Queue-Length Distributions for Multi-Priority Queueing Systems

John N. Daigle
The University of Mississippi
University, MS 38677
wcdaigle@cotton.vislab.olemiss.edu

Matthew Roughan
RMIT University, 110 Victoria St.
Carlton, Vic 3053, AUSTRALIA
matt@serc.rmit.edu.au

Abstract— **The bottleneck in many telecommunication systems has often been modeled by an M/G/1 queueing system with priorities. While the probability generating function for the occupancy distribution of each traffic classes can be readily obtained, the occupancy distributions have been obtainable only rarely. However, the occupancy distribution is of great importance, particularly in those cases where the moments are not all finite.**

We present a method of obtaining the occupancy distribution from the PGF and demonstrate its validity by obtaining the occupancy distributions for a number of cases, including those with regularly varying service time distributions.

I. INTRODUCTION

The bottleneck in telecommunication systems has often been modeled by an M/G/1 queueing system having non-preemptive priority service, where the probability generating function (PGF) for the occupancies of the various traffic classes can be obtained using either classical approaches [17] or Fuhrmann-Cooper decomposition [13]. In principle, moments of the occupancy distributions may be obtained from the PGFs (PGFs). However, several recent studies have shown that such properties as long-range dependence (LRD) [18] and regular variation [6], [10], [15] describe real properties of *real* data. And, if the service times are regularly varying, the moments of the occupancy distribution may be infinite. Hence, there is a real need to calculate the occupancy distribution.

In this paper, we obtain the equilibrium marginal occupancy distribution from PGFs for each class of an M/G/1 system having J priority classes and head-of-the-line (HOL) service. The method, described in Section II, is an extension of the techniques described in [8], [9], and [20], which use discrete Fourier transform techniques and the asymptotic properties of the tail of the distribution for the single-class case.

In Section III, the desired PGFs are obtained from the Laplace-Stieltjes transforms (LSTs) of the waiting time distributions. Section IV gives the asymptotic distributions of the regularly varying and standard cases. We see there that if service times are regularly varying for even one class, the occupancy distributions for all classes are regularly varying. This result demonstrates that the HOL priority system cannot isolate the effects of customers with extreme behavior.

Section V uses the asymptotic results to complete the inversion process, Section VI presents numerical examples, and Section VII presents conclusions.

II. COMPUTATION OF THE OCCUPANCY DISTRIBUTION

Our method of inverting the PGF for the single-class case is now summarized [8]. First the PGF for the occupancy distribution, $\mathcal{F}_{\tilde{n}_j}(\cdot)$, is evaluated at $K+1$ points equally spaced around the unit circle of the complex plane. From the sampled PGF, the inverse fast Fourier transform (IFFT) coefficients, $c_k^{(K)}$ for $0 \le k \le K$, are obtained. The resulting set of coefficients, $c_k^{(K)}$ for $0 \le k \le K$, forms a probability mass function [8][9]; $c_k^{(K)} \ge 0$ and $\sum_{k=0}^{K} c_k^{(K)} = 1$. Furthermore,

$$c_k^{(K)} = p_k + \sum_{m=1}^{\infty} p_{k+m(K+1)}, \tag{1}$$

where p_k is the occupancy probability–that is, $c_k^{(K)}$ is an approximation to p_k corrupted by *alias* terms.

The computational algorithm includes choice of a suitable K, computation of the $c_k^{(K)}$, removal of aliases from $c_k^{(K)}$ to obtain p_k for $0 \le k \le K$, and obtaining the p_k for $k > K$. A summary follows:

1. Initialize: choose initial K, and evaluate the $p_{\tilde{n}_j}(0) = \mathcal{F}_{\tilde{n}_j}(0)$.
2. While error E decreases do
 (a) Evaluate the PGF $\mathcal{F}_{\tilde{n}_j}(\cdot)$, at points $z_k = e^{\frac{-i2\pi k}{K+1}}$ for $0 \le k \le K$.
 (b) Determine $c_o^{(K)} = \frac{1}{K+1}\sum_{k=0}^{K} \mathcal{F}_{\tilde{n}_j}(z_k)$.
 (c) Determine the error measurement E.
 (d) Increase K.
3. Evaluate $c_n^{(K)}$ using the IFFT of the PGF.
4. Evaluate the alias terms $\sum_{m=1}^{\infty} p_{k+m(K+1)}$.
5. Recover the p_n, $0 < n \le K$ by removing the aliased terms from $c_n^{(K)}$ using (1).
6. Compute any number of p_n, $n > K$.

