T-COM/35/1/11698

On the Optimality of Cyclic Transmission in Teletext Systems

M. H. Ammar
J. W. Wong

Reprinted from: IEEE TRANSACTIONS ON COMMUNICATIONS, Vol. COM-35, No. 1, JAN. 1987.

On the Optimality of Cyclic Transmission in Teletext Systems

MOSTAFA H. AMMAR, MEMBER, IEEE, AND J. W. WONG

***Abstract*—Teletext is a one-way information delivery system where pages of information are broadcast to all users in a continuous manner. System response time is an important consideration in the design of teletext systems. One factor contributing to response time is the order in which pages are transmitted. In this paper, we formulate the problem of determining the sequence of page transmissions as a Markovian decision process. Using this formulation we show that, from a response time point of view, a cyclic order of page transmissions is optimal. We also describe two algorithms for designing a teletext broadcast cycle.**

I. Introduction

TELETEXT is an information delivery system where pages of information are broadcast to all users in a continuous manner [1]–[3]. The configuration of a typical teletext system is shown in Fig. 1. In this system, a service computer is connected to the user terminals by a one-way broadcast network. When a page of information is requested by a user, the user terminal examines the broadcast data until the desired page is detected. This page is then captured, stored, and displayed. Note that a request does not propagate beyond the user terminal; such a feature is sometimes described as *pseudointeractive*. Specifics of operational systems may differ; however, they all share the pseudointeractive and information-broadcasting features. In a typical teletext system, the service computer maintains a database, and information pages in the database are updated regularly by service providers. These updates are issued locally or remotely by service provider terminals.

From conceptual and implementation viewpoints, teletext systems represent a very simple approach to organizing information delivery to a large number of users. In this paper, we use the response time experienced by teletext users as our performance measure. Several factors of system design can affect the user response time. Among them are transmission medium speed, amount of processing required at the user terminal, and the efficiency of picture encoding methods. However, these factors are of a "static" nature, since they depend to a large extent on decisions made early in the design process.

Our focus is on a more "dynamic" component of a teletext system which has an impact on the user response time, namely, the sequencing of page transmissions. Teletext system designers have decided, rather arbitrarily, on broadcasting the available information pages in a cyclic manner. In this paper, we formulate the problem of determining the sequence of page transmissions as a Markovian decision process. Using this formulation we show that, within a rather general class of sequencing policies, cyclic transmission is optimal with respect to minimizing the mean response time.

Paper approved by the Editor for Network Protocols of the IEEE Communications Society. Manuscript received July 20, 1985; revised April 17, 1986. This paper was presented at the 24th IEEE Conference on Decision and Control, Fort Lauderdale, FL, December 1985.

M. H. Ammar is with the School of Information and Computer Science, Georgia Institute of Technology, Atlanta, GA 30332.

J. W. Wong is with the Department of Computer Science, University of Waterloo, Waterloo, Ont., Canada N2L 3G1.

IEEE Log Number 8611698.

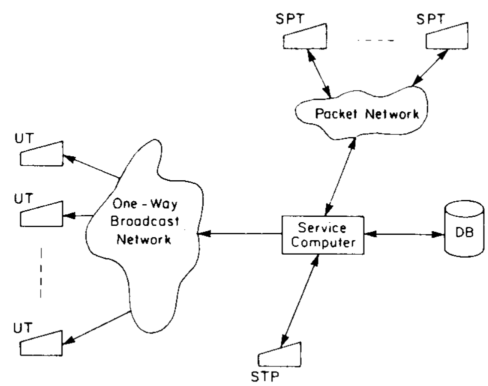

DB - Database
SPT - Service-Provider Terminal
UT - User Terminal

Fig. 1. A typical teletext system.

In recent works [4], [5] Itai and Rosberg discussed the optimal control of a time division multiple access (TDMA) scheme, with the goal of maximizing the system throughput. They showed that an optimal control policy can be found in the class of cyclic policies. An interesting duality relationship exists between TDMA systems and teletext systems. In Section II we describe this duality after a description of the teletext system model which will be used in the subsequent analysis.

The Markovian decision process formulation is given in Section III. The cost function of this process is defined as a limit, and is directly related to the mean response time if it exists. Some characteristics of the optimal policy are obtained in Section IV, with the conclusion that a cyclic policy is optimal among all policies for which the mean response time exists. A sufficient condition for the existence of mean response time is also derived. In Section V, a brief description of two algorithms for designing a teletext broadcast cycle is given. These algorithms are based on the results in [4], [6]. Numerical examples illustrating their response time characteristics are also provided.

