Scheduling in polling systems in heavy traffic

R. Bekker[a], J.L. Dorsman[b,c] *, R.D. van der Mei[a,c], P. Vis[a,c] and E.M.M. Winands[d]

[a]VU University Amsterdam, Amsterdam, Netherlands
[b]Eindhoven University of Technology, Eindhoven, Netherlands
[c]Centre for Mathematics and Computer Science, Amsterdam, Netherlands
[d]University of Amsterdam, Netherlands

ABSTRACT

We consider the classical cyclic polling model with Poisson arrivals and with gated service at all queues, but where the local scheduling policies are not necessarily First-Come-First-Served (FCFS). More precisely, we study the waiting-time performance of polling models where the local service order is Last-Come-First-Served (LCFS), Random-Order-of-Service (ROS) or Processor Sharing (PS). Under heavy-traffic conditions the waiting times turn out to converge to products of generalized trapezoidal distributions and a gamma distribution.

1. INTRODUCTION

A polling system is a multi-queue single-server system in which the server visits the queues in some order to process requests pending at the queues. Polling models find a wealth of applications in areas like computer-communication systems and production systems [2]. In this paper, we study the impact of the local scheduling policy on the waiting-time performance.

It might be natural to assume that the impact of such local scheduling is small, because it only impacts the system performance locally, leaving the amount of time spent outside the targeted queue unaffected. However, [9] illustrates that the impact on the mean waiting time from scheduling within a queue of a polling system can be significant. In many application areas of polling models, such as Bluetooth and 802.11 protocols, scheduling policies at routers and I/O subsystems in web servers, the workloads are known to have high variability and priority-based scheduling could therefore be beneficial. So far, there are only a few papers known where the effect of priority-based scheduling is studied in polling systems (e.g. [1, 3, 9]), whereas the system operation under heavy-traffic (HT) conditions, being the most critical regime from a scheduling point of view, has not been studied at all.

The main results of the paper are closed-form expressions for the (scaled) waiting-time distribution under HT assumptions, explicitly quantifying the impact of the local scheduling policies on the waiting-time performance. The main surprising observation made is that the asymptotic distribution of the waiting time for all of these priority-based scheduling disciplines can be described by the product of a *generalized trapezoidal distribution* and a gamma distribution, both with known parameters, giving fundamental novel insights in the impact of priority-based scheduling on the waiting-time performance.

*The research of J.L. Dorsman is funded in the framework of the STAR-project "Multilayered queueing systems" by the Netherlands Organization for Scientific Research (NWO).

IFIP WG 7.3 Performance 2013, September 24-26, Vienna, Austria.

2. MODEL AND NOTATION

We consider a system of N infinite-buffer queues, $Q_1, \ldots, Q_N$, and a single server that visits and serves the queues in cyclic order. Each queue receives gated service. Customers arrive at Q_i according to a Poisson process with rate λ_i. These customers are referred to as type-i customers. The total arrival rate is denoted by $\Lambda = \sum_{i=1}^{N} \lambda_i$. The service time of a type-i customer is a random variable B_i, with Laplace-Stieltjes transform (LST) $B_i^*(\cdot)$ and finite moments. The kth moment of the service time of an arbitrary customer is denoted by $b^{(k)} = \mathbb{E}[B^k] = \sum_{i=1}^{N} \lambda_i \mathbb{E}[B_i^k]/\Lambda$, $k = 1, 2, \ldots$. The load offered to Q_i is $\rho_i = \lambda_i \mathbb{E}[B_i]$ and the total load offered to the system is equal to $\rho = \sum_{i=1}^{N} \rho_i$. The switch-over time required by the server to proceed from Q_i to Q_{i+1} is an independent random variable S_i with finite mean. Let $S = \sum_{i=1}^{N} S_i$ denote the total switch-over time in a cycle with mean $r := \mathbb{E}[S]$. C_i denotes the cycle time at Q_i, defined as the time between two successive arrivals of the server at queue i, and let $C^*(s)$ denote its LST; it is well-known that $\mathbb{E}[C_i] = r/(1-\rho)$ for each i. The *scheduling policy* determines the *order* in which the customers are served during a visit period at a queue. We consider the following four scheduling policies: FCFS, LCFS, ROS and PS. For policy $P \in \{\text{FCFS, LCFS, ROS, PS}\}$, we denote $i \in I_P$ if Q_i receives scheduling policy P. A necessary and sufficient condition for stability of the system is $\rho < 1$.

