Pseudo-Cyclic Policies for Multi-Queue Single Server Systems

Offer Fabian Hanoch Levy

Department of Computer Science
School of Mathematics
Tel Aviv University

December 1991
Revised: October 1992

Abstract

In this paper we evaluate the waiting time performance of cycle-time guided algorithms in semi-dynamic polling models. In these models knowledge of the system state is used by the server in the process of on-line determination of the visit pattern. Our main focus is on pseudo-cyclic algorithms which achieve *fairness* by visiting every station exactly once in each cycle.

It is shown that in fully symmetric systems the performance of *every* pseudo-cyclic policy is bounded between the performance of the cycle maximization and the cycle minimization strategies. In particular, stochastic dominance is shown with regard to system workload, implying dominance of mean waiting times. In a fluid approximation model the performance ratio between the two bounding policies is derived and shown to be always between 1 and 3/4 under the gated service regime and between 1 and 1/2 under the exhaustive service regime. Under both regimes the ratio approaches 3/4 when the number of stations grows to infinity. Simulation results suggest that a similar performance ratio holds for the general (non-symmetric) stochastic model.

Further we study strategies which are guided by the cycle maximization (minimization) criteria, but which do not constrain themselves to pseudo-cyclic orders. It is shown that depending on the switch-over parameters these more dynamic policies may perform much worse than the pseudo-cyclic schemes.

1 Introduction

Polling systems are models which arise naturally in the modeling of many applications in telecommunications, computer networks, manufacturing etc., in which multiple users compete for the access to a single shared resource. Due to the variety of their applications, these models have been extensively studied in the past two decades. The basic polling system model consists of N stations attended by a single server according to a visit order which may be periodic (polling table), probabilistic (random polling) or dynamic. The service regimes considered range from the traditional exhaustive, gated and limited-1 regimes to various probabilistically limited or global regimes (see, e.g., [2] for further references).

In-spite of the extensive study devoted to polling systems, the topic of efficient operation and optimization received very little attention. In the recently updated polling survey by Takagi [2] out of more than 450 referenced studies only a few address optimization issues. Only very recently researchers started devoting more effort to study the efficient operation of these systems. One way of optimizing a polling model is via the server routing mechanism; such optimization is presented by Boxma, Levy and Weststrate [3,4] who derive efficient operational rules for polling tables. Another way of optimizing the polling model is via the service regime; Levy, Sidi and Boxma [5] compare the efficiency of different service regimes in terms of unfinished work and construct a hierarchy of the different schemes. Dynamic optimization of polling systems has been addressed by Meilijson and Yechiali [6] (for systems with no switch-over periods) and by Hofri and Ross [7] in which the server dynamically decides how long to serve each station.

Browne and Yechiali [1] study a model in which each station is subject to switch-in and swap-out periods and the visit order is *dynamically* determined. They suggest a semi-dynamic routing mechanism in which the server, upon completion of a service period, uses knowledge of the system state in determining the next station to be polled. The visit policy is *pseudo-cyclic* in the sense that at each cycle the server must visit all N stations exactly once but the order of visit within each cycle may change from cycle to cycle. Browne and Yechiali [1] address the optimization of this system where the objective is the minimization (or maximization) of the length of the current pseudo-cycle. Both the gated and exhaustive service regimes are studied and a simple index rule is given by which to choose the next station to visit (the index rule is presented in Section 2.3 below). The suggested index rule may be used to either maximize or minimize the current cycle length. Most of the attention in that work is given to the cycle minimization policy but there is no comparison between the two alternative policies.

This work concentrates on the dynamic polling model in which information about the system state is available before every switch-in epoch. Our goal is to evaluate the waiting time performance of cycle-time guided algorithms and thus derive efficient operational rules for the dynamic model. Since complete waiting time analysis of the general dynamic polling model is not available, we use a restricted model for which the analysis is possible. We first analyze the *Identical Station Fluid Approximation* (ISFA) model (described in Section 2.2) in which we quantify the process of work accumulation under the policies of cycle minimization and maximization. We find that the visit pattern converges in steady state and calculate the incurred waiting times under both policies. We compare the cycle minimization and cycle maximization policies under both the gated and the exhaustive regimes and find that the maximization strategy is better than the minimization strategy.

We then relax some of the model restrictions and analyze the stochastic model with identical stations (IS). In this model we analyze the workload and the mean waiting time under an arbitrary pseudo-cyclic visit policy. We present *stochastic* lower and upper bounds on the workload which imply bounds on the mean waiting time. We show that the cycle maximization strategy yields the lower bound while the cycle minimization strategy yields the upper bound. In the ISFA model the performance ratio of these two policies is shown to approach 3/4 when N goes to infinity. For the unrestricted general (non-symmetric) polling model, we present simulation results comparing the performance of cycle length minimization and maximization policies with that of cyclic polling. Extensive simulations show that the pseudo-cyclic cycle maximization policy is definitely superior to the cycle minimization policy. It is also observed that at medium to high utilizations ($\rho \geq 0.5$), under the gated (exhaustive) regimes, the performance ratio between the two policies ranges between 3/4 (1/2), when the loads are relatively symmetric, to 1, when the loads are relatively asymmetric.

Finally, we examine a polling model with a dynamic routing policy (dropping the pseudo-cycle restriction). For this model under ISFA assumptions we derive the visit pattern convergence at equilibrium under both the minimization and maximization policies. We extend the stochastic workload and mean waiting time bounds to hold for this model as well (under stochastic IS assumptions). We also present simulation results comparing this dynamic routing policy with the pseudo-cyclic routing policy and discuss the trade-off between "fairness" and "efficiency".

The structure of this paper is as follows: In Section 2 we present the general polling model, define the fluid approximation (FA) model and review the criteria for cycle minimization and maximization policies from [1]. In Section 3 we analyze the ISFA model, deriving the converged visit pattern, steady state workload and the waiting times at equilibrium. Section 4 is devoted to bounding the waiting times in the stochastic IS model. In Section 5 we present simulations of the general polling model. In Section 6 we present the model without the pseudo-cycle restriction for which we perform analysis under the ISFA model, extend the waiting time bounds and simulate the general dynamic model in order to compare it with the semi-dynamic model. A summary is provided in Section 7.

Lastly a word about very recent related work. In an independent study, Liu, Nain and Towsley [8] address optimal polling policies. One of their results implies that in fully symmetric polling systems, serving the largest queue (the greedy algorithm) is optimal. This result is similar to our result stating that the cycle maximization policy is stochastically the best policy (in terms of workloads).

2 Model description and preliminaries

Below we describe the general polling model, the fluid approximation model and the criteria for visit order determination.

2.1 The polling system model

We consider a system consisting of N infinite-buffer queues served by a common server. Customers arrive to station i ($i = 1, 2 \ldots N$) according to a Poisson process with rate λ_i; these are called type-i customers. In the analysis of the stochastic IS model we allow for a general independent arrival process distribution. Each type-i customer requires service-time whose duration is an independent random variable distributed as B_i and having first moment b_i.

The server may transit from any station i to any station j ($1 \leq i, j \leq N$) and is not restricted to the transitions $i \rightarrow i + 1$ modulo N as in the cyclic polling model. Upon transition, the server incurs two different switch-over times, *swap-out* which is the time required to leave the current station and *switch-in* which is the time required to enter the next station. Switch-in times for station i are random variables with first moment s_i. Swap-out times for station i are random variables with first moment r_i. The epoch at which the server starts (finishes) serving station i is called a *polling instant* (*departure instant*).