II. Teletext System Model

In our teletext system model, time is slotted. The time to transmit a page of information is assumed to be constant and equal to one slot. This assumption is realistic for certain teletext implementations, e.g., Prestel [2]. Without loss of generality, we will use the slot length as our time unit. Slots are numbered from slot 1. In general, slot l begins at time $l - 1$ and ends at time l.

Let N be the total number of available pages. At the beginning of each slot, a decision is made as to which page to transmit during that slot. Users are assumed to submit independent requests according to a Poisson process with rate

λ. The probability that a request is for page i is assumed to be q_i, $i = 1, 2, \cdots, N$, where $\Sigma_{i=1}^N q_i = 1$. We let $\boldsymbol{q} = (q_1, q_2, \cdots, q_N)$. A request for a page i is satisfied at the end of the first full transmission of page i after the request has been made. If that request arrives during a page i transmission, it has to wait until the next transmission of page i.

For our system we define the following variables:

$X_i(l)$ = number of requests for page i waiting at the beginning of slot l

$b_i(l)$ = number of requests for page i arriving during slot l

$$u_i(l) = \begin{cases} 1 & \text{if page } i \text{ is transmitted during slot } l \\ 0 & \text{otherwise} \end{cases}$$

$Z_i(l)$ = number of slots elapsed since the beginning of the last page i transmission, until the beginning of slot l.

In addition we define the following:

$$\boldsymbol{u}(l) = (u_1(l),\ u_2(l),\ \cdots,\ u_N(l))$$

$$\boldsymbol{u}^{l-1} = (\boldsymbol{u}(1),\ \boldsymbol{u}(2),\ \cdots,\ \boldsymbol{u}(l-1)).$$

System evolution can be described by the following equations:

$$X_i(l+1) = (1 - u_i(l))X_i(l) + b_i(l)$$

$$Z_i(l+1) = 1 + (1 - u_i(l))Z_i(l).$$

We consider the case where exactly one page is transmitted in each slot, and thus we have $\Sigma_{i=1}^N u_i(l) = 1$ for all l. We also assume that, initially, $Z_i(1) = 1$ for all i. This is equivalent to assuming that all pages have been transmitted just prior to the beginning of slot 1. Since we will be considering systems that are run for an infinitely long time, this assumption does not affect the generality of our results.

The teletext system modeled above is equivalent to a (fictitious) time division multiple access (TDMA) system which has N stations and where the packet arrival process to station i is Poisson with rate $q_i\lambda$, $i = 1, 2, \cdots, N$. The number of packets in station i's buffer at the beginning of slot l is represented by $X_i(l)$. In this dual TDMA system, every time a station is given permission to transmit in a slot, all packets present in the station's buffer at the beginning of the slot are transmitted simultaneously.

Duality of queueing systems is a prevalent phenomenon which sometimes leads to interesting results (see, e.g., [7]). Because the dual TDMA system is rather unrealistic, this particular duality is only useful as an intellectual exercise and is not helpful in making the results in [5] directly applicable to our model. However, recognizing this duality led us to the use of an approach similar to that in [5], and consequently, we were able to assert the optimality of cyclic transmission in teletext systems.

III. Markovian Decision Process Formulation

In this section we formulate the problem of sequencing teletext page transmissions as a Markovian decision process. We do this by describing the class of admissible control policies, defining the cost function, and making a formal statement of the problem.

A. Admissible Control Policies

A control policy is a sequence of decisions $\boldsymbol{u}(l)$ for $l = 1, 2, \cdots$. In general each decision can be based on the current system state and the past decisions made. The state of a teletext system at the beginning of slot l can be described by the vector $\boldsymbol{X}(l) = (X_1(l), X_2(l), \cdots, X_N(l))$, which represents the number of requests pending at the beginning of slot l for each page i, $i = 1, 2, \cdots, N$. Due to the one-way nature of the system, however, this state information is not available to the decision maker (i.e., the service computer). Therefore, the decision as to which page to broadcast in slot l cannot be based on $\boldsymbol{X}(l)$. The only other relevant information available to the service computer is contained in the following:

i) the history of decisions made up to the beginning of slot l, given by the vector $\boldsymbol{u}^{l-1}$, and

ii) the probability distribution of $\boldsymbol{X}(l)$ conditioned on past decisions $\boldsymbol{u}^{l-1}$.

