PERIODIC SCHEDULING WITH SERVICE CONSTRAINTS

SHOSHANA ANILY

Faculty of Management, Tel-Aviv University, Tel-Aviv, Israel 69978, anily@post.tau.ac.il

JULIEN BRAMEL

406 Uris Hall, Columbia University, New York, NY 10027, jdb8@columbia.edu

(Received October 1997; revisions received September 1998, December 1998; accepted January 1999)

We consider the problem of servicing a number of objects in a discrete time environment. In each period, we may select an object that will receive a service in the period. Each time an object is serviced, we incur a servicing cost dependent on the time since the object's last service. Problems of this type appear in many contexts, e.g., multiproduct lot-sizing, machine maintenance, and several problems in telecommunications. We assume that at most one object can be serviced in a given period. For the general problem with m objects, which is known to be $\mathcal{NP}$-Hard, we describe properties of an optimal policy; and for the specific case of $m = 2$ objects, we determine an optimal policy.

1. INTRODUCTION

We consider the problem of servicing a set of m objects over the infinite horizon. We assume that decisions concerning when to service an object are made in a discrete time environment, i.e., time is partitioned into periods of, for example, a day, a week, or a month. We assume that the cost of servicing an object is a function of the number of periods since the previous (or next) service to the object. We are interested in finding an optimal policy for servicing all objects that minimizes the average cost per period over the infinite horizon.

Multiproduct lot-sizing clearly falls into this problem framework. In this case, the objects are products, and servicing an object corresponds to ordering the product, or replenishing the product's inventory. The cost of servicing an object (ordering a product) may include a fixed cost for ordering along with the inventory carrying charges for the product over the interval until the next service.

In the case where there are no constraints on the set of objects serviced in a given period, and assuming fairly standard conditions on the cost of service, an optimal policy is easily constructed. Each object is serviced using equidistant (object-specific) time intervals. This optimal interval length between services can be found by minimizing the average service cost per period.

In this paper, we study the periodic scheduling problem with the following *service constraint*: At most one object can be serviced in a given period. This constraint may be because of accounting, space, workforce, or transportation considerations. We give several motivating examples below.

Anily et al. (1997a) consider this same problem in the context of scheduling preventive maintenance of a set of machines. Here the machines are the objects, and servicing a machine means performing maintenance. The service constraint may be because (1) a machine in maintenance cannot be in operation and required production levels prohibit maintaining more than one machine in a given period, and/or (2) there is not enough time, physical space, or other resources to maintain more than one machine in a period. In Anily et al. (1997a), the authors assume that the cost of *operating* a machine in a period is a linear (increasing) function of the number of periods since its last service. That is, they assume no fixed (set-up) cost for actually performing the maintenance. They present a simple algorithm to compute optimal policies for $m = 2$ and a nonpolynomial finite algorithm to find the optimal policy for general m. Because the complexity of this latter algorithm increases fairly fast with m, they also present heuristics and worst-case bounds. To date, it is not clear whether the problem considered in Anily et al. (1997a) is $\mathcal{NP}$-Hard. In Anily et al. (1997b) the same authors consider the problem with three machines. They show that the problem is either solvable to optimality or the optimum can be closely approximated, depending on the machine parameters. In the latter case, the authors provide a heuristic with a guaranteed worst-case bound of $1.0\bar{3}\%$.

Bar-Noy et al. (1997) consider a more general problem where out of the m objects, at most k can be serviced in a given period. In the context of machine maintenance, they assume object-dependent set-up costs as well as operating costs as in Anily et al. (1997a). They show that this problem is $\mathcal{NP}$-Hard for any $1 \leqslant k < m$. Therefore, the problem we study in this paper is $\mathcal{NP}$-Hard for a general number of objects. Bar-Noy et al. (1997) propose approximation algorithms for the case $k = 1$ that are based on properties of the Fibonacci numbers. The proposed heuristics achieve worst-case bounds of 9/8 in the case when there is no fixed cost

Subject classifications: Inventory/production: policies, maintenance/replacement. Mathematics: Convexity. Analysis of algorithms.
Area of review: MANUFACTURING OPERATIONS.

0030-364X/00/4804-0635 $05.00
1526-5463 electronic ISSN

for maintaining a machine, and 1.57 when there is a fixed cost for maintenance. They also prove that a simple greedy algorithm used in Anily et al. (1997a), which was shown empirically to generate close to optimal solutions, has a worst-case bound of 2.

In this paper, we study the problem of periodically servicing m objects under general service cost functions. We assume only that the total cost of servicing an object over t periods, e.g., once in period $p \geqslant 1$ and then not again until period $t + p$, is a convex function of t and is independent of p. One can verify that the problems analyzed in Anily et al. (1997a) and Bar-Noy et al. (1997) fall into this framework. We present general properties that an optimal policy must satisfy for this case. For the case of two objects ($m = 2$), we show that an optimal policy can be easily constructed. Specifically, we show that according to the cost functions there exists an optimal policy whose closed form can be either predetermined or is one of at most four possible simple forms.

In addition to the context of multiproduct lot-sizing and machine maintenance, there are a number of other applications of this same problem. Consider the case where a retailer must collect various products at a number of suppliers. Assume transportation operations are carried out by the retailer (and not the suppliers) who owns a single vehicle. The vehicle cannot visit multiple suppliers on a single day because (1) suppliers are not located close to each other, and/or (2) products cannot be mixed on the vehicle. In addition, such a policy may be attractive from a practical point of view; it balances the workload per day associated with the inventory replenishments and thus necessitates fewer resources to handle these operations.

Several applications in the telecommunications area are presented in Bar-Noy et al. (1997). We describe one here. A database containing m pages is accessed by clients. A broadcast-disk broadcasts only a limited number of these pages in each time period. With each page, there is a probability that a client wishes to access it. A client who wishes to access page i listens to the disk until the end of the time period in which page i is broadcast. When at most one page can be broadcast in a given period, the problem of minimizing the expected time spent by clients is therefore the model described in Anily et al. (1997a).

Related problems are analyzed by Chandrasekaran et al. (1992a, 1992b), Glass (1992, 1994), Hassin and Meggido (1991), Holte et al. (1992), Mok et al. (1989), and Wei and Liu (1983). For example, in Holte et al. (1992), each object has an upper bound on the number of periods between services it receives. In Mok et al. (1989), and Wei and Liu (1983), the exact service intervals are fixed, and the problem is to determine the minimum number of servers needed to form a feasible schedule, where one server can service one object per period.

In the next section, we give some preliminary definitions and notation. In §3, we present properties of an optimal policy for the general problem of m objects. The remainder of the paper is concerned with the case of two objects ($m = 2$). In §4, we consider cases that lend themselves to a straightforward solution method. In §5, we present a set of simple policies. In the last two sections, we consider the case where an optimal solution is not, at first glance, readily found. In §6, we present more specific properties of an optimal solution. In §7, we show that an optimal policy must exist in the set of simple policies presented in §5.

2. PRELIMINARIES

Let $\mathbf{R}_+ \equiv [0, +\infty)$ and $\mathbf{N} \equiv \{1, 2, 3, \ldots\}$. For each object $i = 1, 2, \ldots, m$, we define a general *service cost function* $F_i : \mathbf{N} \mapsto \mathbf{R}_+$. For integer $t \geqslant 1$, this service cost $F_i(t)$ specifies the total cost over t consecutive periods where object i is serviced in the first period of an interval of t periods, and then it is not serviced for the next $t - 1$ periods.

We let $f_i : \mathbf{N} \mapsto \mathbf{R}_+$ denote the respective *average cost function* for object i, which is simply:

$$f_i(t) \equiv \frac{1}{t} F_i(t), \quad \text{for } t \geqslant 1.$$

The average cost function $f_i(t)$ specifies the average cost per period for object i, over the t periods where a service to object i occurs in the first period and no service to object i occurs in the next $t - 1$ periods.

We require that these cost functions satisfy the following properties. For each $i = 1, 2, \ldots, m$:

(**P1**) F_i is convex, i.e., $F_i(t+1) - F_i(t) \leqslant F_i(t+2) - F_i(t+1)$ for all $t \geqslant 1$.
(**P2**) $f_i(t)$ is unbounded as $t \to \infty$.

In the context of lot-sizing, assume product (object) i has a fixed ordering cost of $K_i \geqslant 0$, a linear holding cost of $h_i > 0$ and a demand per period $d_i > 0$, then we have

$$\begin{aligned} F_i(t) &= K_i + (t-1)h_i d_i + (t-2)h_i d_i + \cdots + 2h_i d_i + h_i d_i \\ &= K_i + \frac{h_i d_i}{2} t(t-1), \end{aligned}$$

and $f_i(t) = K_i/t + (h_i d_i/2)(t-1)$. One can verify that in this case (**P1**) and (**P2**) are satisfied.

