Appendix to
Creating Work Breaks From Available Idleness

Xu Sun and Ward Whitt

Department of Industrial Engineering and Operations Research, Columbia University, New York, NY, 10027; {xs2235,ww2040}@columbia.edu

September 9, 2016

Abstract

We show how dynamic priority (DP) rules for assigning available service representatives to arriving customers in customer contact centers can be used to create effective work breaks for the service representatives from naturally available idleness, assuming that the service system is staffed adequately to provide non-negligible idleness. We start by establishing many-server heavy-traffic limits to develop useful approximations for the distributions of server idle times with the customary longest-idle-server-first (LISF) rule and a random-routing (RR) alternative. We show that the pattern of idleness with these rules is totally different but neither produces effective work breaks. We then develop three DP rules and conduct simulation experiments to show that the new DP rules can indeed create effective work breaks from the available idleness. The first DP rule yields unannounced breaks, while the other more refined rules yield announced breaks.

Keywords: customer contact centers, call centers; work breaks; server-assignment rules; many-server queues.

1 Overview

This is an appendix to the main paper, Sun and Whitt [2016]. In §2 we summarize the notation used in the paper, indicating where it is defined and first used. In §3 we elaborate on fitting a truncated normal distribution, which is used in §3 of the main paper. In §4 we elaborate on the simulation methodology. Finally, in §5 we present additional simulation results.

2 List of Abbreviations and Symbols, Plus Location in Text

We now give an overview of the notation in the main paper, indicating where it is first introduced and defined.

LISF	Longest-Idle-Server-First (server assignment rule), §1
RR	Randomized-Routing (server assignment rule), §1
ρ	traffic intensity, §2.1
U_i	interarrival times, §2.1
λ	arrival rate, §2.1
S_i	service times, §2.1
μ	service rate of individual servers, §2.1
s	number of servers, §2.1
$N(t)$	number of customers in the system at time t, §2.1
$B(t)$	number of busy servers at time t, §2.1
$I(t)$	number of idle servers at time t, §2.1
N	steady-state number of customers in the system, §2.1
B	steady-state number of busy servers, §2.1
I	steady-state number of idle servers, §2.1
D	duration of a work break, §2.2
T	interval between successive work breaks, §2.2
β	long-run proportion of time server is on break, §2.2
β_t	long-run proportion of idle time server is on break, §2.2
θ	target duration of a work break, §2.2

Table 1: The notation used in §1 and §2 of the main paper.

Φ	cumulative distribution function (cdf) of a standard normal random variable, §3.1
ϕ	probability density function (pdf) of a standard normal random variable, §3.1
α	steady-state delay probability for the standard $M/M/s$ queueing model, §3.1
ξ	quality-of-service parameter for the square-root staffing rule, §3.1
$N(m, v)$	normal random variable with mean m and variance v, §3.1
V	steady-state idle time, §3.2
$C(t)$	cumulative idleness in an interval $[0, t]$, §3.3
V_c	completed idle time in cycle in progress, §3.3
$M(t)$	completed idle time in cycle in progress, §3.3
$A \equiv A(I)$	random number of arrivals required for RR assignment, §3.4

Table 2: The notation used in §3 of the main paper.

DP	dynamic priority (rule), §4
$DP1(\theta)$	first DP rule (1 for 1 control parameter), §4.1
β_n	Proportion of idle times that are work breaks, §4.1
$DP2(\theta, \tau)$	second DP rule (2 for 2 control parameters), §4.2
p_B	proportion of breaks that are announced, §4.2
p_D	steady-state delay probability, §4.2
Q	steady-state queue-length, §4.2
τ	threshold control for interval between successive work breaks, §4.2
$S_b(t)$	number of servers on break, §4.2
$I_d(t)$	number of idle servers, allowing negative values, §4.2
$G(t)$	gap, §4.2
γ	long-run average gap, §4.2
$DP3(\theta, \tau, \eta)$	third DP rule (3 for 3 control parameters), §4.3
η	Maximum possible number of servers on break at the same time, §4.3
$C(p_B, p_D)$	cost function to choose optimum η, §4.3

Table 3: The notation used in §4 of the main paper.

3 Fitting a Truncated Normal Distribution

In §3.1 of the main paper we observed that we could fit a truncated normal distribution to the steady-state number of idle servers, I. In particular, in equation (3.3) we observed that $I \approx (s - N(m,v))^+$. However, it remains to determine the parameters m and v consistent with the known exact values of $P(I = 0)$, $E[I]$ and $E[I^2]$. We elaborate here.