As in [4] and [5], we argue that basing decisions on the information contained in ii) above makes the information in i) redundant when the optimal policy is considered (see [8]). Furthermore, since the arrival process is assumed to be Poisson, we have, for given $\boldsymbol{u}^{l-1}$,

$$\text{Prob } [X_i(l) = n \mid \boldsymbol{u}^{l-1}] = \frac{(q_i\lambda Z_i(l))^n \exp\ (-q_i\lambda Z_i(l))}{n!}. \quad (1)$$

The $Z_i(l)$'s will serve as sufficient statistics for our decision process, i.e., it is sufficient to base our decisions at the beginning of slot l on the vector $\boldsymbol{Z}(l) = (Z_1(l), Z_2(l), \cdots, Z_N(l))$. Note that the $Z_i(l)$'s are fully specified when $\boldsymbol{u}^{l-1}$ is given.

In general, an admissible policy is one where the decision vector $\boldsymbol{u}(l)$ is based on the state vector $\boldsymbol{Z}(l)$ and l. We allow $\boldsymbol{u}(l)$ to be a probabilistic function of $\boldsymbol{Z}(l)$ and l. A policy is said to be nonrandomized when $\boldsymbol{u}(l)$ is a deterministic function of $\boldsymbol{Z}(l)$ and l. A stationary policy is one where $\boldsymbol{u}(l)$ is independent of l.

B. Cost Function Derivation

The objective of this section is to derive a function that represents the steady-state average cost per unit time of running our teletext system model. In deriving the cost function, we will establish its relationship to the mean response time.

At the end of slot l we incur a unit of cost for each request that was outstanding at the beginning of the slot. Let $w(l)$ be the immediate cost of slot l; we have

$$w(l) = \sum_{i=1}^{N} X_i(l).$$

Let π be a given sequence of page transmissions defined by the decision vectors $\boldsymbol{u}(l)$, $l = 1, 2, \cdots$ (π could be considered as one instance of a randomized policy or the sequence for a nonrandomized policy). We define $V_T(\boldsymbol{q}, \pi)$ to be the total expected cost of running the system until the end of slot T if the request probabilities are given by $\boldsymbol{q}$ and the sequence π is followed. Thus, we have

$$V_T(\boldsymbol{q},\ \pi) = E_\pi \left[\sum_{l=1}^{T} w(l) \right] = \sum_{l=1}^{T} E_\pi[w(l)] \quad (2)$$

where $E_\pi[\cdot]$ is the expected value, given π, taken over all possible request arrival patterns. The expected cost of slot l, $E_\pi[w(l)]$, depends only on the sequence π, and in particular the decisions made up to slot $l - 1$, i.e., $\boldsymbol{u}^{l-1}$. Thus, we have

$$E_\pi[w(l)] = E[w(l) \mid \boldsymbol{u}^{l-1}] = \sum_{i=1}^{N} E[X_i(l) \mid \boldsymbol{u}^{l-1}].$$

From (1), we get

$$E[X_i(l) \mid \boldsymbol{u}^{l-1}] = q_i\lambda Z_i(l). \quad (3)$$

Substituting (3) into (2), we obtain

$$V_T(q, \pi)=\sum_{i=1}^{N}\sum_{l=1}^{T} q_i\lambda Z_i(l).$$

We now define the average cost per unit time of using the sequence π as

$$\bar{V}(q, \pi)=\lim_{T\to\infty}\frac{1}{T}V_T(q, \pi). \quad (4)$$

Note that the limit in (4) may not exist; an example for that to happen is when $q_i > 0$ and π is such that page i is never transmitted.

In our definition of (2), we incur one unit of cost for every full unit of time the request has to wait. It follows that the expected total time spent by requests in system in the time interval $[0, T]$ is given by $V_T(q, \pi) + 0.5\lambda T$ (the $0.5\lambda T$ is added to account for the expected time spent by requests from arrival instant to the beginning of the next slot for the case of Poisson arrivals). The mean number in system in $[0, T]$ as $T \to \infty$ is therefore $\bar{V}(q, \pi) + 0.5\lambda$. If the limit in (4) exists, then by Little's result [9], we have

$$\lambda S(q, \pi)=\bar{V}(q, \pi)+0.5\lambda \quad (5)$$

where $S(q, \pi)$ is the mean response time. On the other hand, if the limit in (4) does not exist, $S(q, \pi)$ is not defined, and we use

$$\bar{V}(q, \pi)=\lim_{T\to\infty}\inf\frac{1}{T}V_T(q, \pi). \quad (6)$$

C. Problem Statement

Following the discussions in Sections III-A and III-B, we formulate a Markovian decision process which is completely defined by the following.

