Dynamic visit-order rules for batch-service polling

Jan van der Wal
Technical University of Eindhoven, The Netherlands
[jan.v.d.wal@tue.nl]

Uri Yechiali
Dept. of Statistics and OR, Tel Aviv University, Israel
[uriy@post.tau.ac.il]

Abstract

We explore server's visit-order policies in non-symmetric polling systems with switch-in and switch-out times, where service is in batches of unlimited size. We concentrate on so-called 'Hamiltonian tour' policies in which, in order to give a *fair* treatment to the various users, the server attends every non-empty queue exactly once during each round of visits (cycle). The server dynamically generates a new visit schedule at the start of each round, depending on the current state of the system (number of jobs in each queue) and on the various non-homogeneous system parameters. We consider three service regimes: Globally Gated, (Locally) Gated, and Exhaustive, and study three different performance measures: (1) Minimizing expected weighted sum of all *sojourn* times of jobs within a cycle. (2) Minimizing the expected length of the *next* cycle, and (3) Maximizing the expected weighted *throughput* in a cycle. For each combination of performance-measure and service-regime we derive characteristics of the optimal Hamiltonian tour. Some of the resulting optimal policies are shown to be elegant index-type rules. Others are the solutions of deterministic NP-Hard problems. Special cases are reduced to Assignment problems with specific cost matrices. The index-type rules can further be used to construct fixed-order, cyclic-type polling tables in cases where dynamic control is not applicable.

1 Introduction

We study optimal server's scheduling policies in heterogeneous polling systems with switch-in and switch-out times, where service is in batches of unlimited size.

We concentrate on *fair* policies, called 'Hamiltonian-tours', where in each new round the server must attend every non-empty queue exactly once, but is free to determine the order of visits.

Polling systems - a set of queues attended by a single server - play a central role in the modeling and analysis of various problems in communications, computer software and hardware, data base systems, manufacturing, maintenance management and road traffic control. Much of the literature deals with occasions where service is given to *one job at a time*. For a comprehensive treatment of such models, as well as description of various applications and extensive literature we refer to Takagi [15–19]. Several surveys dealt with applications and optimization (Levy and Sidi [12]), applications in communications (Grillo [11]), applications in computer networks (Takagi [17]), analysis and optimization (Boxma [5]), and analysis and control (Yechiali [22]). Most of the works considered 'open' polling systems where all jobs are 'transient', i.e., they arrive to, are served by, and leave the system never to return. A few works analyzed 'closed' networks with 'permanent' jobs that are routed from one queue to another, but never leave the system (Altman and Yechiali [2], Dror and Yechiali [8]). Hybrid systems with both permanent and transient jobs were analyzed by Armony and Yechiali [4].

Unlimited batch service models were considered in the literature as applications to videotex, telex and TDMA systems ([3], [9], [13]) as well as for central data base operations ([20]). Ammar and Wong [3] considered a teletext system with N queues, where queue i receives a Poisson stream of requests with rate λ_i. Service times in all queues are *deterministic* (slotted, unit of time each), batches are *unlimited* and there are *no* switching times. The service discipline is (Locally) Gated. Using a homogeneous linear cost function (one unit of cost for each outstanding request at the beginning of a slot) and applying Markov decision process formulation, they showed that the policy that minimizes mean response time is of a cyclic nature, with cycle length of $L \geq N$ slots, in which queue i is visited k_i times, $\sum_{i=1}^{N} k_i = L$. Yet, the problem of finding the *exact* length L was only partially resolved.

Liu and Nain [13] examined the (Locally) Gated and Exhaustive regimes for the case of *zero switching times* and a homogeneous arrival process to all queues, while Dykman et al. [9] indicated (after using Howard's policy-iteration algorithm) that, even with equal and deterministic service requirements and with no switching times, the structure of the optimal policy *could be very complicated*. Van Oyen and Teneketzis [20] formulated a central data base system and an Automated Guided Vehicle (AGV) as a polling system with an infinite-capacity batch server and zero switching times, where the controller observes only the length of the queue at which the server is located. A review termed 'Scheduling with Batching' is presented in [14]. Batches there, however, are *finite* and are defined as the "maximal set of jobs that are scheduled to be served *continuously, one at a time,* on a machine and share a setup". The concentration there is on "classification of problems as polynomially or pseudo-polynomially solvable, binary or unitary (NP-Hard), or open".

Recently, Xia, Michailidis, Bambos and Glynn [21] considered the problem of dynamic allocation of a single server with *finite* batch processing capability to a set of parallel queues. The independent arrival processes are Poisson, but with *equal* rates. The service times of batches are exponentially distributed and *identical*. *Zero* switching times are incurred by the server when moving from one queue to another. They restrict attention to the set of *nonanticipative, nonpreemptive, and nonidling* policies and show that, when buffers are infinite, allocating the server to the longest queue stochastically maximizes the aggregate throughput

of the system.

Our work was motivated by a tape reading problem in a system where large amounts of information are stored on tapes. Irregularly, requests arrive for data on one of these tapes. In order to read the data the tape has to be mounted, read and then dismounted again. If there are more than one requests for a tape, they all can be read in (more or less) the same time, thus suggesting modeling as a batch-service with *unlimited* batch size. Nevertheless, the models we present in this work are *general* and can be applied to all unlimited-batch-service systems mentioned above. Without loss of generality we will use the 'tape language' throughout the paper.

The goal is to establish 'fair' dynamic visit-order policies so as to optimize various performance measures.

Our special fairness approach is to visit the queues in a Hamiltonian way. (See Browne and Yechiali [7].) In each new Hamiltonian cycle only those queues are visited that are nonempty at the start of the cycle, and each such queue is visited *exactly once*. What is to be decided at the beginning of each cycle is the order in which the queues are visited. For a cycle with duration C we define the 'sojourn time of a job J within a cycle' as the total time that J resides in the system during C. We will consider the following performance measures: (1) minimizing the expected (weighted) sum of all *sojourn times* of all jobs within a cycle, (2) minimizing the expected duration of the *next cycle*, and (3) maximizing the expected weighted *throughput* within the cycle. Criterion (2) is a proper one since, given the fairness considerations leading to Hamiltonian tours, and as a result of the unlimited batch service, the duration of the current cycle is *independent* of the order of visits. Thus, one can only influence the duration of the *next cycle*.