The following sections present the pieces of the algorithm for the two possible cases: standard and regularly varying service times. We note in passing that Abate and Whitt [21],[22] present an altenative computational approach.

III. THE PGF OF THE OCCUPANCY DISTRIBUTION

Before proceeding, we review notation. Random variables are indicated using a tilde–service times are denoted by $\tilde{x}$, the lengths of busy periods are denoted by $\tilde{y}$, nonnegative integer-valued random variables are denoted by $\tilde{n}$, and subscripts are used to denote specific service times, busy periods, and numbers. The distribution of a random variable, say $\tilde{x}$, is denoted by $F_{\tilde{x}}(x)$ and its LST is denoted by $F^*_{\tilde{x}}(\cdot)$, the default argument being s. The PGF for the distribution of a nonnegative integer-valued random variable, say $\tilde{n}$, is denoted by $\mathcal{F}_{\tilde{n}}(\cdot)$, the default argument being z.

Class j has priority over class i if $j < i$. With respect to class j, the group of messages having higher priority are referred to as class H, while those messages having lower priority are referred to as class L. We define $\mathcal{X}_L = \{j+1, j+2, \ldots, J\}$ and $\mathcal{X}_H = \{1, 2, \ldots, j-1\}$. The message arrival rates for classes H and L and the corresponding traffic intensities, densities, distributions and LSTs are then defined, with $C \in \{L, H\}$, by

$$\lambda_C = \sum_{i \in \mathcal{X}_C} \lambda_i,$$

$$\rho_C = \lambda_x E[\tilde{x}_C] = \sum_{i \in \mathcal{X}_C} \lambda_i E[\tilde{x}_i],$$

$$F_{\tilde{x}_C}(x) = \frac{1}{\lambda_C} \sum_{i \in \mathcal{X}_C} \lambda_i F_{\tilde{x}_i}(x),$$

$$F^*_{\tilde{x}_C}(s) = \frac{1}{\lambda_H} \sum_{i \in \mathcal{X}_C} \lambda_i F^*_{\tilde{x}_i}(s) \tag{2}$$

where it is understood that if $\mathcal{X}_C = \phi$, then λ_C, $E[\tilde{x}_C]$, ρ_C, $F_{\tilde{x}_C}(x)$, and $F^*_{\tilde{x}_C}(s)$ are all interpreted to be identically zero.

Kleinrock [17, (3.32)] (and many others) gives the LST of the distribution of the waiting time of an arbitrary class j message as

$$F^*_{\tilde{w}_j}(s) = \frac{(1-\rho)G^*_H(s) + \lambda_L\left[1 - F^*_{\tilde{x}_L}\Big(G^*_H(s)\Big)\right]}{s - \lambda_j + \lambda_j F^*_{\tilde{x}_j}\Big(G^*_H(s)\Big)}, \tag{3}$$

where

$$G^*_H(s) = s + \lambda_H - \lambda_H F^*_{\tilde{y}_{HH}}(s). \tag{4}$$

Consider the busy period $\tilde{y}_{HH}$. A busy period starts when a message enters the system during an idle period and ends when there are no more messages to serve. A sub-busy period is a type of restricted busy period. In particular, $\tilde{y}_{ij}$ denotes the length of a sub-busy period of class j messages in which the service time of the first message of the period is drawn from $F_{\tilde{x}_i}$, but the remainder of the dynamics of the busy period are affected by class j traffic only. Then, the random variable $\tilde{y}_{HH}$ denotes the length of a sub-busy period of class H messages started by a class H message, which is statistically identical to an ordinary busy period in an M/G/1 system that has class H messages only. That is, $F^*_{\tilde{y}_{HH}}(s)$ satisfies the well-known functional equation

$$F^*_{\tilde{y}_{HH}}(s) = F^*_{\tilde{x}_H}(s + \lambda_H - \lambda_H F^*_{\tilde{y}_{HH}}(s)), \tag{5}$$

which, except for limited cases, cannot be solved in closed form.