The service policy at queue i is either of:

exhaustive service: The server stays in the station until it is emptied. Note that the service period length is distributed as the sum of n_i independent busy periods in the corresponding $M/G/1$ queueing system (where n_i is the initial number of customers in the station at the polling instant).

gated service: The server serves all (and only) the customers found in the queue at the polling instant. The service period length is distributed as the sum of n_i random variables, the distribution of each is B_i.

We denote by $\rho_i \stackrel{\text{def}}{=} \lambda_i b_i$ the utilization of station i. We also denote four global quantities which are λ, the overall arrival rate ($\lambda \stackrel{\text{def}}{=} \sum_{i=1}^N \lambda_i$), ρ, the overall utilization ($\rho \stackrel{\text{def}}{=} \sum_{i=1}^N \rho_i$), s, the overall switch-in first moment ($s \stackrel{\text{def}}{=} \sum_{i=1}^N s_i$) and r, the overall swap-out first moment ($r \stackrel{\text{def}}{=} \sum_{i=1}^N r_i$).

The time epochs at the termination of a swap-out (and at the beginning of a switch-in) are referred to as *decision points*. At these decision points the server uses knowledge of the system state in determining the next station to visit.

2.2 A fluid approximation model

The fluid approximation approach has been useful in analyzing various queueing systems (see, e.g., Kleinrock [9] pp. 56-62). The special nature of the model is that while keeping ρ_i constant we let $\lambda_i \to \infty$ and $b_i \to 0$. This arrival process possesses an important property: during every time interval Δt, work deterministically flows into station i at rate ρ_i and thus the added amount of work during the time interval is $\rho_i \Delta t$.

An important performance measure is the workload at the different stations. Let u_i^t denote the workload at station i at the beginning of cycle t and u_i the workload at station i at the beginning of an arbitrary cycle at equilibrium. Let x_i^t denote the workload at station i at polling instant in cycle t and x_i the same quantity at equilibrium. By definition, it is obvious that variables u_i^t, u_i, x_i^t and x_i are all non-negative.

2.3 Service order algorithm and criteria

Browne and Yechiali [1] derived a simple index rule for determining an optimal visit order, which minimizes (maximizes) the cycle length. At every decision point the server selects the station with the lowest (highest) value of the following expressions[1]:

$$\begin{array}{ll} \dfrac{n_i b_i + s_i + (1-\rho_i) r_i}{\rho_i} & \text{exhaustive} \\ \dfrac{n_i b_i + (1+\rho_i) s_i + r_i}{\rho_i} & \text{gated} \end{array} \qquad [2.1]$$

Here n_i denotes the number of customers present at station i at the decision epoch (after completion of the swap-out). Note that $n_i b_i$ is the expected amount of work at station i at the decision epoch and thus we can replace $n_i b_i$ by u_i. For the fluid approximation model this substitution is essential since the use of n_i and b_i is meaningless. Thus, in this model

[1] The index rules result from the ratio between the "core" of work and the "growth rate"; see [1] for details.

the index rule is according to the following terms:

$$\begin{aligned} &\frac{u_i + s_i + (1-\rho_i) r_i}{\rho_i} && \text{exhaustive} \\ &\frac{u_i + (1+\rho_i) s_i + r_i}{\rho_i} && \text{gated} \end{aligned} \qquad [2.2]$$

Remark 2.1 *Note that in the FA model since the behavior during the next cycle is fully predictable, all decisions regarding visit pattern can be taken at the beginning of the cycle. Moreover, it is easy to see that exactly the same index rule may be used.*

3 Waiting time analysis of the ISFA model

This section is devoted to the analysis of waiting time in the identical station fluid approximation model. We first analyze the equilibrium visit order pattern for the cycle length minimization and maximization strategies. We analyze the evolution of workload in the system under both strategies for the exhaustive and gated service regimes. We determine the average waiting times and conduct a waiting time comparison of the different policies. Note that in this model all stations are identical and thus we denote: $\tilde{\rho} = \rho_i = \rho/N$, $\tilde{s} = s_i = s/N$ and $\tilde{r} = r_i = r/N$.

3.1 Equilibrium server routing patterns

The following theorem provides a useful tool in the determination of equilibrium routing patterns. The service regime may be gated, exhaustive or a mixture of the two. Note that the theorem is stated for the FA model and in particular it holds for ISFA models.

Theorem 3.1 *In a FA model for every system with mixed service and for any cycle t, if station i is served before station j and $\rho_i \geq \rho_j$ then $u_i^{t+1} \geq u_j^{t+1}$.*

Proof:

Let I_i^{Ex} be an indicator such that

$$I_i^{Ex} = \begin{cases} 1 & \text{if service at station } i \text{ is exhaustive} \\ 0 & \text{otherwise} \end{cases}$$

Without loss of generality, we assume that the service order for cycle t is $(1, 2 \ldots N)$. By definition

$$u_k^{t+1} = \rho_k \left(\sum_{l=k+1}^{N} s_l + \sum_{l=k}^{N} r_l \right) + \rho_k \sum_{l=k+I_k^{Ex}}^{N} x_l^\star \qquad , \qquad [3.1]$$

where $x_l^\star$ denotes the length of the busy period of station l, i.e. it is x_l if the service is gated and $\frac{x_l}{1-\rho_l}$ if it is exhaustive. From here it is easy to see that $u_i^{t+1} - u_j^{t+1} \geq 0$.

□

Theorem 3.2 *In the ISFA model, with either exhaustive or gated service, under the cycle maximization (minimization) strategy the stations are visited always according to the cyclic (alternating) pattern.*

Proof:

For arbitrary time t let us assume without loss of generality that the visit order in cycle t is $(1, 2 \ldots N)$. We go on to execute the cycle and look at the workload of the $t+1$-st cycle. Note that since we consider a fluid approximation model the algorithm is indifferent regarding whether the decisions are taken during the cycle or at the cycle beginning (see Remark 2.1). For simplicity of presentation assume that they are taken at cycle beginning. From Theorem 3.1 and since $\rho_i = \rho_{i+1}$ we have

$$u_i^{t+1} \geq u_{i+1}^{t+1} \qquad i = 1 \ldots N-1$$

Using criteria 2.2 in the identical stations model we see that the stations are ordered according to their workload. Under the cycle maximization strategy the order of cycle $t+1$ is $(1, 2 \ldots N)$ while under the cycle minimization strategy the visit order is $(N, N-1 \ldots 1)$.

□

Note that this result depends neither on t nor on the service discipline and thus we may conclude that for *both* service disciplines the initial service order repeats forever under the cycle length maximization strategy and alternates under the cycle length minimization strategy. Note that though Theorem 3.1 holds for systems with mixed service, Theorem 3.2 can not be extended to such systems since there is a different index rule for different service regimes and so the order of the tables is not maintained as is the case for the uniform service discipline.

3.2 Steady state workload analysis

In this section we find the equilibrium workloads at polling instants under the cycle maximization and minimization policies. To do this we refer to the steady state visit patterns found in the previous section.