We formulate the problem as a polling system with unlimited batch size service and consider Globally Gated, (Locally) Gated and Exhaustive service regimes. Some of the optimization problems turn out to be NP-Hard; some yield simple, elegant, index-type rules; while others are transformed into classical Assignment problems, but with special costs involving the parameters of the problem.

The structure of the paper is as follows. In section 2 we present the model, describe the assumptions and discuss various gating procedures. Minimizing weighted expected *sojourn* times is treated in section 3, where all three service regimes are considered. Section 4 deals with minimizing the expected duration of the *next* cycle, while section 5 aims at maximizing the expected weighted *throughput* in a cycle. The analytical results and their implications are summarized in section 6, and conclusions are discussed in section 7.

2 The model

There is a set of queues (tapes), numbered 1 up to N. Requests (jobs) arrive to queue l at a Poisson rate λ_l. A single server visits the queues according to some order. When a queue is visited all jobs are served simultaneously. This service time is independent of the number of jobs served. The total visit time at queue l consists of three parts: a time R_l to

reach (mount) the tape, a time B_l to read (serve) the tape if the queue is gated, or several repetitions of B_l if the service is exhaustive, and finally, a time D_l to dismount the tape (exit the queue to some 'base' position). The order in which the queues are visited is completely free.

2.1 Hamiltonian cycle approach

In order to give a fair service to all queues, in each cycle the server performs a Hamiltonian tour in which every *nonempty* queue is visited *exactly once*. That is, if at the beginning of the cycle there are $n_l \geq 0$ jobs in queue l, the server visits only those queues for which n_l is *positive*. Thus, even if during the cycle jobs arrive to a queue that was empty at the beginning of the cycle, this queue is not included in the cycle. This process then repeats itself with a new Hamiltonian tour, where the server visits only queues for which the *new* n_l are positive.

We consider several service regimes:

- Globally Gated; i.e., all queues are simultaneously gated (closed) at the start of the cycle;
- (Locally) Gated ; i.e., a queue is gated only when the server arrives; and
- Exhaustive; the server continues serving a visited queue until it becomes empty.

In the two Gated regimes, jobs arriving to a queue after its gate is closed will be served only during the next cycle. The same applies in the exhaustive regime for jobs arriving after the server leaves the queue. For a gated queue l a visit requires a time $R_l + B_l + D_l$. We define $H_l = R_l + B_l + D_l$, $h_l = E[H_l]$, $h_l^{(2)} = E[H_l^2]$ and $\tilde{H}_l(\alpha) = E[e^{-\alpha H_l}]$. Similarly we define r_l, $r_l^{(2)}$, $\tilde{R}_l(\alpha)$, b_l, etc.

Without loss of generality we assume that, at the beginning of the cycle, n_l is positive for queues $l = 1, \ldots, L$ and that $n_l = 0$ for $l = L+1, \ldots, N$. Then, for the Gated cases, the duration of the cycle is $\sum_{l=1}^{L} H_l$, *independent* of the order in which the first L queues are visited. Moreover, for the Exhaustive case as well, since the service is in batches, the duration of the cycle is also *independent* of the order of visits.

3 Minimize weighted expected sojourn times

In this section we consider the criterion of *minimizing expected weighted sum of sojourn times of all jobs in the cycle*. (It is important to note that under batch service, minimizing sojourn times is not always equivalent to minimizing waiting times.) We will analyze the three service regimes introduced before.

3.1 Globally Gated

Under the *Globally Gated regime*, (cf. Boxma, Levy and Yechiali [6]), there is only one thing to look at: the sojourn times of the jobs present at the beginning of the cycle. This follows since the distribution of the cycle time is independent of the visit order, implying that the waiting times of those jobs that arrive during this cycle (and are not served) are not influenced by the order. A visit order (policy) is determined by a permutation $\pi = \{\pi(1), \pi(2, \ldots, \pi(l), \ldots, \pi(L)\}$ of the set $\{1, 2, \ldots, L\}$, where $\pi(l)$ indicates the index of the queue which is l-th to be visited during a Hamiltonian tour. In order to find the optimal order we start with the policy π_1 for which $\pi(l) = l$. The total expected weighted sojourn times for the jobs present at the beginning of the cycle is denoted by $C(\pi_1)$. Then, with w_l denoting the weight associated with the sojourn times of jobs in queue l,

$$C(\pi_1) = \sum_{l=1}^{L} w_l n_l \left((\sum_{j=1}^{l-1} h_j) + r_l + b_l \right) \ .$$

Theorem: The optimal policy is to visit the queues according to a non-increasing order of the index $w_i n_i / h_i$.

Proof: Look at the policy π_1' which has the same order as π_1 except that queues i and $i+1$ are interchanged. That is π_1' is the order $(1, 2, ..., i-1, i+1, i, i+2, ..., L)$. Now, the difference in sojourn 'costs' for these two strategies is

$$C(\pi_1) - C(\pi_1') = w_{i+1} n_{i+1} h_i - w_i n_i h_{i+1} \ .$$

For this to be non-positive, i.e., for strategy π_1 to be at least as good as π_1', we must have

$$\frac{w_i n_i}{h_i} \geq \frac{w_{i+1} n_{i+1}}{h_{i+1}} \ .$$

Applying repeatedly the interchange argument, the optimal visit order is determined by an *index rule*: serve the queues in a non-increasing order of $\frac{w_i n_i}{h_i}$. Q.E.D.

Note that, if all h_i are equal and all w_i are the same, the optimal order is according to the *longest queue first* policy, which is the consequence of the batch servicing.

This result may be compared with the result of Liu and Nain [13], for the (Locally) Gated case, who showed that, for a *fully symmetric* polling system with *zero switching times*, "the so-called Most Customer First policy (in which the server always visits the queue with the largest number of customers) minimizes, in the sense of strong stochastic ordering, the vector of number of customers in each queue whose components are arranged in decreasing order". Note that, in their treatment, the server is free from 'fairness' considerations and is entitled to choose the next queue to visit when exiting a served queue, whereas in our approach the order within each new cycle is *determined at the cycle's beginning*.

3.2 Locally Gated

Gating queue l just before R_l or just before B_l is not essentially different. We consider the case of gating just before B_l. Gating just before R_l then is a special case with $R_l' = 0$ and $B_l' = R_l + B_l$.

