice that $\sqrt{\frac{M}{m+1}} \leq \frac{m+2}{m+1}$ iff $M \leq \frac{(m+2)^2}{m+1}$.
- (c) $\Rightarrow$ (a): if $k = 1$ then (a) follows by inequality (9).
- (b) $\Rightarrow$ (c): set $\rho = \sqrt{\frac{M}{m+1}}$ and consider $S(\rho, M)$. It is not hard to see that $m_2(\rho) = m$. Hence, the policy induced by $S(\rho, M)$ with respect to m is ρ-competitive. But $\rho = r$ and is thus optimal, so $k = 1$.
- (d) $\Rightarrow$ (a): suppose $M \leq \frac{(m+2)^2}{m+1}$. Consider $S(\rho, M)$ where ρ is a variable. For a ρ that solves $m_2(\rho) = m$, it must be that the policy induced by $S(\rho, M)$ with respect to m is ρ-competitive. Hence, $r \leq \rho$. But $m = m_2(\rho) = \frac{M}{\rho^2} - 1 \leq \frac{(m+2)^2}{\rho^2(m+1)} - 1$, so $\rho \leq \frac{m+2}{m+1}$ and (a) holds.

∎

By Lemma 2.10, an alternative way to state Lemma 2.9 is to say that for $m + 1 < M \leq (1 + \frac{1}{m+1})(m + 2)$, it makes sense to change over at most once, and in this case $r^*(m, M) = \sqrt{\frac{M}{m+1}} < 1 + \frac{1}{m+1}$.

Set $r = r^*(m, M)$ and $k = k^*(m, M)$. Since $S^*(m, M)$ has the property that $M_{k+1} = m$, using the closed form (7) for $m_{k+1}(r) = M_{k+1}$ (with $a = M$), we know that r is a real solution of the system

$$
\begin{aligned}
m &= (M + \frac{r^2}{(r-1)^2})r^{-(k+1)} + \frac{k+1}{r-1} - \frac{r^2}{(r-1)^2} && (10) \\
k &= \lceil (m+1)(r-1) \rceil && (11)
\end{aligned}
$$

This system implicitly defines the function $r = r^*(m, M)$, which we desire to characterize. However, the functional dependency of the optimal ratio, r, in m and M is very complicated. For example, when it is known that the optimal policy consists of exactly two changeover thresholds ($k = 2$), then the optimal competitive ratio is

$$
\begin{aligned}
r \ = \ & \left(\frac{M}{2(m+1)} + \frac{1}{27\,(m+1)^3} + \frac{\sqrt{M}\sqrt{27\,Mm^2 + 54\,Mm + 27\,M + 4}\sqrt{3}}{18\,(m+1)^2}\right)^{\frac{1}{3}} + \\
& \left(\frac{M}{2\,(m+1)} + \frac{1}{27\,(m+1)^3} - \frac{\sqrt{M}\sqrt{27\,Mm^2 + 54\,Mm + 27\,M + 4}\sqrt{3}}{18\,(m+1)^2}\right)^{\frac{1}{3}} + \\
& \frac{1}{3\,(m+1)}.
\end{aligned}
$$

Although it may be desirable to find a simple approximation for $r^*(m, M)$, it seems that any accurate approximation will be hard to find, rather complex, and not necessarily informative (for our purposes). We therefore confine ourselves to the much simpler task of finding an asymptotic approximation that will show the nature of $r^*(m, M)$ as the ratio M/m grows. (Then again, such an approximation will be useless for very small competitive ratios.)

Lemma 2.11 *For a fixed m, and sufficiently large M, the optimal competitive ratio (for this restricted version of the problem) is $r = r^*(m, M) = \Theta\left(\frac{\ln M}{\ln \ln M}\right)$.*

Proof. By isolating M in (10), we obtain

$$M = r^{k+1}\left(m + \frac{r^2}{(r-1)^2} - \frac{k+1}{r-1}\right) - \frac{r^2}{(r-1)^2}. \tag{12}$$

Fix m. We assume that $r > 2 + \sqrt{2}$. We will use the following inequalities. The first follows from $r > 2 + \sqrt{2}$:

$$1 < \frac{r^2}{(r-1)^2} < 2. \tag{13}$$

Then from (11) we derive lower and upper bounds on the optimal k

$$(m+1)(r-1) \le k < (m+1)(r-1) + 1 = m(r-1) + r. \tag{14}$$

Using (12), it is not hard to derive the following bounds on M: $r^{k-1} - 2 < M < r^{k+1} - 1$. Further, by applying the bounds on k, (14), we obtain

$$r^{(r-1)m} - 2 < M < r^{r(m+1)-m+1} - 1.$$

Then by applying the natural logarithm to both these inequalities we easily obtain $r = \Theta(\frac{\ln M}{\ln r})$. Applying again the natural logarithm to this asymptotic bound we obtain $r = \Theta(\ln \ln M)$ and the proof is complete. ∎

3 A characterization theorem for simple refusal policies

Having determined a general lower bound for the replacement problem using the restricted adversary, we now step back and examine the general case, in which the flow rate function is entirely unrestricted.

3.1 A subclass of simple refusal policies

We will be focusing on a subclass of *refusal policies* (see Section 1.4 for a definition). A refusal policy in this new subclass has a particularly simple form and is specified by two sequences of real numbers:

- $\{M_i\}_{i=1}^k$ - *changeover thresholds*;
- $\{b_i\}_{i=1}^k$ - *refusal times.*

We require that the changeover threshold sequence is strictly decreasing within the open interval (M, m) and that the refusal time sequence is non-increasing within the time horizon $[0, 1]$. Given these sequences of changeover thresholds and refusal times, the refusal policy is defined as follows: $M_i(t) = M_i$, for $t \leq b_i$, and $M_i(t) = -1$, for $t > b_i$. Thus, up to time b_i the threshold for making the ith changeover is the constant M_i, but after time b_i the online policy refuses to change over for an ith time. By convention, we set $M_0 = M$ and $b_{k+1} = 0$. Here again (as in the discussion in page 11) we can assume that the last threshold, M_k, is greater than m, and then we set $M_{k+1} = m$.

