

Stabilizing Customer Abandonment in Many-Server Queues with Time-Varying Arrivals

Yunan Liu

Department of Industrial Engineering, Room 446, 400 Daniels Hall, North Carolina State University, Raleigh, NC 27695,
yliu48@ncsu.edu

Ward Whitt

Department of Industrial Engineering and Operations Research, Columbia University, New York, NY 10027-6699,
ww2040@columbia.edu

An algorithm is developed to determine time-dependent staffing levels to stabilize the time-dependent abandonment probabilities and expected delays at positive target values in the $M_t/GI/s_t + GI$ many-server queueing model, which has a nonhomogeneous Poisson arrival process (the M_t), general service times (the first GI) and allows customer abandonment according to a general patience distribution (the $+GI$). New offered-load and modified-offered-load approximations involving infinite-server models are developed for that purpose. Simulations show that the approximations are effective. A many-server heavy-traffic limit in the efficiency-driven regime shows that: (i) the proposed approximations achieve the goal asymptotically as the scale increases, and (ii) it is not possible to simultaneously stabilize the mean queue length in the same asymptotic regime.

Key words: staffing; capacity planning; many-server queues; queues with time-varying arrivals; queues with abandonment; infinite-server queues, offered-load approximations; service systems;

History: Submitted on June 3, 2009; revisions (regarded as new submission) on December 1, 2011, and March 28, 2012; accepted July 12, 2012

1. Introduction

In this paper we develop new staffing algorithms for many-server queueing systems with time-varying arrivals, focusing on the challenging case in which service times are relatively long and there is customer abandonment from queue; see Green et al. (2007) for background. Specifically, we develop formula-based algorithms to stabilize the time-dependent abandonment probability and the expected (potential) delay (the delay before starting service for a customer arriving at time t with infinite patience) at any (necessarily related) fixed targets, across a wide range of possible targets, in the $M_t/GI/s_t + GI$ queueing model. This model has a nonhomogeneous Poisson arrival

process (the M_t), a large number s_t of homogeneous servers working in parallel (a function of time t , which is to be determined), independent and identically distributed (i.i.d.) service times with a general distribution (the first GI), unlimited waiting space, the first-come first-served (FCFS) service discipline, and customer abandonment from the queue of customers waiting to start service, with i.i.d. times to abandon having a general distribution (the second GI).

Our results extend Feldman et al. (2008), which introduced a simulation-based *iterative staffing algorithm* (ISA) to stabilize the time-dependent delay probability (the probability that an arrival must wait in queue before starting service). As illustrated by Figure 3 in that paper, the ISA was shown to be remarkably effective at stabilizing the delay probability in the fully Markovian $M_t/M/s_t + M$ model, for all possible quality-of-service (QoS) levels, provided that the arrival rate is suitably large, so that a large number of servers (e.g., 100) is actually required. The target delay probability was allowed to range from 0.1 (high QoS) to 0.9 (low QoS). Indeed, as illustrated by Figures 5 and 6 of that paper, with ISA staffing the delay probability becomes essentially the same as in a corresponding stationary model with constant arrival rate, and thus also showing that the analytically-based modified-offered-load (MOL) approximation (reviewed here in §2) succeeds in stabilizing delays for this $M_t/M/s_t + M$ model. Thus, from the perspective of the delay probability, the effect of the time-varying arrival rate can be eliminated by applying the ISA (or the MOL) to choose the staffing level appropriately.

Since the ISA is based on simulation, it can easily be applied to associated non-Markovian models and even to much more general models. Indeed, experiments showed that the ISA is also effective for stabilizing the delay probability in the more general $M_t/GI/s_t + GI$ model considered here. The ISA also has the advantage of providing automatic verification: Since ISA is based on simulation, we can confirm that ISA achieves its goal in the final simulation results.

Even though the ISA can stabilize the delay probability, it is unable to eliminate the effect of the time-varying arrival rate entirely. When ISA is applied to stabilize the delay probability, the ISA also stabilizes other performance measures to some extent, but as illustrated by Figure 4 of Feldman et al. (2008) and Figure 6 of the accompanying e-companion, significant fluctuations are

seen in the time-varying abandonment probability and average delay with a sinusoidal arrival-rate function when the target QoS is relatively low (the target delay probability is high). An open problem posed in §8 of Feldman et al. (2008) was to develop a way to stabilize the abandonment probability across the full range of target abandonment probabilities.

We address that open problem in this paper. Moreover, we do so with a formula-based algorithm, instead of simulation. Our formula-based algorithm applies to the general $M_t/GI/s_t + GI$ model. Since all performance measures tend to be stabilized together at customary targets with higher QoS, we succeed in obtaining an effective formula-based algorithm for staffing to meet all the standard performance measures at customary targets with higher QoS.

There are two important steps here. The first is to introduce an entirely new *offered-load* (OL) framework, involving two *infinite-server* (IS) queues in series, which we call the *delayed-infinite-server* (DIS) approximating model. (We give background on OL approximations in §2.) The first IS queue represents the waiting room (queue), while the second represents the service facility. The mean number of busy servers in the second IS queue represents the new OL. When the targeted QoS is relatively low (the abandonment probability is high), the OL itself directly yields a good staffing function, called DIS staffing. Moreover, when we use DIS staffing, the entire DIS model serves as a useful approximation for the original $M_t/GI/s_t + GI$ model, so that we obtain useful formulas approximating other performance measures, such as the time-varying mean queue length; see Theorem 1 below. We thus see that the mean queue length cannot be stabilized at the same time as the abandonment probability under higher abandonment-probability targets, and we can quantify the fluctuations in the mean queue length.

We substantiate the good performance of the DIS approximation for heavily loaded systems by establishing a heavy-traffic limit theorem. In particular, we apply our recent heavy-traffic fluid limits for many-server models with time-varying arrival rate and staffing in Liu and Whitt (2012a,b,c) to show that the DIS approximation is asymptotically correct in the overloaded or efficiency-driven (ED) many-server heavy-traffic regime; see Garnett et al. (2002) for background. (Related limits appear in Kang and Ramanan (2010), Kaspi and Ramanan (2011), Mandelbaum et al. (1998),

Whitt (2006), but Mandelbaum et al. (1998) is restricted to the Markovian special case, while the others assume fixed staffing.) As a corollary, we deduce that it is not possible for any algorithm to simultaneously stabilize all performance measures across the full range of target values in the many-server heavy-traffic regime.

Unfortunately, however, the simple DIS approximation does not perform well in the common case in which the target QoS is high (the abandonment probability target is low), which tends to take the system out of the ED regime. In the second step we treat that case by introducing a new *modified-offered-load* (MOL) approximation, which uses the new DIS offered load; we call this the DIS-MOL approximation. We will show that the DIS-MOL approximation is not too dismal!

Paralleling previous MOL approximations in Jennings et al. (1996) and Feldman et al. (2008) (also see Jagerman (1975) and Massey and Whitt (1994, 1997)), our MOL approximation uses steady-state performance formulas from the associated stationary $M/GI/s + GI$ model in a time-varying manner, using an arrival rate determined by the DIS OL. This second step is not immediate either, because the $M/GI/s + GI$ model tends to be intractable. In this step, we apply the approximation for all steady-state performance measures in this model by an associated $M/M/s + M(n)$ model from Whitt (2005), which uses an exponential service time with the same mean and a state-dependent Markovian abandonment process. We conduct simulations to show that this new formula-based DIS-MOL staffing algorithm stabilizes the abandonment probability and the expected waiting time effectively across a wide range of targets. As indicated above, we actually achieve more: Since all performance measures tend to be stabilized together when the target QoS is high, in that case we actually achieve a formula-based staffing algorithm for both the delay probability and the abandonment probability, plus several other performance measures as well.

Here is how the rest of this paper is organized: We start in §2 by reviewing the analytical OL and MOL approaches to the staffing problem. Next in §3 we develop the DIS model and give explicit expressions for all the key performance measures. In §4 we show that the DIS approximation achieves its goal of stabilizing abandonment probabilities and expected delays as the scale increases.

There we also prove that it is impossible to simultaneously asymptotically stabilize all performance functions. In §5 we develop the new MOL approximation for normally loaded systems.

In §6 we perform simulation experiments to validate the approximations, considering the Markovian $M_t/M/s_t + M$ examples with sinusoidal arrival-rate function. We also consider corresponding many-server service systems with non-exponential service times and abandonment times in the e-companion and a longer version available on the authors' web pages. Our simulations add real system constraints including the discretization issues and specified staffing intervals.

In §7 we present extra details for the asymptotic results in §4. Finally, in §8 we draw conclusions. We present additional material in the e-companion and a longer version available on the authors' web pages; the contents are specified in §EC.1.

2. Background on Offered Load Approximations

The general idea of an OL approximation is to initially assume that there are as many resources (here servers) as needed and then see how many are actually used. After we have determined how many servers would be needed if there were no constraint on their availability, we staff to provide the number of servers needed. Since the model is stochastic, the time-dependent number of servers used itself is random, so we must use some deterministic time-dependent partial characterization of this time-dependent distribution, such as the time-dependent mean or that time-dependent mean plus some multiple of the associated time-dependent standard deviation, as in the staffing algorithm proposed by Jennings et al. (1996).

As a consequence, the basic OL approximation is an IS approximation: The time-dependent number in the $M_t/GI/s_t + GI$ system is approximated by the time-dependent number of busy servers in the corresponding $M_t/GI/\infty$ model, having the same arrival processes and service times. This step is effective because the IS model is remarkably easy to analyze. The number of busy servers in the $M_t/GI/\infty$ model at time t has a Poisson distribution with the mean

$$m_0(t) \equiv \int_{-\infty}^t \lambda(s) P(S > t - s) ds, \quad (1)$$

where \equiv denotes “equality by definition” and S is a generic service time; see Eick et al. (1993a,b). (The subscript 0 will be explained later.)

The IS approximation leads to the classical *square-root-staffing* (SRS) formula

$$s_\gamma(t) = m_0(t) + \beta_\gamma \sqrt{m_0(t)}, \quad (2)$$

where γ is the target level of performance, $s_\gamma(t)$ is the required staffing level at time t for that target, β_γ is an associated QoS parameter, and $m_0(t)$ is the mean number of busy servers in the IS model. This mean $m_0(t)$ serves as the appropriate notion of *offered load* (OL) at time t , which is independent of both γ and the abandonment-time cdf. The SRS in (2) is based on a normal approximation for the exact Poisson distribution. If γ is the target delay probability, then β_γ is chosen to satisfy $P(N(0,1) > \beta_\gamma) = \gamma$.

The MOL approximation is a refinement of the OL approximation above, which is needed because in reality there are not infinitely many servers, so that the number in system is not actually so well approximated by a normal distribution. For the $M_t/GI/s_t + GI$ model, the MOL approximation for the performance at time t is the steady state performance in the associated stationary $M/GI/s + GI$ model with constant arrival rate

$$\lambda_0^{MOL}(t) \equiv \frac{m_0(t)}{E[S]}, \quad (3)$$

where $m_0(t)$ is the offered load and S is a random service time. The MOL staffing at time t is the smallest staffing level $s(t)$ such that the performance target is met. Since performance is difficult to analyze in the general $M/GI/s + GI$ model, here we propose using the approximation in Whitt (2005). For further discussion, see §5.

An alternative to using the exact $M/GI/s + GI$ steady-state formula for the delay probability in the MOL approximation is to use a heavy-traffic approximation for it. That leads to a direct application of the SRS, but with a new formula for the QoS parameter β_γ . This approach was used with the delay probability target in §4 of Jennings et al. (1996)) for the $M/M/s$ model and in

Feldman et al. (2008) for the more general $M/M/s + M$ model. For the $M/M/s + M$ model, the Garnett function approximating the delay probability obtained from Garnett et al. (2002) is

$$P(\text{Delay}) \approx \omega(-\beta, \sqrt{\mu/\theta}) \equiv G_1(\beta), \quad (4)$$

where $\omega(x, y) \equiv [1 + h(-xy)/(yh(x))]^{-1}$, $h(x) \equiv \phi(x)/\Phi(x)$, $\phi(x) \equiv (1/\sqrt{2\pi})e^{-x^2/2}$, $\Phi(x) \equiv \int_{-\infty}^x \phi(y)dy$. To stabilize the delay probability at target γ , it suffices to obtain a β_γ by inverting the Garnett function, letting $P(\text{Delay}) \equiv \gamma$.

3. The Delayed-Infinite-Server (DIS) Approximation

Here we use IS models in a new way. Instead of directly replacing the $M_t/GI/s_t + GI$ model by its $M_t/GI/\infty$ counterpart, which ignores the customer abandonment, we represent our model as two IS facilities in series: first the waiting room (or the queue), and then the service facility, as depicted in Figure 1.

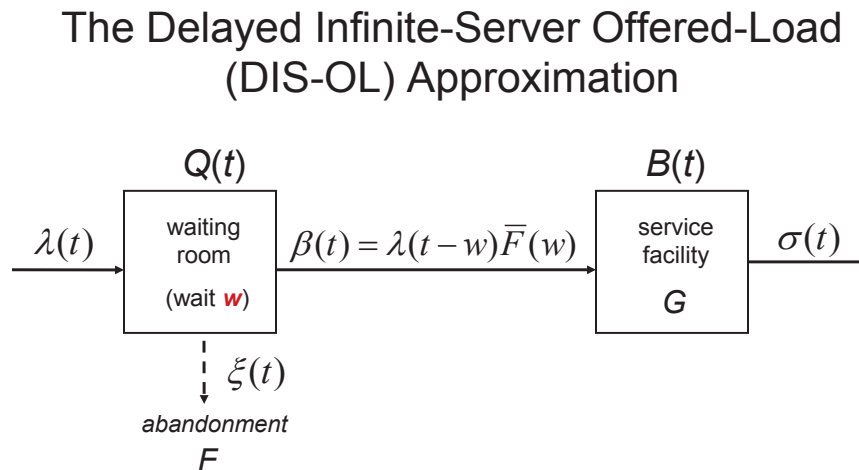


Figure 1 The delayed infinite-server (DIS) approximation for the $M_t/GI/s_t + GI$ queueing model. The contents $Q(t)$ and $B(t)$ are independent Poisson random variables for each t ; the three flows are Poisson processes.

We start by assuming that our goal is to have every arrival that does not elect to abandon wait exactly time w before entering service. To achieve that goal in our approximation, we require that all external arrivals enter the waiting room that has infinite capacity and spend the fixed time w there before they move on to the service facility. (That is done in the approximation, not in the actual system.) While in the waiting room, each customer may abandon instead of entering service, after which the customer is lost. As in the original model, the abandonment times of successive arrivals to the queue are i.i.d. random variables with cumulative distribution function (cdf) F . The resulting model is the approximating DIS model.

In the DIS model with parameter w , the customer always enters service after spending time w in the queue, if the customer has not yet abandoned. That rule is possible because the service facility (the second IS queue) has infinitely many servers. We assume the system starts empty at time 0 and we let the first customer enter service after time w . Thus, for $t \geq w$, customers enter the service facility at rate $\lambda(t-w)\bar{F}(w)$, where $\lambda(\cdot)$ is the arrival-rate function and $\bar{F} \equiv 1 - F$.

Since all arrivals wait precisely the target duration w before entering service, if they do not elect to abandon, the approximate abandonment probability is always $F(w)$. Hence we can initially specify either the target abandonment probability $\alpha \equiv P(Ab)$ or the target delay w . If F is continuous, then there always is a w such that $F(w) = \alpha$ for any given α . If F is also strictly increasing, then $w = F^{-1}(\alpha)$. We assume that F is continuous and strictly increasing. Hence we can work with either α or w in the DIS model.