As indicated, in the Locally Gated regime (similarly to the Globally Gated case), the batch servicing implies that the duration of the cycle, $\sum_{l=1}^{L} H_l$, is not influenced by the order in which the queues are visited. Therefore, we can ignore queues $L+1, \ldots, N$. The difference, compared to the Globally Gated case, is that now we also have to consider new arrivals (during the cycle) to queues 1 to L. Assume again that the visit order is according to policy $\pi_1 = (1, 2, \ldots, L)$.

Consider queue l and define S_l to be the total sojourn time of all queue-l jobs in the cycle. Then S_l is comprised of four terms. The first one is the *sojourn time* of the n_l jobs present at the start of the cycle. The second term is the *waiting time* of all jobs that arrive (during the cycle) before *this* queue is gated. The third term contains the *service time* of the jobs that arrive *before* gating, and the fourth term consists of the *sojourn time*, until the end of the cycle, of those jobs that arrive *after* the queue is gated.

For the second and fourth term we use the following well-known result.

Lemma
Let $X_1, X_2, \ldots$ be the arrival instants of a Poisson process with arrival rate λ, and let Y be a non-negative random variable independent of the process $\{X_i\}$. Let $N(Y)$ be the number of Poisson arrivals in $(0, Y)$. That is, $N(Y) = max\{k | X_k < Y\}$. Then,

$$E[\sum_{j=1}^{N(Y)} (Y - X_j)] = E[\sum_{j=1}^{N(Y)} X_j] = \frac{\lambda}{2} E[Y^2].$$

Combining the above described four terms and following π_1 we have

$$\begin{aligned} E[S_l] = n_l(\sum_{j=1}^{l-1} h_j + r_l + b_l) + \frac{\lambda_l}{2} E[(H_1 + \cdots H_{l-1} + R_l)^2] \\ + \lambda_l(\sum_{j=1}^{l-1} h_j + r_l) b_l + \frac{\lambda_l}{2} E[(B_l + D_l + H_{l+1} + \cdots H_L)^2)] \ . \end{aligned} \tag{1}$$

Thus, the total expected sojourn cost is

$$\begin{aligned} \sum_{l=1}^{L} w_l E[S_l] = \sum_{l=1}^{L} w_l \{ \ n_l(\sum_{j=1}^{l-1} h_j + r_l + b_l) + \frac{\lambda_l}{2} E[(H_1 + \cdots H_{l-1} + R_l)^2] \\ + \lambda_l(\sum_{j=1}^{l-1} h_j + r_l) b_l + \frac{\lambda_l}{2} E[(B_l + D_l + H_{l+1} + \cdots H_L)^2 \ ]\} \ . \end{aligned} \tag{2}$$

The objective function (2) can be simplified by omitting all terms of the form $\sum_{l=1}^{L} \gamma_l$ which are order independent.

We first rewrite the second term in (2):

$$\begin{aligned}
&\sum_{l=1}^{L} \frac{\lambda_l w_l}{2} E[(H_1 + \cdots + H_{l-1} + R_l)^2] \\
&\quad = \sum_{l=1}^{L} \frac{\lambda_l w_l}{2} \left(r_l^{(2)} + 2 r_l \sum_{j=1}^{l-1} h_j + \sum_{j=1}^{l-1} h_j^{(2)} + \sum_{j=1}^{l-1} \sum_{k=1}^{l-1} h_j h_k - \sum_{j=1}^{l-1} h_j^2 \right) .
\end{aligned} \tag{3}$$

Clearly, the first term in the right-hand-side of (3) is order independent.

Now consider the fourth term in (2). We can write it as

$$\sum_{l=1}^{L} \frac{\lambda_l w_l}{2} \{ (b_l^{(2)} + 2 b_l d_l + d_l^{(2)}) + 2(b_l + d_l) \sum_{j=l+1}^{L} h_j + \sum_{j=l+1}^{L} h_j^{(2)} + \sum_{j=l+1}^{L} \sum_{k=l+1}^{L} h_j h_k - \sum_{j=l+1}^{L} h_j^2 \}.$$

Thus, in (2), we can omit not only the terms $\sum_l w_l n_l (r_l + b_l)$, $\sum_l w_l \lambda_l r_l b_l$, $\sum_l \lambda_l w_l r_l^{(2)}/2$, and $\sum_l \lambda_l w_l (b_l^{(2)} + 2 b_l d_l + d_l^{(2)})/2$, but also the combined terms $\sum_l \lambda_l w_l \sum_{j \neq l} h_j^{(2)}/2$, and $\sum_l \lambda_l w_l \sum_{j \neq l} h_j^2/2$.

Writing $h[k,m] = \sum_{j=k}^{m} h_j$, the goal is to find a permutation π that minimizes

$$\begin{aligned}
&\sum_{l=1}^{L} w_{\pi(l)} (n_{\pi(l)} + \lambda_{\pi(l)} b_{\pi(l)}) \cdot h[\pi(1), \pi(l-1)] + \sum_{l=1}^{L} w_{\pi(l)} \lambda_{\pi(l)} \{ r_{\pi(l)} \cdot h[\pi(1), \pi(\ell - 1)] \\
&\qquad + (b_{\pi(l)} + d_{\pi(l)}) \cdot h[\pi(l+1), \pi(L)] + \frac{1}{2} (h[\pi(1), \pi(l-1)])^2 + \frac{1}{2} (h[\pi(l+1), \pi(L)])^2 \}
\end{aligned} \tag{4}$$

This is combinatorially a NP-Hard problem, as we prove in the Appendix.

We now show that the symmetric case leads to a special Assignment Problem, which is solvable in $O(L^3)$ time [1].