3.2 A characterization theorem

We now introduce a characterization theorem that proves very useful for obtaining upper and lower bounds. This theorem provides necessary and sufficient conditions for establishing that a refusal policy $\langle\{M_i\}, \{b_i\}\rangle$ as defined above attains a competitive ratio r. By this theorem, to establish whether or not a particular refusal policy with k thresholds attains a certain competitive ratio, all that is needed is to examine $O(k^2)$ inequalities. This means in particular that given any parameters for the problem (i.e. m and M), one can compute efficiently a numerical approximation of the best refusal policy.

Theorem 3.1 *For each $m \geq 0$ and each $M > m$, let $S = \langle\{M_i\}_{i=1}^k, \{b_i\}_{i=1}^k\rangle$ be a refusal policy with $\{b_i\}$ non-increasing. Then S is r-competitive if and only if the following two conditions hold:*

C1 *for all i and j with $0 \le i \le j \le k$,*

$$M_i b_{j+1} + M_j(1 - b_{j+1}) + j \le r \cdot \mathrm{Min} \begin{bmatrix} M_0, \\ M_{i+1} + 1, \\ M_0 b_{j+1} + m(1 - b_{j+1}) + 1, \\ M_{i+1} b_{j+1} + m(1 - b_{j+1}) + 2 \end{bmatrix};$$

C2 *for all i and j with $0 \le i < j \le k$,*

$$M_i b_j + M_j(1 - b_j) + j \le r \cdot \mathrm{Min} \begin{bmatrix} M_0, \\ M_{i+1} + 1, \\ M_0 b_j + m(1 - b_j) + 1, \\ M_{i+1} b_j + m(1 - b_j) + 2 \end{bmatrix}.$$

Proof. To prove the sufficiency of the above conditions, assume that they are satisfied. Consider the behavior of S against a flow rate function f. For each $0 \le i \le k$, let I_i be the (possibly null) interval of time during which S is paying at the flow rate M_i. Let j be the largest index for which the associated interval is nonempty. Then, throughout the interval I_i, where $i < j$, $f(t) \ge M_{i+1}$. Let the interval I_j be partitioned into two parts, I_j^1 and I_j^2, where $I_j^1 = I_j \cap [0, b_{j+1}]$ and $I_j^2 = I_j \cap (b_{j+1}, 1]$. Then, throughout the interval I_j^1, $f(t) \ge M_{j+1}$ and throughout the interval I_j^2, $f(t) \ge m$. For $i = 0, 1, \ldots j-1$, let α_i denote the length of the interval I_i. Let α_j denote the length of the interval I_j^1 and let α_{j+1} denote the length of the interval I_j^2 (note that $\alpha_j + \alpha_{j+1} > 0$). The cost of the online policy with respect to f, $S(f)$, is $\sum_{i=0}^{j} \alpha_i M_i + \alpha_{j+1} M_j + j$. Let $g(t)$ be the function that is equal to M_{i+1} on I_i, $i = 1, 2, \ldots, j-1$, to M_{j+1} on I_j^1 and to m on I_j^2. Clearly, $g(t) \le f(t)$ for all t, so $\mathrm{OPT}(g) \le \mathrm{OPT}(f)$. We shall prove that the hypotheses of the theorem imply that $S(f) \le r \cdot \mathrm{OPT}(g)$.

Since g is a non-increasing, piecewise-constant function, the optimal policy for the offline player has the property that changeovers occur only at the boundary points between intervals of constancy for g. For each particular choice of changeover points, the offline cost is a linear form, $\mathrm{L}(\alpha_0, \alpha_1, \ldots, \alpha_{j+1})$. It suffices to show that for every choice of j, for every choice of the α_i, and for every such linear form, $S(f) \le r \cdot \mathrm{L}(\alpha_0, \alpha_1, \ldots, \alpha_{j+1})$. For each fixed choice of j, the quantities α_i satisfy the following conditions:

(I1) $\alpha_i \ge 0, \qquad i = 0, 1, \ldots, j;$

(I2) $\sum_{i=0}^{j} \alpha_i = 1 - \alpha_{j+1};$

(I3) $1 - b_{j+1} \ge \alpha_{j+1} \ge 1 - b_j;$

(I4) if $\alpha_j > 0$, then $1 - b_{j+1} = \alpha_{j+1}$.

Regard the $(j+2)$-tuple $(\alpha_0, \alpha_1, \ldots, \alpha_{j+1})$ as a point of a $(j+2)$-dimensional Euclidean space. Then the set of feasible points (possible choices of the α_i's) is determined by the above conditions **I1**-**I4**. Since the (in)equalities in these conditions are linear each inequality fixes a half-space and each equality fixes a hyperplane and the set of feasible points is the intersection of these half spaces and hyperplanes. As this feasible set is a polytope it is completely determined by its corner points and it is hence sufficient to prove that the linear inequality

$$\sum_{i=0}^{j} \alpha_i M_i + \alpha_{j+1} M_{j+1} + j \leq r \cdot L(\alpha_0, \alpha_1, \ldots, \alpha_{j+1}) \tag{15}$$

is satisfied by the corner points of the feasible set. We now claim that in each corner point at most one of the α_i's, $i = 0, 1, \ldots, j$, is non-zero. Consider the system of linear inequalities of conditions **I1** and **I2**. This system has $j+1$ variables. A corner point solution of a system of linear inequalities in $j+1$ variables has the property that it satisfies at least $j+1$ inequalities with equality. Hence in each corner point of the ($(j+2)$-dimensional) feasible set at most one of the α_i's ($i = 0, \ldots, j$) is non-zero. By considering the additional conditions, **I3** and **I4**, we conclude that each corner-point $\alpha = (\alpha_0, \alpha_1, \ldots, \alpha_{j+1})$ of the feasible set satisfies one of the following types of conditions:

1. for some $i \leq j$, $\alpha_i = b_{j+1}$ and $\alpha_{j+1} = 1 - b_{j+1}$;

2. for some $i < j$, $\alpha_i = b_j$ and $\alpha_{j+1} = 1 - b_j$.