We now describe the performance in the approximating DIS model depicted in Figure 1. We start with a targeted waiting time w and the original $M_t/GI/s_t + GI$ model, specified by the arrival-rate function λ , the service-time cdf G and the abandonment-time cdf F . Let S and A be generic service-time and abandonment-time random variables; i.e., $G(x) \equiv P(S \leq x)$ and $F(x) \equiv P(A \leq x)$ for $x \geq 0$. Assume that $E[S] < \infty$. (We do not need to assume that $E[A] < \infty$ because in our approximation abandonments can only occur before time w .) Since F is continuous, F has no point mass at w , i.e., $P(A = w) = 0$, so there is no ambiguity about customer action after waiting in queue for time w .

The approximating model thus becomes a network of two $M_t/GI/\infty$ queues in series. The waiting room has arrival-rate function λ and service times distributed as $T \equiv A \wedge w \equiv \min\{A, w\}$, while the service facility has arrival rate $\lambda(t-w)\bar{F}(w)$, where $\bar{F}(x) \equiv 1 - F(x)$, and the given service-time cdf G . Let F_T be the associated cdf of the truncated random variable T , i.e.,

$$F_T(x) \equiv P(T \leq x) = F(x), \quad 0 \leq x < w, \quad F_T(x) = 1, \quad x \geq w. \quad (5)$$

We see that T has a point probability mass at w , since $P(T = w) = P(A \geq w) = \bar{F}(w)$.

As in Eick et al. (1993a), we assume that the system starts in the infinite past (at $t = -\infty$), with the policy above (customers entering service after waiting w if they have not yet abandoned). With this convention, all processes are defined on the entire real line. That is convenient both for some formulas and for representing the dynamic steady state associated with periodic arrival-rate functions, as in Eick et al. (1993b). If we want the system to start empty at time 0, then we can simply let $\lambda(t) = 0$ for all $t < 0$.

We can split the arrival process into two independent Poisson processes, one for the customers who will eventually be served and the other for the customers who will eventually abandon. Each customer is eventually served with probability $\bar{F}(w)$. We can further modify these two Poisson processes to obtain independent Poisson processes for the process counting customers entering service (at the times they enter service) and the process counting customers abandoning (at the times they abandon). Each can be represented as the departure process from an $M_t/GI/\infty$ queue, which corresponds to the original Poisson arrival process modified by having its points translated to the right by i.i.d. random variables; that is well known to preserve the Poisson property. For the counting process counting customers entering service, the service time is the constant value w ; for the counting process counting customers abandoning, the service time is by the random value $(T|T < w)$. By this construction, we have proved that the arrival process to the service facility is indeed a nonhomogeneous Poisson process with rate $\lambda(t-w)\bar{F}(w)$ at time t . We can thus apply established results for the $M_t/GI/\infty$ model, in particular Theorem 1 of Eick et al. (1993a).

Let $Q(t, y)$ denote the number of customers in queue at time t that have remaining time before abandonment greater than y and let $Q(t) \equiv Q(t, 0)$ denote the total number of customers in queue at time t . The random variable $Q(t, y)$ is depicted in Figure 2, which parallels Figure 1 of Eick et al. (1993a). We put a point at (x, y) in the plane if there is an arrival at x with abandonment time y . Thus $Q(t, y)$ will be the shaded region above the interval $[t - w, t]$ as shown. All arrivals with abandonment times greater than w will be served. All arrivals before time $t - w$ with abandonment times greater than w will have entered service before time t .

Let $B(t, y)$ denote the number of customers in the service facility at time t that have remaining service time greater than y and let $B(t) \equiv B(t, 0)$ denote the total number of busy servers (number of customers in the service facility) at time t . Let $X(t) \equiv Q(t) + B(t)$ be the number of customers in the system at time t . Let $W(t)$ be the potential waiting time at time t , i.e., the virtual waiting time before entering service of a customer with infinite patience arriving at time t .

We now summarize the main structural results for the approximation. For a nonnegative random

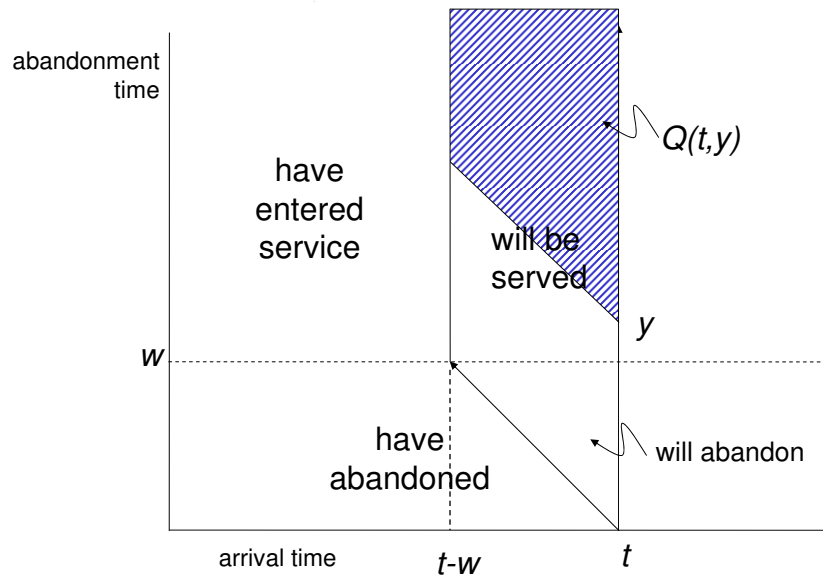


Figure 2 The random variable $Q(t, y)$ representing the queue content at time t that has remaining time before abandonment greater than y , in the DIS approximation.

variable Z with finite mean $E[Z]$ and cdf H , let Z_e denote a random variable with the associated *stationary-excess cdf* (or residual-lifetime cdf) H_e , defined by

$$H_e(y) \equiv P(Z_e \leq y) \equiv \frac{1}{E[Z]} \int_0^y \bar{H}(x) dx, \quad y \geq 0. \quad (6)$$

The moments of Z_e can be easily expressed in terms of the moments of Z via

$$E[Z_e^k] = \frac{E[Z^{k+1}]}{(k+1)E[Z]}, \quad k \geq 1. \quad (7)$$

Both $B(t, y)$ and $Q(t, y)$ are functions of the model parameter α (or $w = F^{-1}(\alpha)$), but in this section we drop the subscript for simplicity.

THEOREM 1. *Consider the DIS approximation for the $M_t/GI/s_t + GI$ model specified above, starting in the distant past with specified delay target (parameter) $w \geq 0$ or abandonment probability target $\alpha = F(w)$. The approximation makes $W(t) = w$ with probability 1 and the probability of abandonment $F(w)$ for all arrivals. Moreover, $Q(t, y_1)$ and $B(t, y_2)$ are independent Poisson random variables for each t and each pair (y_1, y_2) with $y_1 > 0$ and $y_2 > 0$, having means*

$$\begin{aligned} E[Q(t, y)] &= \int_{t-w}^t \lambda(x) \bar{F}(t+y-x) dx = \int_0^w \lambda(t-x) \bar{F}(y+x) dx, \\ E[B(t, y)] &= \bar{F}(w) \int_{-\infty}^t \lambda(x-w) \bar{G}(t+y-x) dx = \bar{F}(w) \int_0^{\infty} \lambda(t-w-x) \bar{G}(y+x) dx. \end{aligned} \quad (8)$$

The total numbers of customers in queue and in service at time t , $Q(t)$ and $B(t)$ respectively, are independent Poisson random variables with means

$$\begin{aligned} E[Q(t)] &= E[Q(t, 0)] = \int_{-\infty}^t \lambda(x) \bar{F}_T(t-x) dx = \int_{t-w}^t \lambda(x) \bar{F}_T(t-x) dx \\ &= E \left[\int_{t-T}^t \lambda(x) dx \right] = E[\lambda(t - T_e)] E[T], \end{aligned} \quad (9)$$

$$\begin{aligned} E[B(t)] &= E[B(t, 0)] = \bar{F}(w) \int_{-\infty}^t \lambda(x-w) \bar{G}(t-x) dx \\ &= \bar{F}(w) E \left[\int_{t-w-S}^{t-w} \lambda(x) dx \right] = \bar{F}(w) E[\lambda(t-w - S_e)] E[S]. \end{aligned} \quad (10)$$

Thus, $X(t)$, the total number of customers in the system at time t is a Poisson random variable

with a mean $E[Q(t)] + E[B(t)]$. The processes counting the numbers of customers abandoning and entering service are independent Poisson processes with rate functions ξ and β , respectively, where

$$\begin{aligned}\xi(t) &= \int_0^w \lambda(t-x) dF(x) = E[\lambda(t-T)|T < w], \\ \beta(t) &= \lambda(t-w)\bar{F}(w).\end{aligned}\tag{11}$$

The departure process (counting the number of customers completing service) is a Poisson process with rate

$$\sigma(t) = \bar{F}(w) \int_0^\infty \lambda(t-w-x) dG(x) = \bar{F}(w)E[\lambda(t-w-S)].\tag{12}$$

Proof. For the most part, these results are a direct application of Theorem 1 of Eick et al. (1993a). Understanding is facilitated by drawing pictures of Poisson random measure. Here we elaborate on only one point: To establish the claim that $Q(t, y_1)$ is independent of $B(t, y_2)$ for each (y_1, y_2) , first observe that the departure process from the queue prior to time t is independent of $\{Q(t, y) : y \geq 0\}$. That implies that the process of customers entering service prior to time t is also independent of $\{Q(t, y) : y \geq 0\}$. But $B(t, y_2)$ depends only on the history of Q up to time t through its own arrival process up to time t . ■

As discussed in Eick et al. (1993a), the last two representations for $E[Q(t)]$ and $E[B(t)]$ in (10) are appealing because they show random time lags, but these random time lags appear inside the expectation in a nonlinear fashion. We get convenient explicit formulas when the arrival rate function λ has special structure, e.g., when λ is sinusoidal or a polynomial, as we show in §EC.2. As in Eick et al. (1993a) and Massey and Whitt (1997), the polynomial arrival rate functions yield useful approximations. We can see directly that the targeted wait of w before starting service increases the random time lags in $E[B(t)]$ by w . This will be negligible if w is small, but not otherwise. As noted in Eick et al. (1993a), the time lag in $E[B(t)]$ involving $w + S_e$ is different from the time lag $w + S$ appearing in the departure rate $\sigma(t)$.

Since the departure process from the service facility is a Poisson process, we see that the approximation extends immediately to yield corresponding approximations for any number of $M_t/GI/s_t + GI$ models in series, in the spirit of Massey and Whitt (1993).

We now indicate how the DIS approximation can be used to specify the staffing function $s_\alpha(t)$ in the original $M_t/GI/s_t + GI$ model in order to achieve target $\alpha \equiv P(Ab)$. For any given target abandonment probability α with the direct DIS approximation, the number of busy servers at time t would be the random variable $B_\alpha(t)$. With the DIS approximation, $B_\alpha(t)$ has a Poisson distribution with mean $m_\alpha(t) \equiv E[B_\alpha(t)]$, for which we give an explicit formula. Hence $B_\alpha(t)$ is approximately normally distributed with both mean and variance equal to $m_\alpha(t)$.

The *simple DIS staffing approximation* is to simply set $s_\alpha(t) = m_\alpha(t)$. We will show that the simple DIS staffing policy is effective when the QoS is low, but not when the QoS is high.

4. Asymptotic Stability

In this section we prove that simple DIS staffing $s_\alpha(t) = m_\alpha(t)$ is effective in stabilizing the abandonment probability and the expected delay at any positive target values α and w related by $\alpha = F(w)$ asymptotically as the scale increases. For that purpose, we apply the many-server heavy-traffic FWLLN established in Liu and Whitt (2012a,b,c). That FWLLN involves a sequence of $G_t/GI/s_t + GI$ queueing models indexed by n . Model n has a general arrival process with time-varying arrival rate $\lambda_n(t) \equiv n\lambda(t)$ (which covers the M_t assumption of the current paper), i.i.d. service times with cumulative distribution function (cdf) G , a time-varying number of servers $s_{n,\alpha}(t) \equiv \lceil ns_\alpha(t) \rceil$ (the least integer above $ns_\alpha(t)$) and customer abandonment from queue, where the patience times of successive customers to enter queue are i.i.d. with general cdf F . The two cdf's G and F are fixed independent of n , and differentiable, with positive probability density functions (pdf's) g and f .

Our scaling of the fixed functions λ and s induces the familiar many-server heavy-traffic scaling. Under that scaling, and under regularity conditions, Liu and Whitt (2012b,c) establish a FWLLN with convergence of the appropriately scaled stochastic processes to deterministic performance functions associated with the fluid model analyzed in Liu and Whitt (2012a). Under this scaling and these regularity conditions, we now show that the DIS staffing achieves stability asymptotically; i.e., the time-dependent mean delay and abandonment probability are stabilized as $n \rightarrow \infty$.

To state the result, let $Q_n(t)$ be the number of customers waiting in queue at time t in the n^{th} queueing model. Let $W_n(t)$ be the corresponding potential waiting time, i.e., the virtual waiting time at time t if there were an arrival at time t , assuming that arrival had unlimited patience. Let $A_n(t)$ be the number of customers that have abandoned in the interval $[0, t]$. Let $A_n(t, u)$ be the number of customers among arrivals in $[0, t]$ that have abandoned in the interval $[0, t + u]$. Let $\bar{Q}_n(t) \equiv n^{-1}Q_n(t)$, $\bar{A}_n(t) \equiv n^{-1}A_n(t)$ and $\bar{A}_n(t, u) \equiv n^{-1}A_n(t, u)$ be the associated FWLLN-scaled processes. (The process $W_n(t)$ is not scaled.) Let $\Lambda(t) \equiv \int_0^t \lambda(s) ds$. Let 1_C be the indicator variable, which is equal to 1 if event C occurs and is equal to 0 otherwise.

Since the arrival processes are nonhomogeneous Poisson processes here, both a functional central limit theorem (FCLT) and a FWLLN hold for the arrival processes, as required in Liu and Whitt (2012b,c). To establish convergence of expected potential waiting times, we assume the the regularity conditions in Liu and Whitt (2012a,b,c) are satisfied, namely, all model parameters (λ , F and G) are piecewisely continuous and differentiable. We assume in addition that the service times have finite second moments.

THEOREM 2. (*asymptotic stability*) *Consider a sequence of $M_t/GI/s_t + GI$ models with the many-server heavy-traffic scaling specified above. Suppose that these systems start empty at time 0, the regularity conditions in Liu and Whitt (2012a,b,c) are satisfied and $E[S^2] < \infty$. Then, with the abandonment-probability target α , under the DIS staffing $s_{n,\alpha}(t) \equiv \lceil n s_\alpha(t) \rceil$, where*

$$s_\alpha(t) = m_\alpha(t) \equiv E[B_\alpha(t)] = \bar{F}(w) \int_0^{t-w} \bar{G}(x) \lambda(t-w-x) dx \cdot 1_{\{t>w\}}, \quad (13)$$

the expected delays and abandonment probabilities are stabilized at their targets w and α with $\alpha = F(w)$ asymptotically as $n \rightarrow \infty$; i.e., for any time b with $w < b < \infty$,

$$\begin{aligned} \sup_{0 \leq t \leq b} \{|\bar{Q}_n(t) - E[Q(t)]|\} &\Rightarrow 0, & \sup_{0 \leq t \leq b} \{|W_n(t) - w|\} &\Rightarrow 0, & E[W_n(t)] &\rightarrow w, & t \geq 0, \\ \sup_{0 \leq t \leq b} \{|\bar{A}_n(t) - A(t)|\} &\Rightarrow 0 & \text{and} & & \sup_{0 \leq t \leq b, w < u < b} \{|\bar{A}_n(t, t+u) - A(t, u)|\} &\Rightarrow 0 & (14) \end{aligned}$$

as $n \rightarrow \infty$, where

$$E[Q(t)] = E[Q(t, 0)] \equiv \int_0^w \lambda(t-x) \bar{F}(x) dx,$$

$$A(t) \equiv \int_0^t \xi(s) ds, \quad \xi(t) \equiv \int_0^w \lambda(t-x)f(x) dx \quad \text{and} \quad A(t, u) \equiv \Lambda(t)\alpha, \quad u > w. \quad (15)$$

REMARK 1. (stabilizing abandonment probabilities, not rates) From (15), we see that the abandonment rate $\xi(t)$ is not stabilized. Instead, the proportion of arrivals arriving in any time interval $[t_1, t_2]$ that eventually abandon before starting service approaches α . That is consistent with our goal to stabilize the abandonment probability, as opposed to the abandonment rate. That is achieved starting empty by not staffing until time w .