The visit pattern of the cycle maximization strategy is cyclic. By balance reasoning we know that

$$E[C_i] = \frac{s+r}{1-\rho} \qquad [3.2]$$

for both the exhaustive and gated service regimes (where C_i is the length of the type-i cycle). Using the cycle length for the exhaustive regime we have:

$$x_j = \tilde{\rho}(1-\tilde{\rho})E[C_j] = \frac{\tilde{\rho}(s+r)(1-\tilde{\rho})}{1-\rho} \tag{3.3}$$

For the gated regime we have:

$$x_j = \tilde{\rho}E[C_j] = \frac{\tilde{\rho}(s+r)}{1-\rho} \tag{3.4}$$

In calculating the equilibrium workloads for the cycle length minimization strategy we use the alternating table visit pattern. In the alternating table we have an up cycle $(1 \ldots N)$ and a down cycle $(N \ldots 1)$. Let x_j (x'_j) denote the polling instant workload of station j during the up (down) cycle. Under the exhaustive regime we have:

$$x_j = (2j-1)\tilde{\rho}(\tilde{s}+\tilde{r}) + \frac{\tilde{\rho}}{1-\tilde{\rho}}\sum_{i=1}^{j-1}(x_i + x'_i) \tag{3.5}$$

$$x'_j = (2N-2j+1)\tilde{\rho}(\tilde{s}+\tilde{r}) + \frac{\tilde{\rho}}{1-\tilde{\rho}}\sum_{i=j+1}^{N}(x_i + x'_i) \tag{3.6}$$

Using balance arguments we have

$$2E[C_i] = \frac{2(s+r)}{1-\rho} \tag{3.7}$$

where $2E[C_i]$ is the mean of the sum of the up cycle and the down cycle. From balance arguments again, we get:

$$\frac{x_j + x'_j}{1-\tilde{\rho}} = \tilde{\rho}2E[C_i] \tag{3.8}$$

From [3.5], [3.6] and [3.8] we have:

$$x_j = \frac{(2j-1-\rho)\tilde{\rho}(\tilde{s}+\tilde{r})}{1-\rho} \tag{3.9}$$

$$x'_j = \frac{(2N-2j+1-\rho)\tilde{\rho}(\tilde{s}+\tilde{r})}{1-\rho} \tag{3.10}$$

Under the gated regime the relations are

$$x_j = (2j-1)\tilde{\rho}(\tilde{s}+\tilde{r}) + \tilde{\rho}\sum_{i=1}^{j-1}(x_i + x'_i) + \tilde{\rho}x'_j \tag{3.11}$$

$$x'_j = (2N - 2j + 1)\tilde{\rho}(\tilde{s} + \tilde{r}) + \tilde{\rho} \sum_{i=j+1}^{N} (x_i + x'_i) + \tilde{\rho} x_j \qquad [3.12]$$

Under the gated regime, [3.7] holds by the same balance reasoning and [3.8] is replaced with:

$$x_j + x'_j = \tilde{\rho} 2E[C_i] \qquad [3.13]$$

From [3.11], [3.12] and [3.13] we have:

$$x_j = \tilde{\rho} x'_j + \frac{(2j - 1 - \rho)\tilde{\rho}(\tilde{s} + \tilde{r})}{1 - \rho} \qquad [3.14]$$

$$x'_j = \tilde{\rho} x_j + \frac{(2N - 2j + 1 - \rho)\tilde{\rho}(\tilde{s} + \tilde{r})}{1 - \rho} \qquad [3.15]$$

From [3.14] and [3.15] we get

$$x_j = \frac{\tilde{\rho}(\tilde{s} + \tilde{r})}{(1 - \rho)(1 + \tilde{\rho})}[(N + \rho) - (N - 2j + 1)] \qquad [3.16]$$

$$x'_j = \frac{\tilde{\rho}(\tilde{s} + \tilde{r})}{(1 - \rho)(1 + \tilde{\rho})}[(N + \rho) + (N - 2j + 1)] \qquad [3.17]$$

3.3 Waiting time analysis

In this section we derive the mean waiting time of an arbitrary customer in the ISFA model. We start with some notation required for the waiting time analysis. Let C_i denote the length of a type-i cycle, i.e. the cycle for station i. The definition of a cycle depends on the service regime in effect: under the gated regime a type-i cycle is the time interval between two consecutive polling instants of station i while under the exhaustive regime it is the time interval between two consecutive departure epochs from station i. Let C_i^R denote the residual time of C_i i.e. the length of the time interval between an arbitrary arrival and the end of the cycle.

It is well known that

$$E\left[C_i^R\right] = \frac{E\left[C_i^2\right]}{2E\left[C_i\right]}. \qquad [3.18]$$

It is also known (see, e.g., [10]) that in a polling system:

$$E\left[W_i\right] = (1 + \rho_i) E\left[C_i^R\right] \qquad \text{gated} \qquad [3.19]$$

$$E\left[W_i\right] = (1 - \rho_i) E\left[C_i^R\right] \qquad \text{exhaustive.} \qquad [3.20]$$

A third relation is

$$E[W] = \sum_{i=1}^{N} \frac{\lambda_i}{\sum_{j=1}^{N} \lambda_j} E[W_i] \qquad [3.21]$$

which in the ISFA model simplifies to

$$E[W] = \frac{1}{N} \sum_{i=1}^{N} E[W_i]. \qquad [3.22]$$

Under the cycle maximization policy, since C_i is deterministic, $E[C_i^2] = E^2[C_i]$ and thus $E[C_i^R] = E[C_i]/2$. Since C_i^R does not depend on i we may find $E[W]$ directly for both service regimes

$$E[W] = \frac{(s+r)(1-\tilde{\rho})}{2(1-\rho)} \qquad \text{Exhaustive} \qquad [3.23]$$

$$E[W] = \frac{(s+r)(1+\tilde{\rho})}{2(1-\rho)} \qquad \text{Gated} \qquad [3.24]$$

Under the cycle minimization policy, we use the alternating table visit pattern for both service regimes. Under the exhaustive regime the cycle is now composed of two sub-cycles as is shown in the following illustration.

$$\ldots i \overbrace{\ldots NN \ldots i}^{C_i'} \overbrace{\ldots 11 \ldots i}^{C_i} \ldots$$

The sub-cycle lengths are found using the following relations

$$E[C_i] = \frac{x_i}{\tilde{\rho}(1-\tilde{\rho})} = \frac{(\tilde{s}+\tilde{r})}{(1-\rho)(1-\tilde{\rho})} \left[(N-\rho) - (N-2i+1)\right] \qquad [3.25]$$

$$E[C_i'] = \frac{x_i'}{\tilde{\rho}(1-\tilde{\rho})} = \frac{(\tilde{s}+\tilde{r})}{(1-\rho)(1-\tilde{\rho})} \left[(N-\rho) + (N-2i+1)\right] \qquad [3.26]$$

Since the cycle length is deterministic we have, again, $E[C_i^R] = E[C_i]/2$ and $E[C_i'^R] = E[C_i']/2$. $E[W_i]$ is found by averaging [3.20] over the sub-cycles:

$$E[W_i] = \frac{E[C_i]}{E[C_i] + E[C_i']}(1-\tilde{\rho})E[C_i^R] + \frac{E[C_i']}{E[C_i] + E[C_i']}(1-\tilde{\rho})E[C_i'^R] = \frac{(\tilde{s}+\tilde{r})\left[(N-\rho)^2 + (N-2i+1)^2\right]}{2N(1-\rho)(1-\tilde{\rho})} \qquad [3.27]$$

$$E[W] = \frac{(\tilde{s}+\tilde{r})\left[3(N-\rho)^2 + (N+1)(N-1)\right]}{6N(1-\rho)(1-\tilde{\rho})} \qquad [3.28]$$