By the definition of the α_i, for the first type of corner point solutions, the online cost is $M_i b_{j+1} + M_j(1 - b_{j+1}) + j$, with $i \leq j$, and the corresponding optimal offline cost is

$$\text{Min} \left[\begin{array}{l} M_0, \\ M_{i+1} + 1, \\ M_0 b_{j+1} + m(1 - b_{j+1}) + 1, \\ M_{i+1} b_{j+1} + m(1 - b_{j+1}) + 2 \end{array} \right].$$

For the second type of corner point solutions, the online cost is $M_i b_j + M_j(1 - b_j) + j$, with $i < j$, and the corresponding optimal offline cost is

$$\text{Min} \left[\begin{array}{l} M_0, \\ M_{i+1} + 1, \\ M_0 b_j + m(1 - b_j) + 1, \\ M_{i+1} b_j + m(1 - b_j) + 2 \end{array} \right].$$

Therefore, by the hypotheses of the theorem, the linear inequality (15) holds for every corner point of the feasible set and we have completed the proof of sufficiency.

To prove that the conditions **C1** and **C2** are necessary, we construct for each condition a flow rate function for which the left-hand side of the condition corresponds to the online cost and the right-hand side (not including the factor r), to the offline cost.

First consider the condition **C1**. Fix any j and $i \leq j$. Let ε be a very small positive real and define f to be the flow rate function such that $f(0) = M$; during the interval $(0, \varepsilon]$, f

decreases continuously to the value $M_{i+1}+\varepsilon$, remains at this value until the time b_{j+1} and then, during the interval $(b_{j+1}, b_{j+1}+\varepsilon]$, decreases continuously to m, and remains at m until time 1. It is easy to see that when ε approaches zero, $S(f)$ approaches $M_i b_{j+1}+M_j(1-b_{j+1})+j$ and $\textsc{opt}(f)$ approaches

$$\operatorname{Min}\left[\begin{array}{l} M, \\ M_{i+1}+1, \\ M b_{j+1}+m(1-b_{j+1})+1, \\ M_{i+1} b_{j+1}+m(1-b_{j+1})+2 \end{array}\right].$$

Therefore, condition **C1** must hold for S to attain a competitive ratio r.

The necessity of condition **C2** is similarly justified. For any i, j with $i<j$, and ε, consider the flow rate function such that $f(0)=M$; during the interval $(0, \varepsilon]$, f decreases continuously to the value $M_{i+1}+\varepsilon$, remains at this value until the time $b_j-\varepsilon$ and then, during the interval $(b_j-\varepsilon, b_j]$, decreases continuously to m, and remains at m until time 1. When ε approaches zero, $S(f)$ approaches $M_i b_j+M_j(1-b_j)+j$ and $\textsc{opt}(f)$ approaches

$$\operatorname{Min}\left[\begin{array}{l} M, \\ M_{i+1}+1, \\ M b_j+m(1-b_j)+1, \\ M_{i+1} b_j+m(1-b_j)+2 \end{array}\right].$$

Hence, condition **C2** must hold for S to attain a competitive ratio r. ∎

4 Upper bounds

Equipped with the characterization theorem of the previous section, we construct and analyze in this section two replacement policies. The first, S^{**}, is strictly optimal in some cases and the second, S^{***}, is approximately optimal.

4.1 Construction of the refusal policy $S^{**}(m, M)$

We now construct a refusal policy and then calculate its competitive ratio. The basis for this construction is the changeover threshold sequence $\{M_i\}$ of the policy $S^*(m, M)$, the optimal policy against the restricted adversary (see Section 2). Recall that $S^*(m, M)=\{M_i\}_{i=1}^k$ where $k=k^*(m, M)$ and

$$\begin{cases} M_0 &= M; \\ M_{i+1} &= \frac{M_i+i}{r}-1, \quad 0 \leq i<k, \end{cases}$$

where $r=r^*(m, M)$ is the lower bound for the problem (or alternatively, the competitive ratio of $S^*(m, M)$ against the restricted adversary). For this r we already know that $k=\lceil(m+1)(r-1)\rceil$ (Lemma 2.8) and that $M_{k+1}=m$ (Lemma 2.7).

What would be a reasonable choice for the refusal times? A simple observation is that it is worthwhile to change over only if the changeover will save the policy at least the penalty it pays to perform it (1). To this end, we define the refusal time b_i (associated with the threshold M_i, $i = 1, 2, \ldots, k$) as the solution of the equation $(1 - b_i)(M_{i-1} - M_i) = 1$. Clearly, the b_i's thus defined are meaningful if and only if b_i lies in the time interval $[0, 1]$, and the sequence $\{b_i\}$ is non-increasing. Using the following two lemmata we establish these facts.

Lemma 4.1 *For all* $1 \leq i \leq k$, $M_{i-1} - M_i > M_i - M_{i+1}$.