For each $i = 1, 2, \ldots, m$, define $f_i^* \equiv \min\{f_i(t) : t \geqslant 1 \text{ and integer}\}$ and let T_i^* denote the set of integers t satisfying $f_i(t) = f_i^*$. The set T_i^* represents those values of t that yield the minimal average cost per period, i.e., without any service constraints it is optimal to service object i once over a sequence of t periods, for any $t \in \mathsf{T}_i^*$. We show below that these definitions are well defined.

Lemma 1. (**P1**) *and* (**P2**) *imply the following, for each* $i = 1, 2, \ldots, m$:

(1) *If t (integer) is a local maximum of f_i, then either $t = 1$ or $f_i(t-1) = f_i(t) = f_i(t+1)$.*
(2) *T_i^* is a nonempty (finite) set of consecutive integers.*
(3) *$f_i(t)$ is strictly decreasing for $1 \leqslant t \leqslant \min(\mathsf{T}_i^*)$ and strictly increasing for $t \geqslant \max(\mathsf{T}_i^*)$. Also $F_i(t)$ is strictly increasing for $t \geqslant \max(\mathsf{T}_i^*)$.*

(4) *Let* $a, b > 0$ *be such that* $a+b \leqslant \min(\mathsf{T}_i^*)$; *then* $F_i(a) + F_i(b) > F_i(a+b)$.

(5) *For* $a+b=\tau$, $F_i(a)+F_i(b)$ *is a nondecreasing function of* $|a-b|$.

(6) *For* $2 \leqslant t \leqslant \min(\mathsf{T}_i^*)$, $F_i(t) - F_i(t-1) < f_i^*$, *and for* $t \geqslant \max(\mathsf{T}_i^*)$, $F_i(t+1) - F_i(t) > f_i^*$.

PROOF. (1) We prove this by showing that for each $i = 1, 2, \ldots, m$, there does not exist a $t \geqslant 2$ such that $f_i(t) \geqslant f_i(t-1)$ and $f_i(t) \geqslant f_i(t+1)$, where at least one of these inequalities is strict. The proof is by contradiction. Assume that there exists a $t \geqslant 2$ with the above property such that $f_i(t) \geqslant f_i(t-1)$ and $f_i(t) > f_i(t+1)$ (the proof of the other case is analogous). By (**P1**), $F_i(t) - F_i(t-1) \leqslant F_i(t+1) - F_i(t)$. Then:

$$t f_i(t) - (t-1) f_i(t-1) \leqslant (t+1) f_i(t+1) - t f_i(t) < (t+1) f_i(t) - t f_i(t) = f_i(t).$$

This implies $f_i(t) < f_i(t-1)$, contradicting our assumption.

(2) By (**P2**) and (1), T_i^* is bounded and therefore nonempty. In addition, (1) implies the set consists of consecutive integers.

(3) To prove the first part, consider three consecutive integers $t, t+1$, and $t+2$, with $t \geqslant 1$. From (**P1**), we must have $2(t+1) f_i(t+1) - t f_i(t) \leqslant (t+2) f_i(t+2)$. This leads to $f_i(t+1) \leqslant [(t+2) f_i(t+2) + t f_i(t)]/(2t+2) \equiv z$. From this expression we see that if $f_i(t+2) < f_i(t)$ then $f_i(t+1) \leqslant z < f_i(t)$. Similarly, if $f_i(t) < f_i(t+2)$ then $f_i(t+1) \leqslant z < f_i(t+2)$. Therefore, the two cases $f_i(t) = f_i(t+1) > f_i(t+2)$ and $f_i(t) < f_i(t+1) = f_i(t+2)$ are impossible. This proves the first part of (3).

For the second part, because $f_i(t)$ is strictly increasing for $t \geqslant \max(\mathsf{T}_i^*)$, then $F_i(t)$ is as well.

(4) Because $a+b \leqslant \min(\mathsf{T}_i^*)$ and $a, b > 0$, then by (3), $f_i(a) > f_i(a+b)$ and $f_i(b) > f_i(a+b)$. Therefore, $a f_i(a) + b f_i(b) > (a+b) f_i(a+b)$.

(5) This follows directly from (**P1**). To see this, assume without loss of generality that $a < b$, for $b = \tau - a$. Then (**P1**) implies: $F_i(a+1) - F_i(a) \leqslant F_i(b) - F_i(b-1)$, which implies $F_i(a+1) + F_i(b-1) \leqslant F_i(a) + F_i(b)$.

(6) Let $t_{\min} \equiv \min(\mathsf{T}_i^*)$, then for $t \leqslant t_{\min}$ (and $t \geqslant 2$):

$$\begin{aligned} F_i(t) - F_i(t-1) &\leqslant F_i(t_{\min}) - F_i(t_{\min}-1) \quad (\text{from } (\mathbf{P1})) \\ &= t_{\min} f_i^* - (t_{\min}-1) f_i(t_{\min}-1) \\ &= (t_{\min}-1)[f_i^* - f_i(t_{\min}-1)] + f_i^* \\ &< f_i^*. \end{aligned}$$

For the other case, let $t_{\max} \equiv \max(\mathsf{T}_i^*)$. Then for $t \geqslant t_{\max}$,

$$\begin{aligned} F_i(t+1) - F_i(t) &\geqslant F_i(t_{\max}+1) - F_i(t_{\max}) \quad (\text{from } (\mathbf{P1})) \\ &= (t_{\max}+1) f_i(t_{\max}+1) - t_{\max} f_i^* \\ &= (t_{\max}+1)[f_i(t_{\max}+1) - f_i^*] + f_i^* \\ &> f_i^*. \quad \square \end{aligned}$$

We remark that we have included condition (**P2**) for the sole purpose of excluding certain pathological cases where an optimal policy is to service a particular object once and then never again. Examples of this occur if, for example, $f_i(t)$ is nonincreasing for all $t \geqslant 1$ (e.g., $f_i(t) = 1/t$), then T_i^* is empty. Similarly, if $f_i(t) = f_i^*$ for all $t \geqslant t'$ for some $t' \geqslant 1$, then $\max(\mathsf{T}_i^*)$ is not finite.

A policy P is a sequence $P = [i_1, i_2, \ldots]$ where $i_k \in \{0, 1, 2, \ldots, m\}$ for $k = 1, 2, \ldots$ denotes the object that is serviced in period k (if $i_k > 0$) or denotes an empty period ($i_k = 0$). An empty period is one in which no object is serviced. For a policy P, let $C(t, P)$ denote the average cost over periods $1, 2, \ldots, t$.

We are interested in policies with bounded average cost. For each such policy P we define

$$C(P) \equiv \limsup_{t \to \infty} C(t, P).$$

A policy is *optimal* if it minimizes $C(P)$. We let $C^* \equiv \min_P C(P)$. In the next section, we show that C^* is well defined.

3. PROPERTIES OF AN OPTIMAL SERVICE POLICY

We define a *cyclic policy* (or simply a *cycle*) to be a schedule of services for all objects over a finite sequence of periods. To create a policy, the cycle is repeated ad infinitum. We use the term *basic cycle* to denote the shortest generating sequence of a cyclic policy. For example, for $m = 3$, a cyclic policy defined by $S = [123123]$ is not basic, but $S = [123]$ is. For a given cycle S, say $S = [1233]$, we say that the cycle uses a four-interval for object 1 (i.e., object 1 is serviced once over a four-interval or every four periods), and uses a four-interval for object 2. For object 3, S uses a one-interval and a three-interval.

A basic cycle S is *minimal* if there does not exist a consecutive subsequence in S servicing all objects which repeats itself. Thus, for $m = 3$, the basic cycle $S = [12323123]$ is not minimal since there is a repetition of the subsequence "123." We note that a minimal cycle is always basic, but not vice versa (as in the example just stated).

It is possible to show, in a similar manner as Anily et al. (1997a), that there exists an optimal solution which is a cyclic policy.

LEMMA 2. *For any policy* P, *there exists a policy* P^* *such that the number of periods between two consecutive services to object* i *is bounded from above by a constant* b_i, *for* $i = 1, 2, \ldots, m$, *and* $C(P^*) \leqslant C(P)$.