State Space:

$$\{Z(l)=(Z_1(l), Z_2(l), \cdots, Z_N(l))\,|\,Z_i(l)=1, 2, \cdots; \; i=1, 2, \cdots, N;\; l=1, 2, \cdots\}.$$

Decision Space:

$$\left\{u(l)=(u_1(l), u_2(l), \cdots, u_N(l))\,|\,u_i(l)=0 \text{ or } 1, \; i=1, 2, \cdots, N \text{ and } \sum_{i=1}^{N}u_i(l)=1, \; l=1, 2, \cdots\right\}.$$

Transition Probabilities:

$$\Pr\,[Z(l+1)=Z(l)+\mathbf{1}-Z_i(l)\mathbf{1}_i\,|\,Z(l),\; u(l)=\mathbf{1}_i]=1, \quad i=1, 2, \cdots, N$$

$$\Pr\,[Z(l+1)\,|\,Z(l),\; u(l)]=0 \quad \text{for all other values of } Z(l+1)$$

where $\mathbf{1} = (1, 1, \cdots, 1)$, and $\mathbf{1}_i = (0, \cdots, 1, \cdots, 0)$ with the 1 in the ith position.

Cost Function:

Expected immediate cost of slot l if decision $u(l)$ is made $=\sum_{i=1}^{N} q_i\lambda Z_i(l)$.

Observe that the immediate cost of a slot is unaffected by the decision taken (or the page transmitted) in that slot. However, it is affected by previous decisions indirectly through the values in $Z(l)$.

IV. Characteristics of the Optimal Policy

We now prove some results that will lead to a useful characterization of the optimal policy. A sequence of page transmissions π can be characterized by the following parameters.

i) Appearance times $l_r^i(\pi)$, $i = 1, 2, \cdots, N$ and $r \geq 1$; $l_r^i(\pi)$ is the index of the slot that contains the rth transmission of page i.

ii) Interappearance gaps $T_r^i(\pi)$, $i = 1, 2, \cdots, N$ and $r \geq 1$; $T_r^i(\pi)$ is the number of slots from the beginning of the rth to the beginning of the $(r + 1)$st transmission of page i.

Note that $T_r^i(\pi) = l_{r+1}^i(\pi) - l_r^i(\pi) = Z_i(l_{r+1}^i)$. In what follows, we obtain an upper bound on interappearance gaps and then discuss the implications of such a bound.

A. An Upper Bound on Interappearance Gaps

In this section, we use the properties of our cost function in (4) or (6) to identify an upper bound for interappearance gaps in an optimal policy. In other words, we show that a number τ exists such that an optimal policy will not allow its state variables $Z_i(l)$, $i = 1, 2, \cdots, N$, to exceed that number.

As a first step in determining the desired upper bound, we prove that a long-run optimal policy contains an infinite number of transmissions of each page.

Lemma 1: For every transmission sequence π where page i is transmitted a finite number of times, there exists another transmission sequence π' (where page i is transmitted infinitely many times), such that

$$\bar{V}(q, \pi')<\bar{V}(q, \pi).$$

Proof: Let $q_{\max} = \max_j \{q_j\}$ and $q_{\min} = \min_j \{q_j\}$. For a, b, c integers, we define the following function:

$$h_j(a, b, c)=q_j\lambda[(b-a)(c-b)].$$

For s_1, s_2, s_3 integers, $0 < s_1 < s_2 < s_3 \leq 2N$, and $\tau \geq 2N$, we have

$$h_j(s_1, s_2, s_3)<q_{\max}\lambda N^2 \quad (7)$$

and

$$h_j(0, s_2, \tau)\geq q_{\min}\lambda(\tau-1). \quad (8)$$

Let l_0 be the slot in which page i is transmitted for the last time in π. Choose $\tau \geq 2N$ such that