In this paper we study HT limits, i.e., the limiting behavior as ρ approaches 1. The HT limits, denoted by $\rho \uparrow 1$, taken in this paper are such that the arrival rates are increased, while keeping both the service-time distributions and the ratios between the arrival rates fixed. For a one-dimensional continuous random variable Y we denote $(1-\rho)Y \to_d \tilde{Y}$ ($\rho \uparrow 1$) if $\lim_{\rho\uparrow 1} \Pr\{(1-\rho)Y < y\} = \Pr\{\tilde{Y} < y\}$ for all y. For each variable x that is a function of ρ, we denote its value *evaluated at* $\rho = 1$ by $\hat{x}$.

Let W_i be the waiting time of a customer at Q_i, defined as the time between the arrival of a customer and the moment at which he enters service. T_i denotes the sojourn time of an arbitrary customer at Q_i, defined as the time between the arrival of a customer and the moment at which he departs from the system. Let the LSTs of W_i and T_i be denoted by $W_i^*(s)$ and $T_i^*(s)$, respectively. It is easy to see that for $i \in I_{FCFS}, I_{LCFS}, I_{ROS}$ we have $T_i^*(s) = W_i^*(s)B_i^*(s)$.

3. ANALYSIS

Let us first consider the distribution of the cycle-time C_i, the time between successive arrivals of the server at queue i. A simple but important observation is that the distribution of C_i is independent of the local scheduling policy, assuming it is work conserving, i.e., the server works at full speed when any work is present at the queue that is being served; note that the policies FCFS, LCFS, ROS and PS under consideration are indeed work conserving.

3.1 Preliminaries

The following result gives a characterization for the limiting behavior of the cycle-time distributions (cf. [7, 8]):

Property 1 (Convergence of the cycle times). *For $i = 1, \dots, N$, $\rho \uparrow 1$,*

$$(1-\rho)C_i \to_d \tilde{\Gamma},$$

where $\tilde{\Gamma}$ has a gamma distribution with parameters

$$\alpha := r\delta\frac{b^{(1)}}{b^{(2)}}, \quad \mu := \delta\frac{b^{(1)}}{b^{(2)}}, \quad \delta := \sum_{i=1}^{N} \hat{\rho}_i(1+\hat{\rho}_i). \tag{1}$$

For FCFS service, the following result gives an expression for the LST of the waiting time W_i in terms of the LST of the cycle time C_i [6]: For $i \in I_{FCFS}$, $Re(s) > 0$, $\rho < 1$,

$$W_i^*(s) = \frac{C_i^*(\lambda_i(1-B_i^*(s)) - C_i^*(s)}{\mathbb{E}\left[C_i\right](s-\lambda_i(1-B_1^*(s))}. \tag{2}$$

Property 1 and (2) imply the following result, characterizing the HT-behavior of the waiting times (cf. [7]):

Property 2 (Asymptotic waiting times for FCFS) *For $i \in I_{FCFS}$, $\rho \uparrow 1$,*

$$(1-\rho)W_i \to_d U_i[\hat{\rho}_i; 1]\tilde{\mathbf{C}}_i, \tag{3}$$

where $U_i[\hat{\rho}_i; 1]$ is uniformly distributed on $[\hat{\rho}_i, 1]$, and where $\tilde{\mathbf{C}}_i$ has a gamma distribution with parameters $\alpha + 1$ and μ, where α and μ are defined in (1).

In the next three subsections we use Properties 1 and 2 to derive asymptotic results for the LCFS, ROS and PS scheduling policy, respectively. Details of the proofs are omitted.