PROOF. Assume that b_i is a constant, which will be specified below. Let P be the given policy and let $\tau(P)$ be the first period in which object i is not serviced during the next $b_i + 1$ periods. We may assume there exists such a period, otherwise there is nothing to prove. In order to construct P^* we will define a sequence of policies P_k, for $k = 0, 1, 2, \ldots$, with $P_0 = P$ and such that $C(P_{k+1}) \leqslant C(P_k)$ and $\tau(P_{k+1}) > \tau(P_k)$, for $k \geqslant 0$.

We show how to construct P_{k+1} from P_k, for $k \geqslant 0$, as follows. Let i be the object that is serviced in period $\tau(P_k)$ and then not serviced until at least period $\tau(P_k) + b_i + 1$.

We will make use of the following observation. For each i, the value

$$F_i(t) - \left(F_i\left(\frac{t}{2} - m\right) + F_i\left(\frac{t}{2} + m\right)\right) \tag{1}$$

can be made arbitrarily large as $t \to \infty$. To see this, let $t_0 \equiv 2[m + \max(\mathsf{T}_i^*)]$ and note that for $t \geqslant t_0$ (i.e., $t/2 - m \geqslant \max(\mathsf{T}_i^*)$), $F_i(t)$ and $f_i(t)$ are strictly increasing according to Lemma 1, part (3). Also note that

$$F_i(t/2 + m) \leqslant m\left(\frac{F_i(t) - F_i(t/2)}{t/2}\right) + F_i(t/2). \tag{2}$$

To see this, consider the line connecting the points $(t/2, F_i(t/2))$ and $(t, F_i(t))$; it has slope $[F_i(t) - F_i(t/2)]/[t/2]$. The right-hand side of (2) is the value along this line at $x = t/2 + m$, which by the convexity of F_i and the fact that $t \geqslant t/2 + m$ must be greater than or equal to $F_i(t/2 + m)$. Thus for $t \geqslant t_0$:

$$\begin{aligned}
&F_i(t) - (F_i(t/2 - m) + F_i(t/2 + m))\\
&\quad \geqslant F_i(t) - \frac{2m}{t}(F_i(t) - F_i(t/2)) - F_i(t/2) - F_i(t/2 - m)\\
&\quad = \left(1 - \frac{2m}{t}\right)(F_i(t) - F_i(t/2)) - F_i(t/2 - m)\\
&\quad = \left(\frac{t/2 - m}{t/2}\right)(F_i(t) - F_i(t/2)) - F_i(t/2 - m)\\
&\quad = \left(\frac{t}{2} - m\right)(2f_i(t) - f_i(t/2) - f_i(t/2 - m))\\
&\quad > 2\left(\frac{t}{2} - m\right)(f_i(t) - f_i(t/2))\\
&\quad = (t - 2m)(f_i(t) - f_i(t/2)).
\end{aligned}$$

We now consider several cases. If $f_i(t) - f_i(t/2)$ is unbounded in t, then it follows that (1) can be made arbitrarily large. Otherwise, if $t[f_i(t) - f_i(t/2)]$ is unbounded in t, then again it follows that (1) can be made arbitrarily large. The only case remaining is when $t[f_i(t) - f_i(t/2)]$ is bounded for all $t \geqslant t_0$. Let us assume that this is the case. We show that f_i must then also be bounded, contradicting (**P2**). Say $t[f_i(t) - f_i(t/2)] \leqslant M$ for some constant M and for all $t \geqslant t_0$. Then

$$f_i(t) \leqslant f_i(t/2) + M/t. \tag{3}$$

Consider $f_i(t)$ at points $t = 2^k$ for integers $k \geqslant \bar{k} \equiv \lceil \log_2 t_0 \rceil$. A recursive application of (3) shows that $f_i(2^k) \leqslant f_i(2^{\bar{k}}) + M$ for all $k \geqslant \bar{k}$. Because f_i is strictly increasing for $t \geqslant t_0$, then it must also be bounded.

Having established that the value of (1) can be made arbitrarily large, we can now prove the claim. Assume b_i is such that $b_i + 1 - 2m$ is even and nonnegative. Let $v_i \equiv \frac{1}{2}(b_i + 1 - 2m) \geqslant 0$.

Renumber the periods so that period $\tau(P_k)$ is period 0. Let I be the interval from periods $v_i + 1$ to $v_i + 2m - 1$. The interval I is $2m - 1$ periods long, and since $v_i + 2m - 1 < b_i + 1$, it does not contain any service to object i. We distinguish between two cases:

Case 1: I has an empty period. Say this empty period is period t with $v_i < t < v_i + 2m$. Let P_{k+1} be an identical policy to P_k, except that we service object i in period t. The savings is therefore $F_i(b_i + 1) - [F_i(t) + F_i(b_i + 1 - t)]$. For $v_i < t < v_i + 2m$, using Lemma 1, part (5) with $a = v_i$ and $b = b_i + 1 - v_i$, and $a = v_i + 2m$ and $b = b_i + 1 - (v_i + 2m)$:

$$\begin{aligned}
&F_i(b_i + 1) - [F_i(t) + F_i(b_i + 1 - t)]\\
&\quad \geqslant F_i(b_i + 1) - \max\{F_i(v_i) + F_i(b_i + 1 - v_i),\\
&\qquad\qquad F_i(v_i + 2m) + F_i(b_i + 1 - v_i - 2m)\}\\
&\quad = F_i(b_i + 1) - [F_i(v_i) + F_i(v_i + 2m)] \quad \text{(by def. of } v_i)\\
&\quad = F_i(b_i + 1) - \left[F_i\left(\frac{b_i + 1}{2} - m\right) + F_i\left(\frac{b_i + 1}{2} + m\right)\right].
\end{aligned}$$

The proof then follows because there exists a large enough b_i for which this is nonnegative. Therefore, the new policy P_{k+1} has $C(P_{k+1}) \leqslant C(P_k)$ and $\tau(P_{k+1}) > \tau(P_k)$.

Case 2: I has no empty periods. Because I is $2m - 1$ periods long, there is an object, say object j, that is serviced at least three times in I. Let t_1, t_2, and t_3 be the periods where object j receives its first three services in I. We construct the policy P_{k+1} by taking P_k and replacing j's second service in I with a service to object i. The difference between the total cost of servicing object j in the policy P_k and the total cost of servicing object j in P_{k+1} is some finite amount

$$Q_j = F_j(t_3 - t_1) - (F_j(t_2 - t_1) + F_j(t_3 - t_2)).$$

If $Q_j \leqslant 0$ then exchange the second service to object j in I with a service to object i. There is no extra cost for deleting the service to object j and the proof follows from Case 1. Therefore we assume $Q_j > 0$. For object i, we save:

$$F_i(b_i + 1) - [F_i(t_2) + F_i(b_i + 1 - t_2)],$$
$$\text{for } v_i + 1 < t_2 < v_i + 2m - 1.$$

Proceeding as in Case 1, there exists a large enough b_i for which this is at least Q_j. Therefore, the new policy P_{k+1} has $C(P_{k+1}) \leqslant C(P_k)$ and $\tau(P_{k+1}) > \tau(P_k)$.

According to the above construction, policies $(P_j)_{j=k}^{\infty}$ coincide in the first $\tau(P_k)$ periods. As $\tau(P_k)$ is monotone increasing, we conclude that a limiting policy P^* exists. By construction, $C(P^*) \leqslant C(P)$. □

We define the *state* of the system at a given period as a vector $(s_1, s_2, \ldots, s_m)$, where $s_i \geqslant 0$ denotes the number of periods since the last service to object i; i.e., if object i is serviced in the period, then $s_i = 0$.

Theorem 1. *There exists an optimal cyclic policy.*

Proof. According to Lemma 2, the number of possible states for object i is bounded from above by a constant $b_i + 1$. Therefore, the total number of possible states, considering the m objects, is bounded by $\prod_{i=1}^{m}(b_i + 1)$. In view of the finiteness of the state space and the stationarity of the model, we conclude that for any optimal policy there exists an alternative optimal policy that is cyclic. □

We will therefore restrict our attention to cyclic policies. We remark that a similar nonpolynomial, but finite, algorithm as in Anily et al. (1997a) can be devised to find the optimal cycle. However, as noted in Anily et al., the running time grows rapidly with the size of the state space.

The following theorem will be useful in subsequent proofs.

THEOREM 2. *Given an optimal cycle* S:

(a) *there exists a consecutive subsequence* S' *of* S, *which defines an optimal minimal cycle, and*

(b) *for each object, the set of service intervals used in* S' *is a subset of the set of service intervals used in* S.