Proof. Under the regularity conditions specified in Liu and Whitt (2012a,b,c), the appropriately scaled versions of the stochastic processes describing the performance in the $G_t/GI/s_t + GI$ queueing models indexed by n converge to corresponding deterministic functions describing the performance of an associated deterministic fluid model having capacity $s(t)$ and fluid input arriving at rate $\lambda(t)$ at time t . The performance of this limiting fluid model is characterized in Liu and Whitt (2012a).

A key assumption in Liu and Whitt (2012a,b,c) is that the fluid model alternates between underloaded intervals and overloaded intervals, with critical loading only holding at isolated points of switching from one regime to another. When we use the DIS staffing in order to stabilize abandonments at a positive target α , we are forcing the system to always operate in an overloaded regime after an initial transient required for the capacity to be filled. Thus, by staffing in that way, we consider the special case in which the system remains overloaded after an initial underloaded interval; i.e., there is a single switching point in the fluid model at $t = w$ (the delay target). In the terminology of Garnett et al. (2002), the limit is for the quality-driven (underloaded) many-server heavy-traffic regime before time w and the efficiency-driven (overloaded) many-server heavy-traffic regime after time w .

In §10 of Liu and Whitt (2012a) it was shown how to staff the fluid model in order to stabilize delays and abandonments. (Abandonment probabilities in the queueing model correspond to proportions of entering fluid that abandon before entering service in the fluid model.) For simplicity, in §10 of Liu and Whitt (2012a) it was assumed that the fluid model starts empty at time 0, and

so we made that assumption, but it was shown how to treat more general initial conditions in §H of the appendix of Liu and Whitt (2012a). When starting empty, the staffing policy to stabilize delays of all fluid to enter service at a target w provides no staffing at all, and thus allows no fluid to enter service, until time w . The stabilizing staffing function after time w is given by (13) above. Given that the delays are stabilized at w , the proportion of arriving fluid at each time to abandon before entering service is $\alpha = F(w)$. Since F is continuous and strictly increasing, it has a unique inverse F^{-1} , so that we could start with α instead of w , and then let $w = F^{-1}(\alpha)$.

Moreover, as discussed in §4 of Liu and Whitt (2012a), there is an intimate connection between the fluid content at time t and the mean of the number of busy servers in an associated IS model. As a consequence, this staffing function for the fluid model coincides with the simple DIS staffing, adjusted appropriately for the scaling factor n ; i.e., in the fluid model, $s_\alpha(t) = E[B_{\alpha,n}(t)]/n$, if we also assume that the fluid model starts empty and we first provide staffing at time w , where w is chosen so that $w = F^{-1}(\alpha)$. Thus we can combine the results above to deduce that DIS staffing achieves its goal asymptotically.

As a consequence, if we use the simple DIS staffing in the fluid model, we succeed in stabilizing the delays in the fluid model. We then apply the FWLLN with that particular staffing function. Then the FWLLN in Liu and Whitt (2012b,c) directly applies the stated limits for $\bar{Q}_n(t)$, $W_n(t)$ and $\bar{A}_n(t)$ in (14). Since all fluid waits exactly w before entering service, if it does not abandon, the same will be true asymptotically for the stochastic model. Thus we get the stated limit for $\bar{A}_n(t, u)$ with $u > w$.

Finally, we apply uniform integrability to get the convergence of means $E[W_n(t)] \rightarrow w$ for each t from the established convergence $W_n(t) \Rightarrow w$, using p. 31 of Billingsley (1999). To obtain the uniform integrability, we show that $\sup_{n \geq 1} E[W_n(t)^2] < \infty$. That is proved in §7.1. ■

From the representation of the DIS approximating mean queue length $E[Q_\alpha(t)]$ in Theorem 1, we can show that it cannot be a constant function with a time-varying arrival rate function. Hence, Theorem 2 together with two additional lemmas below implies the following corollary.

COROLLARY 1. (*the mean queue length is not stabilized asymptotically*) Suppose that the conditions of Theorem 2 hold with the arrival rate function λ not being a constant function. Then, under DIS staffing, the scaled number in queue $n^{-1}Q_n(t)$ converges to the mean DIS queue length, which cannot be a constant function of t after time w . Thus, the DIS staffing function that asymptotically stabilizes the abandonment probability does not asymptotically stabilize the mean queue length.

We use the following lemma to prove Corollary 1. We prove this lemma and the following one in §7.2.

LEMMA 1. (*uniqueness of time-shifted integral*) If

$$m(t) = \int_0^w \lambda(t-x)\bar{F}(x) dx, \quad t \geq w, \quad (16)$$

is a positive constant for all $t \geq w$, then λ is a constant function for $t \geq 0$.

We now observe that the DIS staffing is essentially the only staffing function that can stabilize abandonments and delays.

COROLLARY 2. (*uniqueness*) The DIS staffing in (13) is the unique staffing function that stabilizes abandonment and delays at positive values in the fluid model. Consequently, any other sequence of staffing functions $\{s_n : n \geq 1\}$ that asymptotically stabilizes abandonment and delays must satisfy $n^{-1}s_n(t) \rightarrow s_\alpha(t)$ as $n \rightarrow \infty$.

However, the MOL staffing function for the stochastic system is unique only in the order of $o(n)$, according to the fluid scaling. We use the following lemma to prove Corollary 2. The following lemma shows that there are not multiple staffing functions that produce identical potential waiting time functions or identical abandonment rate functions in the limiting fluid model.

LEMMA 2. (*unique fluid model staffing functions yielding given targeted performance*) For the $G_t/GI/s_t + GI$ deterministic fluid model specified by the model data $(\lambda, s, G, F, b(0, \cdot), q(0, \cdot))$ satisfying the assumptions of Liu and Whitt (2012a) starting empty at time 0, the DIS staffing in (13) is the unique staffing yielding the positive constant target delay w and abandonment $\alpha = F(w)$.

We can combine Corollaries 1 and 2 to deduce the following corollary

COROLLARY 3. (*impossibility of simultaneous stabilization*) *There does not exist a staffing function that can simultaneously stabilize the abandonment probability and the mean queue length at positive targets asymptotically in the many-server heavy-traffic regime.*

5. The DIS-MOL Approximation

Consistent with Theorem 2, in simulation experiments we see that the simple DIS staffing is remarkably effective in stabilizing the abandonment probability and the expected delays in large-scale queueing models with relatively low QoS (high abandonment probabilities and expected delays); see Figure 4 for an example. Unfortunately, however, the DIS approximation does not perform so well for higher QoS (lower abandonment probabilities and expected delays). Such lighter loads tend to move the system from the ED asymptotic regime to the QED asymptotic regime. (With higher QoS, the scale must be extremely large, such as $n = 1000$ or more, before the DIS staffing is effective. Experience indicates that the required scale increases as α decreases.)

Fortunately, for such higher QoS, a new MOL approximation performs remarkably well; see Figure 5 for an example. We now develop it. Let $P_t(Ab)$ be the time-dependent probability that an arrival at time t will eventually abandon. For a stationary model, the offered load can be defined as $\lambda(1 - P(Ab))E[S]$, the rate customers enter service, $\lambda(1 - P(Ab))$, multiplied by the mean service time, $E[S]$. A candidate time-dependent generalization would be the pointwise-stationary approximation for the offered load, $\lambda(t)(1 - P_t(Ab))E[S]$, but $\lambda(t)(1 - P_t(Ab))$ is the rate of arrivals at time t that will eventually enter service; it is not the rate customers actually enter service at time t . By Little's law applied to the service facility in the stationary model, $\lambda(1 - P(Ab))E[S]$ is also the expected number of busy servers in steady state. Experience indicates that the mean number of busy servers in the IS system tends to be a far better analog of the offered load in a nonstationary model.

Thus, to obtain the DIS-MOL approximation for the staffing $s(t)$ at time t to achieve the target

α , we let $w = F^{-1}(\alpha)$ and we look at the steady-state distribution of the stationary $M/GI/s + GI$ model applied with $s = s(t)$ and arrival rate

$$\lambda_{\alpha}^{MOL}(t) \equiv \frac{m_{\alpha}(t)}{E[S](1-\alpha)}. \quad (17)$$

For each fixed time t , we let the DIS-MOL staffing level $s_{\alpha}^{MOL}(t)$ be the least integer staffing level such that the stationary abandonment probability $P(Ab)$ in the $M/GI/s + GI$ model with arrival rate $\lambda = \lambda_{\alpha}^{MOL}(t)$ in (17) and $s = s_{\alpha}^{MOL}(t)$ satisfies $P(Ab) \leq \alpha$.

Since the stationary $M/GI/s + GI$ model itself is quite complicated, we apply the approximation from Whitt (2005), which is based on an associated Markovian $M/M/s + M(n)$ model with state-dependent abandonment rates. Alternatively, since that approach approximates the general service-time distribution by an exponential distribution with the same mean, one can use the exact solution or asymptotic approximations for the associated $M/M/s + GI$ model from Zeltyn and Mandelbaum (2005). Either way, we are exploiting the property that the service-time distribution beyond its mean tends not to matter to much in the stationary $M/GI/s + GI$ model; see Whitt (2005, 2006). In an application this property can be checked with simulation.

REMARK 2. (sensitivity to the service-time distribution beyond its mean) Since the stationary $M/GI/s/0$ loss model and the stationary $M/GI/\infty$ IS model have the celebrated insensitivity property, i.e., since the steady-state performance is independent of the service-time distribution beyond its mean, it is not too surprising that the stationary $M/GI/s + GI$ model should exhibit *approximate* insensitivity to the service-time distribution beyond its mean. However, unlike the performance in these stationary models, the insensitivity is lost when we abandon the stationarity, as was shown in Davis et al. (1995) for the loss model. Similarly, the transient performance in the $M_t/GI/s + GI$ model with time-varying arrival rate is more sensitive to the service-time distribution beyond its mean, as shown by the example in §2 of Liu and Whitt (2012a). That sensitivity is captured in our approximation through the impact of the service-time distribution beyond its mean on the transient performance of the time-varying IS and DIS models.

6. An $M_t/M/s_t + M$ Example with a Sinusoidal Arrival-Rate Function

In this section, we use simulation experiments to show that both the abandonment probability $P_t(Ab)$ and the expected delay $E[W(t)]$ are indeed stabilized (independent of time) in Markovian $M_t/M/s_t + M$ examples for low QoS targets with the DIS algorithm and for all QoS targets with the DIS-MOL algorithm. We show that the new methods also work for non-exponential service and patience distributions in §EC.4 and the longer version available on the authors' web pages.

Our staffing procedure applies to arbitrary arrival-rate functions, because we can apply Theorem 1 to calculate the DIS OL function $m(t)$. Following common practice in the study of time-varying arrival rates, we consider a sinusoidal arrival-rate function

$$\lambda(t) = a + b \cdot \sin(ct), \quad 0 \leq t \leq T. \quad (18)$$

This sinusoidal example is convenient because we can apply Theorem 1 to obtain explicit analytical formulas for the offered load $m_\alpha(t)$, paralleling Eick et al. (1993b). This sinusoidal example also roughly captures the spirit of real systems, as seen in call center data, as in Figure 7 of Feldman et al. (2008). We obtain a concrete many-server example by letting $a = 100$, $b = 20$, $c = 1$ and $T = 20$ in (18). Here we let the individual service rate μ be 1 and the individual abandonment rate θ be 0.5. (Choosing $\mu = 1$ corresponds to measuring time in units of mean service times.)

An important issue for applications is the rate of fluctuation in the arrival-rate function compared to the expected service time. Since a cycle of the sinusoidal arrival-rate function in (18) is $2\pi/c$ and we have set $c = 1$, the length of a cycle is $2\pi \approx 6.3$. Thus there will be slightly more than three cycles during the interval $[0, 20]$; e.g., see Figure 4. If we measure time in hours, then the mean service time is one hour and the full time period is slightly less than one day. Then the three peaks in the arrival rate in Figure 4 occur approximately at 2am, 8am and 2pm. Thus in this example the fluctuation in the arrival rate function are relatively fast compared to the expected service time. From Jennings et al. (1996) and Feldman et al. (2008), we know that the pointwise-stationary approximation does not perform well for this example. That is demonstrated for the abandonment probability and the expected delay in the appendix.

We now turn to DIS staffing. We apply Theorem 1 and Eick et al. (1993b) to calculate the DIS performance functions. Letting $\bar{F}(x) = e^{-\theta x}$, $\bar{G}(x) = e^{-\mu x}$, and $\lambda(t) \equiv 0$ for $t < 0$, we obtain

$$E[B(t)] = \begin{cases} e^{-\theta w} \frac{a}{\mu} - \left(\frac{a}{\mu} - \frac{bc}{\mu^2+c^2}\right)e^{-\theta w-\mu(t-w)} + \frac{be^{-\theta w}}{\mu\sqrt{1+c^2/\mu^2}} \sin[c(t-w) - \phi], & t \geq w \\ 0, & 0 \leq t < w, \end{cases} \quad (19)$$

$$E[Q(t)] = \begin{cases} \frac{a}{\theta}(1 - e^{-\theta w}) + b\sqrt{\frac{x^2+y^2}{\theta^2+c^2}} \sin(ct + \beta - \gamma), & t \geq w \\ \frac{a}{\theta} - \left(\frac{bc}{\theta^2+c^2} - \frac{a}{\theta}\right)e^{-\theta t} + \frac{b}{\sqrt{\theta^2+c^2}} \sin(ct - \gamma), & 0 \leq t < w, \end{cases} \quad (20)$$

where $\phi = \arctan(c/\mu)$, $\gamma = \arctan(a/\theta)$, $\beta = \arctan(y/x)$

$$x = 1 - e^{-\theta w} \cos(cw), \quad y = e^{-\theta w} \sin(cw)$$

From (19), we see that $E[B(t)]$ is asymptotically sinusoidal as $t \rightarrow \infty$, eventually coinciding with the formula for the periodic steady state, given for the more general $M_t/GI/s_t + GI$ model in Theorem EC.1. From (20), we see that $E[Q(t)]$ is sinusoidal when $t \geq w$.