An additional useful result is as follows: let $F^*_{\tilde{x}}(s)$ denote the LST of the distribution of a random period of time, $\tilde{x}$. Suppose observers of this period occur according to an independent Poisson process having rate λ and that the number of such observers is denoted by $\tilde{n}$. Then, the PGF of $\tilde{n}$ is given by

$$\mathcal{F}_{\tilde{n}}(z) = F^*_{\tilde{x}}(\lambda[1-z]). \tag{6}$$

Note that with first-in-first-out (FIFO) queueing, the class j messages left in the queue at the time an arbitrary class j message enters service are precisely those who arrive while that same message is waiting. Also, the arrival process during that message's waiting and service periods is independent of the length of those periods. Further, the number of arrivals that occur during the waiting and service periods are independent. The PGF of the number of class j messages left by an arbitrary departing class j message is then

$$\mathcal{F}_{\tilde{n}_j}(z) = F^*_{\tilde{w}_j}(\lambda_j[1-z]) F^*_{\tilde{x}_j}(\lambda_j[1-z]) \tag{7}$$

where $F^*_{\tilde{w}_j}(s)$ is defined in (3).

In our computational algorithm, we need to evaluate $\mathcal{F}_{\tilde{n}_j}(z)$ and, therefore, $F^*_{\tilde{y}_{HH}}(\lambda_j[1-z_k])$, at points $z_k = e^{-j\frac{2\pi k}{K+1}}$ in the complex plane, but $F^*_{\tilde{y}_{HH}}(\lambda_j[1-z])$ is not given in closed form. The following proposition, upon which a fast-converging iterative procedure is based, is proved in Appendix A:

> **Proposition A:** *For a particular value of z on the unit circle of the complex plane, the value $\nu = F^*_{\tilde{y}_{HH}}(\lambda_j[1-z])$ can always be determined from the expression $\nu_i = F^*_{\tilde{x}_H}(\lambda_j[1-z] + \lambda_H[1-\nu_{i-1}])$ by iteration on i, starting with $\nu_0 = z$.*

IV. Asymptotic tail probabilities

We seek to remove the alias terms of (1) from $c_k^{(K)}$. The alias terms, $p_{k+m(K+1)}$, can be obtained from the asymptotic behavior of the occupancy distribution for sufficiently large K, which we now examine for the standard and regularly varying cases.

A. Asymptotics: standard service times

If all moments of the service-time distributions are finite, then, asymptotically, the probability masses of the occupancy distribution decrease geometrically. The reason for this behavior is readily explained, as for the standard M/G/1 case, through the Laurent series expansion of $\mathcal{F}_{\tilde{n}_j}(z)$ [9].

Since the tail probability masses are geometrically decreasing, for sufficiently large K, there exists a constant, r, such that

$$p_n \simeq p_K r^{(n-K)} \quad \text{for} \quad n > K. \tag{8}$$

Section V specifies an algorithm for computing r. Once r is known, (8) can be used to remove the aliases from the $c_k^{(K)}$.

B. Asymptotics: regularly-varying service times

Next, we derive the asymptotic behavior of the occupancy distribution for the M/G/1 queue with priority when one of the classes of customers has regularly varying service times. The development parallels that of the ordinary M/G/1 system[20].

B.1 Regular variation

Intuitively, regularly varying service times can be described as power-law distributions, e.g. $1 - F_{\tilde{x}}(x) \sim x^{-\alpha}$. The Pareto distribution is a typical example. Several recent studies have shown that properties such as LRD [18] and regular variation [6], [10], [15] describe real properties of *real* data, especially packet traffic.

It has been shown that the tail of the occupancy distribution for the M/G/1 queue with power-law service times is the result of rare arrivals [1] who each demand a single large chunk of work, choking the system's single server. An arriving customer sees a workload dominated by the residual of this high-demand customer. The residual of a power-law lifetime is also power-law with an exponent one lower than the original distribution. Hence, the asymptotic workload distribution takes the form given in [2] Corollary 8.10.4, (from [4])

$$1 - F_{\tilde{w}}(x) \sim \frac{\lambda}{(\alpha-1)(1-\rho)} L(x) x^{-\alpha+1} \tag{9}$$

and the asymptotic occupancy distribution takes the form

$$p_{\tilde{n}}(n) \sim \frac{L(n)\lambda^{\alpha}}{1-\rho} n^{-\alpha}. \tag{10}$$

Thus, if the service times are regularly varying then the occupancy distribution is also regularly varying with one less finite moment than the service-time distribution.