PROOF. We first prove (a). If S is minimal, there is nothing to prove, hence we assume S is not minimal. We describe a method called *pruning* a cycle. The method works as follows: Identify a subsequence of periods, where each object is serviced in the subsequence, such that the subsequence appears (at least) twice in nonoverlapping parts of S. Assume the first (second) occurrence of the subsequence is the interval from period t_1 (t_2) to period t'_1 (t'_2). Consider the cycle formed by the sequence from period $t'_1 + 1$ to period t'_2 (call this cycle S_1) and the cycle formed by the sequence from period $t'_2 + 1$ to the end of the basic cycle and then from period 1 to period t'_1 (call this cycle S_2). It is straightforward to show that the average cost of S is a weighted average of the cost of these two shorter cycles. To see this, let ℓ_i be the length of cycle S_i. Let $C(S)$ represent the average cost per period of a cycle S. Then

$$(\ell_1 + \ell_2)C(S) = \ell_1 C(S_1) + \ell_2 C(S_2).$$

Therefore, because S is optimal, we must have $C(S_1) = C(S_2)$, and the pruned cycle S_1 has the same average cost as S and has only one occurrence of the subsequence. Note that there may be another subsequence that is repeated in S_1. In this case, we can prune this cycle. This procedure will eventually lead to a minimal cycle.

We now prove (b). Given a cycle S, let t_1, t'_1, t_2, t'_2 be as defined above. Note that in S, the state of the system in period t'_1 is identical to the state of the system in period t'_2. Therefore, for each object the set of service intervals used in the cycle S_1 is a subset of the set of service intervals used in S. □

We conclude the following.

COROLLARY 1. *There exists an optimal cycle which is minimal, and hence basic.*

In view of Corollary 1, $C^* = \min_P C(P)$ is well defined. We note the following simple observation.

LEMMA 3. *The sum of the minimums of the average cost functions is a lower bound on the minimal policy cost, i.e.,* $C^* \geqslant \Sigma_{i=1}^m f_i^* \equiv LB$.

The following lemma of Anily et al. (1997b) will be necessary in subsequent proofs.

LEMMA 4. *An optimal basic cycle* S *has the following properties*:

(i) *Extending* S *by* k *periods by inserting any combination of empty periods or services to any set of objects increases the total cost of the basic cycle by at least* kC^*.

(ii) *Removing from* S *any set of* k *periods cannot reduce the total cost of the basic cycle by more than* kC^*.

PROOF. For a solution generated by a basic cycle R, denote by $K(R)$ the total cost of the solution generated by R during one basic cycle. Let T denote the length of the given optimal basic cycle S. Let S' denote a cycle of length $T + k$ derived from S as described in (i). Because S is optimal, we know that $K(S')/(T + k) \geqslant K(S)/T = C^*$, and hence $K(S') - K(S) \geqslant kC^*$, as claimed. Similarly, if S' denotes a basic cycle derived from S as described in (ii), we have $C^* = K(S)/T \leqslant K(S')/(T - k)$, and hence $K(S) - K(S') \leqslant kC^*$. □

4. TWO OBJECTS: THE EASY CASES

We consider the case of two objects, i.e., $m = 2$. In what follows, we use the convention: $i \in \{1, 2\}$ and, once i is defined, we use j to denote the index of the other object (i.e., $j \equiv 3 - i$). Recall that T_i^* is the set of integers t satisfying $f_i(t) = f_i^*$, for $i = 1, 2$.

In this section, we consider several special cases where an optimal cycle can be easily constructed.

4.1. The Case $1 \in \mathsf{T}_1^* \cup \mathsf{T}_2^*$

We first consider the case where $1 \in \mathsf{T}_1^* \cup \mathsf{T}_2^*$. In this case, there are *no empty periods* because it is cheaper to service an object in $\{j : 1 \in \mathsf{T}_j^*\}$ than not to.

We show that an optimal policy must be of one of two simple forms.

THEOREM 3. *If* $1 \in \mathsf{T}_1^* \cup \mathsf{T}_2^*$, *there exists an optimal cycle that is of one of the following forms*:

- *An optimal cycle is of the form* [21…11], *i.e., an optimal cycle consists of a service to object* 2 *and then* $k \geqslant 1$ *services to object* 1.
- *An optimal cycle is of the form* [12…22], *i.e., an optimal cycle consists of a service to object* 1 *and then* $k \geqslant 1$ *services to object* 2.

(*Note that the alternating policy* [12] *falls within both forms.*)

PROOF. Consider an optimal and minimal cycle S. Because S has no empty periods and is minimal, we can assume without loss of generality that S contains a single occurrence of the subsequence "12." As a result the optimal cycle must be of the form $[1\ldots 12\ldots 2]$, i.e., k 1's followed by ℓ 2's ($k\geqslant 1$ and $\ell\geqslant 1$). Assume, by contradiction, that both $k\geqslant 2$ and $\ell\geqslant 2$, i.e., the cycle S contains the subsequence "1122." By using Lemma 1, part (5), switching the second 1 and the first 2 in the subsequence (yielding the subsequence "1212") defines a policy with no worse cost. □

We can now work through and find the exact optimal policy as follows. Suppose first that $1\in \mathsf{T}_1^*$ but $1\notin \mathsf{T}_2^*$; then obviously the optimal cycle is of the form $[21\ldots 11]$. The average cost of such a cycle of length τ is

$$\frac{1}{\tau}(F_2(\tau)+(\tau-2)F_1(1)+F_1(2)) = f_2(\tau)+\frac{1}{\tau}[F_1(2)-2F_1(1)]+F_1(1). \tag{4}$$

From Lemma 1, part (3), for τ large enough, $f_2(\tau)$ is strictly increasing in τ and from (**P2**) it is unbounded. The value $F_1(2)-2F_1(1)$ is nonnegative as $1\in\mathsf{T}_1^*$. Hence an integer τ^* minimizing (4) can be easily found, and it clearly satisfies $\tau^*\geqslant \max(\mathsf{T}_2^*)$. The case $1\notin\mathsf{T}_1^*$ and $1\in\mathsf{T}_2^*$ is analogous.

If $1\in\mathsf{T}_1^*\cap\mathsf{T}_2^*$, then it is not clear a priori what is the exact form of an optimal cycle. However, for the two possible forms, $[211\ldots 1]$ and $[122\ldots 2]$, the best cycle of each form can be found as above. The optimal solution is then obtained by comparing the average cost of these two cycles.

Example: Lot-Sizing. In the context of lot-sizing of two products referred to in §2, we can obtain the following stronger result. Note that $1\in\mathsf{T}_1^*\cup\mathsf{T}_2^*$ corresponds to the case where at least one of the products should be ordered once per period. Let

$$\theta\equiv\tfrac{1}{2}[(K_1-h_1d_1)+(h_2d_2-K_2)].$$

THEOREM 4. *In the case of lot-sizing, if* $1\in\mathsf{T}_1^*\cup\mathsf{T}_2^*$ *then there are three cases*:

- $\theta\leqslant -h_2d_2$. *The optimal cycle is of the form* $[21\ldots 11]$, *i.e., an optimal cycle consists of a replenishment of product* 2 *and then* $k\geqslant 1$ *replenishments of product* 1.
- $-h_2d_2<\theta<h_1d_1$. *The optimal cycle is of the form* $[12]$, *i.e., the alternating policy is optimal.*
- $\theta\geqslant h_1d_1$. *The optimal cycle is of the form* $[12\ldots 22]$, *i.e., an optimal cycle consists of a replenishment of product* 1 *and then* $k\geqslant 1$ *replenishments of product* 2.