3.2 Last-Come-First-Served

The LST for the waiting-time distribution for LCFS is expressed in terms of the cycle-time distributions as follows (cf. [3]): For $i \in I_{LCFS}$, $\rho < 1$, $Re(s) > 0$,

$$W_i^*(s) = \frac{1 - C_i^*(s + \lambda_i(1-B_i^*(s)))}{\mathbb{E}\left[C_i\right](s+\lambda_i(1-B_i^*(s)))}. \tag{4}$$

The following result gives an expression for the asymptotic waiting-time distribution for LCFS service in HT.

Theorem 1 (Asymptotic waiting times for LCFS) *For $i \in I_{LCFS}$, $\rho \uparrow 1$,*

$$(1-\rho)W_i \to_d U_i[0; 1+\hat{\rho}_i]\tilde{\mathbf{C}}_i, \tag{5}$$

where $U_i[0; 1+\hat{\rho}_i]$ is uniformly distributed on $[0, 1+\hat{\rho}_i]$ and $\tilde{\mathbf{C}}_i$ has a gamma distribution with parameters $\alpha + 1$ and μ, where α and μ are defined in (1).

3.3 Random Order of Service

Next we proceed to the Random Order of Service (ROS) local scheduling policy. ROS is represented by ordering marks. Each customer that arrives gets an ordering mark x, a realization from a uniform distribution on $[0, 1]$. When the server arrives at the queue, the gate closes and the customers before the gate are served in order of their marks. It is convenient to condition with respect to x. For $i \in I_{ROS}$, let $W_i(x)$ be the waiting time of a customer in queue i with ordering mark x. Denote the corresponding LST by $W_i^*(s|x)$. In [3], the LST is computed *conditionally on x*, yielding for $i \in I_{ROS}$, $\rho < 1$, $Re(s) > 0$, $x \in [0, 1]$,

$$W_i^*(s|x) = \frac{C_i^*(\kappa_i(s)) - C_i^*(s+\kappa_i(s))}{s\,\mathbb{E}[C_i]}, \tag{6}$$

with $\kappa_i(s) := \lambda_i x(1 - B_i^*(s))$. The next result gives the HT limit of the *conditional* asymptotic scaled delay.

Theorem 2 *For $i \in I_{ROS}$, $\rho \uparrow 1$, $x \in [0, 1]$,*

$$(1-\rho)W_i(x) \to_d U_i[\hat{\rho}_i x; 1+\hat{\rho}_i x]\tilde{\mathbf{C}}_i, \tag{7}$$

where $U_i[\hat{\rho}_i x; 1+\hat{\rho}_i x]$ is uniformly distributed over the interval $[\hat{\rho}_i x, 1+\hat{\rho}_i x]$ and $\tilde{\mathbf{C}}_i$ has a gamma distribution with parameters $\alpha + 1$ and μ, where α and μ are defined in (1).

The following result is useful for unconditioning the result presented in Theorem 2 with respect to the value of x.

Lemma 1 *Let x be a realization of the random variable X having cumulative distribution function $F_X(t) := Pr\{X \le t\}$ and density $f_X(t)$. Assume that the conditional random variable $T|x$ is uniformly distributed on $[a(x), a(x)+1]$ and suppose that $a(0) = 0$, $a(\infty) = \hat{\rho}_i$ and $a(x)$ is increasing in x. Let $a^{-1}(\cdot)$ be the inverse of $a(\cdot)$. Then, the unconditional distribution of $T|x$, denoted by T, has probability density function*

$$f_T(y) = \begin{cases} F_X(a^{-1}(y)) & y \in [0, \hat{\rho}_i) \\ 1 & y \in [\hat{\rho}_i, 1] \\ 1 - F_X(a^{-1}(y-1)) & y \in (1, 1+\hat{\rho}_i]. \end{cases}$$

Combining Theorem 2 and Lemma 1, we obtain the following result.