Since the HOL system has no pre-emption, a class j customer may be blocked by customers of any class. Hence, if even one class has regularly varying service times, we might expect that the tails of the queue lengths for each class would be dominated by the class having regular variation. This expectation is correct and unfortunate, as one of the obvious intentions of dividing the customers into classes is isolation of the effects of one class's "pathological" behavior from the other classes.

The following subsections describe the asymptotic behavior with at least one class having regularly varying service times. For simplicity, we concentrate on the case $\alpha \in (1, 2)$, the most realistic physical case, but note that the results can be generalized.

Before proceeding, we define some terms. We say $h(t)$ *is asymptotically equivalence to* $g(t)$ if $\lim_{t \to t_0} |h(t)|/|g(t)| = 1$, and we write $h(t) \overset{t_0}{\approx} g(t)$. We use t_0^- as the limit as t_0 is approached from below. If $\lim_{t\to t_0} L(xt)/L(t) = 1 \; \forall \; x > 0$, then $L(t)$ is slowly varying at t_0, and $h(t)$ is said to be *regularly varying* at ∞ with exponent p, if

$$h(t) \overset{\infty}{\approx} L(t) t^{p}, \tag{11}$$

where $L(t)$ is slowly varying at ∞. A function $g(t)$ is regularly varying at zero if $g(t) = h(1/t)$ with $h(t)$ regularly varying at ∞. Where t_0 is obvious (.i.e., 0 or ∞), we shall omit it.

B.2 The asymptotics of the workload distribution

We derive here the asymptotics of the workload distribution for the case where at least one class has regularly varying service times. We define the set of classes containing regular variation by $R = \{j = 1, \ldots, J | 1 - F_{\tilde{x}_j}(x) \sim L_j(x) x^{-\alpha_j}\}$, where $L_j(x)$ is slowly varying. For all $j \in R^c$, the complement of R, we define $L_j(x) = 0$. We consider only the case where the α_j are in $(1, 2)$, but in fact only the $\min(\alpha_j)$ need be in the interval.

There are three points through which power-law tails may enter the LST of the workload for customer type j (given in (3) and (4)), and therefore influence class j.

1. **Lower** priority customers, through the term $F^*_{\tilde{x}_L}(\cdot)$.
2. **Same** priority customers, through the term $F^*_{\tilde{x}_j}(\cdot)$.
3. **Higher** priority traffic, through the term $G^*_H(\cdot)$ via $F^*_{\tilde{x}_H}(\cdot)$.

The tails of the distributions $F_{\tilde{x}_H}(x)$ and $F_{\tilde{x}_L}(x)$ are dominated by the $F_k(x)$ which is regularly varying, where $k \in H$or L. If more than one $F_k(x)$ is regularly varying, then the one with the smallest exponent will eventually dominate.

For example, suppose one or more of the classes of priority lower than j is regularly varying, then $F^*_L(x)$ is regularly varying with exponent $\alpha_L = \min(\alpha_k | k \in L)$, and slowly varying function defined by $\lambda_L L_{\tilde{x}_L}(x) = \sum_{i:\alpha_i=\alpha_L} \lambda_i L_i(x)$, *i.e.*

$$1 - F_{\tilde{x}_L}(x) \sim L_{\tilde{x}_L}(x) x^{-\alpha_L}, \tag{12}$$

$$\text{with} \quad L_{\tilde{x}_L}(x) = \frac{1}{\lambda_L} \sum_{i:\alpha_i=\alpha_k} \lambda_i L_i(x). \tag{13}$$

A similar result (with L replaced by H) also holds for the higher priority classes.