PROOF. We show that if $\theta>-h_2d_2$ ($\theta<h_1d_1$) then, in an optimal policy, product 1 (2) cannot be ordered in consecutive periods. This will prove the theorem. The two proofs are analogous, and thus we do only one. Assume that $\theta>-h_2d_2$ and product 1 appears in ℓ, $\ell\geqslant 2$, consecutive periods in an optimal cyclic policy P. (Also assume that ℓ is maximal with this property.) Let P' be the policy resulting from taking policy P and removing one of the ℓ consecutive orders of product 1. In addition, assume the cycle of P is τ periods long and therefore that of P' is $\tau-1$ periods long. Then

$$(\tau-1)C(P')=\tau C(P)-(K_1+\ell h_2d_2),$$

and thus

$$C(P')=C(P)+\frac{C(P)-(K_1+\ell h_2d_2)}{\tau-1}. \tag{5}$$

Now an alternating policy (one where replenishments occur in every period and alternate between the two products) has average cost:

$$\tfrac{1}{2}(K_1+K_2+h_1d_1+h_2d_2)<K_1+2h_2d_2\leqslant K_1+\ell h_2d_2,$$

where the first inequality follows from $\theta>-h_2d_2$. Therefore, $C(P)<K_1+\ell h_2d_2$, for $\ell\geqslant 2$, and combining this with (5), $C(P')<C(P)$, i.e., P cannot be an optimal policy. Hence, product 1 does not repeat in an optimal policy. This completes the proof. □

We can now work through and find the optimal cycle length, which we call τ. If $-h_2d_2<\theta<h_1d_1$, then $\tau=2$, and an alternating policy is optimal. If $\theta\geqslant h_1d_1$, then the optimal cycle consists of one replenishment of product 1 and $\tau-1$ replenishments of product 2. The average cost per period in this cycle is

$$\begin{aligned}&\tfrac{1}{\tau}(K_1+(\tau-1)K_2+h_2d_2+h_1d_1+2h_1d_1+\cdots\\&\quad+(\tau-1)h_1d_1)\\&\qquad=K_2-\frac{h_1d_1}{2}+\frac{K_1-K_2+h_2d_2}{\tau}+\frac{h_1d_1\tau}{2}.\end{aligned}$$

Then the (relaxed) optimal cycle time is

$$\tau^*=\sqrt{\frac{2(h_2d_2+K_1-K_2)}{h_1d_1}},$$

and the best τ is either $\lfloor\tau^*\rfloor$ or $\lceil\tau^*\rceil$. A symmetric formula exists for the other case ($\theta\leqslant -h_2d_2$).

4.2. The Case $1\notin\mathsf{T}_1^*\cup\mathsf{T}_2^*$

The case $1\notin\mathsf{T}_1^*\cup\mathsf{T}_2^*$ is not as straightforward. We deal here with two simple special cases; the remaining, more difficult, cases are dealt with in subsequent sections.

We will need the following definition. Let $\gcd(a,b)$ denote the *greatest common divisor* of the integers $a>0$ and $b>0$. If $\gcd(a,b)=1$, a and b are said to be *relatively prime*.

We first consider the case where there is a $T_1^*\in\mathsf{T}_1^*$ and a $T_2^*\in\mathsf{T}_2^*$ that are not relatively prime. We present below an optimal policy for this case for which the average cost coincides with the lower bound LB (see Lemma 3), proving the optimality of the policy.

THEOREM 5. *If there exist* $T_1^*\in\mathsf{T}_1^*$ *and* $T_2^*\in\mathsf{T}_2^*$ *that are not relatively prime, then there is an optimal cyclic policy whose average cost is LB.*

PROOF. The proof is based on constructing such a policy: Schedule a service to object 1 in period 1 and thereafter every T_1^* periods, and schedule a service to object 2 in period 2 and thereafter every T_2^* periods. Because $\gcd(T_1^*, T_2^*) > 1$, the services to objects 1 and 2 never coincide since $1 + \ell T_1^* = 2 + kT_2^*$ has no solution in integers ℓ and k. The cycle defined by periods $1, 2, \ldots, T_1^* T_2^*$ (which is nonbasic) is a feasible policy. This policy uses only T_1^*-intervals for object 1 and T_2^*-intervals for object 2, and therefore the average cost of the policy is $f_1(T_1^*) + f_2(T_2^*) = \mathrm{LB}$. □

Another simple case occurs when at least one of the functions f_i admits at least two minimizers, i.e., $|\mathsf{T}_1^*| \geqslant 2$, for some $i \in \{1, 2\}$. Assume $1 \notin \mathsf{T}_1^* \cup \mathsf{T}_2^*$, as the case $1 \in \mathsf{T}_1^* \cup \mathsf{T}_2^*$ was dealt with earlier in this section.

THEOREM 6. *If* $|\mathsf{T}_i^*| \geqslant 2$ *for some* $i \in \{1, 2\}$, *and* $1 \notin \mathsf{T}_1^* \cup \mathsf{T}_2^*$, *then there is an optimal cyclic policy whose average cost is LB.*

PROOF. Let $\{T_i^*, T_i^* + 1\} \subseteq \mathsf{T}_i^*$ for some $i \in \{1, 2\}$, and let $T_j^* \in \mathsf{T}_j^*$. We can assume that the conditions of Theorem 5 are not satisfied, i.e., $\gcd(T_i^*, T_j^*) = 1$ and $\gcd(T_i^* + 1, T_j^*) = 1$, otherwise there is nothing to prove. Moreover, we are given that $T_i^* > 1$ and $T_j^* > 1$. We construct an optimal cycle for this case. Service object j in periods $1,\ 1 + T_j^*,\ 1 + 2T_j^*, \ldots$. Service object i in period 2 and then, if period $2 + T_i^*$ is empty, service object i in period $2 + T_i^*$. Then service object i in period $2 + 2T_i^*$ if the period is empty. Continue in this manner until there is some period $2 + kT_i^*$ where object j is serviced for $k \geqslant 1$. Then service object i in the *next* period, i.e., period $2 + kT_i^* + 1$ (which is empty because $T_j^* > 1$). Because there is a repetition of the subsequence $[j, i]$ (first in periods 1 and 2 and then again in periods $2 + kT_i^*$ and $2 + kT_i^* + 1$), the sequence contains a basic cycle as a subsequence. In this basic cycle, object j is serviced every T_j^* periods and object i is serviced every T_i^* periods, except once when it is serviced once over a sequence of $T_i^* + 1$ periods. By using the fact that $f_i(T_i^*) = f_i(T_i^* + 1)$, we see that the average cost of the policy defined by this cycle is LB, proving its optimality. □

5. MINIMUM VIOLATION CYCLES

In this section, we present a set of special cycles that are useful in subsequent proofs and will prove to be very effective.

Due to the results of the previous section, from here on we make the following assumptions:

Assumption Ω.

- For each $i \in \{1, 2, \}$, T_i^* contains only one element, T_i^*.
- For each $i \in \{1, 2\}$, $T_i^* \geqslant 2$.
- T_1^* and T_2^* are relatively prime.

These are conditions for which an optimal policy has not yet been established. Under Assumption Ω there may be empty periods; the other cases were analyzed in §4.

We number the objects so that $T_1^* \leqslant T_2^*$. Assumption Ω then implies $T_2^* \geqslant 3$.

Given that $T_1^* > 1$ and $T_2^* > 1$ are relatively prime, it is clear that a cycle cannot be constructed where object i is serviced every T_i^* periods, for $i = 1, 2$. Therefore, in any cycle there will be occasions where some object i is always serviced once over a number of periods different from T_i^*, for $i = 1$ and/or 2. We call such a service interval a *violation*. We define a special set of cycles, which we call *minimum violation* basic cycles.

Let $i \in \{1, 2\}$ and $z_1, z_2 \in \{-1, +1\}$. There are four such minimum violation cycles, which are constructed as follows. The cycles depend on the values of T_1^* and T_2^*; however, we will omit this in the notation. We construct the minimum violation cycle, called $S_i(z_i)$. To this end, define

$$k_i^* \equiv \min\{k \geqslant 0 : (kT_i^* + (T_i^* + z_i)) \bmod (T_j^*) = 0\}, \tag{6}$$

and

$$m_i \equiv (k_i^* T_i^* + (T_i^* + z_i))/T_j^*. \tag{7}$$

The cycle $S_i(z_i)$ is $\tilde{T}_i(z_i) \equiv m_i T_j^*$ periods long. The index i denotes the object for which we have a violation, i.e., the object that is not serviced exclusively using T_i^*-intervals. More specifically, object i is serviced k_i^* times using an interval of length T_i^* and once using an interval of length $T_i^* + z_i$. Object j is always serviced once every T_j^* periods, a total of m_i times. The exact service policy $S_i(z_i)$ depends on the value z_i. More precisely, we specify the services for $m_i T_j^*$ consecutive periods starting in period $(1 + z_i)/2$:

- service object i in periods $1 + z_i + kT_i^*$, for $k = 0, 1, \ldots, k_i^*$, and
- service object j in periods $1 + \ell T_j^*$, for $\ell = 0, 1, \ldots, m_i - 1$.

(Observe that if $z_i = +1$, the specification of the cycle is from period 1 to period $m_i T_j^*$; if $z_i = -1$ it is from period 0 to period $m_i T_j^* - 1$.)

First, we prove that the minimum violation cycles are feasible and minimal.