$$q_{\min}(\tau-1)>q_{\max}N^2. \quad (9)$$

Let $l_1 = l_0 + \tau$. In the transmission sequence π, some page, say j, is transmitted at least three times during the time interval $[l_0, l_1]$. This is because τ, the number of slots in the interval, is larger than or equal to $2N$, and only $N - 1$ different pages are transmitted. Let these three times be t_1, t_2, and t_3 such that

$$l_0<t_1<t_2<t_3\leq l_1=l_0+\tau.$$

Consider another sequence π_1 which is constructed from π by replacing the transmission of page j in slot t_2 of π by a transmission of page i. The reduction in the expected immediate cost over the time interval $[l_0, l_1]$ using π_1 instead of π is derived in [11] and given by

$$\Delta=h_i(0, t_2-l_0, l_1-l_0)-h_j(t_1-l_0, t_2-l_0, t_3-l_0). \quad (10)$$

Using (7) and (8) into (10) and the fact that $l_1 - l_0 = \tau$, we get

$$\Delta>\lambda[q_{\min}(\tau-1)-q_{\max}N^2].$$

Since τ is chosen according to (9) and $\lambda > 0$, we have $\Delta > 0$.

The above procedure can be repeated infinitely many times, because after each replacement we can still identify a slot where page i is transmitted for the last time. Furthermore, each replacement results in a positive improvement in the immediate cost of an interval of length τ. We can thus construct a sequence π' where page i appears infinitely many times from π such that

$$\bar{V}(q, \pi') < \bar{V}(q, \pi).$$

This concludes the proof of the lemma.

We learn from Lemma 1 that it is suboptimal to stop transmitting a page altogether, regardless of how low its request probability is. In the next lemma we show that the cost of a transmission sequence where page i appears infinitely many times may be improved by splitting interappearance gaps of lengths greater than τ into gaps of lengths less than or equal to τ.

Lemma 2: For every q there exists a τ [given by (9)] such that for every transmission sequence π with $\max_r \{T_r^i(\pi)\} > \tau$, there exists another transmission sequence π' such that

$$\max_r \{T_r^i(\pi')\} \le \tau$$

and

$$\bar{V}(q, \pi') \le \bar{V}(q, \pi).$$

Proof: Let r_0 be such that $T_{r_0}^i(\pi) > \tau$. Following the proof of Lemma 1, there exists a slot s in the time interval $[l_{r_0}^i(\pi), l_{r_0+1}^i(\pi)]$ such that a reduction in the immediate cost of the interval can be obtained by transmitting page i in slot s. Using successive refinement steps as above, we can construct a sequence π' from π where $\max_r \{T_r^i(\pi')\} \le \tau$. Furthermore, since each refinement step will imply a reduction in expected immediate cost over an interval, we get

$$\bar{V}(q, \pi') \le \bar{V}(q, \pi).$$

Note that the average cost per unit time for π' is strictly less than that for π only if we can apply the refinement step above infinitely many times.

B. Optimality of Cyclic Policies

Define C to be the set of all cyclic transmission sequences, i.e.,

$$C = \{\pi | u(l) = u(l + L)\}$$

for some integer $L < \infty$, called the cycle length. We now prove the following theorem.

Theorem: There exists an optimal nonrandomized stationary policy defined by a transmission sequence $\pi^* \in C$.

Proof: A direct consequence of Lemma 2 is that we can improve the performance of a policy by replacing every instance of the policy where the state $Z_i(l)$ is larger than τ with another instance where $Z(l)$ is always an element of the set

$$\{Z = (Z_1, Z_2, \cdots, Z_N) | Z_i \le \tau, i = 1, 2, \cdots, N\}.$$

We are thus dealing with a finite state Markovian decision process. From [12], we know that an optimal nonrandomized stationary policy exists. For a nonrandomized stationary policy, the decision vector $u(l)$ is a deterministic function of the state $Z(l)$ and is independent of l. This implies that if the states at times l and l' are equal, then the states at times $l + m$ and $l' + m$ are also equal for all m. Since the number of states is finite, some state must repeat. Thus, the optimal nonrandomized stationary policy is also cyclic and is defined by a sequence π^* where $\pi^* \in C$. The cycle length is the smallest integer L such that $Z(l) = Z(l + L)$. This completes the proof of the theorem.