We first compare the DIS-MOL staffing function $s_\alpha^{MOL}(t)$ to the DIS staffing function, which




Figure 3 A comparison of the OL function $m_\alpha(t)$ and the DIS-MOL staffing function $s_\alpha^{MOL}(t)$ for the $M_t/M/s_t + M$ example with sinusoidal arrival rate for abandonment probability targets $\alpha = 0.1$ and 0.001 .

is the OL $m_\alpha(t)$. In Figure 3, we plot $\{s_\alpha^{MOL}(t) : 0 \leq t \leq T\}$ and $\{m_\alpha(t) : 0 \leq t \leq T\}$ ($T = 20$) with two abandonment probability targets: $\alpha = 0.1$ and $\alpha = 0.001$. These two targets are extreme cases. The first one ($\alpha = 10\%$) represents a low QoS; the second one ($\alpha = 0.1\%$) represents a high QoS. Figure 3 shows that $\{s_\alpha^{MOL}(t) : 0 \leq t \leq T\}$ and $\{m_\alpha(t) : 0 \leq t \leq T\}$ coincide when the QoS is low. Therefore, in order to stabilize abandonment under low QoS, we can just staff the system with the OL function $m_\alpha(t)$, given in Theorem 1. However, when the QoS is high, Figure 3 shows that the DIS-MOL staffing function is significantly above the OL function, which makes the MOL refinement necessary. Our experiments indicate that the DIS-MOL staffing is never less than the OL (DIS staffing).

We now use simulation experiments to show that DIS-MOL approximation achieves the desired time-stable performances under all QoS targets. First we evaluate the simple DIS approximation at low QoS targets (where it is nearly identical to DIS-MOL). Figure 4 shows simulation estimates of key performance measures with target abandonment probability α for $0.05 \leq \alpha \leq 0.20$. (Additional details about the simulation estimates are given in §EC.3.) Figure 4 shows that both the abandonment probability and the expected delay are stabilized at the targets α and $w = F^{-1}(\alpha) = -1/\theta \log(1 - \alpha)$. Moreover, as predicted, the queue-length processes are not stabilized; they agree closely with the formulas in (20). All these quantities were estimated by performing multiple (5000) independent replications under the staffing function $s_\alpha(t) = m_\alpha(t) = E[B(t)]$ in (19).

Figure 4 shows that the simple DIS approximation works quite well for α between 0.05 and 0.20, and associated expected delays ranging from 0.1 to 0.45. However, at least 5% abandonment may not be acceptable. For higher QoS targets, the DIS-MOL approximation is needed.

In Figure 5, we again plot the expected queue lengths, abandonment probabilities, delay probabilities and expected delays, using the DIS-MOL approximations with relatively low abandonment probability targets $0.005 \leq \alpha \leq 0.02$. The DIS-MOL approximation works remarkably well after the initial transient period.

In the DIS model, since all customers are required to wait in queue for w before entering service, unlike the expected queue length, we are not able to produce an approximate delay probability,

because the DIS approximation predicts that every customer should be delayed. Of course, in the original stochastic model, the delay probability is always less than 1. The delay probability increases as the abandonment probability α increases. When α is big enough, e.g., $\alpha = 0.2$, the delay probability is nearly 1, as shown in Figure 4; when α is as small as 0.005, the delay probability goes down to 0.2, as shown in Figure 5. Figure 4 and 5 show that the delay probabilities are stabilized when α is small and are sinusoidal otherwise. This phenomenon is consistent with Figure 4 in Feldman et al. (2008), which shows the abandonment probabilities when the delay probability is the target. There, both were stabilized in with high QoS, but the abandonment probabilities were not stabilized under low QoS (high delay probability targets).




Figure 4 A comparison of the simple DIS approximation with simulation: Estimated time-dependent expected queue lengths, abandonment probabilities, delay probabilities and expected delays for the $M_t/M/s_t + M$ example with sinusoidal arrival under the OL Staffing, with low QoS ($\alpha = 0.05, 0.10, 0.15, 0.20$).

We conclude this section by remarking that the DIS-MOL approximation is consistent with a square-root-staffing (SRS) formula. Paralleling (2), the candidate new SRS formula for the $M_t/GI/s_t + GI$ model based on the DIS OL with abandonment target α would be

$$s_\alpha(t) = m_\alpha(t) + \beta_\alpha \sqrt{m_\alpha(t)}. \quad (21)$$

Formula (21) differs from (2) by using the DIS OL $m_\alpha(t)$ instead of the standard OL $m_0(t)$ in (1). In §EC.5 we verify that the DIS-MOL approximation is indeed consistent with the SRS staffing in (21); i.e., we show that the DIS-MOL staffing function $s_\alpha^{MOL}(t)$ has the form of the SRS formula (21). Following Feldman et al. (2008), for an abandonment probability target α , we let $D_\alpha(t) \equiv s_t^{MOL} - m_\alpha(t)$ be the difference between the DIS-MOL staffing and the OL functions, and let $\beta_\alpha(t) \equiv D_\alpha(t)/\sqrt{m_\alpha(t)}$ be the *implicit QoS function* for DIS-MOL. Figure EC.4 shows that $\beta_\alpha(t) \approx \beta_\alpha$, independent of t , where β_α is a nonnegative decreasing function of α . However, it remains to find




Figure 5 Estimated time-dependent expected queue lengths, abandonment probabilities, delay probabilities and expected delays for the $M_t/M/s_t + M$ example with sinusoidal arrival under the DIS-MOL staffing with relatively high QoS ($\alpha = 0.005, 0.01$ and 0.02).

a simple formula for the new QoS constant β_α . In §EC.5 we show results of fitting it to simulation data.

7. Additional Proofs for §4

7.1. A Bound to Justify Uniform Integrability in Theorem 2

To justify the uniform integrability used to imply convergence in moments $E[W_n(t)] \rightarrow w$ given the established convergence in distribution $W_n(t) \Rightarrow w$, we show that $\sup_{n \geq 1} E[W_n(t)^2] < \infty$, where $W_n(t)$ is the potential waiting time at time t in model n ; see p. 31 of Billingsley (1999).

LEMMA 3. *Under the assumptions of Theorem 2, $\sup_{n \geq 1} E[W_n(t)^2] < \infty$.*

Proof. It suffices to use a crude upper bound. Thus we simplify. First, we bound the arrival process in model n , $N_n(t)$, above by a stationary Poisson process over the interval $[0, T]$ by letting the fluid arrival rate be $\lambda_{bd} \geq \lambda(t)$ for all t in a bounded interval $[0, T]$. Similarly, we bound the staffing function in the fluid model below by a positive capacity s_{bd} with $0 < s_{bd} \leq s(t)$, $0 \leq t \leq T$, which is possible by Assumption 11 of Liu and Whitt (2012a). We also assume that there is no abandonment. These modifications can be shown to increase the process $W_n(t)$ in sample path stochastic order, as in Whitt (1981).

With those simplifications, we have a sequence of classical $M/GI/s/\infty$ models, with arrival rate $n\lambda_{bd}$ and staffing level $\lceil ns_{bd} \rceil$, $n \geq 1$. We now focus on these models, without changing the notation. Let $W_n^a(k)$ be the actual delay of the k^{th} arrival and observe that $W_n(t) \leq W_n(N_n(t)+) \leq W_n^a(N_n(t)) + S$, where $N_n(t)$ is the homogeneous Poisson arrival process with rate $n\lambda_{bd}$ and S is a generic service time independent of $W_n^a(N_n(t))$. Next we bound the given $M/GI/s$ models with the customary FCFS service discipline above by the associated $M/G/s$ model assigning the customers to servers in a cyclic or round robin order. In particular, we next apply the stochastic bounds in Wolff (1977, 1987) for the moments to deduce that

$$E[(W_n^a(N_n(t)))^2] \leq E[W_n^{a,c}(N_n(t))^2],$$

where the additional superscript c on the right side denotes the cyclic service discipline.

With the cyclic service discipline, the model is equivalent to separate $GI/GI/1$ models. However, the arrival process at each server in model n has Erlang $E_{\lceil s_{bd}n \rceil}$ interarrival times, which change with n . This will not be a serious difficulty, because we can relate these arrival processes back to the original Poisson arrival process. Next, the upper bound can be further bounded above by the sum of all service times assigned to that single server up to time T , i.e.,

$$E[W_n^{a,c}(N_n(t))]^2 \leq E \left[\left(\sum_{i=1}^{N_n^c(T)} S_i \right)^2 \right], \quad 0 \leq t \leq T, \quad (22)$$

where $N_n^c(T)$ is the number of arrivals assigned to that individual server, which is a renewal process with Erlang interarrival times. Since every $\lceil ns_{bd} \rceil^{\text{th}}$ arrival in the original system is assigned to this server, $N_n^c(T) \leq (N_n(T)/s_{bd}n) + 1$. Combining these results, we get

$$E[(W_n^a(N_n(t)))^2] \leq E \left[\left(\sum_{i=1}^{\lceil N_n(T)/s_{bd}n \rceil + 1} S_i \right)^2 \right], \quad 0 \leq t \leq T, \quad (23)$$

where $N_n(t)$ is a Poisson process with rate $\lambda_{bd}n$. Hence,

$$E[N_n(T)/s_{bd}n] = \frac{\lambda_{bd}n}{s_{bd}n} = \frac{\lambda_{bd}}{s_{bd}} < \infty \quad \text{and} \quad \text{Var}(N_n(T)/s_{bd}n) = \frac{\lambda_{bd}n}{(s_{bd}n)^2} \leq \frac{\lambda_{bd}}{s_{bd}^2} < \infty,$$

from which the desired uniform bound follows, using formulas for compound Poisson random variables and the condition that $E[S^2] < \infty$. ■

7.2. Proofs of Two Lemmas in §4

We now prove the two lemmas used with Theorem 2 to prove Corollary 2.

Proof of Lemma 1. Let $m_w \equiv \int_0^w \bar{F}(x) dx$ and $p_w(x) \equiv \bar{F}(x)/m_w$, $0 \leq x < w$. Then the function m in (16) can be expressed as an integral weighted average on the interval $[w, \infty)$ with respect to the positive pdf p_w , namely,

$$m(t) = m_w \int_0^w \lambda(t-x)p_w(x) dx, \quad t \geq w.$$

Now extend the pdf p and the arrival rate function λ to the entire real line by letting $p_w(x) = 0$ if $x \leq 0$ or if $x \geq w$ and $\lambda(t) = 0$ for $t < 0$. Then $m(t) = c$ for all $t \geq w$ if and only if the convolution integral

$$m_{c,w}(t) \equiv m_w \int_{-\infty}^{\infty} \lambda_{c,w}(t-x)p_w(x) dx = 0 \quad \text{for all } t \geq 0, \quad (24)$$

where $\lambda_{c,w}(t) \equiv \lambda(t+w) - c$ for all $t \geq 0$. In particular, $m_{c,w}$ is the convolution of the function $\lambda_{c,w}$ with respect to the density p , which is decreasing and positive on its interval of support, $[0, w]$; i.e., $m_{c,w} = p_w \star \lambda_{c,w}$, as on p. 143 of Feller (1971). Now let

$$\hat{p}_w(s) \equiv \int_{-\infty}^{\infty} e^{-st} p_w(t) dt$$

for positive real s , and similarly for the other functions. Exploiting basic properties of transforms of convolutions, we get $\hat{m}_{c,w}(s) = \hat{\lambda}_{c,w}(s) \hat{p}_w(s)$ for all positive real s . If $m_{c,w}(t) = 0$ for all t , then necessarily $\hat{m}_{c,w}(s) = 0$ for all positive real s . Since $\hat{p}_w(s) > 0$ for all positive real s , we deduce that necessarily $\hat{\lambda}_{c,w}(s) = 0$ for all $s > 0$. Since $\hat{\lambda}_{c,0}(s) = e^{sw} \hat{\lambda}_{c,w}(s)$, where $\hat{\lambda}_{c,0}(s)$ coincides with the ordinary Laplace transform of $\lambda_{c,w}(t)$, we see that $\tilde{\lambda}_{c,0}(s) = 0$ for all positive real s . By §VII.6 of Feller (1971), that implies that $\lambda_{c,w}(t) = 0$ for all t . ■

Proof of Lemma 2. Since the fluid model starts empty, in order to have all fluid experience delay of w , no staffing at all can be provided until time w . After time w the staffing must be as in (13) in order to achieve maintain the target delay w . In turn, that fixed delay must be obtained to provide the fixed abandonment proportion $\alpha = F(w)$.

To elaborate, there must be a first time that the staffing deviates from the DIS staffing. We can trace the implications of a change in the staffing function in the fluid model, referring to results in Liu and Whitt (2012a). First, a change in the staffing function s necessarily changes the rate fluid enters service $b(\cdot, 0)$. To see that, first observe that a change in s necessarily changes \hat{a} in (19) (of Liu and Whitt (2012a), like all references in this proof) to an associated \bar{a} . In particular, (19) implies that the function \hat{a} is a monotone function of the derivative s' . If we increase s' over some interval $[0, u]$, then, \hat{a} necessarily increases over $[0, u]$. From the monotonicity of the fixed point equation for the rate fluid enters service in (18), we can apply Theorem 2 to deduce that the $b(\cdot, 0)$ increases as well, and similarly for a decrease. Hence, a change in the staffing function s forces a change in the rate fluid enters service, $b(\cdot, 0)$, which is monotone in s' . That change in the function $b(\cdot, 0)$ in turn forces a change in the BWT w by virtue of the ODE in Theorem 3. However, here an increase in $b(t, 0)$ forces a decrease in $w(t)$. Thus the first change of staffing changes the BWT

$w(t)$. The change in the BWT w forces corresponding changes in the fluid density in queue $q(t, x)$ by Corollary 2, the PWT $v(t)$ by Theorem 5, and in the abandonment rate function α by (7). The change in $b(\cdot, 0)$ produces changes in the fluid density in service $b(t, x)$ via (15) and the service completion rate σ via (9). Since we are concerned with stabilizing the PWT $v(t)$ at w , the change in $v(t)$ implies that the DIS staffing is the unique staffing that achieves the stabilizing goals. ■

8. Conclusions

We have developed a systematic formula-based procedure to stabilize the abandonment probability and the expected delay in an $M_t/GI/s_t + GI$ model with a time-varying arrival rate, across a wide range of performance targets, by providing an appropriate staffing function $s(t)$. The first step in §3 involves the delayed infinite-server (DIS) model with two $M_t/GI/\infty$ IS models in series. In this model (but not in the actual system) each customer arrives at the first IS system (representing the queue) and stays there for a fixed time $w = F^{-1}(\alpha)$ (the target). Customers may abandon from this first queue. If they do not, then they proceed to the second IS system (representing the service facility). Since the number of busy servers in the service facility, $B(t)$, is a Poisson random variable for each t , it is easy to analyze. Its mean $m_\alpha(t) \equiv E[B(t)]$ as a function of the abandonment probability target α is the offered-load function used in this paper. We gave explicit formulas for $m_\alpha(t)$ in §3 and §EC.2.

We found that the DIS mean $m_\alpha(t)$ itself provides an excellent staffing function for low Quality-of-Service (QoS) targets. Indeed, in §4 we proved that it achieves the stabilizing goal asymptotically as the scale increases. However, to obtain a staffing function that works for all QoS levels, we developed a new modified-offered-load (MOL) approximation in §5, obtaining our overall DIS-MOL approximation. As in previous MOL approximations, our MOL approximation exploits the associated stationary model with a constant arrival rate depending on the appropriate offered load. To treat the steady-state of the stationary $M/GI/s + GI$ model, we use an approximation developed in Whitt (2005), which is based on an associated Markovian $M/M/s + M(n)$ model with state-dependent abandonment rates.

Our simulation experiments have shown that the DIS-MOL approximation not only stabilizes expected delays and abandonment probabilities, but it also describes other performance measures, e.g., the expected queue length $E[Q(t)]$. As in Feldman et al. (2008), we find that other performance measures are stabilized to a great extent, but not fully (across a wide range of performance targets). That was illustrated in Figures 4 and 5. Indeed, in Corollary 3 we showed that it is not possible to simultaneously stabilize all performance measures asymptotically in the efficiency-driven many-server heavy-traffic regime. However, just as in Feldman et al. (2008), we find that DIS-MOL simultaneously stabilizes all standard performance measures with higher QoS targets, when the system is in the quality-and-efficiency driven regime.