LEMMA 5. *Under assumption* Ω, *for* $i \in \{1, 2\}$ *and* $z_i \in \{-1, +1\}$, *the minimum violation cycle* $S_i(z_i)$ *is feasible and minimal.*

PROOF. To prove feasibility, we have to show that there does not exist two overlapping services (services in the same period) to objects 1 and 2 in $S_i(z_i)$. Assume by contradiction that there do exist two overlapping services. This means that there exists an ℓ $(0 \leqslant \ell \leqslant m_i - 1)$ and a k $(0 \leqslant k \leqslant k_i^*)$ such that

$$\ell T_j^* + 1 = 1 + z_i + kT_i^*,$$

or equivalently

$$\ell T_j^* = kT_i^* + z_i. \tag{8}$$

We can write (8) as

$$((k-1)T_i^* + T_i^* + z_i) \bmod (T_j^*) = 0.$$

Note that $k \neq 0$, otherwise $\ell T_j^* = z_i$, which is impossible. But by the definition of k_i^* in (6), we obtain that $k_i^* \leqslant k-1$, contradicting our assumption that $k \leqslant k_i^*$.

To prove the minimality of $S_i(z_i)$, we have to show that there does not exist any repeating subsequence (with services to both objects) in $S_i(z_i)$. Observe that the last service in the cycle $S_i(z_i)$ is in period $\tau \equiv \max\{k_i^* T_i^* + 1 + z_i, (m_i - 1)T_j^* + 1\}$. So, if there is a repeating subsequence it must be between periods $(1+z_i)/2$ and τ. However, within this interval object 1 (object 2) is feasibly serviced at equidistant intervals of T_1^* (T_2^*). Because T_1^* and T_2^* are relatively prime, the length of this interval is less than $T_1^* T_2^*$, i.e., $\tau - ((1+z_i)/2) + 1 < T_1^* T_2^*$. When object k is feasibly serviced every T_k^* periods, for $k = 1, 2$, over an interval of length less than $T_1^* T_2^*$ (and T_1^* and T_2^* are relatively prime), it is easily verified that there is no repeating subsequence servicing both objects. $\square$

6. TWO OBJECTS: PROPERTIES OF AN OPTIMAL CYCLE

In this section, we prove a number of properties of an optimal cycle.

LEMMA 6. *Under assumption* Ω, *any optimal basic cycle has the following properties*:

(a) *object* 1 *is never serviced in three consecutive periods, and*

(b) *object* 2 *is never serviced in two consecutive periods.*

PROOF. Let S be an optimal and basic cycle.

Assume that S does not satisfy (a), i.e., say object 1 is serviced in periods 1, 2, and 3. Eliminating the service in period 2 and leaving this period empty saves $2F_1(1) - F_1(2)$, which is strictly positive from the fact that $T_1^* \geqslant 2$ and Lemma 1, part (4) (with $a = b = 1$). This contradicts the optimality of S.

Assume that S does not satisfy (b). Say object 2 is serviced in periods 1 and 2. Let t_1 (t_2) be the last (first) period before (after) period 1 in which object 1 is serviced. Note that $t_1 \leqslant 0$ and $t_2 \geqslant 3$. Let $p = t_2 - t_1$ be the length of the corresponding service interval for object 1. We prove (b) by applying Lemma 4, distinguishing between three cases:

Case i: $p > T_1^*$. In this case, remove period 2. The total savings is

$$
\begin{aligned}
&F_2(1) + F_1(p) - F_1(p-1)\\
&\quad = f_2(1) + F_1(p) - F_1(p-1)\\
&\quad > f_2(T_1^*) + F_1(p) - F_1(p-1)\\
&\quad > f_2(T_1^*) + f_1^* \geqslant C^*,
\end{aligned}
$$

contradicting Lemma 4, part (ii). (The first inequality follows from Lemma 1, part (3) and $T_1^* \geqslant 2$; the second inequality follows from Lemma 1, part (6); and the last inequality from the fact that $f_2(T_1^*) + f_1^*$ represents the average cost of a feasible policy where each of the objects is serviced every $T_1^* > 1$ periods.)

Case ii: $p < T_1^*$. In this case, we add an empty period in between period 1 and 2. The total increase in cost is

$$F_2(2) - F_2(1) + F_1(p+1) - F_1(p) < f_2^* + f_1^* \leqslant C^*,$$

contradicting Lemma 4, part (i). (The first inequality follows from Lemma 1, part (6), $T_2^* \geqslant 3$ and $p < T_1^*$. The second inequality follows from the fact that $f_2^* + f_1^*$ is a lower bound on the optimal average cost.)

Case iii: $p = T_1^*$. In this case, we add T_1^* periods in between period 1 and period 2 as follows: After servicing object 2 in period 1 we add $T_1^* + t_1 - 2$ empty periods, then a period with a service to object 1 and then $1 - t_1$ empty periods. It is easily verified that the intervals for object 1 starting in period t_1 and the following one are each T_1^* periods long. The total increase in cost is

$$
\begin{aligned}
&F_1(T_1^*) + F_2(T_1^* + 1) - F_2(1)\\
&\quad = T_1^* f_1^* + [F_2(T_1^* + 1) - F_2(T_1^*)]\\
&\qquad + [F_2(T_1^*) - F_2(T_1^* - 1)] + \cdots + [F_2(2) - F_2(1)]\\
&\quad < T_1^* f_1^* + T_1^* f_2^* \leqslant T_1^* C^*,
\end{aligned}
$$

contradicting Lemma 4, part (i). (The first inequality follows from Lemma 2.1, part (6), and $T_1^* + 1 \leqslant T_2^*$.) $\square$

We show that there exists an optimal cycle where each object is serviced using at most two interval lengths which are consecutive integers. Let $\mathscr{P}^*$ be the set of optimal basic cyclic policies generated by a minimal cycle. According to Corollary 1, $\mathscr{P}^*$ is not empty.

LEMMA 7. *Under Assumption* Ω, *there exists a policy in* $\mathscr{P}^*$ *that has each object serviced using at most two interval lengths that are consecutive integers.*

PROOF. Given an arbitrary policy $P \in \mathscr{P}^*$, if the claim of the lemma is not satisfied, we will perform a series of simple changes to the policy's generating optimal, minimal cycle that does not increase the policy's average cost. The resulting cycle uses, for each object, only service intervals whose lengths are two consecutive integers. If this cycle is not minimal, a minimal cycle of strictly shorter length can be constructed by using the pruning technique of Theorem 2. In addition, by part (b) of Theorem 2, for each object, the set of service intervals have the desired property as well.

Suppose there exists an optimal policy $P \in \mathscr{P}^*$, generated by a minimal cycle S, that uses two different service intervals for object $i \in \{1, 2\}$, say τ_i and $\tau_i + \Delta$ for $\Delta \geqslant 2$. Note that if $i = 2$ then $\tau_i \geqslant 2$ according to Lemma 6, part (b).

We start by assuming that S contains two adjacent intervals for object i, the first of length τ_i and the other of length $\tau_i + \Delta$, $\Delta \geqslant 2$. We will perform a shift of the services to object i that changes the τ_i-interval and the ($\tau_i + \Delta$)-interval into a ($\tau_i + 1$)-interval and a ($\tau_i + \Delta - 1$)-interval. This technique maintains feasibility, does not affect the services to object j, does not affect the total length of the cycle, and does not increase the total (or average) cost (according to Lemma 1, part (5)). Renumber the periods, if necessary, so that object i is serviced in periods $1, \tau_i + 1$ and $(\tau_i + 1) + (\tau_i + \Delta)$. Therefore, period $k \equiv \tau_i + 1$ is the beginning of the $(\tau_i + \Delta)$-interval. Note that the next service to object i, after the one in period k, is not until period $k + \tau_i + \Delta \geqslant k + 3$. If period $k + 1$ is empty, then a simple shift of the service to i currently in period k to period $k + 1$ achieves the desired result. Otherwise, if period $k + 1$ has a service to object j, then by the minimality of S all services to object i in S, except the service to object i in period k, are followed by an empty period. We claim that we can assume without loss of generality that period $k + 2$ is empty. To prove this claim, note that because $k + \tau_i + \Delta \geqslant k + 3$, if there is a service in period $k + 2$, it must be to object j. Also, in such a case, $j \neq 2$ because otherwise object $j = 2$ is serviced in two consecutive periods ($k + 1$ and $k + 2$), contradicting Lemma 6, part (b). Therefore, $i = 2$ and $j = 1$. Note that object 2 is not serviced in period $k + 3$ (because $k + \tau_2 + \Delta \geqslant k + 4$). Therefore, consider object 1. It is serviced in period $k + 1$, and assume that it is also serviced in period $k + 2$. By Lemma 6, part (a), object 1 cannot be serviced in period $k + 3$. In this case, period $k + 3$ is empty, and the service to object $j = 1$ in period $k + 2$ can be moved to the empty period $k + 3$ at no additional cost (according to Lemma 1, part (5)). Therefore, we can assume period $k + 2$ is empty. Now shift all services to object i in S forward by one period except for the service in period k, which is shifted forward two periods (to period $k + 2$). This accomplishes the desired result (i.e., we have transformed the τ_i and $\tau_i + \Delta$ intervals into a $(\tau_i + 1)$- and a $(\tau_i + \Delta - 1)$-interval).