3.2.1 A symmetric case

Suppose $b_l = b$, $d_l = d$, $r_l = r$, so that $h_l = h$ (this might be the case when all tapes are similar). One can verify that the objective can be simplified to finding the permutation π that minimizes

$$\begin{aligned}
&\sum_{l=1}^{L} w_{\pi(l)} (n_{\pi(l)} + \lambda_{\pi(l)} b) h \cdot l + \sum_{l=1}^{L} w_{\pi(l)} \lambda_{\pi(l)} \{ [2rh - (L+2)h^2] \cdot l + h^2 \cdot l^2 \} \\
&= \sum_{l=1}^{L} \alpha_{\pi(l)} \cdot l + \sum_{l=1}^{L} \beta_{\pi(l)} \cdot l^2 ,
\end{aligned}$$

with $\alpha_{\pi(l)} = w_{\pi(l)}(n_{\pi(\ell)} + \lambda_{\pi(\ell)}b)h + w_{\pi(l)}\lambda_{\pi(l)}[2rh - (L+2)h^2]$ and $\beta_{\pi(l)} = w_{\pi(l)}\lambda_{\pi(l)}h^2$. It follows that this problem can be formulated and readily solved as an Assignment problem with costs $C_{ij} = C_{\pi(i),j} = \alpha_{\pi(i)}j + \beta_{\pi(i)} \cdot j^2$ for $i, j = 1, 2, \ldots, L$.

3.3 Exhaustive regime

In the Exhaustive regime, servicing in queue l is repeated until the queue is empty. This means that if during the service time B_l one or more new jobs arrive (with probability $\int_0^\infty (1 - e^{-\lambda_l t}) dP(B_l \le t) = 1 - \tilde{B}_l(\lambda_l)$) then, after completing the present service time, a new service time is started. This can be repeated several times in a geometric fashion.

Let E_l denote this geometric sum of service times in queue l, with repetition rate $1 - \tilde{B}_l(\lambda_l)$, mean e_l, second moment $e_l^{(2)}$ and Laplace transform $\tilde{E}_l(\alpha)$. Then

$$e_l = \frac{b_l}{\tilde{B}_l(\lambda_l)} \ .$$

The Laplace transform $\tilde{E}_l(\alpha)$ is derived as follows:
For given values $B_{l1}, B_{l2}, \cdots$, all distributed as B_l, we have $E_l = \sum_{j=1}^{k} B_{lj}$ with probability $\prod_{j=1}^{k-1}(1 - e^{-\lambda_l B_{lj}})e^{-\lambda_l B_{lk}}$, $k = 1, 2, \cdots$.
Thus,

$$\begin{aligned}
\tilde{E}_l(\alpha) &= \sum_{k=1}^{\infty} E\left[e^{-\alpha(\sum_{j=1}^{k} B_{lj})} \prod_{j=1}^{k-1} (1 - e^{-\lambda_l B_{lj}}) e^{-\lambda_l B_{lk}} \right] \\
&= \tilde{B}_l(\alpha + \lambda_l) \sum_{k=1}^{\infty} \left[\tilde{B}_l(\alpha) - \tilde{B}_l(\alpha + \lambda_l) \right]^{k-1} \\
&= \frac{\tilde{B}_l(\alpha + \lambda_l)}{1 - \tilde{B}_l(\alpha) + \tilde{B}_l(\alpha + \lambda_l)} \ .
\end{aligned}$$

By differentiation we get,

$$e_l^{(2)} = \frac{1}{\tilde{B}_l(\lambda_l)} \left(b_l^{(2)} + 2b_l e_l + 2e_l \tilde{B}_l^{(')}(\lambda_l) \right) \ .$$

Let G_l be the total visit time to queue l: $G_l = R_l + E_l + D_l$. Then, for the weighted sojourn time criterion there is only a little difference between the Gated and the Exhaustive cases. If in the analysis of the Gated model we replace H_l, h_l and $h_l^{(2)}$ by G_l, g_l and $g_l^{(2)}$, respectively, and we add the term for the sojourn times of the jobs served during the repetitions (on the average $\lambda_l e_l$ jobs, each requiring one residual and one normal service time), then for policy π_1 we get (cf. (2))

$$\sum_{l=1}^{L} w_l E[S_l] = \sum_{l=1}^{L} w_l \ \{ \ n_l(\sum_{j=1}^{l-1} g_j + r_l + b_l) + \frac{\lambda_l}{2} E[(\sum_{j=1}^{l-1} G_j + R_l)^2] + \lambda_l(\sum_{j=1}^{l-1} g_j + r_l) b_l$$
$$+ \lambda_l e_l(\frac{b_l^{(2)}}{2b_l} + b_l) + \frac{\lambda_l}{2} E[(D_l + \sum_{j=l+1}^{L} G_j)^2] \ \} \ . \tag{5}$$

The structure of this expression looks similar to equation (2), but in fact is even more involved than in the (Locally) Gated case. Even for the symmetric case, when all R_l, B_l and D_l are the same, the $\tilde{B}_l(\lambda_l)$ terms cause us to no longer have an Assignment problem as a solution of (5).

3.4 Extension: modifying the Hamiltonian tour

A possibly improved procedure (within the framework of Hamiltonian tours) is the following. Suppose the visit order dictated by the performance measure and the gating procedure is $1, 2, ..., L$. Then, after exiting queue 1, the rule can be re-applied with regard to the *remaining* $L-1$ queues, taking into account the new values of the n_l's. Accordingly, the next queue to be visited will be determined, continuing in this manner until the last queue (among the original L) will be serviced. Note, however, that such a modified procedure will violate the main purpose of the Globally Gated regime aimed at *not* letting jobs arriving during the current cycle to be serviced *before* jobs that have arrived during the previous cycle. Considering for a moment a Markov Decision Process approach, which aims at acting optimally at each step, one can apply a new the visit-order rule (for example, the nonincreasing $w_l n_l / h_l$ policy) after each visit of a queue, doing it in a continuous routine while neglecting the Hamiltonian tour restriction and generating a dynamic, though erratic, polling visit table (for more about polling tables see Yechiali [22]).

4 Minimize next cycle duration

Since in all three service regimes the duration of each Hamiltonian cycle is independent of the order of visits, a criterion which seems to look a bit further ahead is to *minimize the expected duration of the next Hamiltonian cycle.* Besides short cycles being clearly good, this criterion also implies that we are trying to serve new arrivals already in the first cycle.
A possibly different way of looking one cycle ahead is to *minimize the expected number of nonempty queues at the beginning of the next cycle.* These two problems are, as we will see, closely related.