Under the gated regime, the sub-cycle definition differs

$$\ldots \overbrace{i \ldots NN \ldots}^{C_i'} \overbrace{i \ldots 1, 1 \ldots}^{C_i} i \ldots$$

The sub-cycle lengths are found from the following relations

$$E[C_i] = \frac{x_i}{\tilde{\rho}} = \frac{(\tilde{s}+\tilde{r})}{(1-\rho)(1+\tilde{\rho})}\left[(N+\rho)-(N-2i+1)\right] \qquad [3.29]$$

$$E[C_i'] = \frac{x_i'}{\tilde{\rho}} = \frac{(\tilde{s}+\tilde{r})}{(1-\rho)(1+\tilde{\rho})}\left[(N+\rho)+(N-2i+1)\right] \qquad [3.30]$$

As is in the exhaustive system, $E[C_i^R] = E[C_i]/2$, $E[C_i'^R] = E[C_i']/2$ and thus the mean waiting times are:

$$E[W_i] = \frac{(\tilde{s}+\tilde{r})\left[(N+\rho)^2+(N-2i+1)^2\right]}{2N(1-\rho)(1+\tilde{\rho})} \qquad [3.31]$$

$$E[W] = \frac{(\tilde{s}+\tilde{r})\left[3(N+\rho)^2+(N+1)(N-1)\right]}{6N(1-\rho)(1+\tilde{\rho})} \qquad [3.32]$$

Remark 3.1 *Knowing that the visit order consists of an alternating table, it is possible to derive the mean waiting time of an arbitrary customer in the* **stochastic** *IS model thus yielding results [3.28] and [3.32] as special cases.*

3.4 Waiting time comparison

This section is devoted to a waiting time comparison of the two strategies examined. As the variance of the elevator pattern cycle-time is greater than that of the cyclic pattern, we expect that the cycle length maximization strategy would yield lower waiting times.

Let R_{Ex} be the ratio $E[W_{max}]/E[W_{min}]$ in the exhaustive regime

$$R_{Ex} = \frac{3(N-\rho)^2}{3(N-\rho)^2+(N+1)(N-1)} \qquad [3.33]$$

Note that R_{Ex} obeys

$$1/2 \leq R_{Ex} \leq 1 \qquad [3.34]$$

We see that the maximization strategy is better than the minimization strategy for all $N > 1$. Note that for $N = 1$ both strategies coincide in theory and practice.

Figure 1:

For the gated regime, define R_{Ga} similarly to R_{Ex}. Thus we have:

$$R_{Ga} = \frac{3(N+\rho)^2}{3(N+\rho)^2 + (N+1)(N-1)} \quad [3.35]$$

$$3/4 \leq R_{Ga} \leq 1 \quad [3.36]$$

Again, we see that the maximization strategy is better than the minimization strategy for all $N > 1$ while they coincide for $N = 1$.

For large systems we have for both service regimes

$$Lim_{N\to\infty} \frac{E[W_{max}]}{E[W_{min}]} = \frac{3}{4} \quad [3.37]$$

Remark 3.2 *As suggested by Boxma [11], expression [3.37] can be derived by analyzing a continuous polling system. For simplicity of illustration, let us denote this system by a unit parameter circle on which we arbitrarily place three points of interest: cycle beginning, server location and arbitrary customer arrival. This situation is described in Figure 1 in which the cycle beginning was placed at the top of the circle, c denotes the customer location and s the server location.*

From the figure it is easy to see that under the cycle length maximization strategy, in order to arrive to the customer, the server will require either one third time unit (on the average), or two thirds time units (on the average) depending on its moving direction. Thus, on the average it will require 1/2 time unit. In contrast, under the cycle length minimization strategy it will require either one third or three thirds time units, leading to an average requirement of 2/3 time units.

Comparing the exhaustive and the gated service regimes we have for any ρ and for every

$N = 1, 2 \ldots$:

$$R_{Ex} \leq R_{Ga} \qquad [3.38]$$

In the next section we show that the performance of these two policies form upper and lower bounds of the performance of any pseudo-cyclic policy. This concludes that the performance difference between various pseudo-cyclic policies is relatively small.

4 Stochastic bounding of pseudo-cyclic policies (IS model)

The previous section was devoted to the analysis of mean waiting time in an ISFA model. In this section we analyze and bound the mean waiting time of an arbitrary customer in the stochastic IS model. In this part of the work we do not require Poisson arrivals but the arrival process may be that of a general independent distribution (customers inter-arrival times are independent of each other and of the system state). We are interested in the mean delay of an arbitrary policy in the family of pseudo-cyclic ("fair") policies. Since it is difficult to analyze the mean waiting time of an arbitrary policy we bound its performance by that of two extreme policies. In fact, we show that the cycle length maximization strategy yields the lower bound while the cycle length minimization strategy yields the upper bound. In this sense, one may say that the cycle maximization strategy and the cycle minimization strategy are the best and worst in the family of pseudo-cyclic policies.

We start by defining the family of pseudo-cyclic visit order policies under consideration

$$\Pi = \{\pi \mid \pi = (\pi_1, \pi_2, \ldots),\ \forall k = 1, 2 \ldots,\ \forall i = 1 \ldots N, i \in \{\pi_{kN+1}, \pi_{kN+2} \ldots \pi_{kN+N}\}\}$$

This is the family of finite or infinite visit order policies, which result from the concatenation of various permutations of $1 \ldots N$. If a certain policy is finite then it is a member of Π simply by repeating it forever.

Below we prove that the amount of work in any pseudo-cyclic policy is stochastically higher (lower) than in the cycle maximization (minimization) policy. This implies that the mean waiting time under an arbitrary pseudo-cyclic policy is higher (lower) than that of the cycle maximization (minimization) policy. The structure of the proof is as follows: We first show that ordering the stations for cycle length maximization (minimization) is equivalent to ordering them by decreasing (increasing) workload. We then show that the workload is stochastically minimized (maximized) when visiting stations according to decreasing (increasing) workload. Finally, Little law is used to derive the corresponding relations between the mean waiting times.

We start with some notation. Let $\pi_{max} \in \Pi$ denote the visit pattern resulting from the cycle length maximization strategy and $\pi_{min} \in \Pi$ denote the pattern of the minimization strategy. Let us also denote by $U^{(t)}$ the system workload at time t and by $U_i^{(t)}$ the workload of queue i at time t.

Proposition 4.1 *Let n_i and n_j be the number of customers at station i and station j respectively at a decision point. Under the cycle length maximization (minimization) strategy station i will be visited before station j iff $n_i \geq n_j$ ($n_i \leq n_j$).*

The proof follows immediately from the definition of the decision criteria quantities in the IS model (Eq. 2.2).