Note that when $\alpha_L \in (1, 2)$ it has been shown [2], [3] that if $F_{\tilde{x}_L}(x)$ obeys (12) then

$$F^*_{\tilde{x}_L}(s) \sim 1 - \frac{s}{\mu_L} + \frac{L_{\tilde{x}_L}(1/s)\Gamma(2-\alpha_L)}{\alpha_L - 1} s^{-\alpha_L}, \tag{14}$$

where $1/\mu_L = E[\tilde{x}_L]$. Again, we may replace the indices L in the above relationship by j or H. However, $F^*_{\tilde{x}_H}(\cdot)$ appears in the PGF of interest through $G^*_H(s)$. Unlike in Section II, numerical evaluation of $F^*_{\tilde{y}_{HH}}(\cdot)$ is not sufficient; we require a closed form. However, from de Meyer and Teugals ([2], p.388), we have the following result for the M/G/1 busy period:

Theorem 1 (de Meyer and Teugals) For a stable M/G/1 queue of traffic intensity ρ, regularly varying function $L(x)$ and $\alpha \geq 1$, the following are equivalent:

$$\begin{array}{lll} \text{(i)} & 1 - F_{\tilde{x}}(x) & \overset{x\to\infty}{\sim} & L(x)x^{-\alpha}, \\ \text{(ii)} & 1 - F_{\tilde{y}}(x) & \overset{x\to\infty}{\sim} & (1-\rho)^{-\alpha-1} L(x) x^{-\alpha}, \end{array}$$

where $F_{\tilde{x}}(x)$ is the distribution function of the service times, and $F_{\tilde{y}}(x)$ is the distribution function of the busy-period.

Furthermore, Little's result and elementary renewal theory tell us that for a stable queueing system $E[\tilde{y}] = E[\tilde{x}]/(1-\rho)$. Hence for $\alpha_H \in (1, 2)$, the asymptotic form of the LST of the busy period of class H customers is

$$\begin{aligned} F^*_{\tilde{y}_{HH}}(s) \sim \quad & 1 - \tfrac{s}{\mu_H(1-\rho_H)} + \tfrac{\Gamma(2-\alpha_H)}{(\alpha_H-1)(1-\rho_H)} \\ & L_{\tilde{x}_H}\left(\tfrac{1}{s}\right) \left(\tfrac{s}{1-\rho_H}\right)^{\alpha_H}. \end{aligned} \tag{15}$$

Combining (12)-(15), we find the following:

Proposition B: *Given the HOL priority system described above, $\rho_H + \rho_j < 1$, $\alpha_{\mathcal{R}} = \min_j \alpha_j \in (1,2)$, $\mathcal{R} = \{j | \alpha_j = \alpha_{\mathcal{R}}\}$, and $\lambda_{\mathcal{R}} L_{\mathcal{R}}(x) = \sum_{i\in\mathcal{R}} \lambda_i L_{\tilde{x}_i}(x)$, the LST of the workload distribution obeys*

$$F^*_{\tilde{w}_j}(s) \sim \quad 1 - \tfrac{\Gamma(2-\alpha_{\mathcal{R}})}{(\alpha_{\mathcal{R}}-1)(1-\rho_j-\rho_H)} \lambda_{\mathcal{R}} L_{\mathcal{R}}\left(\tfrac{1}{s}\right) \left(\tfrac{s}{1-\rho_H}\right)^{\alpha_{\mathcal{R}}-1}. \tag{16}$$

The above proposition, which is proved in Appendix B, can be transformed to give the following asymptotic workload

$$1 - F_{\tilde{w}_j}(x) \sim \frac{(1-\rho_H)^{-\alpha_{\mathcal{R}}+1}}{(\alpha_{\mathcal{R}}-1)(1-\rho_j-\rho_H)} \lambda_{\mathcal{R}} L_{\mathcal{R}}(x) x^{-\alpha_{\mathcal{R}}+1}. \tag{17}$$

by using the following theorem from [12, pp.445-6]:

Theorem 2 (Tauberian) If $0 < p < \infty$, and a positive measure concentrated on $(0, \infty)$ defined by the nondecreasing function $H(t)$ has LST $\tilde{H}(s)$, then

$$H(t) \overset{\infty}{\sim} \frac{L(t)}{\Gamma(p+1)} t^p \quad \Leftrightarrow \quad \tilde{H}(s) \overset{0}{\sim} L\left(\frac{1}{s}\right) s^{-p},$$

where $L(t)$ is a slowly varying function at infinity.

Compare (17) to (9). The obvious similarity results from the same cause–the queue is "clogged" by a single rare customer.