Consider again the strategy π_1 which serves the queues in the order 1, 2 up to L. We will compute for this strategy the expected duration of the next cycle. In order to do so we need the probability that a specific queue will be empty in the beginning of the next cycle. Define

$$p_l = Prob\{Queue\ l\ is\ empty\ at\ the\ start\ of\ the\ next\ cycle\}$$

We emphasize again that for the Globally and Locally Gated regimes, as well as for the Exhaustive case, the duration of the present Hamiltonian cycle does not depend on the visit order within the cycle. Consider first queues which are not visited in this cycle. The probability that queue l, $l > L$, is empty in the beginning of the next cycle is independent of the order. Formally, defining for the Gated regimes $H^{(L)} = \sum_{j=1}^{L} H_j$, we have

$$p_l = E[e^{-\lambda_l H^{(L)}}] = \tilde{H}^{(L)}(\lambda_l) \ , \ l > L \ .$$

In the Exhaustive regime we get a similar expression with $G^{(L)} = \sum_{l=1}^{L} G_l$.

4.1 Globally Gated

When simultaneously gating at the beginning of the cycle, the p_l for $l = 1, \ldots, L$ also satisfy $p_l = \tilde{H}^{(L)}(\lambda_l)$. Hence, in this case, all orders are stochastically identical (and optimal).

4.2 Locally Gated

Similarly to the sojourn times criterion, the gating is just before B_l .
Then, for $l \leq L$, p_l is the probability that there are no arrivals into queue l during the remaining duration of the present cycle $B_l + D_l + H_{l+1} + \cdots + H_L$. Hence,

$$p_l = \tilde{B}_l(\lambda_l)\tilde{D}_l(\lambda_l) \prod_{j=l+1}^{L} \tilde{H}_j(\lambda_l) \ .$$

The expected duration of the next cycle under π_1 is

$$\sum_{l=1}^{N}(1-p_l)h_l = \sum_{l=1}^{N} h_l - \sum_{l=1}^{L} \tilde{B}_l(\lambda_l)\tilde{D}_l(\lambda_l)h_l \prod_{j=l+1}^{L} \tilde{H}_j(\lambda_l) - \sum_{l=L+1}^{N} \tilde{H}^{(L)}(\lambda_l)h_l \ . \qquad (6)$$

We immediately observe that, in contrast to the 'minimizing sojourn time' performance measure, the above expression is *independent* of the number of jobs present at the start of the cycle.

The first and third terms in (6) are order independent, so we consider

$$V(\pi_1) = \sum_{l=1}^{L} \tilde{B}_l(\lambda_l)\tilde{D}_l(\lambda_l)h_l \prod_{j=l+1}^{L} \tilde{H}_j(\lambda_l)$$

and wish to find a permutation π that maximizes

$$V(\pi) = \sum_{l=1}^{L} \tilde{B}_{\pi(l)}(\lambda_{\pi(l)})\tilde{D}_{\pi(l)}(\lambda_{\pi(l)})h_{\pi(l)} \prod_{j=l+1}^{L} \tilde{H}_{\pi(j)}(\lambda_{\pi(l)})$$

In general this expression is hard to maximize. Therefore, we consider two special cases:

1. All λ_l are equal, say $\lambda_l = \lambda$,
2. All H_l are identical.

4.2.1 Case 1: all λ_l equal

Assume $\lambda_l = \lambda$ for all l. As previously, compare two strategies π_1 and π_1' which differ only in the order in which queues i and $i+1$ are visited. In π_1 the order is i, $i+1$ and in π_1' the order is reversed.

Comparing $V(\pi_1)$ and $V(\pi_1')$, the contributions for queues l less than i and larger than $i+1$ are the same. We only have to compare the expressions for i and $i+1$ in the two orders.

It follows that π_1 is better than π_1' if, simplifying the notation to $a_j := \tilde{H}_j(\lambda)$ and $c_l = \tilde{B}_l(\lambda)\tilde{D}_l(\lambda)h_l$,

$$V(\pi_1) - V(\pi_1') = (c_i a_{i+1} + c_{i+1} - c_{i+1}a_i - c_i) \prod_{l=i+2}^{L} a_l > 0 \ .$$

That is, if

$$\frac{c_{i+1}}{1-a_{i+1}} - \frac{c_i}{1-a_i} > 0.$$

Therefore, if all arrival rates are equal, the optimal visit schedule follows a nondecreasing order of the index

$$\frac{c_l}{1-a_l} = \frac{\tilde{B}_l(\lambda)\tilde{D}_l(\lambda)h_l}{1-\tilde{H}_l(\lambda)}.$$

4.2.2 Case 2: all $R_l,\ B_l,\ D_l$ are identical

If all R_l, B_l, D_l, and therefore H_l, are identical, writing $f_l = \tilde{B}_l(\lambda_l)\tilde{D}_l(\lambda_l)h$ and $g_l = \tilde{H}(\lambda_l)$, the expression for $V(\pi_1)$ simplifies to

$$V(\pi_1) = \sum_{l=1}^{L} f_l g_l^{L-l} \ .$$

Thus, similarly to the consideration in Section 3, if queue i is in position j in the visit order, then its 'contribution' is $f_i g_i^{L-j}$. Therefore, in order to maximize $V(\pi)$, one has to solve an Assignment problem with rewards $C_{ij} = C_{\pi(i),j} = f_{\pi(i)} g_{\pi(i)}^{L-j}$.

Note that this case is more involved than case 1 above, but essentially simpler than the general case in which the terms in the objective function depend not only on $\pi(i)$ and j but also on the order of the other queues.

4.3 Exhaustive regime

The time to visit a queue now consists of $G_l = R_l + E_l + D_l$ instead of $H_l = R_l + B_l + D_l$. Due to the Exhaustive discipline (and under π_1) we have $p_l = \tilde{D}_l(\lambda_l)\prod_{j=l+1}^{L}\tilde{H}_j(\lambda_l)$ for $l \le L$. Thus, writing $g_l = r_l + e_l + d_l$ the expression equivalent to $V(\pi_1)$ in Section 4.2 is

$$\sum_{l=1}^{L} \tilde{D}_l(\lambda_l)(r_l + e_l + d_l) \prod_{j=l+1}^{L} \tilde{R}_j(\lambda_l)\tilde{E}_j(\lambda_l)\tilde{D}_j(\lambda_l) \ .$$

This leads to essentially the same optimization problem as for the Locally Gated variant.