We next turn to show that the system workload is minimized when more heavily loaded stations are visited earlier. We do this by referring to an arbitrary visit pattern $\pi 0 \in \Pi$. Assume that the system operates under $\pi 0$ and follows an arbitrary sample path and that at t' the server is facing the *i-th* decision point; Further assume that the station to be visited next ($\pi 0_i$) has less customers present than the station to be visited afterwards ($\pi 0_{i+1}$). We then define an alternative visit pattern $\pi 1 \in \Pi$ as follows:

$$\pi 1_j = \begin{cases} \pi 0_j & j < i \text{ } or \text{ } \pi 0_j \neq \pi 0_i, \pi 0_{i+1} \\ \pi 0_{i+1} & \pi 0_j = \pi 0_i \text{ } and \text{ } j \geq i \\ \pi 0_i & \pi 0_j = \pi 0_{i+1} \text{ } and \text{ } j \geq i \end{cases}$$

In other words, $\pi 1$ is the visit pattern resulting from $\pi 0$ in which after the $i-1$-*st* index all occurrences of the $\pi 0_i$-*th* and the $\pi 0_{i+1}$-*st* stations are interchanged (Note that the interchanged stations must belong to the same pseudo-cycle). We say that $\pi 0$ is taken by server 0 and $\pi 1$ is taken by server 1. We construct the following sample path for $\pi 1$: The realizations of the arrival, service demand and switch-time processes are interchanged between the two stations from time t' on. We also define the service requirement of the n customers present at queues $\pi 0_i$ and $\pi 0_{i+1}$ at t'. Let the service demand of the l-th customer ($l = 1, 2 \ldots n$) served under $\pi 1$ equal that of the l-th customer served under $\pi 0$. Using these definitions we next show that under $\pi 1$ with the constructed sample path the system workload is lower than under $\pi 0$ with the original sample path.

For the sake of presentation let us distinguish between three classes of customers. A customer is an *old customer* if he arrived before t' but was served afterwards. A customer is a *new customer* if he arrives after t'. The customers who arrive and are served before t' comprise the third customer group. Let us also denote by old and new work the service demand of the old and new customers respectively. Let us denote by τ_k^j the *k-th* polling instant ($k = 0, 1 \ldots$) under strategy πj ($j = 0, 1$) and let t_k^j similarly denote the departure instance. Let N_k^j (O_k^j) denote the new (old) workload consumed during the *k-th* service period ($k = 0, 1 \ldots$) under πj (where $j = 0, 1$). Let us denote by $S_0(t)$ the accumulated switch-time under $\pi 0$ up to time t and by $S_1(t)$ the same quantity for $\pi 1$.

In the following lemma we show that $t_k^0 \leq t_k^1$ for $k = 0, 1 \ldots$. We compare the performance of the server under $\pi 0$ and the original sample path with that of the server under $\pi 1$ with the constructed sample path. The proof of the lemma uses some ideas from Lemma 1 in [5] with the appropriate modifications.

Lemma 4.1 *For $k = 0, 1 \ldots$ and $n = 1 \ldots N$*

$$
\begin{array}{ll}
(a) & t_k^0 \leq t_k^1 \\
(b1) & \sum\limits_{\substack{l:\pi 0_l = \pi 0_i \\ i \leq l \leq i+k}} N_l^0 \leq \sum\limits_{\substack{l:\pi 1_l = \pi 1_i \\ i \leq l \leq i+k}} N_l^1 \\
(b2) & \sum\limits_{\substack{l:\pi 0_l = \pi 0_{i+1} \\ i \leq l \leq i+k}} N_l^0 \leq \sum\limits_{\substack{l:\pi 1_l = \pi 1_{i+1} \\ i \leq l \leq i+k}} N_l^1 \\
(b3) & \sum\limits_{\substack{l:\pi 0_l = n \\ i \leq l \leq i+k}} N_l^0 \leq \sum\limits_{\substack{l:\pi 1_l = n \\ i \leq l \leq i+k}} N_l^1 \qquad n \neq \pi 0_i, \pi 0_{i+1} \\
(b4) & \sum\limits_{l=0}^{k} O_l^0 \leq \sum\limits_{l=0}^{k} O_l^1
\end{array}
$$

Proof:

Without loss of generality we may assume $i = 0$. By definition the proof trivially holds for $k < 0$. The proof for $k \geq 0$ is by way of induction.

Induction basis:

For $k = 0$ we know that both servers entered the stations at the same time since the switch processes are swapped. Upon entry the servers are faced with the old customers (which existed at the decision point) and possibly also new customers which arrived during the switch-in period. By definition, the new customers are identical in both interchanged stations. Since server 1 is facing more old work, he will take longer to consume it and turn to the new customers. Thus, $t_0^0 \leq t_0^1$ under both the gated and the exhaustive regimes.

The proof of $(b1)$ depends on the service regime; Under the gated regime, the servers will serve exactly the same new customers arriving during the switch-in period. Under the exhaustive regime the amount of consumed new work may differ between the two. We have already established that server 1 starts attending the new customers after server 0. If any new customers arrive in the time interval between the time server 0 started consuming new work and the time server 1 does, then clearly server 1 will consume more new work by the end of the current busy period. If there are no arrivals in that time interval then the amount of consumed new work is identical. The proofs of $(b2)$ and $(b3)$ are immediate since no service was rendered in all stations other than $\pi 0_0$ and $\pi 1_0$.

The proof of $(b4)$ follows immediately from the fact that the amount of old work consumed by both servers is identical to that present in the stations at t'. Therefore the lemma holds for $k = 0$.

Induction step

We assume now the lemma holds for k, we prove it holds for $k+1$.

We first prove (a) by way of contradiction. For the contrary, assume $t^0_{k+1} > t^1_{k+1}$. Let us denote

$$W^1_{\pi 1_{k+1}}(t^1_{k+1}) = t^1_{k+1} - t' - \left[\sum_{\substack{l:\pi 1_l \neq \pi 1_{k+1} \\ 0 \le l \le k}} N^1_l + \sum_{l=0}^{k} O^1_l + S_1(t^1_{k+1}) - S_1(t') \right] \qquad [4.1]$$

$$W^0_{\pi 0_{k+1}}(t^1_{k+1}) = t^1_{k+1} - t' - \left[\sum_{\substack{l:\pi 0_l \neq \pi 0_{k+1} \\ 0 \le l \le k}} N^0_l + \sum_{l=0}^{k} O^0_l + S_0(t^1_{k+1}) - S_0(t') \right] \qquad [4.2]$$

which is the accumulated consumed new work of the stations $\pi 1_{k+1}$ and $\pi 0_{k+1}$ under $\pi 1$ and $\pi 0$ respectively. From the induction hypothesis (a) and since the switch-times are identical, server 0 is actively serving at time t^1_{k+1}. From the induction hypothesis $(b1 - b4)$, we know that $W^0_{\pi 0_{k+1}}(t^1_{k+1}) \geq W^1_{\pi 1_{k+1}}(t^1_{k+1})$. Under the gated regime, $W^1_{\pi 1_{k+1}}(t^1_{k+1})$ is the total amount of new work that arrived to station $\pi 1_{k+1}$ under $\pi 1$ up to time τ^1_{k+1}. By the definition of $\pi 1$ and the sample paths, this is also the amount of new work that arrived at station $\pi 0_{k+1}$ under $\pi 0$ up to the same time. By the same reasoning we may conclude that $W^0_{\pi 0_{k+1}}(t^1_{k+1})$ is the total amount of work that arrived at station $\pi 0_{k+1}$ under $\pi 0$ up to time τ^0_{k+1}. Since $\tau^0_{k+1} \leq \tau^1_{k+1}$ while $W^0_{\pi 0_{k+1}}(t^1_{k+1}) \geq W^1_{\pi 1_{k+1}}(t^1_{k+1})$ we get a contradiction. The proof for the exhaustive regime is done by replacing the polling instants τ^j_l with the departure instants t^j_l. Hence $t^0_{k+1} \leq t^1_{k+1}$ and we have proved (a) for $k+1$.