In §6 we showed that a modification of the classical square-root-staffing (SRS) formula in (2) can be applied for staffing to meet abandonment-probability targets provided that we use the DIS offered-load function $m_\alpha(t)$. However, in general, it remains to determine an appropriate QoS parameter β_α . In §EC.5 we developed an explicit approximation formula for the QoS parameter β_α for the Markovian $M_t/M/s_t + M$ special case; see (EC.1). Its special linear separable structure reveals how performance depends on the model parameters.

We have demonstrated that the DIS-MOL approximation is remarkably effective by performing simulation experiments for both Markovian and non-Markovian models with sinusoidal arrival-rate functions, when the arrival rates are not too small (around 100 with mean service $ES = 1$). In general, the performance of DIS-MOL tends to improve as the scale (arrival rate and number of servers) increases. We have also considered both smaller and larger systems, in particular, for average arrival rates ranging from $s = 20$ to $s = 1000$. The performance of DIS-MOL is spectacular for $s = 1000$ and still reasonable for $s = 20$.

While conducting the simulation experiments, we considered several discretization issues: how to convert a continuous staffing function into an integer-valued staffing function; what is the consequence of agents being required to finish their current services when called to leave; how does the size of a fixed-staffing period affect this approach (discussed in the e-companion).

Much work remains to be done in the future. For example, it remains to establish supporting theory for the DIS-MOL approximation, paralleling Massey and Whitt (1994). So far, we only can conjecture that the DIS-MOL staffing is never less than the DIS staffing. We also need asymptotic results supporting the excellent performance of the DIS-MOL approximation under a wide range of targets. We conjecture that it is asymptotically correct in the QED many-server heavy-traffic regime (in a meaningful way, e.g., that $\sqrt{n}P_t^n(Ab) \rightarrow \alpha$ as $n \rightarrow \infty$, independent of t , where α_n is the target in model n , which is required to satisfy $\sqrt{n}\alpha_n \rightarrow \alpha$ as $n \rightarrow \infty$). It also remains to stabilize performance measures in multi-class multi-pool systems and in systems with different service disciplines.

Acknowledgments.

This research was supported by NSF grants DMI-0457095, CMMI-0948190 and CMMI-1066372.

References

- Billingsley, 1999. *Convergence of Probability Measures*, second ed., Wiley, New York.
- Brown, L., N. Gans, A. Mandelbaum, A. Sakov, H. Shen, S. Zeltyn, L. Zhao. 2005. Statistical analysis of a telephone call center: a queueing-science perspective. *J. Amer. Statist. Assoc.* **100** 36–50.
- Davis, J. L., W. A. Massey, W. Whitt. 1995. Sensitivity to the service-time distribution in the nonstationary Erlang loss model. *Management Sci.* **41** 1107–1116.
- Eick, S. G., W. A. Massey, W. Whitt. 1993a. The physics of the $M_t/G/\infty$ queue. *Oper. Res.* **41** 731–742.
- Eick, S. G., W. A. Massey, W. Whitt. 1993b. $M_t/G/\infty$ queues with sinusoidal arrival rates. *Management Sci.* **39** 241–252.
- Feldman, Z., A. Mandelbaum, W. A. Massey, W. Whitt. 2008. Staffing of time-varying queues to achieve time-stable performance. *Management Sci.* **54** 324–338.
- Feller, W. 1971. *An Introduction to Probability Theory and its Applications*, second ed., Wiley, New York.
- Garnett, O., A. Mandelbaum, M. Reiman. 2002. Designing a call center with impatient customers. *Manufacturing Service Oper. Management* **4** 208–227.
- Green, L. V., P. J. Kolesar, W. Whitt. 2007. Coping with time-varying demand when setting staffing requirements for a service system. *Production and Operations Management* **16** 13–39.

- Jagerman, D. L. 1975. Nonstationary blocking in telephone traffic. *Bell System Tech. J.* **54** 625–661.
- Jennings, O. B., A. Mandelbaum, W. A. Massey, W. Whitt. 1996. Server staffing to meet time-varying demand. *Management Sci.* **42** 1383–1394.
- Kang, W., K. Ramanan. 2010. Fluid limits of many-server queues with reneging. *Ann. Appl. Prob.* **20** 2204–2260.
- Kaspi, H., K. Ramanan. 2011. Law of large numbers limits for many-server queues. *Ann. Appl. Prob.* **21** (2011) 33–114.
- Liu, Y., W. Whitt. 2012a. The $G_t/GI/s_t + GI$ many-server fluid queue. *Queueing Systems* **71** 405–444.
- Liu, Y., W. Whitt. 2012b. A many-server fluid limit for the $G_t/GI/s_t + GI_t$ queueing model experiencing periods of overloading. *Operations Research Letters*, doi:10.1016/j.orl.2012.05.010
- Liu, Y., W. Whitt. 2012c. Many-server heavy-traffic limit for queues with time-varying parameters. Columbia University, NY. <http://www.columbia.edu/~ww2040/allpapers.html>
- Mandelbaum, A., W. A. Massey, M. I. Reiman. 1998. Strong approximations for Markovian service networks. *Queueing Systems* **30** 149–201.
- Massey, W. A., W. Whitt. 1993. Networks of infinite-server queues with nonstationary Poisson input. *Queueing Systems* **13** 183–250.
- Massey, W. A., W. Whitt. 1994. An analysis of the modified offered load approximation for the nonstationary Erlang loss model. *Ann. Appl. Probabil.* **4** 1145–1160.
- Massey, W. A., W. Whitt. 1997. Peak congestion in multi-server service systems with slowly varying arrival rates. *Queueing Systems* **25** 157–172.
- Whitt, W. 1981. Comparing counting processes and queues. *Adv. Appl. Prob.* **14** 207–220.
- Whitt, W. 2005. Engineering solution of a basic call-center model. *Management Sci.* **51** 221–235.
- Whitt, W. 2006. Fluid models for multiserver queues with abandonments. *Operations Research* **54** 37–54.
- Wolff, R. W. 1977. An upper bound for multi-channel queues. *J. Appl. Prob.* **14** 884–888.
- Wolff, R. W. 1987. An upper bound for multi-channel queues. *J. Appl. Prob.* **24** 547–551.
- Zeltyn S., A. Mandelbaum. 2005. Call centers with impatient customers: many-server asymptotics of the $M/M/n+G$ queue. *Queueing Systems* **51**, 361–402.

This page is intentionally blank. Proper e-companion title page, with INFORMS branding and exact metadata of the main paper, will be produced by the INFORMS office when the issue is being assembled.

E-Companion

EC.1. Overview

This e-companion has five more sections. In §EC.2 we extend §3 by giving explicit DIS performance formulas in structured special cases, when the arrival-rate function is sinusoidal and quadratic. In §EC.3 we indicate how the simulation estimates were made and give confidence intervals. In §EC.4 we present results of simulation experiments showing that the DIS-MOL staffing algorithm and the approximations for other performance measures are effective for $M_t/GI/s_t + GI$ models with non-exponential service-time and abandonment-time distributions. In §EC.5 we relate the DIS-MOL staffing for the special case of the $M_t/M/s_t + M$ model to a square root staffing formula. In §EC.6 we consider several discretization issues and real-world constraints for the DIS-OL and DIS-MOL staffing functions. First the staffing levels must be integer-valued, so they must be rounded. Second, when the staffing levels decrease, we do not remove a server until one completes the service in progress. Throughout we assume that server assignments can be switched when a server leaves, so that only the total number of servers matters. As a consequence, when the number of servers is scheduled to decrease by one when all servers are busy, then a server leaves at the time of the next service completion. In addition, we have considered specified staffing intervals, e.g., requiring that staffing changes can only be made on the half hour.

Additional material appears in a longer version available on the authors' web pages. In §EC.7, we supplement previous results in Jennings et al. (1996) and Feldman et al. (2008) showing that the pointwise stationary approximation does not perform well when the mean service time is relatively long. In §EC.8 we present the results from additional simulation experiments to show that the DIS-MOL approximation is effective for stabilizing the abandonment probability and the expected waiting time in $M_t/GI/s_t + GI$ models with non-exponential service-time and abandonment-time distributions. In §EC.9 we compare our SRS approximation for stabilizing the abandonment probability to the associated SRS formula developed by Feldman et al. (2008) for the delay probability, based on the Garnett function in (4). In §EC.10 we report simulation experiments for larger and

smaller $M_t/M/s_t + M$ systems, specifically for the same sinusoidal arrival rate function in (18) for average arrival rates $a = 20$, $a = 50$ and $a = 1000$; the main paper considered the case $a = 100$. Finally, to put the staffing issue in perspective, in §EC.11 we study the sensitivity of the performance to a change of a single server. To do so, we give numerical results for the stationary $M/M/s + M$ model. These show that errors in meeting the target may be partly due to the impact of a single server.

EC.2. DIS Approximations in Structured Special Cases

Since many service systems have daily cycles, it is natural to consider sinusoidal and other periodic arrival-rate functions, as was done in Jennings et al. (1996), Feldman et al. (2008). When we do so, we can apply Eick et al. (1993b) to obtain explicit formulas for the performance measures in that setting. For periodic arrival processes, it is natural to focus upon the dynamic steady state, which prevails because we have started the system empty at time $t = -\infty$. The position within the cycle is determined by simply defining the arrival-rate function over the entire real line, and assuming that it is periodic. In that way, time 0 corresponds to a definite place within a cycle.

THEOREM EC.1. *Consider the DIS-OL approximation for the $M_t/GI/s_t + GI$ model with sinusoidal arrival-rate function $\lambda(t) = a + b \cdot \sin(ct)$ and delay target w . Then $Q(t)$ and $B(t)$ are independent Poisson random variables having sinusoidal means*

$$\begin{aligned} E[Q(t)] &= E[T](a + \gamma \sin(ct - \theta)), \\ &\text{for } \theta \equiv \arctan[\phi_1(T)/\phi_2(T)], \quad \gamma \equiv b\sqrt{\phi_1(T)^2 + \phi_2(T)^2}, \\ E[B(t)] &= \bar{F}(w)E[S](a + \tilde{\gamma} \sin[c(t - w) - \tilde{\theta}]), \\ &\text{for } \tilde{\theta} \equiv \arctan[\phi_1(S)/\phi_2(S)], \quad \tilde{\gamma} \equiv b\sqrt{\phi_1(S)^2 + \phi_2(S)^2}, \end{aligned}$$

where $\phi_1(X) \equiv E[\sin(cX_e)]$, $\phi_2(X) \equiv E[\cos(cX_e)]$ for a nonnegative random variable X .

It is easy to see that the extreme values of $E[Q(t)]$ and $E[B(t)]$ occur at $t_Q = t_\lambda + \theta/c$ and $t_B = t_\lambda + \tilde{\theta}/c$, where $t_\lambda = \pi/2\gamma + n\pi/\gamma$ for n integer are times at which the extreme values of $\lambda(t)$ occurs. And their extreme values are $Q(t_Q) = E[T](a \pm \gamma)$ and $B(t_B) = \bar{F}(w)E[T](a \pm \tilde{\gamma})$. As shown

in Eick et al. (1993b), much nicer formulas are obtained in the special case of exponential service times, but we do not display them.

As discussed in §4 of Eick et al. (1993a), Massey and Whitt (1997) and §4 of Green et al. (2007), for general arrival-rate functions, it is convenient to use simple approximations stemming from linear and quadratic approximations formed by Taylor series approximations. These generate simple time lags and space shifts.

THEOREM EC.2. *Consider the DIS-OL approximation for the $M_t/GI/s_t + GI$ model with quadratic arrival-rate function $\lambda(t) = a + bt + ct^2$ and delay target w . Then $Q(t)$ and $B(t)$ are independent Poisson random variables having quadratic means*

$$E[Q(t)] = E[T]\lambda(t - E[T_e]) + cE[T]Var[T_e],$$

$$E[B(t)] = \bar{F}(w)E[S]\lambda(t - w - E[S_e]) + c\bar{F}(w)E[S]Var[S_e]$$

The quadratic case shows a deterministic time shift which corresponds to interchanging the order of expectation in Theorem 1, and a space shift that is the variance term multiplied by a constant. We obtain expressions for linear arrival-rate functions by simply letting the quadratic coefficient be $c = 0$.

EC.3. More about the Simulation Experiments

We now provide more details about the experiment with the $M_t/M/s_t + M$ example in §6.

EC.3.1. Estimating Performance Measures

We first describe our estimation procedure for the following time-dependent performance measures: (i) the mean potential waiting time, $E[W(t)]$, (ii) the abandonment probability, $P_t(Ab)$, (iii) the delay probability, $P_t(Delay)$, and (iv) the mean queue length, $E[Q(t)]$.

We estimated these performance measures in a time interval $[0, T]$ with $T = 20$. First, for $E[W(t)]$, we kept track of all customer arrivals in each sample path. For a customer n , we keep track of the arrival time, A_n , and the time that the customer enters service, E_n . Therefore, one value for this

sample path is $(t, \hat{W}(t)) = (A_n, E_n - A_n)$. Of course, this customer may have already abandoned by time E_n . Since we are interested in the *potential waiting time*, assuming infinite patience, we keep track of the time that the customer would enter service even after it abandons; i.e., our procedure includes the behavior of virtual customers.

We divide the interval $[0, 20]$ into subintervals or “bins.” The bin size for $E[W(t)]$ is 0.1. For the performance measures $P_i(Ab)$ and $P_i(Delay)$, we used a larger bin size because there are less sample points. In particular, we used a bin size of 0.25. At each sample path, for the interval $[0, 0.25]$ (the procedures remain the same for other intervals $[0.25, 0.5], \dots, [19.75, 20]$), we obtained an estimate by computing the ratios $A(0, 0.25)/N(0, 0.25)$ and $D(0, 0.25)/N(0, 0.25)$, where $A(t, t + \Delta t)$ is the total number of abandonments among arrivals in interval $[t, t + \Delta t]$, $D(t, t + \Delta t)$ is the number of delays encountered by arrivals in interval $[t, t + \Delta t]$, and $N(t, t + \Delta t)$ is the total number of arrivals in $[t, t + \Delta t]$. For the mean queue length, $E[Q(t)]$, we used a bin size 0.05. In particular, we sampled the queue length once every 0.05 units of time.

EC.3.2. Confidence Intervals

We estimated the above performance measures by generating independent replications. For the example in §6 we used 5000 independent sample paths. To verify the statistical precision of our estimates, we also provide the 95% confidence intervals. Let $z_{0.05} \equiv 1.96$ be the 95% percentile of the standard normal distribution, the confidence interval is $[\bar{\mu} - 1.96 S_n/\sqrt{n}, \bar{\mu} + 1.96 S_n/\sqrt{n}]$, where $\bar{\mu} \equiv (1/n) \sum_{i=1}^n X_i$ is the sample mean and $S^2 \equiv 1/(n-1) \sum_{k=1}^n (X_i - \bar{\mu})^2$ is the sample variance. (Since the sample size is so large, the t distribution associated with unknown variance is approximately normal.) Let $\epsilon_{P(AB)}$ and $\epsilon_{E[W]}$ be the half length of the confidence intervals. These are shown in the Table EC.3.2 for different targets α^* and w^* . Since the performance is stabilized, these apply approximately at each t . The values in Table EC.3.2 are averages over $t \in [0, T]$. The plotted performance functions appear even smoother, because there is dependence among successive values.