Observe that a cyclic policy is defined by a single transmission sequence. For cyclic policies, the mean response time exists [6], [11]. It follows from (5) that the limit in (4) also exists. We thus conclude that among all policies for which the mean response time exists, a cyclic policy is optimal as far as minimizing the mean response time is concerned.

In general, it may be difficult to determine the existence of mean response time for a given policy. However, a sufficient condition for the mean response time to exist can be derived. For a given transmission sequence π, the interappearance gaps (or $T_r^i(\pi)$'s) can be viewed as a discrete time stochastic process for page i where r is the discrete time index. A sufficient condition for the existence of mean response time is that the mean and second moment of the interappearance gap (denoted by $E[T^i(\pi)]$ and $E[(T^i(\pi))^2]$, respectively) exist for all i, i.e., with probability 1,

$$E[T^i(\pi)] = \lim_{k\to\infty} \frac{1}{k} \sum_{r=1}^{k} T_r^i(\pi) \tag{11}$$

and

$$E[(T^i(\pi)^2] = \lim_{k\to\infty} \frac{1}{k} \sum_{r=1}^{k} (T_r^i(\pi))^2. \tag{12}$$

The sufficient condition can be proved as follows. When (11) and (12) exist, the mean response time of a request for page i is given by the residual life [10] of the interappearance gap it encounters plus the transmission time of page i, or

$$S_i(q, \pi) = \frac{E[(T^i(\pi))^2]}{2E[T^i(\pi)]} + 1. \tag{13}$$

This implies that the mean response time over all requests exists and given by $S(q, \pi) = \sum_{i=1}^{N} q_i S_i(q, \pi)$.

As a final remark, the mean response time exists for a policy if all sequences generated by that policy satisfy the sufficient condition mentioned above.

V. Design of Teletext Broadcast Cycles

In the last section, we presented a Markovian decision process formulation for teletext page transmissions, and established an upper bound on interappearance gaps in an optimal policy. It was shown that in deriving the optimal policy, one need only consider a finite state space. However, for typical teletext systems, the state space can be very large. This makes the application of solution techniques for Markovian decision problems, such as those in [13], [14], not practical. In this section, we address the issue of designing a teletext transmission cycle with near-optimal mean response time.

In general, for a given cycle of length L, the following parameters can be identified.

i) Appearance frequencies k_i, $i = 1, 2, \cdots, N$; k_i is the number of appearances of page i in the cycle. We require that $k_i \ge 1$ and $\sum_{i=1}^{N} k_i = L$.

ii) Interappearance gaps T_r^i, $r = 1, 2, \cdots, k_i$ and $i = 1, 2, \cdots, N$; T_r^i is the number of slots between the beginning of the rth and the $(r + 1)$st appearance of page i in the cycle for $r < k_i$. $T_{k_i}^i$ is the number of slots between the beginning of the k_ith appearance in the current cycle and the beginning of the first appearance in the next cycle (see Fig. 2).

In [6], it was shown that the mean response time over all requests is given by

$$S = \frac{1}{2L} \sum_{i=1}^{N} q_i \sum_{r=1}^{k_i} (T_r^i)^2 + 1. \tag{14}$$

It was shown in [6] that a lower bound for S occurs at

$$T_r^i = L/k_i \qquad \text{for all } r \tag{15}$$

and

$$k_i/k_j = \sqrt{q_i}/\sqrt{q_j}. \tag{16}$$

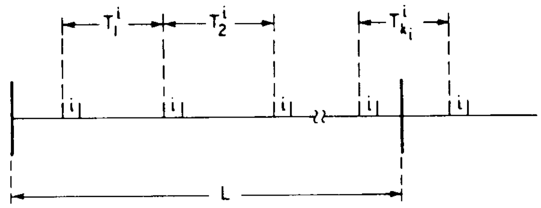

Fig. 2. A teletext broadcast cycle.

This lower bound is given by

$$S \geq \frac{1}{2}\left(\sum_{i=1}^{N} \sqrt{q_i}\right)^2 + 1. \tag{17}$$

Our discussion of cycle design is based on the following criteria.

i) The cycle length L must not exceed L^* (L^* is related to the amount of memory available to store the cycle).

ii) The mean response time S is minimized.

Our approach is to design a near-optimal cycle for each cycle length L in the range $N \leq L \leq L^*$, and adopt the cycle with the best mean response time. In what follows, we describe two algorithms that can be used to design a cycle of length L. In our descriptions, the cycle positions are numbered $0, 1, \cdots, L-1$, and we assume without loss of generality that $q_1 \geq q_2 \geq \cdots \geq q_N$.