Proof. Let $C = M + \frac{r^2}{(r-1)^2}$. Using formula (7) with $a = M$ and $m_i(r) = M_i$ we obtain

$$\begin{aligned} M_{i-1} - M_i &= \frac{C}{r^{i-1}} + \frac{i-1}{r-1} - \frac{r^2}{(r-1)^2} - \frac{C}{r^i} - \frac{i}{r-1} + \frac{r^2}{(r-1)^2} \\ &= C\left(\frac{1}{r^{i-1}} - \frac{1}{r^i}\right) + \frac{i-1-i}{r-1} \\ &= C\frac{r-1}{r^i} - \frac{1}{r-1} \\ &= \left(M + \frac{r^2}{(r-1)^2}\right) \cdot \frac{r-1}{r^i} - \frac{1}{r-1} \\ &= \frac{[M(r-1)^2 + r^2](r-1)}{r^i(r-1)^2} - \frac{1}{r-1} \\ &= \frac{M(r-1)^2 + r^2 - r^i}{r^i(r-1)}. \end{aligned} \tag{16}$$

Hence,

$$\begin{aligned} \frac{M_{i-1} - M_i}{M_i - M_{i+1}} &= \frac{r^{i+1}(r-1)}{r^i(r-1)} \cdot \frac{M(r-1)^2 + r^2 - r^i}{M(r-1)^2 + r^2 - r^{i+1}} \\ &= r \cdot \frac{M(r-1)^2 + r^2 - r^i}{M(r-1)^2 + r^2 - r^{i+1}} > 1, \end{aligned}$$

for every $r > 1$. ∎

Lemma 4.2 *For all* $1 \leq i \leq k$, $M_{i-1} - M_i > 1$.

Proof. By Lemma 4.1 the difference $M_{i-1} - M_i$ is decreasing with i. It is thus sufficient to prove the claim for $i = k$; that is, to prove that $M_{k-1} - M_k \geq 1$.

Using the identities $M_{k-1} = r(M_k + 1) - k + 1$ and $M_k = r(M_{k+1} + 1) - k = r(m+1) - k$, and the bound $k < r(m+1) - m$ (for $k = k^*(m, M)$) we obtain

$$\begin{aligned} M_{k-1} - M_k &= r(M_k + 1) - k + 1 - M_k \\ &= M_k(r-1) + r - k + 1 \end{aligned}$$

$$
\begin{aligned}
&= (r(m+1)-k)(r-1)+r-k+1\\
&> (r(m+1)-(r(m+1)-m))(r-1)+r-(r(m+1)-m)+1\\
&= 1.
\end{aligned}
$$

∎

From Lemma 4.2 it follows that $b_i \in (0,1)$. Further, by Lemma 4.1 it follows that the sequence $\{b_i\}$ is strictly decreasing. Hence, the b_i's are well defined refusal times. By convention we take $b_{k+1} = 0$ and thus we have

$$
b_i = \begin{cases} 1 - \frac{1}{M_{i-1}-M_i}, & 0 \le i \le k; \\ 0, & i = k+1. \end{cases}
$$

Denote this resulting refusal policy, comprised of the changeover threshold sequence $\{M_i\}$ and the above refusal time sequence $\{b_i\}$, by $S^{**}(m, M)$. Somewhat surprisingly (as will be shown), $S^{**}(m, M)$ attains the competitive ratio $r^*(m, M)$, when $m = 0$ or $M \le \frac{(m+2)^2}{m+1}$. As $r^*(m, M)$ is the lower bound for the problem, $S^{**}(m, M)$ is optimal in these cases. The first result of this section will be devoted to proving one of these bounds. In this rest of this section unless otherwise specified r and k will denote $r^*(m, M)$ and $k^*(m, M)$, respectively.

4.2 Optimality of $S^{**}(m, M)$ when $k = 1$

In this section we prove that $S^{**}(m, M)$ is optimal when $k = k^*(m, M) = 1$; i.e. when $S^{**}(m, M)$ consists of one threshold. An equivalent assumption to $k = 1$, is that $\sqrt{\frac{M}{m+1}} \le \frac{m+2}{m+1}$ (see Lemma 2.10).

Theorem 4.1 *For any m and M with $\sqrt{\frac{M}{m+1}} \le \frac{m+2}{m+1}$, $S^{**}(m, M)$ attains a competitive ratio $r^*(m, M) = \sqrt{\frac{M}{m+1}}$ and hence, it is optimal.*

Proof. As $r = r^*(m, M) = \sqrt{\frac{M}{m+1}}$ (see Lemma 2.9), $S^{**}(m, M)$ consists of one changeover threshold M_1, with a corresponding refusal time b_1. $M_0 = M$, $M_2 = m$, and we can write $S^{**}(m, M)$ explicitly:

$$
\begin{aligned}
M_1 &= M_0/r - 1\\
&= \sqrt{M(m+1)} - 1;\\
b_1 &= 1 - 1/(M_0 - M_1)\\
&= 1 - 1/(M - \sqrt{M(m+1)} + 1)\\
&= \frac{M - \sqrt{M(m+1)}}{M - \sqrt{M(m+1)} + 1}.
\end{aligned}
$$

We shall use theorem 3.1 to prove that $S^{**}(m, M)$ attains the ratio r. Since for these M and m, r is also a lower bound for the problem, it will follow that $S^{**}(m, M)$ is optimal. It remains to confirm that the conditions of the theorem hold.

Since b_2 is set to zero, proving that condition **C1** holds reduces to proving that the following three inequalities hold, corresponding to the three possible values of the pair (i, j) (i.e. $(0, 0)$, $(0, 1)$ and $(1, 1)$).

$$M \leq r \cdot \min \{M, M_1 + 1, Mb_1 + m(1 - b_1) + 1, M_1 b_1 + m(1 - b_1) + 2\}; \tag{17}$$

$$M_1 + 1 \leq r \cdot \min\{M, M_1 + 1, m + 1, m + 2\}. \tag{18}$$

$$M_1 + 1 \leq r \cdot \min\{M, m + 1, m + 2\}. \tag{19}$$

Consider first inequalities (18) and (19). In both these inequalities $m + 1$ is the minimum of the right-hand side (by assumption, $m+1 < M$) so it remains to prove that $M_1+1 \leq r(m+1)$. But since $m = M_2$, both inequalities follow by the identity $M_2 = \frac{M_1+1}{r} - 1$.