Since N has a Poisson distribution in the $M/GI/\infty$ model, we approximate the conditional distribution of I given $I > 0$ by a truncated normal distribution. Thus, for the $M/GI/s$ model, we approximate by

$$B \approx N(m,v) \wedge s, \quad \text{and} \quad I \approx (s - N(m,v))^+, \tag{3.1}$$

where the mean m and variance v can be obtained by solving the equations

$$\begin{aligned} P(I=0) &\approx P((N(m,v) \geq s) \approx \alpha \\ E[(I)^k] &\approx (1-\alpha)E[(s-N(m,v))^k | N(m,v) < s] \quad \text{for} \quad k = 1,2, \end{aligned} \tag{3.2}$$

where

$$\begin{aligned} E[s - N(m,v)|N(m,v) < s] &= s - E[N(m,v)|N(m,v) < s] \qquad (3.3) \\ E[(s-N(m,v))^2|N(m,v) < s] &= s^2 - 2sE[N(m,v)|N(m,v) < s] + E[N(m,v)^2|N(m,v) < s] \end{aligned}$$

and

$$\begin{aligned} E[N(m,v)|N(m,v) < s] &= \sqrt{v}E[N(0,1)|N(0,1) < (s-m)/\sqrt{v}] + m \\ E[N(m,v)^2|N(m,v) < s] &= vE[N(0,1)^2|N(0,1) < (s-m)/\sqrt{v}] \\ &\quad +2m\sqrt{v}E[N(0,1)|N(0,1) < (s-m)/\sqrt{v}] + m^2, \end{aligned} \tag{3.4}$$

with explicit formulas given, e.g., in Proposition 18.3 of Browne and Whitt [1995]. After solving for m and v, we obtain the tail probability approximation

$$P(I > x) \approx P(N(m,v) < s - x) = \Phi((s-x-m)/\sqrt{v}) = \Phi^c((x-(s-m))/\sqrt{v}), \quad x \geq 0, \tag{3.5}$$

where $\Phi^c(x) \equiv 1 - \Phi(x)$.

It is natural to calculate the parameter pairs (m,v) by doing an exhaustive search in the neighborhood of the overall mean and variance $(E[I], Var(I)) = (1-\rho)s, \rho s(1-\alpha))$. Of course, the approximation is asymptotically correct if s is very large for $\rho < 1$. Then we may use the associated QD MSHT approximation in which $\alpha \approx 0$.

With exhaustive search in mind, we observe that we can apply (3.2) and (3.3) to rewrite the two moment equations as

$$\begin{aligned} E[N(m,v)|N(m,v)<s] &= s(\rho-\alpha)/(1-\alpha) \\ E[N(m,v)^2|N(m,v)<s] &= \rho s+\rho^2 s^2-[\alpha(1-\rho^2)s^2]/(1-\alpha). \end{aligned} \tag{3.6}$$

Note that the formulas in (3.6) are correct for $\alpha = 0$. To find the appropriate m and v, we can calculate the righthand sides and then compute for (m, v) near $(E[B], Var(B)) = (\rho s, \rho s(1-\alpha))$ and find where the two equations in (3.6) are satisfied.

We now illustrate this algorithm for the base case in §2.3 of the main paper. In particular, we consider the $M/M/s$ model with LISF, $\lambda = 90$, $\mu = 1$ and $s = 100$.

4 Simulation Methodology

The simulation results for the idle time distribution in the $M/M/s$ model with the LISF and RR routing rules and model parameters $s = 100$, $\mu = 1$, $\lambda = 90$ and $\rho = 0.9$ are reported in §3 of the main paper; e.g., see the histgrams for LISF and RR in Figure 1 of the main paper. These were based on 100 i.i.d. replications of the $M/M/100$ system observed over a time interval of length $10,000$ after a warmup period of length 100 to allow the system that started empty to approach steady state. Idle time data were collected from all 100 servers. We used all observed idle times that started before time $9,980$, allowing a final interval of length 20 to avoid terminal end effects.

As usual, the first two moments $m_k \equiv E[V^k]$, $k = 1, 2$, are estimated as sample averages within each replication. Within each replication, the estimated variance is $\sigma^2 = m_2 - m_1^2$. Then the overall estimates $\bar{m}_1$ and $\bar{\sigma}^2$ are estimated as the sample averages of the 100 values. Moreover, 95% confidence intervals are estimated for the overall averages in the usual way based on a sample of 100 i.i.d. observations. We examine the sample distribution of the 100 values to verify that the approach is reasonable.