Theorem 3 (Asymptotic waiting times for ROS) *For $i \in I_{ROS}$, $\rho \uparrow 1$,*

$$(1-\rho)W_i \to_d U_i\tilde{\mathbf{C}}_i$$

where U_i has a trapezoidal distribution with density function

$$f_{U_i}(y) = \begin{cases} y/\hat{\rho}_i & y \in [0, \hat{\rho}_i) \\ 1 & y \in [\hat{\rho}_i, 1] \\ (1+\hat{\rho}_i - y)/\hat{\rho}_i & y \in (1, 1+\hat{\rho}_i], \end{cases}$$

and where $\tilde{\mathbf{C}}_i$ has a gamma distribution with parameters $\alpha + 1$ and μ, where α and μ are defined in (1).

3.4 Processor sharing

When x is the amount of work that a tagged customer brings into the system, we define the waiting time $W_i(x)$ as the sojourn time *minus* the service time x of the tagged customer. The following result expresses the LST of $W_i(x)$ in terms of the cycle times (cf. [3]): For $i \in I_{PS}$, $\rho < 1$, $Re(s) > 0$, $x > 0$,

$$W_i^*(s|x) = \frac{C_i^*(\zeta_i(s,x)) - C_i^*(s + \zeta_i(s,x))}{s\,\mathbb{E}[C]}, \tag{8}$$

where $\zeta_i(s,x) := \lambda_i(1 - \varphi_i(s,x))$, with $\varphi_i(s,x)$ the LST of $\min(B_i, x)$.

For PS, the most interesting performance metrics are the conditional and unconditional sojourn times, denoted $T_i(x)$ and T_i. Note that under HT scalings, the limiting distributions $W_i(x)$ and $T_i(x)$ are the same for all x, and so are the unconditional waiting and sojourn times W_i and T_i.

Theorem 4 (Conditional sojourn time for PS) *For $i \in I_{PS}$, $\rho \uparrow 1$, $x > 0$, we have*

$$(1-\rho)T_i(x) \to_d U_i(x)\tilde{\mathbf{C}}_i, \tag{9}$$

where $U_i(x)$ is uniformly distributed over the interval $[a_i(x), 1 + a_i(x)]$, with $a_i(x) := \hat{\lambda}_i\,\mathbb{E}[\min(B_i, x)]$ and $\tilde{\mathbf{C}}_i$ has a gamma distribution with parameters $\alpha + 1$ and μ, where α and μ are defined in (1).

Then, using Lemma 1, we obtain the following result.

Theorem 5 (Unconditional sojourn time for PS) *For $i \in I_{PS}$, $\rho \uparrow 1$, we have*

$$(1-\rho)T_i \to_d U_i\tilde{\mathbf{C}}_i,$$

where U_i has a generalized trapezoidal distribution with density function

$$f_{U_i}(y) = \begin{cases} F_{B_i}(a_i^{-1}(y)) & y \in [\hat{\lambda}_i m, \hat{\rho}_i) \\ 1 & y \in [\hat{\rho}_i, 1 + \hat{\lambda}_i m] \\ 1 - F_{B_i}(a_i^{-1}(y-1)) & y \in (1 + \hat{\lambda}_i m, 1 + \hat{\rho}_i], \end{cases}$$

where $a_i(x) = \hat{\lambda}_i\,\mathbb{E}[\min(B_i, x)]$ and m is the lowest possible value of B_i.

To illustrate this, suppose B_i is a uniformly distributed random variable on the interval $[a_i, b_i]$. Then it follows after some calculus that $a_i^{-1}(y) = b_i - \sqrt{(1 - y/\hat{\rho}_i)(b_i^2 - a_i^2)}$, and that the density function of U_i is

$$f_{U_i}(y) = \begin{cases} 1 - \frac{\sqrt{(1-y/\hat{\rho}_i)(b_i^2 - a_i^2)}}{b_i - a_i} & y \in [\hat{\lambda}_i a_i, \hat{\rho}_i) \\ 1 & y \in [\hat{\rho}_i, 1 + \hat{\lambda}_i a_i] \\ \frac{\sqrt{(1-(y-1)/\hat{\rho}_i)(b_i^2 - a_i^2)}}{b_i - a_i} & y \in (1 + \hat{\lambda}_i a_i, \hat{\rho}_i + 1]. \end{cases}$$

Figure 1 illustrates $f_{U_i}(y)$ for the case $a_i = 0.1340$, $b_i = 1.8660$, $\hat{\lambda}_i = 0.077$ and thus $\hat{\rho}_i = 0.077$, and the case $a_i = 0.2679$, $b_i = 3.7321$, $\hat{\lambda}_i = 0.2315$ and thus $\hat{\rho}_i = 0.4615$.