This proves that any two consecutive service intervals for object i either are the same length or their lengths differ by exactly one period. Assume now that S contains an interval, call it I_0, of length τ_i for object i, then $z > 0$ intervals of length $\tau_i + 1$ followed by an interval, call it I_2, of length $\tau_i + 2$. We show below that this cycle can be transformed, at no additional cost, to a cycle in which the τ_i-interval and the $(\tau_i + 2)$-interval are adjacent, which was dealt with above. Let period 0 be the period that is at the beginning of I_0, so that object i is serviced in period 0.

We first describe a shifting technique that reduces z by one each time and does not change the cost of the solution.

As long as $z > 0$, we do the following. Let period $k = \tau_i + z(\tau_i + 1)$ be the period at the beginning of I_2, when object i is serviced. If there is no service to either object in period $k + 1$, then move the service to object i currently in period k to period $k + 1$. The cost of the solution does not change since the number of intervals of each type does not change; a $(\tau_i + 2)$-interval becomes a $(\tau_i + 1)$-interval and a $(\tau_i + 1)$-interval becomes a $(\tau_i + 2)$-interval and in between the τ_i-interval and the $(\tau_i + 2)$-interval there are $(z - 1)$ $(\tau_i + 1)$-intervals. Otherwise, if there is a service in period $k + 1$, it must be a service to object j because the next service to object i is not until period $k + \tau_i + 2 \geqslant k + 3$. Because the cycle is minimal, there cannot be a repeating subsequence. Therefore, all services to object i in the cycle, except for the one in period k, are followed by an empty period. In addition, period $k + 2$ can be assumed to be an empty period for the same reasons as detailed above (when analyzing the $(\tau_i + \Delta)$-interval). Shift all services to object i forward by one period, except for the service in period k, which is shifted forward two periods (to period $k + 2$). This reduces z by one without changing the total number of intervals of each type used, for each object. This procedure can be repeated until $z = 0$ which was dealt with above. □

Let $\mathscr{P}^{**} \subseteq \mathscr{P}^{*}$ denote those policies of $\mathscr{P}^{*}$ where object i is serviced using at most two interval lengths which are consecutive integers, for $i \in \{1, 2\}$. Lemma 7 shows that $\mathscr{P}^{**}$ is nonempty. The following lemma shows that there exists an optimal policy in $\mathscr{P}^{**}$ where the only service intervals used to service object i are either T_i^* and $T_i^* - 1$ or T_i^* and $T_i^* + 1$, for $i \in \{1, 2\}$.

LEMMA 8. *Under assumption* Ω, *there exists an optimal policy* $P \in \mathscr{P}^{**}$ *where for object* $i \in \{1, 2\}$ *the service intervals used are of length either* $T_i^* - 1$ *and* T_i^* *or* T_i^* *and* $T_i^* + 1$.

PROOF. According to assumption Ω, $T_i^* \geqslant 2$ for each $i = 1, 2$. From Lemma 7, there exists an optimal policy $P \in \mathscr{P}^{**}$ such that the set of service intervals for object i consists of either a single value τ_i or two values τ_i and $\tau_i + 1$.

By definition, $f_i(t)$ is minimized at T_i^*. Let $T_i^{(2)}$ be the second best minimizer of $f_i(t)$. (Note assumption Ω implies $f_i(T_i^*) < f_i(T_i^{(2)})$.) By Lemma 1, part (3), without loss of generality we can assume that $T_i^{(2)}$ is either $T_i^* - 1$ or $T_i^* + 1$. Suppose there exists an optimal policy P for which the claim of this lemma does not hold for at least one object, say object $i \in \{1, 2\}$. We construct a new policy P' that services object j every T_j^* periods and services object i in intervals of at most two lengths: T_i^* and $T_i^{(2)}$. Let $z_i = T_i^{(2)} - T_i^*$, and construct the minimum violation cycle $S_i(z_i)$. The average cost of object j in P' is $f_j(T_j^*)$, which is a lower bound on the average cost for object j, and therefore is not higher than that cost in P. The average cost of object i is a weighted average of $f_i(T_i^*)$ and $f_i(T_i^{(2)})$ which is strictly better than *any* weighted average of any $f_i(x)$ and $f_i(x + 1)$ where $x \geqslant T_i^* + 1$ or $x \leqslant T_i^* - 2$. □

The policy proposed in the proof of Lemma 8 has one object serviced using its optimal interval length (e.g., every T_j^* periods), and the other object is serviced using two types of intervals: its optimal interval length (T_i^* periods) and its second best interval length ($T_i^{(2)}$ periods). That is, the only time we deviate from the lower bound LB is when we service object i after $T_i^{(2)}$ periods instead of T_i^* periods. At what rate this occurs depends on the *length* of the cycle $S_i(T_i^{(2)} - T_i^*)$. It is this issue that we deal with in the remainder of this paper. It is interesting to note that it may be optimal to use T_i^* along with a value other than the second best value. For example, say $T_1^* = 4$ and $T_2^* = 5$ with $T_1^{(2)} = 5$. The cycle formed by using 4- and 5-intervals for object 1 and 5-intervals for object 2 leads to the cycle [12000]. However, using 3- and 4-intervals for object 1 along with 5-intervals for object 2 leads to the cycle [120010201002100]. The latter cycle may have a smaller average cost.

Lemmas 7 and 8 have established that there exists an optimal policy that is generated by a minimal basic cycle with the following particular structure, which we call canonical structure.

DEFINITION 1. Under assumption Ω, a minimal cycle S is *canonical* if there exist $z_1, z_2 \in \{-1, +1\}$ and integers $p_1, p_2, q_1, q_2 \geqslant 0$ such that object i is serviced (in S) p_i times using an interval of length T_i^* and q_i times using an interval of length $T_i^* + z_i$, for $i = 1, 2$.

We say a canonical cycle S is generated by $(z_1, p_1, q_1, z_2, p_2, q_2)$. Any canonical cycle clearly satisfies $p_1 T_1^* + q_1(T_1^* + z_1) = p_2 T_2^* + q_2(T_2^* + z_2)$. Also because T_1^* and T_2^* are relatively prime, a canonical cycle must have $q_1 + q_2 > 0$.

We now derive several additional properties of canonical cycles.

LEMMA 9. *Under assumption* Ω, *in a canonical cycle the number of services to object* 2, *i.e.*, $p_2 + q_2$, *is bounded as follows*:

$$p_2 + q_2 \leqslant \begin{cases} T_1^* - 1, & \text{if } q_1 = 0, \\ T_1^*, & \text{if } q_1 > 0. \end{cases}$$

PROOF. Let S be a canonical cycle. First note that $T_2^* > T_1^*$ and $|z_i| = 1$ (for $i = 1, 2$) imply that any interval between services to object 1 in S contains *at most* one service to object 2. Arbitrarily number the services to object 2 in S: $1, 2, \ldots, p_2 + q_2$. Let $n_k \geqslant 1$ be the number of periods since the last service to object 1 when object 2 receives service number k, for $k = 1, 2, \ldots, p_2 + q_2$. Because S is minimal, each n_k must be a distinct integer.

If $q_1 = 0$, then $n_k \in \{1, 2, \ldots, T_1^* - 1\}$, for each k; hence $p_2 + q_2 \leqslant T_1^* - 1$. If $q_1 > 0$, there may exist (if $z_1 = +1$) at least one $(T_1^* + 1)$-interval for object 1 in S and therefore $n_k \in \{1, 2, \ldots, T_1^*\}$, for each k; hence $p_2 + q_2 \leqslant T_1^*$. □

The following lemma characterizes the average cost of canonical cycles. It will be useful in the next section. Define $\delta_i(z_i) \equiv (T_i^* + z_i)[f_i(T_i^* + z_i) - f_i(T_i^*)] > 0$, for $i = 1, 2$. (The strong inequality follows from assumption Ω. The other cases were solved in §4.)