Next we prove that inequality (17) holds. We will consider four cases, that correspond to the four expressions inside the "min" operator in (17). The first inequality, $M \leq rM$, trivially holds. The second inequality, $M \leq r(M_1+1)$, follows by the identity $M_1 = M_0/r - 1$. Next we want to check that the following inequality holds:

$$M \leq r \cdot (Mb_1 + m(1 - b1) + 1). \tag{20}$$

Substituting r with $M/(M_1 + 1)$ and b_1 with $1 - 1/(M - M_1)$, inequality (20) becomes $(M - M_1)^2 \geq M - m$. Then, by substituting M_1 with $\sqrt{M(m + 1)} - 1$ it becomes equivalent to

$$M + 1 - \sqrt{M(m + 1)} - \sqrt{M - m} \geq 0.$$

The left hand side is a non-increasing function of m and thus attains its minimum value, zero, when $m = M - 1$.

The last inequality from (17) to check is

$$M \leq r \cdot (M_1 b_1 + m(1 - b_1) + 2).$$

Using again the substitutions $r = M/(M_1 + 1)$ and $b_1 = 1 - 1/(M - M1)$ this inequality becomes

$$(M - M_1)(M_1 - m - 1) \leq (M - M_1 - 1)(M_1 - m),$$

which is equivalent to $M - M_1 \geq M_1 - m$ that clearly holds (see Lemma 4.1).

Lastly, we have to verify that condition **C2** holds. Since $0 \leq i < j \leq 1$, we want to show that $Mb_1 + M_1(1 - b_1) + 1 \leq r \cdot \min\{M, M_1 + 1, Mb_1 + m(1 - b_1) + 1, M_1 b_1 + m(1 - b_1) + 2\}$ holds. Consider the left hand side. Using the identity $b_1 = 1 - 1/(M - M_1)$, we will show that it is exactly M.

$$Mb_1 + M_1(1 - b_1) + 1 \quad = \quad M(1 - \frac{1}{M - M_1}) + \frac{M}{M - M_1} + 1$$

$$
\begin{aligned}
&= M - \frac{M}{M - M_1} + \frac{M_1}{M - M_1} + 1 \\
&= M + \frac{M_1 - M}{M - M_1} + 1 \\
&= M - 1 + 1 = M.
\end{aligned}
$$

Thus, **C2** is identical to (17), and the proof is therefore complete. ∎

4.3 Optimality of $S^{**}(m, M)$ when $m = 0$

Similar to the proof of Theorem 4.1, we can prove optimality of $S^{**}(m, M)$ for all values of M whenever $m = 0$. As in the proof of Theorem 4.1, the idea is to verify that the two conditions of Theorem 3.1 hold with respect to $S^{**}(m, M)$. Nevertheless, in this case the proof we know of is extremely laborious and and involves lengthy "bookkeeping". We state the following theorem without a proof.[3]

Theorem 4.2 *Let $m = 0$. For each $M > 1$, $S^{**}(m, M)$ is $r^*(m, M)$-competitive and hence optimal.*

4.4 Is $S^{**}(m, M)$ always optimal?

In the previous sections we learned that $S^{**}(m, M)$ is optimal in two special cases. This was established by showing that $S^{**}(m, M)$ attains the lower bound in these special cases. Unfortunately, $S^{**}(m, M)$ does not attain the lower bound ratio $r^*(m, M)$ for all values of m and M. For example, for $M = 100$, $m = 3.625$, with $r^*(3.625, 100) = 2$ (see Example 2.1), it can be verified that $S^{**}(3.625, 100)$ violates some of the conditions of Theorem 3.1, and therefore, by the 'only if' direction of this theorem, it cannot attain a competitive ratio of 2. However, for this example it can be shown (using Theorem 3.1) that $S^{**}(3.625, 100)$ is approximately 2.0351-competitive. Similarly, we worked out various numerical examples in which $S^{**}(m, M)$ cannot attain a competitive ratio $r^*(m, M)$. For each such example, using theorem 3.1, we also calculated a number $\alpha \geq 1$ such that $S^{**}(m, M)$ is $[\alpha \cdot r^*(m, M)]$-competitive. Table 1 summarizes some of these examples. The table shows α as a function of M, with a fixed $m = 2$. We conjecture that $\lim_{M\to\infty} \alpha(M) = 1$. In other words, we conjecture that $r^*(m, M)$ is indeed the optimal competitive ratio for the replacement problem and that $S^{**}(m, M)$ is near-optimal.

4.5 A weak upper bound for $S^{**}(m, M)$ when $m > 0$

For the case $m > 0$, we now prove that for sufficiently large $r = r^*(m, M)$, $S^{**}(m, M)$ is r^2-competitive. In fact, we shall prove what appears to be a stronger statement - that a *time-independent* version of $S^{**}(m, M)$ is r^2-competitive. In contrast to a refusal policy, a

[3] A proof of Theorem 4.2 can be found in [6].

M	10	20	30	50	100	1000	10^9	10^{20}
α	1	1	1.00341	1.01065	1.00812	1.00281	1.0000105	$1 + .668 \times 10^{-9}$

Table 1: α as a function of M for a fixed $m = 2$

time-independent policy makes its changeovers without considering the time remaining (see Section 1.4). Intuitively, time-independent policies should be weaker than refusal policies. Given m and M, we construct a time-independent version of $S^{**}(m, M)$ by modifying the refusal times such that $b_i = 1$, $i = 1, \ldots, k$ (as before, $b_{k+1} = 0$). The performance of this resulting policy is, in fact, identical to the policy $S^*(m, M)$ from section 2, so we denote it by $S^*(m, M)$. The reason that we would like to view $S^*(m, M)$ as a (degenerate) refusal policy is that we shall use theorem 3.1 (that assumes a refusal policy) to derive an upper bound on its competitive ratio. The proof of the following theorem appears in Appendix A

Theorem 4.3 *Let $m > 0$ and $M > m + 1$ be any numbers such that $r = r^*(m, M) \geq 1.619$. Then, $S^*(m, M)$ is $[r \cdot \min\{r, m + 2\}]$-competitive.*

For any flow rate function f, $S^{**}(m, M)(f) \leq S^*(m, M)(f)$. This easily follows from the choice of the b_i's: if for some j, at time $t > b_j$ S^{**} refuses to change over to the threshold M_j but S^* does, the penalty (1) paid by S^* is greater than the savings in flow payments corresponding to this changeover. Further, since the b_i's are non-increasing, the same argument applies to any subsequent changeovers that may be performed by S^* (and will be "refused" by S^{**}). Hence, we have

Corollary 4.4 *For sufficiently large $r = r^*(m, M)$, $S^{**}(m, M)$ is $[r \cdot \min\{r, m+2\}]$-competitive.*

4.6 Approximately optimal policy for $m > 0$

We now present a stronger upper bound for the general case where m is positive. We shall identify a policy that for every positive m achieves a competitive ratio that for sufficiently large M, is within a constant factor of $r^*(m, M)$ (the constant is independent of M).