The first algorithm was developed by the authors [6]. A brief description of this algorithm is given below.

Algorithm 1:

1) Select integer k_i's such that $\Sigma_{i=1}^N k_i = L$ and that k_i/k_j is as close to $\sqrt{q_i}/\sqrt{q_j}$ as possible for all i, j [see (16)].

2) For $i = 1$ to N, select integer T_r^i such that $\Sigma_{r=1}^{k_i} T_r^i = L$, and that T_r^i is as close to L/k_i as possible for all r [see (15)].

3) Assign page 1 to cycle positions: $0, T_1^1, T_1^1 + T_2^1, \cdots, \Sigma_{r=1}^{k_1-1} T_r^1$.

4) For $i = 2$ to $N - 1$, assign cycle positions to page i with the objective of matching the interappearance gaps obtained in step 2. Note that this is not always possible because one or more cycle positions needed for page i may have been assigned to pages with lower indexes already (see [6] for more details).

5) Assign page N to the remaining free positions in the cycle.

The second algorithm is based on the golden ratio policy proposed in [5] for TDMA systems. This policy is inspired by the use of the golden ratio ($\phi^{-1} = (\sqrt{5} - 1)/2 = 0.618\ldots$) in the multiplicative hash function to distribute keys as uniformly as possible over a hash table [15]. Our version of the golden ratio algorithm is described below.

Algorithm 2:

1) Same as step 1 of Algorithm 1.

2) For $i = 1$ to N

For $r = \Sigma_{j=1}^{i-1} k_j$ to $\Sigma_{j=1}^{i} k_j - 1$

$$b_r := r\phi^{-1} - \lfloor r\phi^{-1} \rfloor$$

$$c_r := i$$

3) Sort b_r, $r = 0, 1, \cdots, L-1$ in ascending order, i.e., find the permutation $\sigma = \{\sigma(0), \sigma(1), \cdots, \sigma(L-1)\}$ of the numbers $\{0, 1, \cdots, L-1\}$ such that $b_{\sigma(0)} \leq b_{\sigma(1)} \leq \cdots \leq b_{\sigma(L-1)}$.

4) For $r = 0$ to $L - 1$, assign cycle position r to page with index $c_{\sigma(r)}$.

We now present numerical examples to show the performance characteristics of the two algorithms mentioned above. Our examples are based on the following selection of parameter values: $N = 100$, $L^* = 1000$, and q_i is assumed to follow Zipf's law [16], i.e., if $q_1 \geq q_2 \geq \cdots \geq q_N$, then $q_i = c/i$ where c is a normalizing constant. The results are shown in Fig. 3. We observe that the best cycle obtained by Algorithm 1 has length $L = 743$. The overall mean response time of this cycle is $S = 34.4$ which is only 0.3 percent higher than the lower bound given in (17). The best cycle obtained by Algorithm 2 has length $L = 148$. The mean response time in this case is $S = 36.2$ which is 5.5 percent higher than the lower bound. Our results indicate that both algorithms yield cycles with good response time performance, but Algorithm 1 produces a better cycle. This can be explained by the fact that Algorithm 1 results in cycles with interappearance gaps closer to the ideal values given by (15).

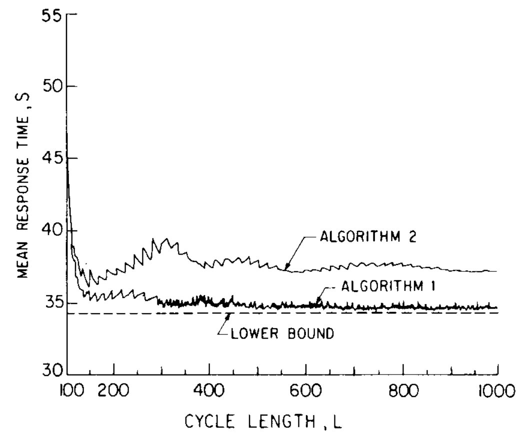

Fig. 3. Overall mean response time versus cycle length.