LEMMA 10. *Under assumption* Ω, *let* P *denote the policy generated by a canonical cycle of length* $T = p_i T_i^* + q_i(T_i^* + z_i)$, *for* $i = 1, 2$; *then*:

$$C(P) = LB + \frac{q_1 \delta_1(z_1) + q_2 \delta_2(z_2)}{T}.$$

PROOF.

$$\begin{aligned} C(P) &= \frac{1}{T} \sum_{i=1}^{2} [p_i T_i^* f_i(T_i^*) + q_i(T_i^* + z_i) f_i(T_i^* + z_i)] \\ &= \frac{1}{T} \sum_{i=1}^{2} [p_i T_i^* (f_i(T_i^*) - f_i(T_i^*)) \\ &\qquad + q_i(T_i^* + z_i)(f_i(T_i^* + z_i) - f_i(T_i^*))] \\ &\quad + \sum_{i=1}^{2} f_i(T_i^*) \\ &= \frac{1}{T} \sum_{i=1}^{2} [q_i(T_i^* + z_i)(f_i(T_i^* + z_i) - f_i(T_i^*))] + \text{LB} \\ &= \frac{1}{T} \sum_{i=1}^{2} q_i \delta_i(z_i) + \text{LB}. \quad \square \end{aligned}$$

7. TWO OBJECTS: THE OPTIMALITY OF MINIMUM VIOLATION CYCLES

In this section, we combine the properties described in the previous sections to prove our main result: *There exists a minimum violation cycle that is optimal.*

We show that for any canonical cycle S generated by $(z_1, p_1, q_1, z_2, p_2, q_2)$, the average cost of the better of the two minimum violation cycles $S_1(z_1)$ or $S_2(z_2)$ is no more than the average cost of S. Let $C(S)$ denote the average cost per period of a policy generated by S.

LEMMA 11. *Under assumption* Ω, *given any canonical cycle* S *generated by* $(z_1, p_1, q_1, z_2, p_2, q_2)$, *we have*

$$\min\{C(S_1(z_1)), C(S_2(z_2))\} \leqslant C(S).$$

PROOF. Recall that $\delta_i(z_i) \equiv (T_i^* + z_i)[f_i(T_i^* + z_i) - f_i(T_i^*)] > 0$, for $i = 1, 2$, and

$$\tilde{T}_1(z_1) = m_1 T_2^* = k_1^* T_1^* + T_1^* + z_1, \tag{9}$$

$$\tilde{T}_2(z_2) = m_2 T_1^* = k_2^* T_2^* + T_2^* + z_2, \tag{10}$$

where m_i and k_i^* are as defined in §5. Let T denote the length of the canonical cycle S, where:

$$T = p_1 T_1^* + q_1(T_1^* + z_1), \tag{11}$$

and

$$T = p_2 T_2^* + q_2(T_2^* + z_2). \tag{12}$$

Because $S_i(z_i)$ is a canonical cycle, in view of Lemma 10, for $i = 1, 2$, the average cost of the policy defined by $S_i(z_i)$ is $\text{LB} + \delta_i(z_i)/\tilde{T}_i(z_i)$. The average cost of S (of length T) is $\text{LB} + (1/T)(q_1 \delta_1(z_1) + q_2 \delta_2(z_2))$. Therefore, to prove that the cheaper (in terms of average cost) of the two cycles $S_1(z_1)$ and $S_2(z_2)$ is at least as good as S, we need to show that

$$\frac{q_1 \delta_1(z_1) + q_2 \delta_2(z_2)}{T} \geqslant \min\left\{\frac{\delta_1(z_1)}{\tilde{T}_1(z_1)}, \frac{\delta_2(z_2)}{\tilde{T}_2(z_2)}\right\}. \tag{13}$$

To prove (13), we investigate the relationship between the values $(z_1, p_1, q_1, z_2, p_2, q_2)$, defining the canonical cycle S and the numbers k_1^*, k_2^*, m_1 and m_2.

One can verify (using (9)–(12)) that $T_1^*(p_1 - q_1k_1^* - q_2m_2) = T_2^*(p_2 - q_1m_1 - q_2k_2^*)$, and because T_1^* and T_2^* are relatively prime it follows that there exists an integer ℓ^* such that

$$p_1 = q_1k_1^* + q_2m_2 + \ell^* T_2^*, \tag{14}$$

and

$$p_2 = q_1m_1 + q_2k_2^* + \ell^* T_1^*. \tag{15}$$

We now show that Lemma 9 implies $\ell^* \leqslant 0$. To see this, consider both cases: $q_1 = 0$ and $q_1 > 0$. If $q_1 = 0$ then Lemma 10 implies $0 \leqslant p_2 < T_1^*$. This with (15) implies $\ell^* \leqslant 0$. If $q_1 > 0$, then Lemma 10 implies $0 \leqslant p_2 \leqslant T_1^*$. This with (15) and $m_1 > 0$ imply $\ell^* \leqslant 0$.

We can now prove (13). Hence,

$$\begin{aligned}
T &= p_1T_1^* + q_1(T_1^* + z_1)\\
&= (q_1k_1^* + q_2m_2 + \ell^* T_2^*)T_1^* + q_1(T_1^* + z_1)\\
&\quad \text{(from (14))}\\
&\leqslant (q_1k_1^* + q_2m_2)T_1^* + q_1(T_1^* + z_1) \quad (\text{because } \ell^* \leqslant 0)\\
&= (q_1k_1^* + q_2m_2)T_1^* + q_1(m_1T_2^* - k_1^*T_1^*) \quad \text{(from (9))}\\
&= q_1m_1T_2^* + q_2m_2T_1^*\\
&= q_1\tilde{T}_1(z_1) + q_2\tilde{T}_2(z_2)\\
&= \frac{q_1\delta_1(z_1)}{\delta_1(z_1)/\tilde{T}_1(z_1)} + \frac{q_2\delta_2(z_2)}{\delta_2(z_2)/\tilde{T}_2(z_2)}\\
&\quad (\text{because } \delta_1(z_1) > 0 \text{ and } \delta_2(z_2) > 0)\\
&\leqslant \frac{q_1\delta_1(z_1) + q_2\delta_2(z_2)}{\min\{\delta_1(z_1)/\tilde{T}_1(z_1), \delta_2(z_2)/\tilde{T}_2(z_2)\}},
\end{aligned}$$

which is precisely (13). □

We can now conclude the following.

THEOREM 7. *Under assumption Ω, an optimal basic cycle exists in the set of four minimum violation cycles.*

PROOF. The result follows because Lemmas 7 and 8 establish the fact that there exists an optimal minimal cycle that is canonical, and Lemma 11 shows that, given any canonical cycle, there always exists a minimum violation cycle of equal or inferior average cost per period. □

We conclude therefore that under Assumption Ω, an optimal solution to the problem can be found by simply calculating the average cost of each of the four minimum violation basic cycles $S_1(+1)$, $S_1(-1)$, $S_2(+1)$, and $S_2(-1)$, and choosing the one with minimum average cost.

REFERENCES

Anily, S., C. A. Glass, R. Hassin. 1997a. The scheduling of maintenance service. *Discrete Appl. Math.* **82** 27–42.

——,——,——. 1999. Scheduling maintenance service to three machines. *Annals of Oper. Res.* Special Vol. on *Combinatorial Optimisation*: *Recent Advances in Theory and Practice.*

Bar-Noy, A., R. Bhatia, J. Naor, B. Schieber. 1997. Minimizing service and operations costs of periodic scheduling. Working Paper, Electrical Engineering Dept., Tel-Aviv University.

Chandrasekaran, R., A. Mukhopadhyay, P. Yang. 1992a. Optimal ordering policies for a capacity constrained two-item inventory system. Working Paper, The University of Texas at Dallas.

——,——,——. 1992b. Optimal ordering policies for a capacity constrained multi-item inventory system. Working Paper, The University of Texas at Dallas.

Glass, C. A. 1992. Feasibility of scheduling lot-sizes of three products on one machine. *Management Sci.* **38** 1482–1494.

——. 1994. Feasibility of scheduling lot-sizes of two frequencies on one machine. *Euro. J. Oper. Res.* **75** 354–364.

Hassin, R., N. Meggido. 1991. Exact computation of optimal inventory policies over an unbounded horizon. *Math. Oper. Res.* **16** 534–546.

Holte, R., L. Rosier, I. Tulshinsky, D. Varvel. 1992. Pinwheel scheduling with two distinct numbers. *Theoretical Comput. Sci.* **100** 105–135.

Mok, A., L. Rosier, I. Tulshinsky, D. Varvel. 1989. Algorithms and complexity of the periodic maintenance problem. *Microprocessing and Microprogramming* **27** 657–664.

Wei, W., C. Liu. 1983. On a periodic maintenance problem. *Oper. Res. Lett.* **2** 90–93.