B.3 Asymptotic behavior of the PGF

As in Section 2, we use (7) to derive the PGF for the number of customers. We construct $F^*_{\tilde{w}_j}(s) F^*_{\tilde{x}_j}(s)$, noting that $F^*_{\tilde{x}_j}(s)$ has terms with powers of $s \geq 1$ only, and, except for a constant, makes no contribution to the regularly varying terms. We then substitute $s = \lambda_j(1-z)$, note the slowly varying $L_{\mathcal{R}}$, and get

$$\mathcal{F}_{\tilde{n}_j}(z) \sim \quad 1 - \tfrac{\Gamma(2-\alpha_{\mathcal{R}})}{(\alpha_{\mathcal{R}}-1)(1-\rho_j-\rho_H)} \left(\tfrac{\lambda_j}{1-\rho_H}\right)^{\alpha_{\mathcal{R}}-1} \lambda_{\mathcal{R}} L_{\mathcal{R}}\left(\tfrac{1}{1-z}\right) (1-z)^{\alpha_{\mathcal{R}}-1}, \tag{18}$$

B.4 Asymptotic occupancy distribution

We now transform (18) to obtain the asymptotic occupancy behavior using the following Tauberian Theorem [12, p. 447].

Theorem 3: If $\{h_n\}$, $n = 0, 1, 2\ldots$ is an ultimately monotone positive sequence with generating function $H^*(z)$ that converges for $0 \leq z < 1$, and $0 \leq p < \infty$, then

$$h_n \overset{\infty}{\sim} \frac{L(n)}{\Gamma(p)} n^{p-1} \;\Leftrightarrow\; H^*(z) \overset{1^-}{\sim} L\left(\frac{1}{1-z}\right) \frac{1}{(1-z)^p},$$

where L is slowly varying at infinity.

Since $\alpha_{\mathcal{R}} \in (1,2)$, we may rearrange (18) so that $p > 0$. Thus,

$$\frac{1-\mathcal{F}_{\tilde{n}_j}(z)}{1-z} \sim \quad \tfrac{\Gamma(2-\alpha_{\mathcal{R}})}{(\alpha_{\mathcal{R}}-1)(1-\rho_j-\rho_H)} \left(\tfrac{\lambda_j}{1-\rho_H}\right)^{\alpha_{\mathcal{R}}-1} \lambda_{\mathcal{R}} L_{\mathcal{R}}\left(\tfrac{1}{1-z}\right) (1-z)^{\alpha_{\mathcal{R}}-2}. \tag{19}$$

The left-hand side of (19) is simply the LST of $1 - P_{\tilde{n}_j}(n) = \sum_{i=n+1}^{\infty} p_{\tilde{n}_j}(i)$. Applying the Tauberian theorem, we get

$$1 - P_{\tilde{n}_j}(n) \sim \frac{\lambda_{\mathcal{R}} L_{\mathcal{R}}(n)}{(\alpha_{\mathcal{R}}-1)(1-\rho_j-\rho_H)} \left(\frac{\lambda_j}{(1-\rho_H)n}\right)^{\alpha_{\mathcal{R}}-1} \tag{20}$$

The following lemma, proved in [20] is then used to derive the asymptotic form of the distribution

Lemma 1: Suppose that $\bar{U}_n$ has ultimately monotone difference $u_n = \bar{U}_{n-1} - \bar{U}_n$, for $n > 0$. If $\bar{U}_n \sim L(n) n^p$ with $p < 0$ then $u_n \sim -p \frac{\bar{U}_n}{n}$.

The result of applying the lemma is the asymptotic form of the occupancy distribution needed to remove the aliases from the $c_k^{(K)}$:

$$p_{\tilde{n}_j}(n) \sim \frac{\lambda_{\mathcal{R}} L_{\mathcal{R}}(n)}{1-\rho+\rho_L} \left(\frac{\lambda_j}{1-\rho_H}\right)^{\alpha_{\mathcal{R}}-1} n^{-\alpha_{\mathcal{R}}}. \tag{21}$$

V. Removal of aliases

We now describe algorithms to compute complete occupancy distributions for the standard and the regularly varying cases, respectively.