The case where all λ_l are equal is the same as before, implying that the optimal visit schedule is determined by the nondecreasing order of

$$\frac{\tilde{D}_l(\lambda)g_l}{1-\tilde{H}_l(\lambda)} = \frac{\tilde{D}_l(\lambda)(r_l+e_l+d_l)}{1-\tilde{R}_l(\lambda)\tilde{E}_l(\lambda)\tilde{D}_l(\lambda)} \ .$$

However, for case 2 (all H_l equal) there is a significant difference in the meaning of the assumption that $R_l + E_l + D_l$ is the same for all l. It implicitly implies that if the H_l are the same and the E_l are the same, then all λ_l are the same as well. Hence, in this case, all queues are stochastically identical and all visit orders are optimal.

4.4 Number of nonempty queues

As was indicated, the problem of minimizing the expected number of nonempty queues at the beginning of the next cycle is not very different from the problem of minimizing the expected duration of the next cycle . The only difference is that in the expression for the objective function given in (6) the terms h_l disappear because the expected number of nonempty queues equals $\sum_{l=1}^{N}(1-p_l)$. Thus, the structure of the optimization problem is exactly the same.

5 Maximizing weighted throughput in a cycle

A third objective for optimization is to maximize the expected weighted throughput during the cycle. This objective may seem to be more interesting from the viewpoint of the system operator than from the viewpoint of the requester. However, the more jobs are served in the current cycle, the fewer jobs are left to wait until the next cycle. Thus, this objective is an interesting one for the requesters as well. For the Globally Gated case the cycle throughput is order independent. Also, as we will see, the structure of the optimization problem is the same for the Locally Gated and for the Exhaustive regimes.

Note again that, for all service disciplines, the duration of the cycle is independent of the order in which the queues are visited, implying that one can ignore the queues $L+1$ up to N.

5.1 Locally Gated

Let π_1 be again the order $1, 2, \cdots, L$. Define $M(\pi_1)$ to be the expected weighted throughput during the cycle under order π_1. Then, gating just before B_l,

$$M(\pi_1) = \sum_{l=1}^{L} w_l n_l + \sum_{l=1}^{L} w_l \lambda_l (\sum_{j=1}^{l-1} h_j + r_l)$$

Now, similarly to what we have seen previously, the optimal visit order is an index rule and is independent of n_l: visit the queues in a nondecreasing order of

$$\frac{w_l \lambda_l}{h_l} \ .$$

Note that this rule does not affect the long run weighted throughput, which is fixed and equal to $\sum_l w_l \lambda_l$, but focuses on *early* weighted throughput.

Remark. If all w_l are the same and all λ_l are equal, then the optimal order is in a non-increasing order of h_l, enabling accumulation of as many jobs as possible in queues to be visited later in the cycle. Similarly, if all w_l and h_l are equal, the queues are visited in a non-decreasing order of λ_l, again in order to generate as many jobs as possible in the most active queues.

5.2 Exhaustive regime

The result here is very similar to the one derived for the Locally Gated case. The corresponding expression for $M(\pi_1)$ changes into

$$M(\pi_1) = \sum_{l=1}^{L} w_l n_l + \sum_{l=1}^{L} w_l \lambda_l (\sum_{j=1}^{l-1} g_j + r_l + e_l) \ .$$

Therefore, the optimal visit schedule follows a non-decreasing order of

$$\frac{w_l \lambda_l}{g_l} \ .$$

6 Summary of Results

Table 1 below summarizes the analytical results considering the optimal policies (visit-order rules) for the various combinations of Performance-Measure and Gating-Procedure. The criterion 'Minimizing weighted expected sojourn times (during a Hamiltonian tour)' leads to rules involving the values of the n_l's, i.e., the number of requests present at the start of the cycle in each queue. This implies that the visit order will change from one cycle to another as a result of the dynamic evolution of the system. Furthermore, it may enable one to modify the visit order for the remaining queues in the cycle after each time the server exits a queue. One can possibly exercise a one-step look-ahead procedure (following the relevant rule) and apply it repeatedly without being confined by the Hamiltonian tour restriction. The objectives 'Minimize expected duration of next cycle' and 'Maximizing expected weighted throughput' lead to Hamiltonian procedures that do *not* involve the n_l's. Thus, if we try to modify the visit order within a cycle, we'll come up with the same visit order that has been determined at the start of the cycle. However, one can use rules which are λ_l dependent to determine static, fixed-order, cycles in cases where dynamic control is not applicable.

Table 1: Optimal policies for the various combinations of Performance Measure and Gating Procedure

Gating Procedure Performance Measure	Globally Gating	Locally Gating	Exhaustive
Minimize Weighted Expected Sojourn Times [Section 3]	**Index rule:** decreasing $\frac{w_\ell n_\ell}{h_\ell}$	Hard Combinatorial Problem involving the n_ℓ's **Special case:** $h_\ell = h$; distinct λ_ℓ's: Assignment Problem with costs c_{ij}'s involving the n_ℓ's	Hard Combinatorial Problem involving the n_ℓ's
Minimize Expected Duration of Next Cycle [Section 4]	All Policies are Stochastically Equal	Hard Combinatorial Problem *not* involving the n_ℓ's **Special case 1:** $\lambda_\ell = \lambda$ **Index rule:** increasing $\frac{\tilde{B}_\ell(\lambda)\tilde{D}_\ell(\lambda)h_\ell}{1-\tilde{H}_\ell(\lambda)}$ **Special case 2:** $h_\ell = h$ Assignment Problem with rewards c_{ij}'s *not* involving the n_ℓ's	Hard Combinatorial Problem *not* involving the n_ℓ's **Special case 1:** $\lambda_\ell = \lambda$ **Index rule:** increasing $\frac{\tilde{D}_\ell(\lambda)g_\ell}{1-\tilde{G}_\ell(\lambda)}$ **Special case 2:** $h_\ell = h$, $e_\ell = e$ Fully Symmetric, all Policies are Stochastically Equal
Maximize Expected Weighted Throughput in a Cycle [Section 5]	All Policies are Stochastically Equal	**Index rule:** increasing $\frac{w_\ell \lambda_\ell}{h_\ell}$	**Index rule:** increasing $\frac{w_\ell \lambda_\ell}{g_\ell}$

7 Conclusions

We have considered the problem of finding dynamic visit order rules for a polling system with unlimited batch servicing, where service times are independent of the batch sizes. We adopt a 'fair' dynamic Hamiltonian cycle approach where in each new cycle a new visit order is determined, based on the dynamic evolution of the system. Within the cycle we have considered various performance measures and various gating procedures. Some of the Hamiltonian cycle problems lead to elegant solutions in the form of an index rule. Some lead to an Assignment problem, while others result in combinatorial hard problems. Furthermore, the n_l-dependent rules may be re-applied within a cycle to modify it, or as a one-step look-ahead repeated procedure. The non n_l-dependent index rules can be used to construct fixed-order cyclic polling tables in cases where static rules are required.