α^*	0.2	0.15	0.1	0.05	0.02	0.01	0.005
$\bar{\mu}_{P(Ab)}$	0.20	0.156	0.0899	0.0547	0.0202	0.0098	0.0051
$\epsilon_{P(Ab)}$	0.0045	0.0042	0.0034	0.0029	0.0012	0.0008	0.00056
$w^* = F^{-1}(\alpha^*)$	0.45	0.033	0.21	0.11	0.041	0.020	0.010
$\bar{\mu}_{E[W(t)]}$	0.449	0.343	0.194	0.129	0.0388	0.0183	0.0093
$\epsilon_{E[W(t)]}$	0.0186	0.0162	0.0119	0.0099	0.0054	0.0037	0.0027

Table EC.1 Confidence intervals for estimates of $E[W(t)]$ and $P_t(Ab)$ at different targets α^* and w^* based on 5000 independent replications.




Figure EC.1 Time-dependent confidence interval for $E[W(t)]$ and $P_t(Ab)$ with $\alpha^* = 0.2$.

To illustrate the consequence, in Figure EC.1 we also plot the variance envelope (i.e., confidence interval) for $\alpha = 0.2$. The red dashed curves are the targets, the solid light blue curves are means (i.e., mid points of the confidence intervals) and the dashed blue curves are the confidence intervals.

EC.4. Non-Exponential Service and Patience Distributions

The DIS-MOL approximation also performs well for the more general $M_t/GI/s_t + GI$ model with non-exponential service and patience distributions. We illustrate by displaying results from a simulation experiment for the $M_t/H_2/s_t + E_2$ model.

Let A and S be the generic abandon and service times. Here we fix their means $E[A] \equiv 1/\theta$ and $E[S] \equiv 1/\mu$ with $\mu \equiv 1$ and $\theta \equiv 0.5$. For the Erlang distribution, we consider Erlang-2 (E_2), the sum of two i.i.d. exponential random variables, which has squared coefficient of variation (SCV) $c_X^2 \equiv Var(X)/E[X]^2 = 1/2$. For the Hyper-exponential (H_2) distribution, X is a mixture of two




Figure EC.2 Simulation estimates of expected queue length, abandonment probability and expected delay for the $M_t/H_2/s_t + E_2$ under the DIS-OL staffing, with the system heavily loaded ($0.05 \leq \alpha \leq 0.20$).

independent exponential random variables, i.e., X has complementary cdf $1 - G(x) = p \cdot e^{-\lambda_1 x} + (1 - p) \cdot e^{-\lambda_2 x}$. We choose $p = 0.5(1 - \sqrt{0.6})$, $\lambda_1 = 2p/E[X]$, $\lambda_2 = 2(1 - p)/E[X]$, which yields $c_X^2 = 4$ (and balanced means).

We use the same sinusoidal arrival rate function in (18) with the same parameters: $a = 100$, $b = 20$, $c = 1$. Just as for the Markovian example in §6, here we plot the simulation estimations of




Figure EC.3 Simulation estimates of expected queue length, abandonment probability and expected delay for the $M_t/H_2/s_t + E_2$ under the DIS-MOL staffing, with the system lightly loaded ($0.01 \leq \alpha \leq 0.05$).

the expected queue lengths, abandonment probabilities and expected delays for different abandon probability targets α in a 20-hour day for the $M_t/D/s_t + H2$. Just as before, our estimates are based on 5000 independent replications. Estimates of the confidence intervals appear in Table EC.4.

α^*	0.2	0.15	0.1	0.05	0.02	0.01	0.005
$\bar{\mu}_{P(Ab)}$	0.1904	0.1408	0.0931	0.0506	0.0191	0.0093	0.0041
$\epsilon_{P(Ab)}$	0.0024	0.0022	0.0020	0.0015	0.0011	0.0009	0.0007
$w^* = F^{-1}(\alpha^*)$	0.825	0.684	0.532	0.356	0.215	0.149	0.104
$\bar{\mu}_{E[W(t)]}$	0.79	0.64	0.49	0.32	0.201	0.138	0.093
$\epsilon_{E[W(t)]}$	0.0245	0.0221	0.0191	0.0153	0.0127	0.0103	0.0092

Table EC.2 Confidence intervals for estimates of $E[W(t)]$ and $P_t(Ab)$ of the $M_t/H2/s_t + E_2$ model, at different targets α^* and w^* based on 5000 independent replications.

EC.5. A New Square-Root Staffing (SRS) Formula

At the end of §6, we introduced the new SRS formula (21) for the $M_t/GI/s_t + GI$ model based on the DIS OL with abandonment target α . Formula (21) differs from (2) by using the DIS OL $m_\alpha(t)$ instead of the standard OL $m_0(t)$ in (1).

We now verify that the DIS-MOL approximation is indeed consistent with the SRS staffing in (21); i.e., we show that the DIS-MOL staffing function $s_\alpha^{MOL}(t)$ has the form of the SRS formula (21). Following Feldman et al. (2008), for an abandonment probability target α , we let $D_\alpha(t) \equiv s_t^{MOL} - m_\alpha(t)$ be the difference between the DIS-MOL staffing and the OL functions, and let $\beta_\alpha(t) \equiv D_\alpha(t)/\sqrt{m_\alpha(t)}$ be the *implicit QoS function* for DIS-MOL. In Figure EC.4, we plot the implicit QoS function $\beta_\alpha(t)$ for different targets α for the Markovian example with sinusoidal arrival rate in §6. Figure EC.4 shows that, just as in Feldman et al. (2008), except for an initial transience, $\beta_\alpha(t)$ is independent of time and decreasing in α .

Figure EC.4 shows that there exists a QoS function β_α , which is positive, strictly decreasing in α with the other parameters fixed, that makes the SRS function perform as well as DIS-MOL in stabilizing the abandonment probability and the expected waiting time. For all $\alpha \geq 0.1$, we found that $\beta_\alpha \approx 0$, consistent with Figure 3 and our observation more generally, that the DIS-MOL staffing always exceeds the DIS staffing. However, it remains to determine an explicit formula for




Figure EC.4 The implicit QoS function $\beta_\alpha(t) \equiv (s_\alpha(t) - m_\alpha(t))/\sqrt{m_\alpha(t)}$ for different abandonment probability target α , in the $M_t/M/s_t + M$ example, where $a = 100$, $b = 20$, $c = 1$, $\theta = 0.5$, $\mu = 1$.

the appropriate function β_α (even though it could be obtained numerically from plots like those in Figure EC.4).

We did statistical fitting to determine an approximation for the QoS function β_α to use in (21) for the $M_t/M/s_t + M$ model with a sinusoidal arrival rate function like (18) as a function the delay target α and the model parameters $(\lambda(t), \mu, \theta)$. First, we found that the dependence upon α tends to be linear in $-\log(\alpha)$. Second, we found that it sufficed to consider the impact of $\lambda(t)$ via its time average $\bar{\lambda}$. These fitting experiments lead us to propose the following linear approximation (in $-\log(\alpha)$) as an engineering solution:

$$\beta(\alpha, \bar{\lambda}, \theta, \mu) \approx K \cdot (-\log(\alpha)) + C(\bar{\lambda}/\mu, \theta/\mu), \quad (\text{EC.1})$$

$$\text{where } C(\bar{\lambda}/\mu, \theta/\mu) \equiv -K \log(\sqrt{\bar{\lambda}/\mu}) + H(\theta/\mu),$$

where constant $K = 0.43$, $\bar{\lambda}$ is the average arrival rate, and $H(x)$ is the linear function $H(x) = 0.2822x - 0.0133$. Figure EC.5 shows both β_α obtained from the DIS-MOL approximation and its linear approximation (EC.1), with different $\bar{\lambda}/\mu$ and θ/μ . We fix $\mu = 1$. Since the arrival is




Figure EC.5 Comparison of the linear approximation in (EC.1) for β_α to the value of β_α obtained from the DIS-MOL approximation.

sinusoidal, $a = \bar{\lambda}$ is the average arrival rate. This approximation is simple because (i) it is linear in $-\log(\alpha)$ when the other parameters are fixed, (ii) and the dependence on $\bar{\lambda}/\mu$ and θ/μ is separable. Figure EC.5 shows that linear approximation (EC.1) is good when a is not too small ($a > 10$) so that we are in the many-server regime, and when θ/μ is not too big ($\theta/\mu < 4$), where the model begins to behave as a loss model.

EC.6. Discretization Issues and Real-World Constraints

In this section we consider several discretization issues and real-world constraints for staffing functions. First, our basic method produces continuous staffing levels, but we have to staff at integer levels. To be conservative, we could always round up, but to obtain greater accuracy, we round to the nearest integer. That does not affect the performance much.

When the staffing schedule calls for decreasing the number of servers, we impose a real system constraint: We do not release any server until that server completes the service in process. However, we make an additional assumption that greatly helps achieve the required staffing level: We assume

that server assignments can be switched when a server is scheduled to leave. In other words, if a specific server is scheduled to depart, then any other available server can complete the service. That means that we must only wait until the next service completion by *any* server when all servers are busy and the staffing level is scheduled to decrease by one.

This realistic feature makes the actual staffing level always lie above the integer-valued scheduled staffing level, as shown in Figure EC.24. As a consequence, the system will be slightly overstaffed when the staffing function $s(t)$ decreases. However, this effect tends to be insignificant when there are a large number of servers, provided we allow server switching as assumed above. (The number of servers ranges from 70 to 120 in the Markovian example considered in §6.) If the service-time distribution is exponential and all servers are busy when the number of servers is called to decrease, the lack-of-memory property of the exponential distribution implies that the time lag between the actual jump-down time of $s(t)$ and the scheduled one is $1/n\mu$, if there are currently n servers in the system.

We can propose methods to counter this effect. In particular, during periods of planned staffing decrease, we can deliberately set the staffing level a bit lower to anticipate for that lag in releasing servers. One direct approach is to re-schedule a new staffing function which takes into consideration these time lags since they can easily be quantified. In §6 we don't have to modify the DIS and DIS-MOL staffing functions for the $M_t/M/s_t + M$ example with sinusoidal arrival since this effect is insignificant, as shown in Figure 4 and 5.

It is also interesting to consider another realistic constraint for the staffing function of the DIS-MOL approximation: Many real service systems have staffing periods with constant staffing levels. For example, it is common to only change staffing levels every half hour or every hour. We now consider this case. It is evident that the smaller the staffing period is, the less accuracy we should lose because of the requirement that the staffing function remain constant over each staffing interval.

In Figure EC.6 (and Figures EC.25 and EC.26 in the longer version), we plot the abandonment




Figure EC.6 Simulation estimates for the $M_t/M/s_t + M$ example with sinusoidal arrival under DIS-MOL staffing with fixed staffing period $d = 30$ minutes and 2 hours, $\alpha = 0.1$.

probabilities, expected delays with abandonment probability target $\alpha = 0.1, 0.05$ and 0.02 , respectively for two cases: fixed staffing periods of 2 hours and 30 minutes. In particular, we change the staffing once every 30 minutes (or 2 hours) by setting the constant number of servers in each period equal to the midpoint of the proposed continuous staffing function. When the staffing period is long (2 hours), fluctuations brought by the crudeness of the discreteness of the staffing functions are unavoidable. However, these performance measures are still quite stable except during the initial

transient periods, which is caused by the steep rise of the DIS-MOL staffing function.

The rest of this e-companion only appears in the authors' longer online version.

EC.7. The Pointwise Stationary Approximation (PSA)

A standard approximation for queues with time-varying arrival rates is the pointwise-stationary approximation (PSA); see Green and Kolesar (1991) and Whitt (1991). The general idea is to use steady-state performance measures associated with the corresponding stationary model in a time-varying way. For the $M_t/GI/s_t + GI$ model considered here, that means using the steady-state performance measures associated with the stationary $M/GI/s + GI$ model computed at time t as if the constant arrival rate λ in the stationary model is always $\lambda(t)$, the value of the time-varying arrival rate at time t , and the constant staffing s in the stationary model is always $s(t)$, the value of the time-varying staffing at time t . When the staffing is to be determined, only the arrival rate is used, and the staffing at time t is determined to meet the requirement, using the steady-state performance measures associated with the stationary $M/GI/s + GI$ model, with this special arrival rate.

The PSA is well known to be effective when service times are relatively short, but not when the service times are relatively long. Examples were given in Jennings et al. (1996) and Feldman et al. (2008). We give an example here in the context of the example in §6. We let the abandon probability target α be 0.1. That corresponds to a target delay of $w = (-1/\theta) \log(1 - \alpha) \approx 0.2107$.

Figure EC.7 shows that simulation estimates of the expected delay and abandonment probability (the solid lines) under PSA are significantly different from their targets (the dashed lines). One can see that the system alternates between overstaffed and understaffed periodically, although the time average over a day is close. Figure EC.7 shows a time lag between the peaks of the arrival rate and the performance measures.

EC.8. More Experiments with Non-Exponential Distributions

We now complement §EC.4 by presenting additional simulation results to show that the DIS-MOL approximation is effective for non-Markovian $M_t/GI/s_t + GI$ models with non-exponential service and patience distributions.




Figure EC.7 Simulation estimates of the abandonment probability and expected delay with PSA staffing for the $M_t/M/s_t + M$ example with the sinusoidal arrival rate, $\alpha = 0.1$ and $w = -(1/\theta) \log(1 - \alpha) = 0.21$.

In particular, we consider three models: (i) deterministic service times and hyper-exponential abandonment times ($M_t/D/s_t + H2$); (ii) Erlang service distribution and exponential abandonment times ($M_t/E2/s_t + M$); (iii) Hyper-exponential service times and exponential abandonment times ($M_t/H2/s_t + M$), where the E_2 and H_2 distributions are the same as before.

We use the same sinusoidal arrival rate function in (18) with the same parameters: $a = 100$, $b = 20$, $c = 1$. Just as for the Markovian example in §6, here we plot the simulation estimations of the expected queue lengths, abandonment probabilities and expected delays for different abandon probability targets α in a 20-hour day for the $M_t/D/s_t + H2$ (in Figures EC.8 and EC.9), $M_t/H2/s_t + M$ (in Figures EC.10 and EC.11) and $M_t/E2/s_t + M$ (in Figures EC.12 and EC.13) models. As before, the figures show that the DIS-MOL approximation consistently performs well.

The plots here are based on 500 independent sample paths in each case. Paralleling Table EC.3.2, Table EC.8 below shows the corresponding confidence intervals in Figures 4 and 5. These serve as rough estimates of the confidence intervals in these new plots.

α^*	0.2	0.15	0.1	0.05	0.02	0.01	0.005
$\bar{\mu}_{P(Ab)}$	0.20	0.156	0.091	0.054	0.020	0.010	0.0055
$\epsilon_{P(Ab)}$	0.014	0.013	0.011	0.0091	0.0037	0.0025	0.0017
$w^* = F^{-1}(\alpha^*)$	0.45	0.033	0.21	0.11	0.041	0.020	0.010
$\bar{\mu}_{E[W(t)]}$	0.44	0.34	0.19	0.13	0.038	0.019	0.010
$\epsilon_{E[W(t)]}$	0.059	0.051	0.038	0.032	0.017	0.012	0.0082

Table EC.3 Confidence intervals for estimates of $E[W(t)]$ and $P_t(Ab)$ at different targets α^* and w^* based on 500 independent replications.




Figure EC.8 Simulation estimates of expected queue length, abandonment probability and expected delay for the $M_t/D/s_t + H2$ under the DIS-OL staffing, with the system heavily loaded ($0.05 \leq \alpha \leq 0.20$).