For the second part of the lemma $(b1 - b4)$ we define $W^j_0(t)$ to be the total amount of new work that arrived at station $j = \pi 0_{k+1}$ up to time t under $\pi 0$. Let us denote by $I(j)$ the following index substitution mechanism:

$$I(j) = \begin{cases} \pi 1_0 & j = \pi 0_0 \\ \pi 1_1 & j = \pi 0_1 \\ j & \text{otherwise} \end{cases}$$

Let $W^{I(j)}_1(t)$ denote the amount of new work that arrived at station $I(j)$ up to time t. Using these definitions under the gated regime we have,

$$\sum_{\substack{l:\pi 0_l=j\\ 0\le l\le k+1}} N_l^0 = W_0^j(\tau_{k+1}^0) \le W_1^{I(j)}(\tau_{k+1}^1) = \sum_{\substack{l:\pi 1_l=I(j)\\ 0\le l\le k+1}} N_l^1 \qquad [4.3]$$

where the first equality follows from the fact that by time τ_{k+1}^0 server 0 will exhaust all the new work that arrived up to that time. The second equality, similarly, follows the same reasoning for server 1 and the inequality follows since $\tau_{k+1}^0 \le \tau_{k+1}^1$ and from the definition of $W_0^j(t)$ and $W_1^{I(j)}(t)$. For stations other than $\pi 0_{k+1}$ and $\pi 1_{k+1}$, $(b1), (b2)$ and $(b3)$ trivially hold for $k+1$.

The proof of $(b4)$ follows from the fact that for every station j $(j = 1 \ldots N)$ apart from the two interchanged stations, the old workload is identical under $\pi 0$ and $\pi 1$. Since the old work in both interchanged stations is consumed during the $k = 1$ visit, we have:

$$\sum_{l=0}^{k} O_l^0 = \sum_{l=0}^{k} O_l^1 \qquad \forall k \ge 1$$

and thus the proof of the lemma is concluded.

$\square$

Next, we define stochastic dominance and, based on the lemma establish the stochastic dominance of the workload under $\pi 0$ over the workload under $\pi 1$.

Definition 4.1 *Let X^1 and X^2 be two real valued random variables, we say that X^1 is stochastically smaller than X^2, denoted $X^1 \le_{\mathrm{st}} X^2$, if $P(X^1 \le x) \ge P(X^2 \le x)$ for any real x.*

Theorem 4.1 *Let $\pi 0 \in \Pi$ be an arbitrary policy. Let t' be the epoch of the i-th decision point. If $U_{\pi 0_i}^{(t')} \le U_{\pi 0_{i+1}}^{(t')}$ and $\lceil i/N \rceil = \lceil (i+1)/N \rceil$, then there exists a policy $\pi 1$ such that $\forall t \ge 0, \quad U_{\pi 1}^{(t)} \le_{\mathrm{st}} U_{\pi 0}^{(t)}$.*

Proof:

Obviously, $\pi 1$ is the policy defined just before Lemma 4.1.

Consider policy $\pi 0$ with the original sample path and policy $\pi 1$ with the constructed sample path. For $k = 1 \ldots i-1$ the theorem holds by definition. For $k \ge i$ we have shown in Lemma 4.1 that $t_k^0 \le t_k^1$. By definition of the sample paths we have $\tau_k^0 \le \tau_k^1$ for $k = 1, 2 \ldots$. By definition, $\forall t \ge 0$, $S_0^{(t)} \ge S_1^{t)}$ which implies that $\forall t \ge 0$, $U_0^{(t)} \ge U_1^{t)}$. Since the statistics of all stations are the same we may now use sample path mapping considerations to conclude the proof.

$\square$

We now come to the main theorem of this section stating that the cycle maximization and the cycle minimization policies yield respectively the lowest and highest mean waiting times over all pseudo-cyclic schemes. Let us denote by $E[W \mid \pi]$ the mean waiting time of an arbitrary customer under policy π.

Theorem 4.2 *In an IS model,*

$$\forall \pi \in \Pi \qquad E[W \mid \pi_{max}] \leq E[W \mid \pi] \leq E[W \mid \pi_{min}]$$

Proof:

Proposition 4.1 shows that the cycle length maximization (minimization) strategy polls the stations according to decreasing (increasing) workloads. Theorem 4.1 implies that ordering the stations in decreasing (increasing) order of workload minimizes (maximizes) the mean workload. Application of Little law concludes the proof.

□

Remark 4.1 *Note that the theorem holds for all IS models with a general independent arrival process.*

Combining Theorem 4.2 and the analysis of the ISFA model, we may now bound the performance of all pseudo-cyclic policies in the ISFA model.

Theorem 4.3 *In an ISFA model for any pseudo-cyclic policy $\pi \in \Pi$ we have*

$$\begin{array}{ll} 1/2E[W \mid \pi_{min}] \leq E[W \mid \pi] \leq E[W \mid \pi_{min}] & \text{exhaustive} \\ 3/4E[W \mid \pi_{min}] \leq E[W \mid \pi] \leq E[W \mid \pi_{min}] & \text{gated} \end{array}$$

Where $E[W \mid \pi_{min}]$ is given in [3.28] and [3.32].

5 The general (non-symmetric) polling model

In the previous sections we focused our analysis on several models of identical stations. In this section we remove all restrictions and examine the general non-identical stations polling model. Unfortunately, the general model is very difficult to analyze. In view of this difficulty, we resort to simulations. We simulate runs of the two strategies (cycle length minimization and maximization) and of cyclic polling. In the examination, all three strategies are tested over the same sample path realization (with 20 different paths). The distribution of all random variables in the simulation is exponential and the service regime is gated with various parameters.

The test cases were chosen as to cover a wide variety of parameter settings. Table 1 presents the 9 test cases. In the test cases $b_i = 1$ for all settings.

Table 2 presents the simulation results. The first two columns describe the cases (corresponding to Table 1). Columns $3 - 5$ present the mean waiting time and its standard deviation under the cycle maximization, cyclic and cycle minimization policies respectively. The last column presents the average of the mean waiting time ratio between the cycle maximization and the cycle minimization policies (averaged over the 20 repetitions).

Num.	Description	N	Station #	ρ_i	s_i	r_i
1	IS low ρ	10	$1 \ldots 10$	0.01	1.0	1.0
2	IS high ρ	10	$1 \ldots 10$	0.09	1.0	1.0
3	Unequal ρ_i	2	1	0.1	1.0	1.0
	2 stations		2	0.8	1.0	1.0
4	Unequal ρ_i	10	$1 \ldots 9$	0.01	1.0	1.0
	10 stations		10	0.8	1.0	1.0
5	Unequal s_i	2	1	0.25	1.0	1.0
	2 stations		2	0.25	10.0	1.0
6	Unequal s_i	10	$1 \ldots 9$	0.05	1.0	1.0
	10 stations		10	0.05	10.0	1.0
7	Unequal s_i	10	$1 \ldots 5$	0.05	1.0	1.0
	half & half		$6 \ldots 10$	0.05	10.0	1.0
8	Unequal r_i	10	$1 \ldots 5$	0.05	1.0	1.0
	half & half		$6 \ldots 10$	0.05	1.0	10.0
9	Unequal ρ_i	2	1	0.8	1.0	1.0
	and s_i		2	0.1	10.0	1.0

Table 1: Simulation test cases.