The policy we consider is a time-independent policy, $\{M_i\}_{i=1}^k$, where the sequence of changeover thresholds, $\{M_i\}$, is defined by the following recurrence relation. For each $\rho > 1$, set $k = \lfloor \rho \rfloor$. Then we define

$$\begin{cases} M_0 & \stackrel{\text{def}}{=} M; \\ M_{i+1} & \stackrel{\text{def}}{=} \frac{M_i + k}{\rho} - 1, \quad \text{integer } i \geq 1. \end{cases} \tag{21}$$

The following lemma states that for every positive m, the sequence defined by (21) decreases below m after a finite (possibly more than k) number of steps.

Lemma 4.3 *Fix $m > 0$. Then, there exists j such that $M_j < m$.*

Proof. It can be shown (for example, by induction on j) that

$$M_j = \frac{M}{\rho^j} + \frac{\rho - k}{(\rho-1)\rho^j} + \frac{k-\rho}{\rho-1}. \tag{22}$$

Since $k \leq \rho$ and $\rho > 1$, for every non-negative j,

$$\frac{\rho - k}{(\rho-1)\rho^j} + \frac{k-\rho}{\rho-1} \leq 0.$$

ρ^j increases with j so M_j decreases with j. Further, since $\lim_{j\to\infty} \frac{M}{\rho^j} = 0$, $\lim_{j\to\infty} M_j \leq 0$. ∎

In fact, by a straightforward refinement of the proof of Lemma 4.3, we can obtain the following, stronger statement.

Lemma 4.4 *For sufficiently large ρ, the sequence $\{M_i\}$ strictly decreases below m within $k+1$ steps.*

Call a ρ for which $M_{k+1} \leq m$ and $M_k > m$ *good.* Having Lemma 4.4 it is easy to see that good ρ exist. For each ρ, each $m > 0$ and each $M > m+1$, let $S^{***}_\rho(m, M)$ denote the policy $\{M_i\}$ (as defined by (21)). Now, by considering $S^{***}_\rho(m, M)$ as a (degenerate) time-dependent policy, we can apply Theorem 3.1 to prove that $S^{***}_\rho(m, M)$ is ρ-competitive for all good ρ and almost all values of M.

Lemma 4.5 *For any good ρ, any $m > 0$, and any $M \geq \max\{m+1, \frac{\lfloor\rho\rfloor}{\rho-1}\}$, $S^{***}_\rho(m, M)$ is ρ-competitive.*

Proof. First we note that the assumption $M > \frac{\lfloor\rho\rfloor}{\rho-1} = \frac{k}{\rho-1}$, implies that $M_1 + 1 - M \leq 0$, which means that $\min\{M_i+1, M\} = M_i+1$. We will use this fact later. Let us now specialize the three conditions of Theorem 3.1 to the case where $b_{k+1} = 0$ and $b_i = 1$, $i = 1, 2, \ldots, k$. That is, when the (degenerate) refusal policy is a time-independent policy. For $r = \rho$, the two conditions of Theorem 3.1 reduce to the following condition.

C1' $0 forall \leq i \leq j \leq k$,

$$M_i + j \leq \rho \cdot \min\{M_0, M_{i+1} + 1\};$$

However, by the definition of the M_i's, for all i, $M_i + k = \rho(M_{i+1}+1)$, so under the assumption that $M_1 + 1 \leq M$, condition **C1'** readily holds. Hence, by Theorem 3.1, $S^{***}_\rho(m, M)$ is ρ-competitive. ∎

Fix $m > 0$. For each (sufficiently large) M define $\rho(M)$ to be the minimum (infimum) good ρ. The next lemma establishes that the growth of $\rho(M)$ with M is asymptotically the same as the growth of $r^*(m, M)$.

Lemma 4.6 *For a fixed m, $\rho(M) = \Theta\left(\frac{\ln M}{\ln\ln M}\right)$.*

Proof. The proof is very similar to the proof of Lemma 2.11 and hence will only be sketched. For the analysis below abbreviate $\rho = \rho(M)$. We first isolate M in the inequality $M_k > m$ (ρ is good) to obtain

$$M > \rho^k \left(m + \frac{k-\rho}{\rho^k(\rho-1)} + \frac{\rho-k}{\rho-1} \right). \tag{23}$$

Assuming that ρ is greater than, say 2, it can be easily verified that the sum of the two right most terms inside the parentheses is always in $(0,1)$. In addition, we have the bounds $\rho - 1 < k \leq \rho$. Hence, $M > m\rho^k > m\rho^{\rho-1}$, so $\rho = O\left(\frac{\ln M}{\ln \rho}\right)$. Similarly, from $M_{k+1} \leq m$ (ρ is good) we have $M \leq (m+1)\rho^{k+1} \leq (m+1)\rho^{\rho+1}$. Therefore, $\rho = \Omega\left(\frac{\ln M}{\ln \rho}\right)$. Hence, a bootstrap derivation such as the one used in the proof of Lemma 2.11 completes the proof. ∎

By considering the lower bound from Section 2 we therefore obtain that for sufficiently large M, $S^{***}_{\rho(M)}(m,M)$ is approximately optimal; that is, S^{***} is $\rho(M)$-competitive and there exists a constant c such that and $\rho(M) \leq c \cdot r^*(m,M)$.