EC.9. Comparison with Feldman et al. (2008)

Feldman et al. (2008) showed that the delay probability can be stabilized under a different SRS staffing function, namely,

$$s_0(t) = m_0(t) + \beta_\gamma \sqrt{m_0(t)}, \quad (\text{EC.2})$$

where $m_0(t)$ is the offered-load (OL) function of the $M_t/GI/\infty$ IS model, given in (1); i.e., $m_0(t) = E[\lambda(t - S_e)]E[S]$, where $\lambda(t)$ is the arrival rate and S_e is a random variable with the stationary-excess distribution associated with the service time S as in (6). In other words, $m_0(t)$ is just our

$m_\alpha(t)$ in (8) with $\alpha = w = 0$.

In Feldman et al. (2008), the QoS parameter β_γ is a function of the delay probability target γ . Moreover, this β_γ can be approximated by the Garnett function given in (4). In order to stabilize the delay probability at target γ , it suffices to obtain a β_γ by inverting the Garnett function, letting $P(\text{Delay}) \equiv \gamma$.

We now compare the performance of our DIS-MOL approximation (the dashed lines) to the performance of the delay-probability offered-load (P(Delay)-OL) staffing (EC.2) (the solid lines) in Figure EC.14. We plot simulation estimates for expected delay $E[W]$, abandonment probability $P(\text{Ab})$ and delay probability $P(\text{Delay})$ of the Markovian model considered in §6 under the staffing functions of these two approximations, with abandonment probability targets $\alpha = 0.1, 0.05, 0.02$. We first treat the time averages of simulation estimates of $P(\text{Delay})$ obtained from DIS-MOL as the delay target γ (the bold numbers in Table EC.9) and we invert function G_1 in (4) with γ to obtain β_γ , then we use P(Delay)-OL staffing function (EC.2).

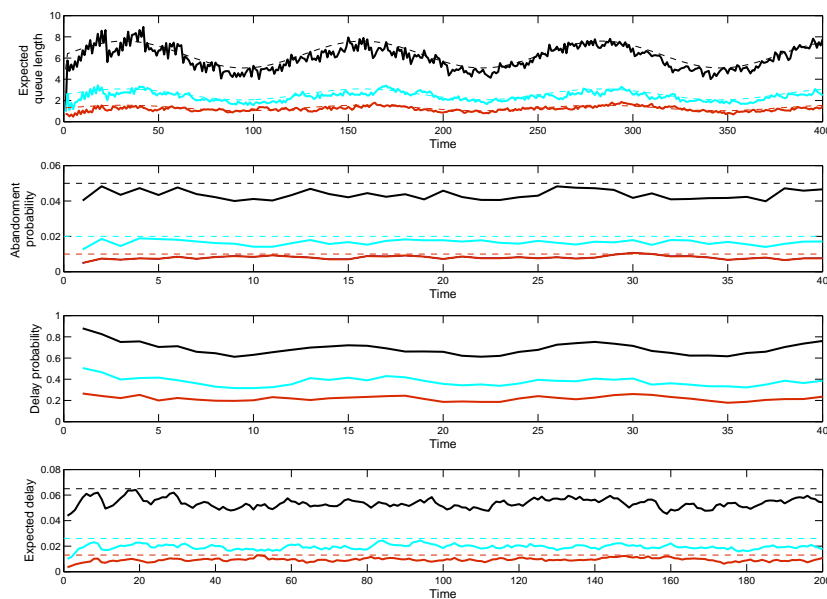


Figure EC.9 Simulation estimates of expected queue length, abandonment probability and expected delay for the $M_t/D/s_t + H2$ under the DIS-MOL staffing, with the system lightly loaded ($0.01 \leq \alpha \leq 0.05$).




Figure EC.10 Simulation estimates of expected queue length, abandonment probability and expected delay for the $M_t/H2/s_t + M$ under the DIS-OL staffing, with the system heavily loaded ($0.05 \leq \alpha \leq 0.20$).




Figure EC.11 Simulation estimates of expected queue length, abandonment probability and expected delay for the $M_t/H2/s_t + M$ under the DIS-MOL staffing, with the system lightly loaded ($0.01 \leq \alpha \leq 0.05$).




Figure EC.12 Simulation estimates of expected queue length, abandonment probability and expected delay for the $M_t/E2/s_t + M$ under the DIS-OL staffing, with the system heavily loaded ($0.05 \leq \alpha \leq 0.20$).

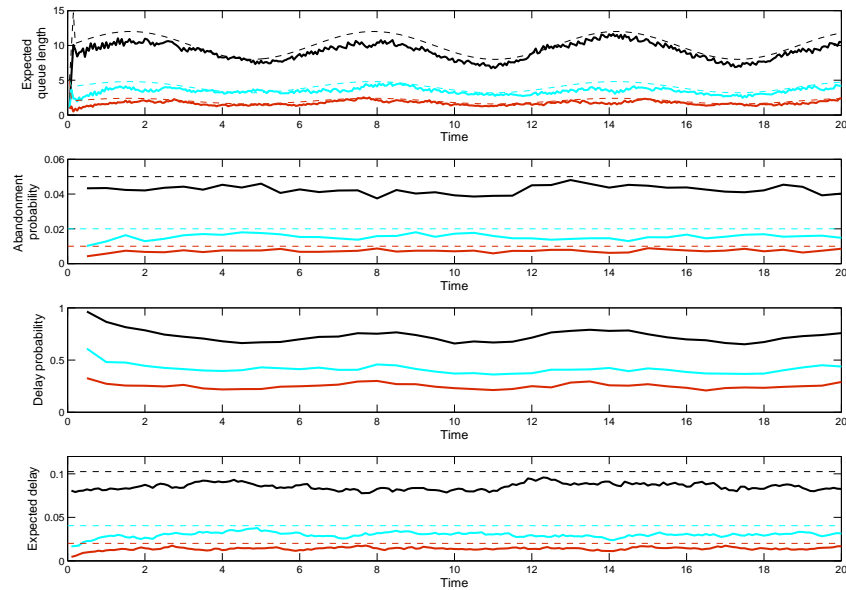


Figure EC.13 Simulation estimates of expected queue length, abandonment probability and expected delay for the $M_t/E2/s_t + M$ under the DIS-MOL staffing, with the system lightly loaded ($0.01 \leq \alpha \leq 0.05$).

Figure EC.14 shows that the DIS-MOL staffing function works better in stabilizing $P(Ab)$ and $E[W]$, while the P(Delay)-OL staffing function works better in stabilizing $P(Delay)$, when $\alpha = 0.1, 0.05$, which is consistent with the different targets used: Here we aim to stabilize $P(Ab)$ and $E[W]$; Feldman et al. (2008) aim to stabilize $P(Delay)$. Figures 3 and 4 in Feldman et al. (2008) show that the delay-probability can be stabilized over a wide range of targets, but in doing so, the abandonment probability is not stabilized under heavy loads. Here, with out different method aimed at the abandonment probability, we find that it can be stabilized over a wide range of targets, but in doing so, the delay probability is not stabilized under heavy loads. From Figure EC.14 we see that when $\alpha = 0.02$ both methods work well. This observed “duality” is consistent with the




Figure EC.14 Simulation estimates for the $M_t/M/s_t + M$ example with sinusoidal arrival (18), under two staffing functions: (i) DIS-MOL staffing $s_\alpha^{MOL}(t)$ (the dashed lines); (ii) P(Delay)-OL staffing $s_0(t)$ in (EC.2) (the solid lines). In (EC.2) we obtain β_γ by inverting function G_1 in (4) letting delay probability target γ be the time average of $P(Delay)$ under DIS-MOL. The abandonment probability targets are $\alpha = 0.1, 0.05, 0.02$ (light straight lines), the other parameters are fixed: $a = 100$, $b = 20$, $c = 1$, $\theta = 0.5$, $\mu = 1$.

heavy-traffic theory. The limits in Garnett et al. (2002) are for the quality-and-efficiency-driven (QED) regime. As α increases, we tend to shift to the different efficiency-driven (ED) regime, as in Whitt (2004, 2006). In §4 we saw that our DIS method is asymptotically correct in ED regime, but then our DIS approaches the OL as α decreases, so it should be consistent with Feldman et al. (2008) under lighter loads.

In Table EC.9, we give more detailed measurements. Here for each curve, we compute its time average, which characterizes how far it is from its target, and the difference between its maximum and minimum which quantifies the magnitude of its fluctuation. Here Both (max-min) and time average are obtained by sampling $t \in [3, 20]$, to avoid initial transient. In Table EC.9, measurements under column 'DIS-MOL' and 'Garnett 1' correspond to the DIS-MOL and P(Delay)-OL staffing (EC.2). In Table EC.9, the abandonment probability targets are $\alpha = 0.1, 0.05, 0.02$; the other parameters are fixed: $a = 100$, $b = 20$, $c = 1$, $\theta = 0.5$, $\mu = 1$.

We see that as the abandonment probability target α decreases, the relative error, i.e., (max-min)/average, increases. The reason is that, as α goes to 0, the error becomes more and more sensitive to the discreteness of the staffing function. In other words, adding or removing a single server affects the performance more when α is smaller.

			Garnett 2	Garnett 1	DIS-MOL
$\alpha = 0.1$	$P(Ab)$	Time avg.	0.111	0.101	0.101
		Target - time avg.	-0.011	-6.9E-04	-0.0013
		Max - Min	0.044	0.041	0.013
		(Max-Min)/avg.	0.399	0.407	0.129
	$E[W]$	Time avg.	0.237	0.213	0.213
		Target - time avg.	-0.026	-0.002	-0.0022
		Max - Min	0.110	0.093	0.029
		(Max-Min)/avg.	0.464	0.437	0.137
	$P(Delay)$	Time avg.	0.944	0.928	0.929
		Target - time avg.	-0.0148	0.0004	0
		Max - Min	0.032	0.044	0.083
		(Max-Min)/avg.	0.034	0.048	0.089
$\alpha = 0.05$	$P(Ab)$	Time avg.	0.050	0.044	0.046
		Target - time avg.	9.6E-05	0.0056	0.0042
		Max - Min	0.022	0.019	0.014
		(Max-Min)/avg.	0.432	0.430	0.303
	$E[W]$	Time avg.	0.102	0.090	0.093
		Target - time avg.	9.8E-04	0.013	0.010
		Max - Min	0.053	0.041	0.026
		(Max-Min)/avg.	0.520	0.454	0.284
	$P(Delay)$	Time avg.	0.726	0.692	0.696
		Target - time avg.	0.011	0.0038	0
		Max - Min	0.091	0.068	0.114
		(Max-Min)/avg.	0.125	0.098	0.163
$\alpha = 0.02$	$P(Ab)$	Time avg.	0.020	0.019	0.018
		Target - time avg.	-3.1E-04	0.0021	0.0016
		Max - Min	0.011	0.0083	0.0073
		(Max-Min)/avg.	0.52	0.46	0.40
	$E[W]$	Time avg.	0.040	0.035	0.036
		Target - time avg.	7.7E-05	0.0054	0.0042
		Max - Min	0.018	0.019	0.017
		(Max-Min)/avg.	0.45	0.54	0.46
	$P(Delay)$	Time avg.	0.437	0.408	0.412
		Target - time avg.	0.013	0.004	0
		Max - Min	0.072	0.095	0.103
		(Max-Min)/avg.	0.164	0.233	0.251

Table EC.4 Simulation estimates for the $M_t/M/s_t + M$ example with sinusoidal arrival rate in (18), under three staffing functions: (i) DIS-MOL staffing $s_\alpha^{MOL}(t)$ (column 'DIS-MOL'); (ii) P(Delay)-OL staffing $s_0(t)$ in (EC.2), where we obtain β_γ by inverting function G_1 in (4), letting the target γ be the time-average $P(Delay)$ under DIS-MOL shown in bold numbers (column 'Garnett 1'); (iii) the same SRS formula in (EC.2), but with β_α instead of β_γ , where β_α is obtained by inverting function G_2 in (EC.3) with target α (column 'Garnett 2').

We now explain the second Garnett approximation in Table EC.9. In Garnett et al. (2002), in




Figure EC.15 The QoS parameter β_α as a function of $-\log \alpha$ in the DIS-MOL approximation (red), in approximation (EC.3) (blue), and its linear approximation based on (EC.1) (black dashed) for the $M_t/M/s_t + M$ model with sinusoidal arrival. The other parameters fixed: $a = 100$, $b = 20$, $c = 1$, $\theta = 0.5$, $\mu = 1$.

addition to the heavy-traffic approximation for $P(\text{Delay})$, there is also a heavy-traffic approximation for the abandonment probability in the $M/M/s + M$ model:

$$P(\text{Ab}) \approx \left[1 - \frac{h(\beta\sqrt{\mu/\theta})}{h(\beta\sqrt{\mu/\theta} + \sqrt{\theta/(N\mu)})} \right] \cdot \omega(-\beta, \sqrt{\mu/\theta}) \equiv G_2(\beta), \quad (\text{EC.3})$$

where $N \equiv \lambda/\mu + \beta\sqrt{\lambda/\mu}$, and λ is the constant arrival rate for the $M/M/s + M$ model.

Instead of inverting G_1 in (4) with a delay probability target γ (which we here obtain from the average delay under DIS-MOL), we can invert G_2 in (EC.3) to obtain β_α for a given abandonment probability target α , and apply P(Delay)-OL staffing in (EC.2) based on the OL $m_0(t)$ with QoS parameter β_α . The simulation estimates are under column 'Garnett 2' of Table EC.9. Table EC.9 shows that 'Garnett 2' performs nearly the same as 'Garnett 1.' This is consistent with the results here and in Feldman et al. (2008): In order to stabilize the abandonment probability under heavy loads, it is necessary to use our DIS offered load $m_\alpha(t)$ instead of the direct OL $m_0(t)$.

Figure EC.4 shows that there exists a function β_α , which is strictly decreasing in α with the

other parameters fixed, that makes the SRS function perform as well as DIS-MOL in stabilizing the abandonment probability and the expected waiting time. We now proceed to study this function. Figure EC.15 shows the values of β_α with $0.001 \leq \alpha \leq 0.4$ in natural log scale (the red curve) for the Markovian example with sinusoidal arrival rate function with $a = 100$, $b = 20$, $c = 1$. Figure EC.4 shows that the function β_α is approximately a linear function of $-\log(\alpha)$ (thus exponential with α) for small α ($\alpha < 0.1$). When α is relatively large ($\alpha > 0.1$), β_α is nearly constant (close to 0) which is consistent with Figure 3, i.e., the simple DIS-OL approximation works well when the system is in heavy load.

In Figure EC.15 we also plot the approximation (EC.3) in Garnett et al. (2002) (the blue curve). It is seen that DIS-MOL is close to (EC.3) for $0.001 < \alpha < 0.1$, where the system may be regarded as being in the QED regime.

We also want to characterize how β_α depends on the other model parameters. First, experiments show that β_α depends upon the arrival rate function $\lambda(t)$ only through its time average $\bar{\lambda} \equiv (1/T) \int_0^T \lambda(t) dt$. Next, we can eliminate μ , i.e., we can let it be 1 by choosing appropriate measurement units. That is equivalent to dividing $\bar{\lambda}$ and θ by μ in general.

Therefore, we consider $\beta(\alpha, \bar{\lambda}, \theta, \mu) = \beta(\alpha, \bar{\lambda}/\mu, \theta/\mu, \mu) = \beta(\alpha, \bar{\lambda}, \theta)$, which is independent of μ , as a function of the abandonment probability target α , the average scaled arrival rate $\bar{\lambda}/\mu$, and the scaled abandonment rate θ/μ . Therefore, we were led to propose the approximation given in (EC.1). We had initially proposed a third-order polynomial for $H(\cdot)$, namely, $H(x) = p_1 x^3 + p_2 x^2 + p_3 x + p_4$ with $p_1 = 0.0024$, $p_2 = -0.012$, $p_3 = 0.28$, $p_4 = 0.011$. However, a referee persuaded us that the linear approximation for H was just as good.