We now do a rough analysis to estimate the statistical precision for the estimate of the mean $E[V]$. First, because the mean service time was 1 and the mean idle time was $(1-\rho)/\rho = 0.1111$, the mean service cycle was 1.1111. Hence, each server has about $14,980/1.111 \approx 13,400$ or 1.34×10^4 idle times per replication. For all 100 servers, that produces about 1.34×10^6 idle times per replication. Given that the mean and variance are $E[V] = 0.111$ and $Var(V) = 0.01$ by Theorem 3.2 and Example 3.2, acting as if the idle times are mutually independent (which we know they are not), the sample mean $\bar{m}_1$ would have sample variance about $Var(\bar{m}_1) \approx (0.01)/1.34 \times 10^6 = 75 \times 10^{-8}$ and the standard error would be $\bar{s} \equiv \sqrt{Var(\bar{m}_1)} \approx 8.6 \times 10^{-4} = 0.00086$. Allowing for positive dependence, we might conservatively anticipate a sample variance of about $4-5$ times greater or $400 \times 10^{-8} = 4 \times 10^{-6}$, from which we get the estimated standard error $\bar{s} = 2 \times 10^{-3} = 0.002$ within each replication.

Finally, 100 replications leads to $Var(\bar{m}_1) \approx 4 \times 10^{-8}$ with standard error $\bar{s} \equiv \sqrt{Var(\bar{m}_1)} \approx 2 \times 10^{-4} = 0.0002$. Thus, we would anticipate 95% confidence intervals for $E[V] = 0.11111$ based on $n = 100$ i.i.d. replications of about

$$E[V] \pm \frac{1.96\bar{s}}{\sqrt{n}} = 0.11111 \pm \frac{1.96(0.0002)}{10} \approx 0.11111 \pm 0.00004 \quad \text{or} \quad [0.11107, 0.11115].$$

The way to estimate the variance $Var(V)$ is less straightforward. As indicated above, we estimate the variance within each replication as $\bar{\sigma}^2 = \bar{m}_2 - \bar{m}_1^2$. We thus obtain 100 estimates of the variance, one of each replication. We then estimate the overall variance as the sample average of these, and estimate the CI assuming that these are Gaussian distributed with unknown variance. We examine the distribution of these sample variance estimates.

Finally, we estimate the distribution of V by constructing a histogram based on all the data.

5 More Simulation Results

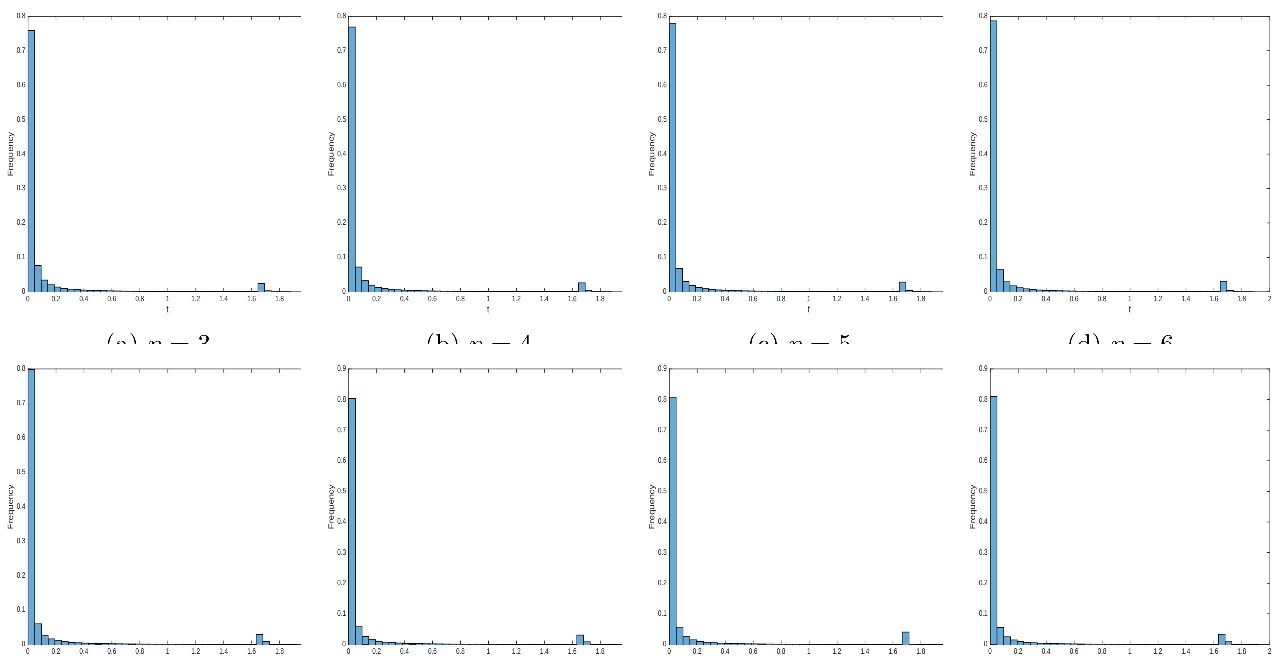

(a) $\eta = 3$ (b) $\eta = 4$ (c) $\eta = 5$ (d) $\eta = 6$

(e) $\eta = 7$ (f) $\eta = 8$ (g) $\eta = 9$ (h) $\eta = 10$

Figure 1: Histogram of the idle times with rule DP3($\theta = 5/3, \tau = 20, \eta$) estimated from simulation