Example 4.1 It can be shown that for $M = 100$, $m = 3.625$, $\rho(M) = 2.905$. This means that $S^{***}_{2.905}(3.625, 100)$, with two changeover thresholds ($k = 2$), attains a competitive ratio of 2.905. For comparison, recall from previous examples that the lower bound for such m and M is $r^*(3.625, 100) = 2$, and that $S^{**}(3.625, 100)$ is 2.035-competitive.

5 Future work

This work is the first that studies the online replacement problem from the perspective of competitive analysis. Not surprisingly, we leave many unresolved questions. In this section we shall mention a few.

An intriguing question is whether or not one can obtain strictly tight bounds for this problem. A sensible starting point to investigate this question would be to try refining the time-independent policy $S^{***}_\rho(m,M)$ (e.g. consider a refusal policy based on $S^{***}_\rho(m,M)$). Also, it would be of interest to obtain tighter bounds on the performance of $S^{**}(m,M)$ for $m > 0$ and $M > (1 + \frac{1}{m+1})(m+2)$.

Our problem formulation is still somewhat simplistic to accurately model various real-life applications. We now point out several possible extensions that will lead towards more realistic models. A relatively simple extension (but probably, quite technical) would be the introduction of interest rates (or discount factors); i.e. to measure all costs by their present value. In our model, we assume that the changeover penalty is the same for all flow changeovers. It would be of interest to extend our results to the case where each flow rate is offered with a possibly different changeover penalty.

It is reasonable that one could obtain strictly better performance by allowing only discrete flow rate sequences. For example, instead of considering the continuous time horizon $[0,T]$, one could formulate the problem in terms of n, the length of the flow rate *sequence*. In this case, the sequence cannot vary (and in particular, drop) in an arbitrarily small time interval. It is probably the case that the lower bound for this discrete problem variant is an increasing function of n (approaching the lower bound in this paper). However, for practical purposes, n is bounded from above and is typically rather small.

There are various examples of online problems in which one obtains a dramatic improvement of the competitive ratio by allowing the online player to use randomization (e.g. [3, 7]). A further intriguing question is whether or not randomized replacement policies can improve the (optimal) deterministic performance obtained here.

Acknowledgments

We thank Allan Borodin, Vincent Feltkamp, Faith Fich, Sergiu Hart, Nati Linial, Motty Perry, Charlie Rackoff, Dan Rosen, Hisao Tamaki, Asher Wolinsky and mainly Brenda Brown for many helpful discussions and useful comments.

References

[1] R. Bellman. Equipment replacement policy. *Journal of the Society for Industrial and Applied Mathematics*, 3:133–136, 1955.

[2] R.E. Bellman and S.E. Dreyfus. *Applied Dynamic Programming*. Cambridge University Press, Princeton, New Jersey, 1962.

[3] A. Borodin, N. Linial, and M. Saks. An optimal on-line algorithm for metrical task system. *Journal of the ACM*, 39(4):745–763, 1992 1987.

[4] C. Derman. On optimal replacement rules when changes of state are markovian. In R. Bellmen, editor, *Mathematical Optimization Techniques*. 1963.

[5] S. Dreyfus. A generalized equipment study. *Journal of the Society for Industrial and Applied Mathematics*, 8:425–435, 1960.

[6] R. El-Yaniv. *On-Line Algorithms and Financial Decision Making*. PhD thesis, Department of Computer Science, University of Toronto, January 1994.

[7] A. Fiat, R. M. Karp, M. Luby, L. A. McGeoch, D. D. Sleator, and N. E. Young. Competitive paging algorithms. Technical Report CMU-CS-88-196, Carnegie Mellon University, Pittsburgh, PA, 1988.

[8] R.L. Graham. Bounds for certain multiprocessing anomalies. *SIAM Journal on Computing*, 17:263–269, 1969.

[9] G. Hadley. *Non Linear and Dynamic Programming.* Addison–Wesley Publishing Company, Inc., 1964.

[10] D.S. Johnson. *Near-Optimal Bin Packing Algorithms.* PhD thesis, Massachusetts Institute of Technology, 1973.

[11] D.S. Johnson, A. Demers, J.D. Ullman, M.R. Garey, and R.L. Graham. Worst-case performance bounds for simple one-dimensional packing algorithms. *SIAM Journal on Computing*, 3:299–325, 1974.

[12] E. Sheshinski and Y. Weiss. Optimum pricing policy under stochastic inflation. In E. Sheshinski and Y. Weiss, editors, *Optimal Pricing, Inflation, and the Cost of Price Adjustment.* The MIT Press, Cambridge, Massachusetts; London England, 1993.

[13] D.D. Sleator and R.E. Tarjan. Amortized efficiency of list update and paging rules. *Communication of the ACM*, 28:202–208, 1985.

[14] D.J Wilde and C.S. Beighther. *Foundations of Optimization.* Prentice–Hall, Inc., Englewood Cliffs, N.J., 1967.

[15] A.C. Yao. New algorithms for bin packing. *Journal of the ACM*, 27:207–227, 1980.

A Proof of Theorem 4.3

We apply Theorem 3.1 to calculate an upper bound on the competitive ratio of $S^*(m, M)$. In particular we show that it attains a competitive ratio αr, where α is a constant smaller than $\min(r, m+2)$ to be determined later. In other words, we show that the inequalities of conditions **C1** and **C2** of Theorem 3.1 hold (but with a competitive ratio αr instead of r). The way we go about this task is to *assume* that the inequalities in **C1** and **C2** hold, and from each inequality we derive some constraints (lower bounds) on α.

Throughout the analysis, recall that $S^*(m, M)$ has a degenerate refusal time sequence; that is, for all $1 \le i \le k$, $b_i = 1$, and $b_{k+1} = 0$.