Acknowledgement

We thank Professor Arie Tamir for very helpful discussions regarding the NP-Hard property.

References

[1] R.K. Ahuja, T.L. Magnanati and J.B. Orlin, Network Flows, Theory, Algorithms and Applications, Prentice Hall, Englewood Cliffs, NJ, (1993).

[2] E. Altman and U. Yechiali, Polling in a Closed Network, Probability in the Engineering and Informational Sciences, Vol. 8, No. 3, 327–343, (1994).

[3] M.H. Ammar and J.W. Wong, "On the Optimality of Cyclic Transmission in Teletext Systems", IEEE Transactions on Communications, Vol. COM-35, No. 1, pp. 68–73, (1987).

[4] R. Armony and U. Yechiali, Polling Systems with Permanent and Transient Customers, Stochastic Models, Vol. 15, No. 3, 395–427, (1999).

[5] O.J. Boxma, Analysis and optimization of polling systems, in: Queueing, Performance and Control in ATM, eds. J.W. Cohen & C.D. Pack, North-Holland, Amsterdam, pp. 173–183, (1991).

[6] O. Boxma, H. Levy and U. Yechiali, Cyclic Reservation Schemes for Efficient Operation of Multiple-Queue Single-Server Systems, Annals of Operations Research 36, pp. 187–208, (1991).

[7] S. Browne and U. Yechiali, Dynamic Priority Rules for Cyclic-Type Queues, Adv. in App. Prob. 21, pp. 432–450, (1989).

[8] H. Dror and U. Yechiali, Closed Polling Models with Failing Nodes, Queueing Systems, Vol. 35, 55–81 (2000).

[9] H.D. Dykeman, M.H. Ammar and J.W. Wong, "Scheduling Algorithms for Videodex Systems under Broadcast Delivery", in Proceedings of the International Conference on Communications (ICC'86), pp. 1847–1851, (1986).

[10] M.R. Garey and D.S. Johnson, Computers and Intractability: A Guide to the Theory of NP-Completeness, W.H. Freeman, New York, NY, (1979).

[11] D. Grillo, Polling mechanism models in communications systems – some applications and examples, in: Stochastic Analysis of Computer and Communication Systems, ed. H. Takagi, North-Holland, Amsterdam, pp. 659–698, (1990).

[12] H. Levy, M. Sidi, Polling systems: applications, modeling, and optimization, IEEE Trans. Commun. 38 pp. 1750–1760, (1990).

[13] Z. Liu and P. Nain, "Optimal Scheduling in Some Multiqueue Single-Server Systems", IEEE Transactions on Automatic Control, Vol. 37, No. 2, pp. 247–252, (1992).

[14] C.N. Potts and M.Y. Kovalyov, "Scheduling with Batching: A Review", European Journal of Operational Research, Vol. 120, pp. 228–249, (2000).

[15] H. Takagi, Analysis of polling systems, MIT Press, Cambridge, Mass. (1986).

[16] H. Takagi, Queueing analysis of polling systems: an update, in: Stochastic Analysis of Computer and Communication Systems, ed. H. Takagi, North-Holland, Amsterdam, pp. 267–318, (1990).

[17] H. Takagi, Application of polling models to computer networks, Comp. Netw. ISDN Syst. 22 pp. 193–211, (1991).

[18] H. Takagi, Queueing analysis of polling models: Progress in 1990–1994, in: Frontiers in Queueing: Models and Applications in Science and Engineering, ed. J.H. Dshalalow, CRC Press, Buca Raton, pp. 119–146, (1997).

[19] H. Takagi, Queueing Analysis of Polling models, ACM Comput. Surveys 20, pp. 5–28, (1998).

[20] M.P. Van Oyen and D. Teneketzis, Optimal Batch Service of a Polling System under Partial Information, Method and Models in OR, Vol. 44, No. 3, pp. 401–419, (1996).

[21] C.H. Xia, G. Michailidis, N. Bambos and P.W. Glynn, Optimal Control of Parallel Queues with Batch Service, Probability in the Engineering and Informational Sciences, Vol. 16, No. 3, pp. 289-307, (2002).

[22] U. Yechiali, Analysis and control of polling systems, in: Performance Evaluation of Computer and Communication Systems, eds. L. Donatielo & R. Nelson, Springer-Verlag, pp. 630–650, (1993).

8 Appendix: The Scheduling Problem (4) is NP-Hard

We wish to show that the scheduling-optimization problem (4) of finding a permutation π so as to

$$\text{(A1)}\quad \textit{Minimize}\ \left\{\sum_{l=1}^{L}\omega_{\pi(l)}(n_{\pi(l)}+\lambda_{\pi(l)}b_{\pi(l)})\cdot(\sum_{j=1}^{l-1}h_{\pi(j)}) + \sum_{l=1}^{L}\omega_{\pi(l)}\lambda_{\pi(l)}\left[r_{\pi(l)}(\sum_{j=1}^{l-1}h_{\pi(j)})+(b_{\pi(l)}+d_{\pi(l)})(\sum_{j=l+1}^{L}h_{\pi(j)})+\frac{1}{2}(\sum_{j=1}^{l-1}h_{\pi(j)})^2+\frac{1}{2}(\sum_{j=l+1}^{L}h_{\pi(j)})^2\right]\right\}$$

is NP-Hard.

To prove this assertion we consider a special instance of the problem where, for all ℓ, $b_\ell = 0$, $d_\ell = r_\ell$ and $w_\ell = 1$. Then problem (A1) is reduced to minimizing

$$\text{(A2)}\quad f(\pi)=\sum_{l=1}^{L}n_{\pi(l)}\left(\sum_{j=1}^{l-1}h_{\pi(j)}\right)+\frac{1}{2}\sum_{l=1}^{L}\lambda_{\pi(l)}\left[\left(\sum_{j=1}^{l-1}h_{\pi(j)}\right)^2+\left(\sum_{j=l+1}^{L}h_{\pi(j)}\right)^2\right]$$

It is known [10] that the following *Partition* problem is NP-Complete:

Input: $a_1, a_2, \ldots, a_k$ are positive integers ($k \geq 2$).