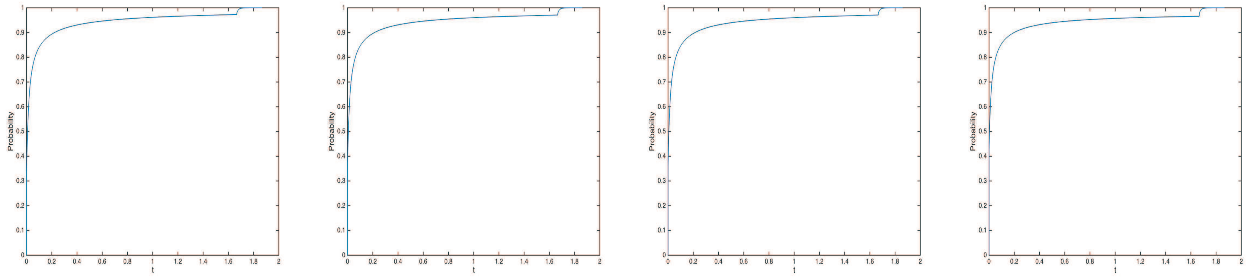

(a) $\eta = 3$ (b) $\eta = 4$ (c) $\eta = 5$ (d) $\eta = 6$

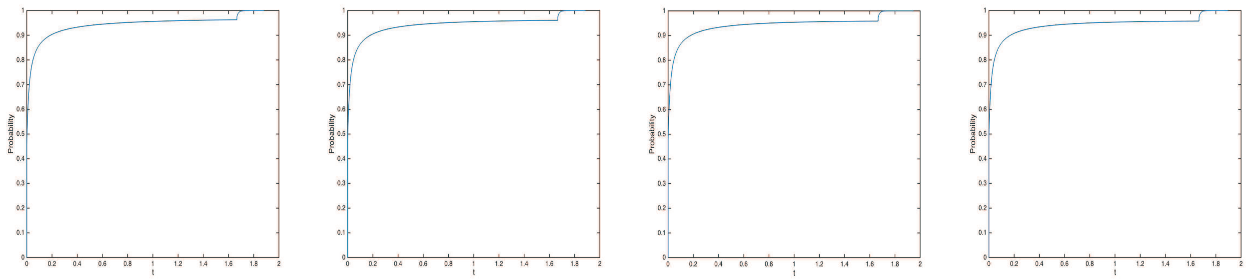

(e) $\eta = 7$ (f) $\eta = 8$ (g) $\eta = 9$ (h) $\eta = 10$

Figure 2: Empirical CDF of the idle-time with rule DP3($\theta = 5/3, \tau = 20, \eta$) estimated from simulation