Consider first condition **C1** for the case $j = k$. In this case, since $b_{j+1} = b_{k+1} = 0$, the minimum of the four operands of the min operator on the right hand side of **C1** is $m + 1$. Hence we should prove that $M_k + k \le r(m + 1)$ which surely holds with equality since $m = M_{k+1}$ and M_{k+1} equals, by definition, $(M_k + k)/r - 1$.

For all $0 \le j \le k - 1$, since $b_{j+1} = 1$, the minimum on the right hand side of **C1** is $M_{i+1} + 1$ and the left hand side of **C1** reduces to $M_i + j$. Similarly, for all $1 \le j \le k$, since $b_j = 1$, the minimum on the right hand side of **C2** is $M_{i+1} + 1$. Therefore, since the left hand side of **C2** reduces to $M_i + j$, for any $j \le k - 1$, **C1** is identical to **C2**.

Therefore, it remains to calculate the minimal $\alpha \ge 1$ for which

G1 for $0 \le i \le j \le k - 1$, $M_i + j \le \alpha r(M_{i+1} + 1)$;

G2 for $0 \le i \le k - 1$, $M_i + k \le \alpha r(M_{i+1} + 1)$.

Notice that when $i = j$, the inequality **G2** follows immediately (with $\alpha = 1$) by the identity $M_{j+1} = \frac{M_j + j}{r} - 1$.

For the rest of the analysis we need a few identities that will now be derived. First, using the identity $k = \lceil (m + 1)(r - 1) \rceil$ we write $k = (m + 1)(r - 1) + x$, with $0 \le x < 1$. Using this representation of k, the fact that $M_{k+1} = m$ and the identity $M_i = r(M_{i+1} + 1) - i$, it is not hard to verify the following identities that express a few of the smaller changeover thresholds in terms of m, r and x.

$$
\begin{aligned}
M_k &= m + 1 - x; \\
M_{k-1} &= m + r + 2 - x(r + 1); \\
M_{k-2} &= m + r^2 + 2r + 3 - x(r^2 + r + 1); \\
M_{k-2} &= m + r^3 + 2r^2 + 3r + 4 - x(r^3 + r^2 + r + 1);
\end{aligned}
$$

Consider **G1** (with $i < j$). $j \le k - 1 < (m + 1)(r - 1)$. Using the identity $M_{i+1} = \frac{M_i + i}{r} - 1$ we have

$$
\begin{aligned}
\alpha r M_{i+1} - M_i &= \alpha r(\frac{M_i + i}{r} - 1) - M_i \\
&= M_i(\alpha - 1) + \alpha i - \alpha r.
\end{aligned}
$$

Therefore, together with the upper bound on j we learn that it is sufficient to find (the minimal) α that satisfies $(m+1)(r-1) \le (\alpha-1)M_i + \alpha i$.

Since α and i are nonnegative, it is sufficient to find an α that satisfies

$$(m+1)(r-1) \le (\alpha-1)M_i. \tag{24}$$

Since the M_i are decreasing, it is sufficient to consider the minimum M_i which is M_{k-2}. Using the above identity for M_{k-2} we have $M_{k-2} > m+r+2$. Then, for

$$\alpha_1 \stackrel{\text{def}}{=} 1 + \frac{(m+1)(r-1)}{m+r+2},$$

we have an equality in (24). It is easy to check that $\alpha_1 < \min(r, m+2)$.

Consider **G2**. By the same argument as above we obtain that it is sufficient to find (the minimal) α that satisfies

$$k \le (\alpha-1)M_i + \alpha i. \tag{25}$$

We cannot solve this inequality in the same way we solved **G1**; that is, we cannot "sacrifice" the term αi in the right hand side of (25). The reason is that this time i can take values as large as $k-1$ and the right hand side can be very small. Therefore, we "pump-out" of (25) more constraints that correspond to large values of i, and then, after ensuring that M_i cannot be too small, we solve (25) exactly as we solved **G1**.

First, for the case $i = k-1$ we solve the equation $k = (\alpha-1)M_{k-1} + \alpha(k-1)$ for α to obtain

$$\alpha \ge \alpha(x) \stackrel{\text{def}}{=} \frac{mr+2r-rx+1}{mr+2r-rx}.$$

It is easy to see that $\alpha'(x)$ is positive, and therefore, $\alpha(1) = \frac{mr+r+1}{mr+r}$ maximizes $\alpha(x)$. But it is also evident that $\alpha(1) < \alpha_1$, and therefore, it is included in the previous constraint.

for $i = k-2$, the same procedure results in a new constraint:

$$\alpha_2 \stackrel{\text{def}}{=} \frac{mr+2r+2}{mr+2r}.$$

It can be verified that $\alpha_2 \le r$, for each $r \ge 1.619$, and that $\alpha_2 < m+2$.

Lastly, we can assume that $i \le k-3$. This time we solve $k = (\alpha-1)M_{k-3}$ for α to obtain

$$\alpha \ge \beta(x) \stackrel{\text{def}}{=} \frac{mr+4r+3+r^3+2r^2-x(r^3+r^2+r)}{mr^3+2r^2+3r+4-x(r^3+r^2+r+1)}.$$

By differentiation, it can be shown that $\beta(x)$ is increasing with x. Therefore, our last constraint α_3, is simply the value $\beta(1)$; namely,

$$\alpha_3 \stackrel{\text{def}}{=} \frac{mr+3r+3+r^2}{m+2r+3+r^2}.$$

It is easy to check that $\alpha_3 < m+2$, and that $\alpha_3 \le r$ for any $r \ge 1.175$.

Therefore, for $r \ge 1.619$, $\max \alpha_i \le \min(r, m+2)$, and the proof is complete. Note that in general, we have actually proven that $S^*(m, M)$ attains a competitive ratio αr where $\alpha = \max(\alpha_1, \alpha_2, \alpha_3)$. ∎