Output: For $\sum_{j=1}^{k} a_j = A$, is there a subset $S \subseteq \{1, 2, \ldots, k\}$ such that $\sum_{j\in S} a_j = A/2$?

That is, the process of finding an answer 'Yes' or 'No' to the above output question is NP-Hard.

We start with transforming the Partition problem into the polling-optimization scheduling problem by taking $L = k + 1$; $n_l = 1$ for $l = 1, \ldots, k + 1$; $h_l = a_l$ for $l = 1, 2, \ldots, k$; $h_{k+1} = h_L = 1$; $\lambda_l = 1$ for $l = 1, 2, \ldots, k$; $\lambda_{k+1} = \lambda_L = M = 4(k + 1)A^2$.

We now prove the following:

Claim: The Output for the *Partition* problem is 'Yes' if and only if there exists a permutation $\overline{\pi}$ of $\{1, 2, \ldots, k + 1\}$ such that $f(\overline{\pi})$, the objective value of the respective scheduling-optimization problem, satisfies

$$\text{(A3)}\quad f(\overline{\pi}) \leq (k+1)A+\frac{k}{2}A^2+\frac{MA^2}{4}=B$$

where $M = 4(k + 1)A^2$.

Proof: Assume that there exists a partition and construct a permutation $\overline{\pi}$ satisfying $f(\overline{\pi}) \leq B$.

Suppose that the set $S \subseteq \{1, 2, \ldots, k\}$ is a Partition. That is, $\sum_{j \in S} a_j = \sum_{j \in S} h_j = A/2$. Let $|S| = m$. Let $\overline{\pi}$ be a permutation of $\{1, 2, \ldots, k+1\}$ satisfying

$$\begin{aligned} &\overline{\pi}(j) \in S \quad \text{for } j = 1, 2, \ldots, m \quad (m \leq k-1) \\ &\overline{\pi}(m+1) = L \\ &\overline{\pi}(j) \in \{1, 2, \ldots, k\} - S \quad \text{for } j = m+2, \ldots, k+1 \ . \end{aligned}$$

Substituting for $\overline{\pi}$ in $f(\cdot)$ we obtain

$$\begin{aligned} f(\overline{\pi} = &\sum_{l=1}^{L} \left(\sum_{j=1}^{l=1} h_{\overline{\pi}(j)} \right) + \frac{1}{2} \sum_{\substack{l=1 \\ l \neq m+1}}^{L} \left[\left(\sum_{j=1}^{l-1} h_{\overline{\pi}(j)} \right)^2 + \left(\sum_{j=\ell+1}^{L} h_{\overline{\pi}(j)} \right)^2 \right] \\ &+ \frac{1}{2} \lambda_L \left[\left(\sum_{j=1}^{m} h_{\overline{\pi}(j)} \right)^2 + \left(\sum_{j=m+2}^{L} h_{\overline{\pi}(j)} \right)^2 \right] \\ \leq\ & L \cdot A + \frac{1}{2}(L-1) \left(\sum_{j=1}^{k} h_j \right)^2 + \frac{1}{2} M \left[\left(\sum_{j \in S} a_j \right)^2 + \left(\sum_{j \notin S} a_j \right)^2 \right] \\ =\ & L \cdot A + \frac{L-1}{2} A^2 + \frac{1}{2} M \left[(\frac{A}{2})^2 + (\frac{A}{2})^2 \right] = L \cdot A + \frac{L-1}{2} A^2 + \frac{M}{4} A^2 = B \end{aligned} \tag{7}$$

The inequalities follow since, giving $h_j = a_j$ positive integers for $j = 1, 2, \ldots, L-1$, and $h_L = 1$, $\sum_{j=1}^{l-1} h_{\overline{\pi}(j)} \leq \sum_{j=1}^{L-1} a_j = A$ for every $l = 2, 3, \ldots, L$.

Next, suppose that the answer to the Output problem is 'No', that is, there is no partition such that $\sum_{j \in S} a_j = A/2$. We will show that for *any* permutation π, $f(\pi) > B$. Consider an arbitrary permutation π of $\{1, 2, \ldots, k+1 = L\}$ and suppose that $\pi(m) = k+1 = L$ for some m. Then

$$f(\pi) \geq \frac{1}{2} \lambda_{\pi(m)} \left[\left(\sum_{j=1}^{m-1} h_{\pi(j)} \right)^2 + \left(\sum_{j=m+1}^{L} h_{\pi(j)} \right)^2 \right]$$

Since there is no partition, $\sum_{j=1}^{m-1} h_{\pi(j)} \equiv x \neq \frac{A}{2}$, and $\sum_{j=m+1}^{L} h_{\pi(j)} \equiv y \neq \frac{A}{2}$, where $x + y = \sum_{j=1}^{k} a_j = A$.

Hence, since x and y are integers, $x^2 + y^2 \geq \left(\frac{A}{2} - \frac{1}{2}\right)^2 + \left(\frac{A}{2} + \frac{1}{2}\right)^2$ (this follows from the convexity and symmetry of the program $\min\{x^2 + y^2\}$ s.t. $x + y = A; x, y \geq 0$). Therefore,

$$\begin{aligned} f(\pi) &\geq \frac{1}{2} M \left[\left(\frac{A}{2} - \frac{1}{2} \right)^2 + \left(\frac{A}{2} + \frac{1}{2} \right)^2 \right] \\ &= \frac{M}{2} \left(\frac{A^2}{2} + \frac{1}{2} \right) = \frac{M}{4} A^2 + \frac{M}{4} = \frac{M}{4} A^2 + (k+1) A^2 \\ &= \frac{M}{4} A^2 + \frac{k}{2} A^2 + \left(\frac{k}{2} + 1 \right) A^2 > \frac{M}{4} A^2 + \frac{k}{2} A^2 + (k+1) A = B \end{aligned} \tag{8}$$

The last inequality follows since, for $k \geq 2$ and A positive integer satisfying$A \geq 2$, $\left(\frac{k}{2}+1\right) A^2 > (k+1)A$. Q.E.D.

To summarize, solving our polling scheduling-optimization problem is equivalent to solving the Partition problem, which is NP-Hard.