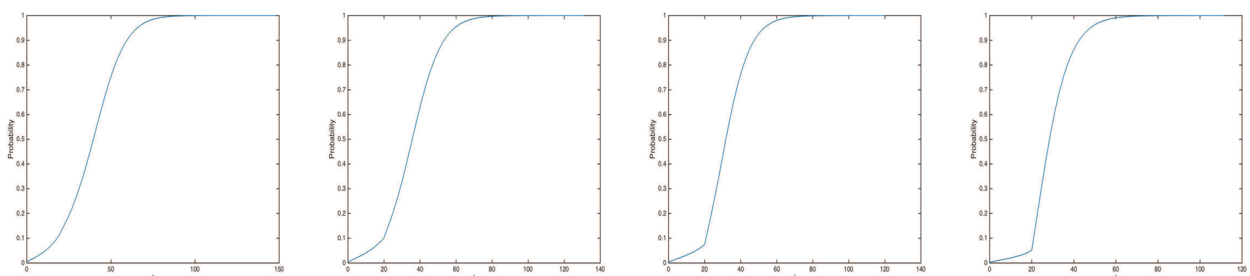

(a) $\eta = 3$ (b) $\eta = 4$ (c) $\eta = 5$ (d) $\eta = 6$

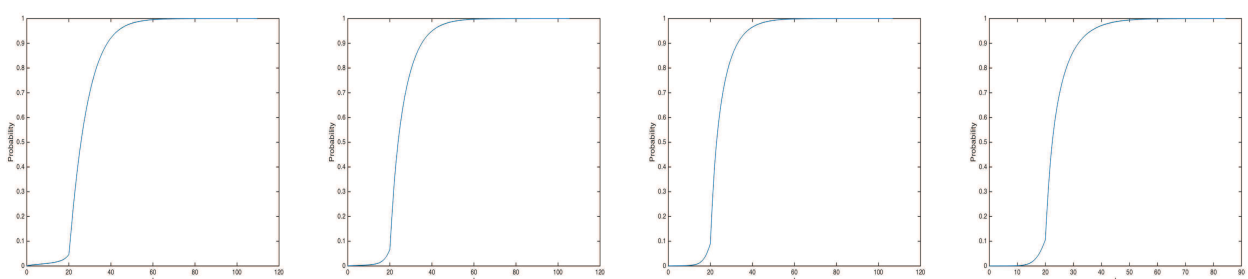

(e) $\eta = 7$ (f) $\eta = 8$ (g) $\eta = 9$ (h) $\eta = 10$

Figure 3: Empirical CDF of the inter-break-time distribution with rule DP3($\theta = 5/3, \tau = 20, \eta$) estimated from simulation

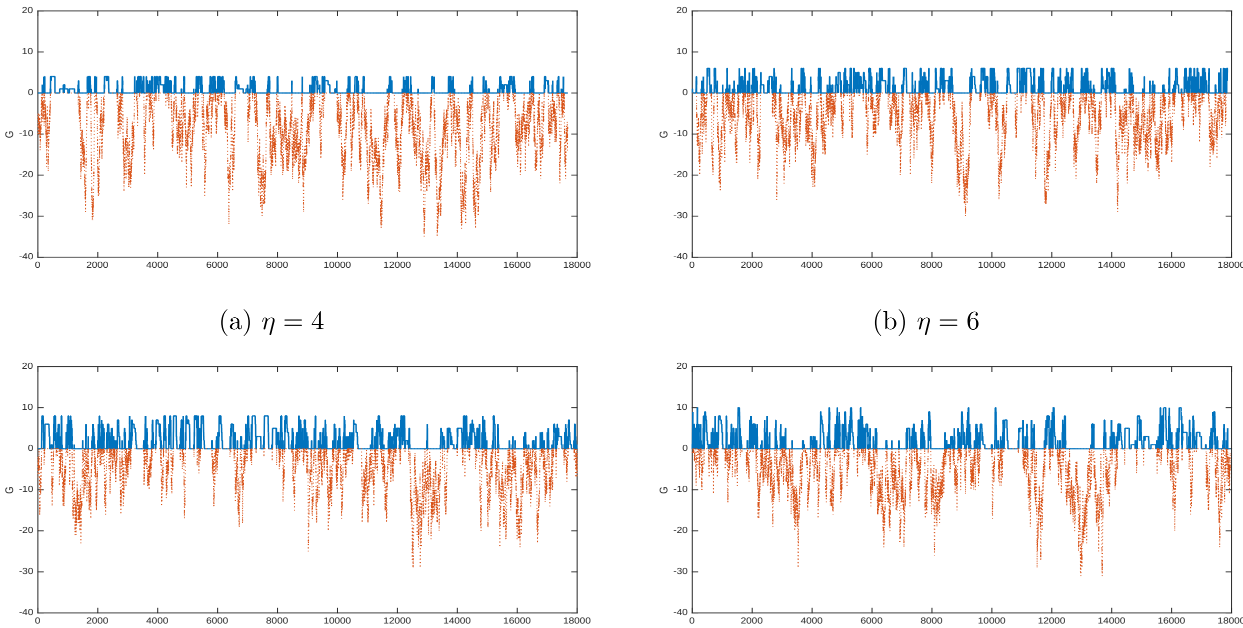

(a) $\eta = 4$

(b) $\eta = 6$

(c) $\eta = 8$

(d) $\eta = 10$

Figure 4: Sample paths of $G(t) \equiv S_b(t) - I(t)$, with rule DP3 as a function of η for $\theta = 5/3$ and $\tau = 20$

References

d

S. Browne and W. Whitt. Piecewise-linear diffusion processes. In J. Dshalalow, editor, *Advances in Queueing*, pages 463–480. CRC Press, Boca Raton, FL, 1995.

X. Sun and W. Whitt. Creating work breaks from available idleness. Columbia University, http://www.columbia.edu/∼ww2040/allpapers.html, 2016.