Some remarks are due. First, an interesting observation is that in HT the influence of the scheduling policy only manifests itself in the distribution of U_i, while the parameters of the gamma distribution do not depend on the scheduling policy. This observation is in line with the fact that the cycle-time and queue-length distributions at polling instants are stochastically identical among different local scheduling policies, as observed in Section 3. Second, there is a remarkable difference in sensitivity with respect to the service-time distributions. On the one hand, for policies like FCFS, LCFS and ROS (see Property 2 and Theorems

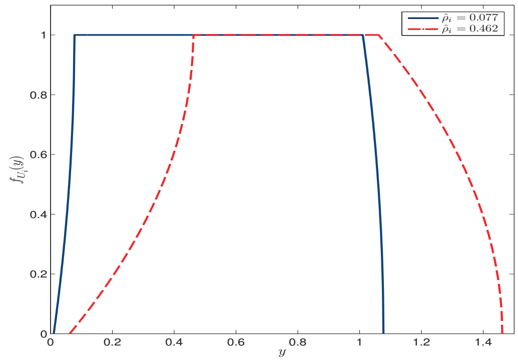

Figure 1: Density function of U_i for uniform distribution in case of PS scheduling

1 and 3) the asymptotic waiting-time distribution is of the form $U_i[a_i; b_i]\Gamma$, where the a_i and b_i depend on the service-time distribution at Q_i only through its *mean*, whereas for PS the asymptotic sojourn-time distribution does depend on the complete service-time distributions (Theorem 5). Third, the HT results form an excellent basis for the development of waiting-time and sojourn time approximations for stable systems: For $\rho < 1$, $i = 1, \ldots, N$, $x > 0$,

$$Pr\{T_i > x\} \approx Pr\{U_i\Gamma > x(1-\rho)\}, \tag{10}$$

where U_i and Γ are characterized in Theorems 1 to 5 above. Refinement of the approximations can also be obtained along the lines of [4]. Initial simulation experiments show promising results. Lastly, the HT asymptotics are easily extendable to renewal arrivals, following the lines described in [5].

4. REFERENCES

[1] Boon, M.A.A., Adan, I.J.B.F. and Boxma, O.J. (2010). A polling model with multiple priority levels. *Performance Evaluation* **67**, 468-484.

[2] Boon, M.A.A., Van der Mei, R.D. and Winands, E.M.M. (2011). Applications of polling systems. *Surveys in Operations Research and Management Science* **16**, 67-82.

[3] Boxma, O.J., Bruin, J. and Fralix, B. (2009). Sojourn times in polling systems with various service disciplines. *Performance Evaluation* **66**, 621–639.

[4] Dorsman, J.L., Van der Mei, R.D and Winands, E.M.M. (2011). A new method for deriving waiting-time approximations in polling systems with renewal arrivals. *Stochastic Models* **27**, 318–332.

[5] Olsen, T.L. and Van der Mei, R.D. (2005). Periodic polling systems in heavy-traffic: renewal arrivals. *Operations Research Letters* **33**, 17-25.

[6] Takagi, H. (1986). *Analysis of Polling Systems* (MIT Press, Cambridge, MA).

[7] Van der Mei, R.D. (1999). Distribution of the delay in polling systems in heavy traffic. *Performance Evaluation* **31**, 163-182.

[8] Van der Mei, R.D. (2007). Towards a unifying theory on branching-type polling systems in heavy traffic. *Queueing Systems* **57**, 29–46.

[9] Wierman, A., Winands, E.M.M. and Boxma, O.J. (2007). Scheduling in polling systems. *Performance Evaluation* **64**, 1009-1028.