Num.	Description	Max. policy		Cyclic policy		Min. policy		Max/Min	
		$E[W]$	σ_w	$E[W]$	σ_w	$E[W]$	σ_w	$E[f]$	σ_f
1	IS low ρ	6.63	6.13	11.84	7.41	20.01	10.05	0.33	0.002
2	IS high ρ	112.79	75.10	115.68	71.24	145.35	98.71	0.77	0.014
3	Unequal ρ_i 2 stations	40.79	19.88	41.06	20.03	42.76	21.01	0.95	0.005
4	Unequal ρ_i 10 stations	153.27	53.53	154.07	52.85	157.75	55.81	0.97	0.009
5	Unequal s_i 2 stations	33.24	19.00	33.26	19.01	35.56	23.01	0.94	0.008
6	Unequal s_i 10 stations	40.49	29.49	44.07	27.83	53.01	33.27	0.76	0.009
7	Unequal s_i half & half	117.09	72.56	120.30	69.84	148.51	95.82	0.79	0.007
8	Unequal r_i half & half	20.39	14.19	22.57	13.69	26.22	15.04	0.78	0.008
9	Unequal ρ_i and s_i	179.52	64.52	179.66	64.58	187.68	71.64	0.95	0.005

Table 2: Simulation results for the general polling model with pseudo-cyclic routing

An important observation that can be made is that the cycle maximization strategy is better than the cycle minimization strategy for all variations of the model.

A very surprising observation is that the bound and limits proven for the ISFA model seem to hold in general for a wide variety of parameters in a stochastic non-symmetric system. Apart for the first test case, in all other test cases the ratio was greater than or equal to 3/4. In each of the 8 cases $\rho \geq 1/2$ apart for the first case where $\rho = 0.1$. These observations lead us to conjecture that the bounds and limits of the ISFA model waiting times may be valid (under a wide range of parameters) in the general model as well. Specifically the conjectures regarding the mean waiting time of an arbitrary pseudo-cyclic policy π are:

$$1/2E[W \mid \pi_{min}] \leq E[W \mid \pi] \leq E[W \mid \pi_{min}] \qquad \text{exhaustive}$$

$$3/4E[W \mid \pi_{min}] \leq E[W \mid \pi] \leq E[W \mid \pi_{min}] \qquad \text{gated}$$

We may also make the following observations:

- The (relative) performance difference decreases when the system utilization increases. The same observation was made in the ISFA model since in lightly loaded systems, the probability of empty queues is relatively high. In such systems under the cycle minimization strategy the server first passes through the empty stations thus performing needless switch-over periods. In contrast, the cycle maximization strategy has a lower probability of encountering an empty station. When the system load increases the probability of empty stations decreases and the performance gap narrows.

- The effects of variations in switch-in and swap-out time on the performance differences are similar to each other. This is intuitive due to their similar nature.

6 Dynamic server routing

In the analysis in the previous sections we considered visit policies which force the server to visit every station exactly once in each cycle. That restriction is an inherent part of cycle length guided policies. In this section we present a modification to the model analyzed so far. This modification allows us to relax restrictions regarding visit order fairness and to handle a larger family of visit patterns. Since we have shown that maximizing cycle length is equivalent to a greedy strategy which will visit the station with the heaviest expected load (in the IS model) we may remove the pseudo-cycle restriction and examine the dynamic greedy strategies. This modification affects only the controversial fairness factor since a station may be polled several times before another is polled once.

6.1 ISFA model analysis

In this section we analyze the visit pattern generated by the server in the ISFA model under the dynamic minimization and maximization strategies. The decision criteria used is 2.2. Under the dynamic minimization strategy the server polls the station with the lowest criterion value while under the dynamic maximization it polls the station with the highest value.

Let us begin with the analysis of the dynamic maximization policy. Since the server no longer serves in cycles, let u_i^t denote the workload of station i at decision point t. We may easily find the visit pattern convergence with the following lemma

Lemma 6.1 *In an ISFA model under the dynamic maximization policy, if j is the station selected by the server at decision point t (which implies $u_j^t \geq u_i^t$ for $i = 1 \ldots N$) then $u_j^{t+1} \leq u_i^{t+1}$ for $i = 1 \ldots N$.*

Proof:

Proof follows directly from the definition of u_i^{t+1} which depends on the service policy employed. Under the exhaustive regime we have

$$u_i^{t+1} = \begin{cases} \tilde{\rho}\tilde{r} & i = j \\ u_i^t + \tilde{\rho}(\tilde{s} + \tilde{r}) + \frac{\tilde{\rho}}{1-\tilde{\rho}} u_j^t & i \neq j \end{cases} \qquad [6.1]$$

Under the gated regime we have

$$u_i^{t+1} = \begin{cases} \tilde{\rho}(u_j^t + \tilde{r}) & i = j \\ u_i^t + \tilde{\rho}(\tilde{s} + u_j^t + \tilde{r}) & i \neq j \end{cases} \quad [6.2]$$

□

The lemma shows that once a station is polled it becomes the one with the lowest criterion value in the next decision point. Since work flow is constant and identical to all stations, the station will be polled again only after all other $N-1$ stations were polled, thus yielding a cyclic visit pattern. The steady-state analysis for the ISFA model under the cycle maximization strategy of both workload and waiting time is, therefore, valid for this dynamic strategy as well.

The analysis of the dynamic minimization policy is done in a similar manner, with the following Lemma

Lemma 6.2 *In an ISFA model under the dynamic minimization policy, if j is the station selected by the server at decision point t (which implies $u_j^t \leq u_i^t$ for $i = 1 \ldots N$) then $u_j^{t+1} \leq u_i^{t+1}$ for $i = 1 \ldots N$.*

Proof:

The proof follows directly from Lemma 6.1 since under the dynamic minimization strategy the station visited is the one with the lowest criterion value.

□

From the lemma we see that once a station is polled, it will be polled again forever. In this case the visit pattern generated by the server degenerates to a repetition of a single station. Since the other $N-1$ stations are never polled the system is unstable and an arbitrary customer suffers an indefinite delay.

6.2 IS model waiting time bounds

In Section 4 we presented stochastic bounds for the waiting times of an arbitrary pseudo-cyclic policy. In this section we show that these bounds may be easily extended to the dynamic stochastic IS model.

Let us denote by Π^d the family of all infinite visit patterns. In particular $\Pi \subset \Pi^d$. Let $\pi^d_{max} \in \Pi^d$ denote the visit pattern resulting from the dynamic maximization strategy and let π^d_{min} denote the pattern of the dynamic minimization strategy.

The following theorem states that the dynamic maximization and minimization policies yield respectively the lowest and highest mean waiting times over all dynamic schemes.

Theorem 6.1 *In a dynamic IS model*

$$\forall \pi \in \Pi^d \qquad E[W \mid \pi^d_{max}] \leq E[W \mid \pi] \leq E[W \mid \pi^d_{min}]$$

Proof:

The proof is identical to that of the pseudo-cyclic IS model with a minor modification. In the definition of $\pi 1$ used in Lemma 4.1 one may drop the restriction stating that i and $i+1$ must be in the same pseudo-cycle. After dropping this restriction, all steps of the proof are valid for this model as well.

□

6.3 Waiting time simulation runs

In this section we offer simulation runs comparing the cycle length maximization strategy with the dynamic maximization policy. The test cases are presented in Table 1. In the simulation both pseudo-cyclic and dynamic maximization policies were simulated over the same arrival and service demand sample paths. The switch-over periods were generated randomly. All random processes were simulated with the exponential distribution function. The service regime was gated.

Table 3 presents the simulation results. The first two columns correspond to Table 1. Column 3 presents the mean waiting time of pseudo-cyclic maximization and its standard deviation. Column 4 presents the mean and standard deviation of the dynamic maximization. The last column presents the mean of the performance ratio between the dynamic and the pseudo-cyclic maximization policies and its standard deviation.