In Figure EC.5 we plot both β_α obtained from the DIS-MOL approximation and its linear approximation (EC.1), with different $\bar{\lambda}/\mu$ and θ/μ . Here we still use the Markovian example with sinusoidal arrival considered in §6. We fix $\mu = 1$. Since the arrival is sinusoidal, $a = \bar{\lambda}$ is the average arrival rate. Figure EC.5 shows that linear approximation (EC.1) is good when a is not too small ($a > 10$) so that we are in the many-server regime, and when θ/μ is not too big ($\theta/\mu < 4$), where the model begins to behave as a loss model.

EC.10. Small Arrival Rate and Number of Servers

In this final section we investigate how the DIS-OL and DIS-MOL approximations perform for smaller systems, when the average arrival rate (and thus the number of servers) is small. We also consider a very large system to give an overall perspective on scale.

We again consider the Markovian $M_t/M/s_t + M$ example with sinusoidal arrival rate ($\lambda(t) = a + b \cdot \sin(c \cdot t)$). As before, we use the parameters $\mu = 1$ and $\theta = 1/2$. In Figure EC.16-EC.23, we plot simulation estimates with parameters $a = 20, 50, \text{ and } 1000, b = 0.2a, c = 1$. In each case, we consider the system heavily and lightly loaded.

When $a = 20$ and 50 , the number of servers is around 20 and 50 . Performance measures (expected delay and abandonment probability) have bigger fluctuations than the case $a = 100$. But when $a = 1000$, all performance measures coincide with their targets nearly perfectly.




Figure EC.16 Simulation estimates for the $M_t/M/s_t + M$ example with sinusoidal arrival rate function in (18) when $a = 20$, $b = 4$, $c = 1$. The system is heavily loaded ($\alpha = 0.05, 0.10, 0.15, 0.20$).

For the smaller system with $a = 20$, Figures EC.16 and EC.18 show that the realized abandonment probabilities are stabilized at a level somewhat below the target. It is natural to wonder if that is due to the effect of a single server (and so is actually the correct staffing) or if instead it is a genuine error due to the smaller scale.

In order to address this question, we consider the simple refinement of using one server less, i.e., we replace $s(t)$ by $s(t) - 1$. Indeed, that can be achieved by rounding down instead of rounding up. The results of this modification are shown in Figures EC.17 and EC.19. We see that the performance sometimes exceeds the target with staffing by $s(t) - 1$, but it is actually much closer. Such a modification is easy to consider in applications when verifying the results by simulation.




Figure EC.17 Simulation estimates for the $M_t/M/s_t + M$ example with sinusoidal arrival rate function in (18) with staffing $s(t) - 1$, when $a = 20$, $b = 4$, $c = 1$. The system is heavily loaded ($\alpha = 0.05, 0.10, 0.15, 0.20$).

We also remark that we see that a single server does have a large impact. We explore this issue further in §EC.11.

EC.11. Sensitivity of Performance to a Single Server

In order to evaluate the performance of alternative staffing algorithms, it is useful to understand the impact of a single server. We should not expect more precision than the effect of one server. We could staff too high (so that the performance is better than the target), but one server less might fail to satisfy the target constraint. In that case, we have actually achieved our goal, even though the simulation shows some error.

In order to gain insight into the sensitivity of the abandonment probability $P(Ab)$ to a single server, we give exact numerical results for the stationary $M/M/s + M$ model in Table EC.11. In particular, we do a sensitivity analysis showing the impact of a single server as a function of the target α and the arrival rate λ . We consider the stationary $M/M/s + M$ model with abandonment probability targets $\alpha = 0.2, 0.1, 0.01$ and 0.005 , and arrival rates $\lambda = 20, 100$ and 1000 . The other parameters are fixed: $\mu = 1, \theta = 0.5$.

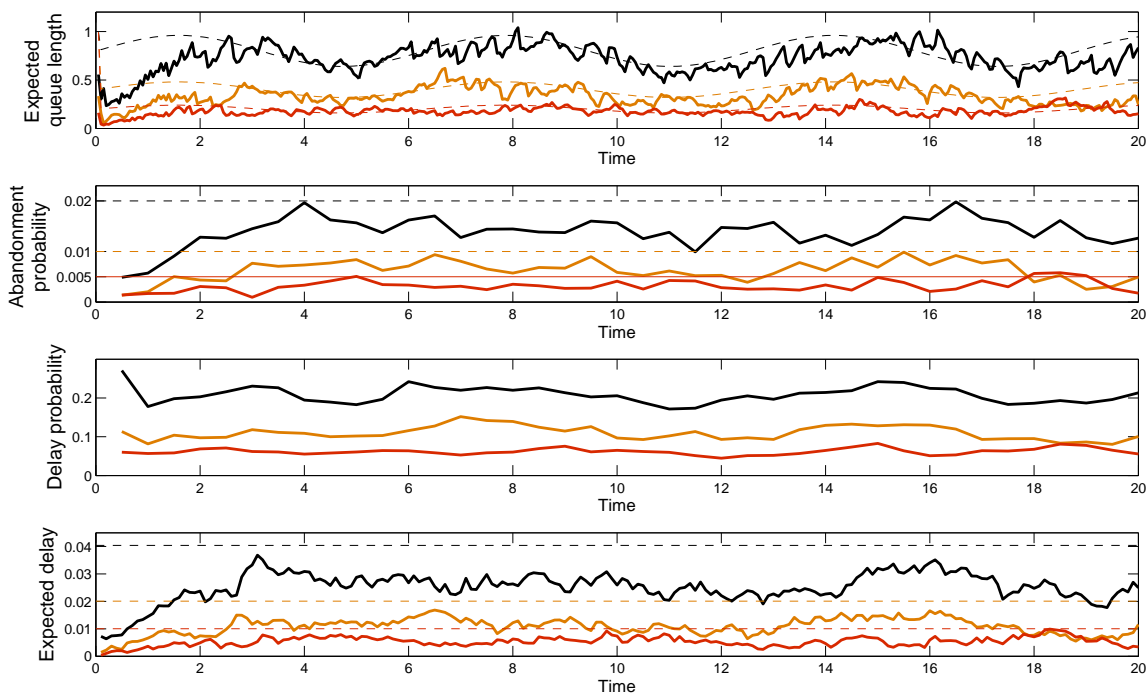


Figure EC.18 Simulation estimates for the $M_t/M/s_t + M$ example with sinusoidal arrival rate function in (18) when $a = 20, b = 4, c = 1$. The system is lightly loaded ($\alpha = 0.005, 0.01, 0.02$).

Table EC.11 shows how the errors depend on the two parameter λ and α . As we would expect, the sensitivity increases as λ , and thus the number of servers, decreases. Under heavy loads, we can predict the quantitative impact from the heavy-traffic approximation

$$P(Ab) \approx (\lambda - s\mu)/\lambda;$$

see Whitt (2004).

The impact of α depends on how we look at it. The absolute error increases as α increases. As α increases, we get more and more overloaded. Then the formula above applies. But under lower α , there has to be more slack, so that one server can make less difference.

However, On the other hand, Table EC.11 shows that the relative error as a function of α actually goes the other way. It does not change so much.




Figure EC.19 Simulation estimates for the $M_t/M/s_t + M$ example with sinusoidal arrival rate function in (18) with staffing $s(t) - 1$, when $a = 20$, $b = 4$, $c = 1$. The system is lightly loaded ($\alpha = 0.005, 0.01, 0.02$).




Figure EC.20 Simulation estimates for the $M_t/M/s_t + M$ example with sinusoidal arrival rate function in (18) when $a = 50$, $b = 10$, $c = 1$. The system is heavily loaded ($\alpha = 0.05, 0.10, 0.15, 0.20$).

EC.12. More on Discretization Issues and Real-World Constraints

In this section we present extra detail for §EC.6. First, Figure EC.24 shows the resulting integer-valued staffing function, fitting to the nearest integer.

Next in Figures EC.25 and EC.26 we plot the abandonment probabilities, expected delays with abandonment probability target $\alpha = 0.05$, and 0.02 , respectively for two cases: fixed staffing periods of 2 hours and 30 minutes.

References

Green, L. V., P. J. Kolesar. 1991. The pointwise stationary approximation for queues with nonstationary arrivals. *Management Sci.* **37** 84–97.

Whitt, W. 1991. The pointwise stationary approximation for $M_t/M_t/s$ queues is asymptotically correct as the rates increase. *Management Sci.* **37** 307–314.

Whitt, W. 2004. Efficiency-driven heavy-traffic approximations for many-server queues with abandonments. *Management Sci.* **50** 1449–1461.




Figure EC.21 Simulation estimates for the $M_t/M/s_t + M$ example with sinusoidal arrival rate function in (18) when $a = 50$, $b = 10$, $c = 1$. The system is lightly loaded ($\alpha = 0.005, 0.01, 0.02$).




Figure EC.22 Simulation estimates for the $M_t/M/s_t + M$ example with sinusoidal arrival rate function in (18) when $a = 1000$, $b = 200$, $c = 1$. The system is heavily loaded ($\alpha = 0.05, 0.10, 0.15, 0.20$).




Figure EC.23 Simulation estimates for the $M_t/M/s_t + M$ example with sinusoidal arrival rate function in (18) when $a = 1000$, $b = 200$, $c = 1$. The system is lightly loaded ($\alpha = 0.005, 0.01, 0.02$).




Figure EC.24 The discretized staffing function and the actual staffing function.

α	λ	Staffing level s	$P_s(Ab)$	$P_{s-1}(Ab)$	Absolute error $P_{s-1}(Ab) - P_s(Ab)$	Relative error Abs.err./ $P_s(Ab)$
0.2	20	17	0.1681	0.2095	0.415	0.2467
	100	81	0.1901	0.2001	0.01	0.0524
	1000	801	0.199	0.2000	0.001	0.005
0.1	20	19	0.0997	0.1312	0.0315	0.3165
	100	91	0.0945	0.1034	0.0088	0.0933
	1000	901	0.099	0.1000	1.00E-03	0.0101
0.01	20	26	0.0072	0.0112	0.004	0.5624
	100	108	0.0088	0.0106	0.0018	0.2027
	1000	1001	0.01	0.0105	4.67E-04	0.0468
0.005	20	27	0.0045	0.0072	2.70E-03	0.5984
	100	111	0.0049	0.006	0.0011	0.2228
	1000	1015	0.0049	0.0052	2.75E-04	0.0557

Table EC.5 Sensitivity analysis for the stationary $M/M/s + M$ model, with abandonment probability target $\alpha = 0.2, 0.1, 0.01$ and 0.005 , arrival rate $\lambda = 20, 100$ and 1000 . The other parameters are fixed: $\mu = 1, \theta = 0.5$.

$P_s(Ab)$ is the steady-state abandonment probability with s servers.




Figure EC.25 Simulation estimates for the $M_t/M/s_t + M$ example with sinusoidal arrival under DIS-MOL staffing with fixed staffing period $d = 30$ minutes and 2 hours, $\alpha = 0.05$.

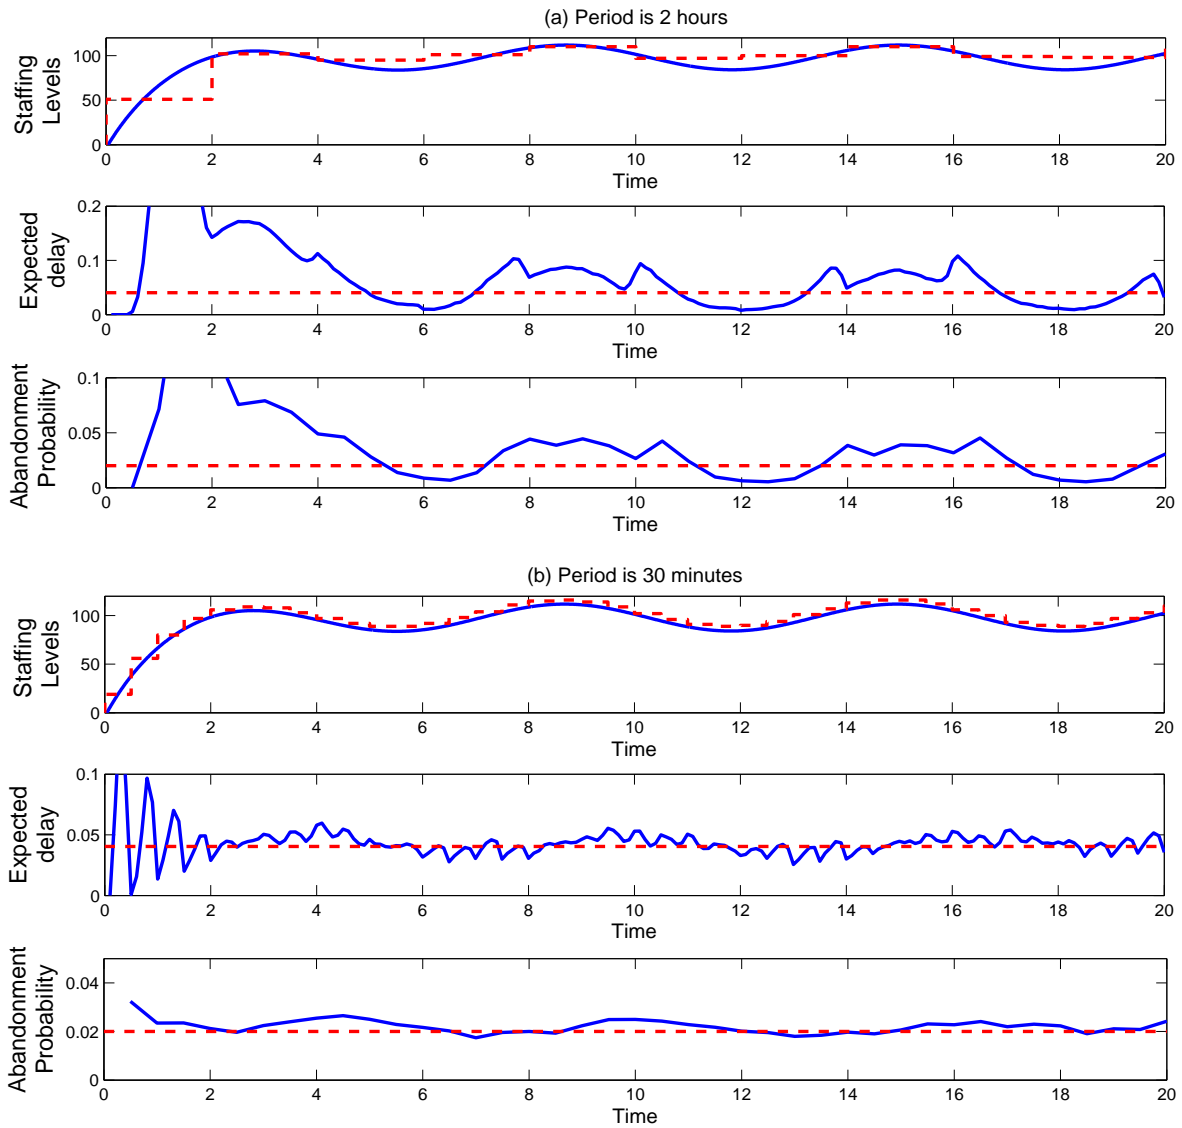


Figure EC.26 Simulation estimates for the $M_t/M/s_t + M$ example with sinusoidal arrival under DIS-MOL staffing with fixed staffing period $d = 30$ minutes and 2 hours, $\alpha = 0.02$.