VI. CONCLUSION

In this paper, we have formulated the problem of sequencing page transmissions in a teletext system as a Markovian decision process, and shown that among all policies for which the mean response time exists, a cyclic policy is optimal as far as minimizing the mean response time is concerned. We have also described two algorithms for designing teletext broadcast cycles, and illustrated by numerical examples that both algorithms yield cycles with near-optimal mean response time.

REFERENCES

[1] J. R. Storey, A. Vincent, and R. Fitzgerald, "A description of the broadcast Telidon system," *IEEE Trans. Consumer Electron.*, vol. CE-26, Aug. 1980.

[2] P. Mothersole, "Teletext and Viewdata: New information systems using the domestic television receiver," *Proc. IEE*, vol. 126, Dec. 1979.

[3] J. Gecsei, *The Architecture of Videotex Systems.* Englewood Cliffs, NJ: Prentice-Hall, 1983.

[4] Z. Rosberg, "Optimal decentralized control in a multiaccess channel with partial information," *IEEE Trans. Automat. Contr.*, vol. AC-28, Feb. 1983.

[5] A. Itai and Z. Rosberg, "A golden ratio control policy for a multiple-access channel," *IEEE Trans. Automat. Contr.*, vol. AC-29, Aug. 1984.

[6] M. H. Ammar and J. W. Wong, "The design of teletext broadcast cycles," *Perform. Eval.*, vol. 5, Nov. 1985.

[7] M. H. Ammar and S. B. Gershwin, "Equivalence relations in queueing models of manufacturing networks," in *Proc. 19th IEEE Conf. Decision Contr.*, Albuquerque, NM, Dec. 1980.

[8] C. Striebel, "Sufficient statistics in the optimum control of stochastic systems," *J. Math. Anal. Appl.*, vol. 12, 1965.

[9] J. D. C. Little, "A proof of the queueing formula $L = \lambda W$," *Oper. Res.*, vol. 9, May 1961.

[10] L. Kleinrock, *Queueing Systems, Vol. 1: Theory.* New York: Wiley-Interscience, 1975.

[11] M. H. Ammar, "Performance analysis of information systems using broadcast delivery," Ph.D. dissertation, Dep. Elec. Eng., Univ. Waterloo, Waterloo, Ont., Canada, 1985.
[12] C. Derman, "On sequential decisions and Markov processes," *Management Sci.*, vol. 9, Jan. 1962.
[13] —, *Finite State Markovian Decision Processes*. New York: Academic, 1970.
[14] R. A. Howard, *Dynamic Programming and Markov Processes*. Cambridge, MA: M.I.T. Press, 1960.
[15] D. E. Knuth, *The Art of Computer Programming*, Vol. 1. Reading, MA: Addison-Wesley, 1973.
[16] G. K. Zipf, *Human Behaviour and the Principle of Least Effort*. Reading, MA: Addison-Wesley, 1949.

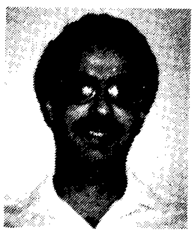

Mostafa H. Ammar (S'83–M'85) received the S.B. and S.M. degrees in electrical engineering and computer science from the Massachusetts Institute of Technology, Cambridge, in 1978 and 1980, respectively, and the Ph.D. degree in electrical engineering from the University of Waterloo, Waterloo, Ont., Canada, in 1985.

From 1980 to 1982, he worked at Bell-Northern Research, Ottawa, Ont., first as a Member of the Scientific Staff, and then as Manager of Data Network Planning. Currently, he is an Assistant Professor at the School of Information and Computer Science, Georgia Institute of Technology, Atlanta. His research interests are in the area of the performance analysis of communication networks and services.

Dr. Ammar is a member of the Associaton for Computing Machinery, Eta Kappa Nu, and the Association of Professional Engineers of the Province of Ontario.

★

J. W. Wong received the B.S. degree in engineering and the M.S. and Ph.D. degrees in computer science from the University of Califronia at Los Angeles in 1970, 1971, and 1975, respectively.

Since 1975, he has been with the University of Waterloo, Waterloo, Ont., Canada, where he is currently a Professor of Computer Science. From September 1981 to August 1982, he was a Visiting Scientist at the IBM Zurich Research Laboratory, Zurich, Switzerland, working on performance analysis of local area networks. His research interests include computer networks, communication protocols, and performance evaluation.

Dr. Wong has served on program committees of several conferences on computer communications; he was Program Chairman of the IEEE INFOCOM '84 Conference.