Num.	Description	Pseudo-cyclic policy		Dynamic policy		Dyn/PsCyc	
		$E[W]$	σ_w	$E[W]$	σ_w	$E[f]$	σ_f
1	IS low ρ	6.63	6.13	3.12	2.66	0.47	0.007
2	IS high ρ	110.62	73.20	112.05	76.47	1.01	0.027
3	Unequal ρ_i 2 stations	40.85	20.12	41.57	19.22	1.02	0.028
4	Unequal ρ_i 10 stations	153.36	52.28	197.40	53.63	1.30	0.076
5	Unequal s_i 2 stations	33.05	18.96	44.85	27.29	1.36	0.039
6	Unequal s_i 10 stations	40.77	29.83	113.56	84.05	2.79	0.109
7	Unequal s_i half & half	117.33	72.56	157.21	116.12	1.34	0.030
8	Unequal r_i half & half	20.40	14.15	51.51	63.59	2.53	0.064
9	Unequal ρ_i and s_i	183.42	65.09	197.48	61.45	1.08	0.037

Table 3: Simulation results for the general polling model, pseudo-cyclic and dynamic routing policies

It is a common concept in queueing theory that by relaxing “fairness” restrictions one may improve efficiency. In Table 3 we see that this is not the case in this model. When all stations are identical and the utilization is low, the dynamic strategy is better than the pseudo-cyclic strategy. However, when the utilization increases or there is a difference between the stations, the pseudo-cyclic strategy is better. These observations can be explained as follows.

In a model where stations differ in their switch-over periods, the dynamic strategy yields mean waiting times up to $2\frac{1}{2}$ times those of the pseudo-cyclic strategy. The reason for the difference may be illustrated by the following example. Consider a two station model where both stations are identical in all aspects except for switch-in times. The first station requires a long switch-in period while the second requires a much shorter one. When the number of customers in both stations is similar both the dynamic and pseudo-cyclic maximization policies will favor the first station. While the pseudo-cyclic strategy is limited to a single visit before polling the second station, the dynamic strategy may poll the first station several times. Thus the accumulated switch-time of the dynamic strategy is much greater than that of the pseudo-cyclic strategy. The effect does not occur when we compare pseudo-cyclic minimization with pseudo-cyclic maximization since both strategies will eventually perform the same amount of switching in every pseudo-cycle.

The visit order criterion 2.2 refers to the mean expected amount of work divided by the average work consumption rate. The result is that long switch-over periods increase the likelihood of a station to be polled. Intuitively one would conjecture that the less idle time "wasted" by the server, the more efficient is the routing strategy. Keeping the conjecture in mind together with the simulation results presented in Table 3, we believe that the routing criteria in the general model should include a *trade-off* between existing workload and expected switch-time, however it is beyond the scope of this work to suggest other criteria.

The second observation made was that in a model where the system utilization is low, the dynamic strategy is better. It is in this kind of model that the dynamic nature of the routing process plays a key role. In many decision points most stations are empty while others have a small number of customers. By allowing the server to disregard the empty stations and serve only where service is needed we achieve much better performance (the waiting times of the dynamic strategy are twice as good). When system utilization increases both strategies will follow along a similar visit pattern and performance is similar. The small performance difference between the two is caused by decision points where several stations have the same criterion value. In such cases the pseudo-cyclic restriction may be an advantage. One may overcome this performance difference by keeping track of previous visits for both servers and in a decision point where the server must choose between several equal criterion value stations, force the server to choose the station whose elapsed time since the previous visit is the longest. With this modification the performance of the dynamic strategy is expected to be slightly better than that of the pseudo-cyclic strategy due to the reasons mentioned above.

7 Summary and future work

Our goal in this work has been to derive efficient (in terms of waiting times) operational rules for semi-dynamic polling systems. We focused our attention on cycle minimization and cycle maximization strategies and considered other policies as well.

In the identical station fluid approximation model we presented complete analysis of work accumulation and equilibrium workload and waiting times. We showed that the cycle maximization strategy yields lower waiting times than the cycle minimization strategy. In one-station systems both strategies are identical while as the number of stations grows to infinity, the performance ratio between the two increases up to 3/4.

In the stochastic identical station model we showed that the cycle maximization and the cycle minimization strategies form upper and lower bounds on any pseudo-cyclic policy. These bounds are in the sense that the workload in any policy is stochastically lower (higher) than that in the cycle minimization (maximization) policy. This stochastic dominance of workload translates to similar dominance of the mean delays.

For the general (asymmetric) polling model we conducted a comparison (via simulation) between three policies: 1) Cycle maximization, 2) Cycle minimization and 3)Cyclic visit order. The results show that the cycle maximization policy is the best while the cycle minimization strategy is the worst among the three.

We finally considered policies which do not require pseudo-cyclic orders. We showed that in the ISFA model the strategy which is guided by the cycle maximization criterion was not affected by the relaxation of the pseudo-cycle restriction; in contrast the strategy which is guided by the cycle minimization criterion ends-up serving a single station and becomes unstable with infinite delays in all but one station. We have extended the pseudo-cyclic bounds of the semi-dynamic models to hold for this model as well. In simulation runs we compared the pseudo-cyclic maximization policy with the dynamic maximization policy and observed that in this model "fairness" may lead to efficiency.

It is an open question whether the stochastic dominance result may be extended to non-identical station models. We conjecture that similar results may be derived for systems with identical switch-over periods and identical service time distributions. However, such results may require more sophisticated analysis. Another important open question is how to incorporate cost factors into these dynamic optimization problems.

References

[1] S. Browne and U. Yechiali (1987), "Dynamic Priority Rules for Cyclic-Type Queues", Adv. Appl. Prob., 21, pp. 432, 1989.

[2] H. Takagi (1990), "Queueing Analysis of Polling Models", In: H. Takagi (ed.), Stochastic Analysis of Computer and Communication Systems (North-Holland Publ. Cy., Amsterdam).

[3] O. J. Boxma, H. Levy and J. A. Weststrate (1990), "Optimization of Polling Systems", in: *Performance '90*, eds. P.J.B King, I. Mitrani and R.J. Pooley (North-Holland, Amsterdam), pp. 349-361, September 1990 (also Technical report BS-R8932, CWI April 1991).

[4] O. J. Boxma, H. Levy and J. A. Weststrate (1991), "Efficient Visit Frequencies for Polling Tables: Minimization of Waiting Costs", *Queueing Systems*, vol. 9 (1991), pp. 133-162.

[5] H. Levy, M. Sidi and O. J. Boxma (1990) "Dominance Relations in Polling Systems", *Queueing Systems*, vol. 6 (1990), pp. 155-172.

[6] I. Meilijson and U. Yechiali (1977), "On Optimal Right-Of-Way Policies at a Single-Server Station When Insertion of Idle Times is Permitted", Stochastic Processes and Their Applications, 6, pp. 25, 1977.

[7] M. Hofri and K. W. Ross (1987), "On the Optimal Control of Two Queues with Server Setup Times and Its Analysis", *SIAM J. on Comp.*, 16, pp. 399, 1987.

[8] Z. Liu, P. Nain, D. Towsley, (1991), "On optimal polling policies", *Queueing Systems*, **11**, 59–83, July 1992.

[9] L. Kleinrock (1976), "Queueing Systems, Vol II", New York: John Wiley & Sons.

[10] H. Takagi (1986), "Analysis of Polling Systems", MIT Press.

[11] O. J. Boxma, Private communication.