An algorithm for using IFFT techniques to obtain the occupancy probabilities for the M/G/1 system is fully described in Daigle [8], and the major results from that reference are repeated here. For a given K, compute $c_k^{(K)}$. As noted in (1), p_k obeys

$$p_k = c_k^{(K)} - \sum_{m=1}^{\infty} p_{k+m(K+1)} \tag{22}$$

for $0 \leq k \leq K$. Thus, we must estimate the alias terms $p_{k+m(K+1)}$ and removing them from $c_k^{(K)}$ to get p_k.

A. Standard case

In the standard case, the occupancy distribution decays geometrically with some as yet unspecified decay constant r. The coefficients $c_k^{(K)}$ may be used to compute the decay constant for the geometrically decreasing tail probabilities. To this end $p_0 = \mathcal{F}_{\tilde{n}_j}(0)$ is evaluated. Next, one computes the geometric decay rate for the tail probabilities, r, from the following formula:

$$r_0 \approx \left(c_0^{(K)} - p_0\right) / c_K^{(K)}. \tag{23}$$

The value of r may also be estimated using the $c_n^{(K)}$ by computing $r_n^{(K)} = c_n^{(K)} / c_{n-1}^{(K)}$ for large n, *e.g.* $n > 3K/4$. By comparing the difference $E_r = \max_{3K/4<n\leq K} \left\{ \text{abs}\left(r_0 - r_n^{(K)}\right) / r_0 \right\}$, we may assess the accuracy of r_0. The accuracy of r_0 as an estimate for r depends on the correct choice of K and so we determine the best value of K by increasing it iteratively and comparing the new estimated decay constant with that obtained from the smaller value of K. The best value of K varies from problem to problem, generally being quite small for standard problems. The method of comparing results and choosing the appropriate value of K is discussed further in [8].

Once K and r_0 are found, the sum term of (22) may be calculated and removed from the $c_k^{(K)}$ to give p_k for $0 \le k \le K$.[1] Specifically,

$$\sum_{m=1}^{\infty} p_{k+m(K+1)} = \left(c_0^{(K)} - p_0\right) r_0^n. \tag{24}$$

Finally, one computes any desired number of remaining probabilities according to

$$p_n = p_K {r_0}^{(n-K)} \quad \text{for} \quad n > K. \tag{25}$$

B. Regularly varying case

This case uses exactly the principles described in the previous section. But, in this case we can predict the asymptotic behavior exactly and therefore need not measure a decay constant. Since the algorithm is described in detail in [20], we describe it here only briefly.

For simplicity, consider $L_{\mathcal{R}}(t) = L_{\mathcal{R}}$. Then, from (21),

$$p_{\tilde{n}_j}(n) \sim B n^{-\alpha_{\mathcal{R}}}, \text{ where } B = \frac{\lambda_{\mathcal{R}} L_{\mathcal{R}}}{1 - \rho + \rho_L} \left(\frac{\lambda_j}{1 - \rho_H}\right)^{\alpha_{\mathcal{R}} - 1}. \tag{26}$$

Recall that $c_n^{(K)}$ consist of p_n plus tail terms. For K large enough, $p_{K+1}, p_{K+2}, \ldots$, are well approximated by the power-law form, and the sum of aliased terms may be evaluated. Thus from (26) we get

$$\sum_{m=1}^{\infty} p_{k+m(K+1)} \overset{K \to \infty}{\approx} B(K+1)^{-\alpha_{\mathcal{R}}} \zeta\left(\alpha_{\mathcal{R}}, \tfrac{n+K+1}{K+1}\right) \tag{27}$$

where $\zeta(\alpha, q)$ is a generalized Riemann zeta-function [14, Equation 9.521]. When α is real and greater than 1 the function is well approximated [19] by the lower bound

$$\zeta(\alpha, q) \quad \approx \quad \sum_{n=0}^{N} \frac{1}{(q+n)^\alpha} + \frac{(q+N+1)^{1-\alpha}}{\alpha - 1}, \tag{28}$$

where the positive error term $\epsilon(N)$ satisfies $\epsilon(N) < (q+N)^{-\alpha}$. We have an analytic expression for the value of B, but it is useful to directly estimate it from Equations (1) and (27) with $n = 0$. Using $p_0 = \mathcal{F}_{\tilde{n}_j}(0)$, the estimate is

$$\hat{B}(K) \quad = \quad \frac{(c_0^{(K)} - p_0)(K+1)^{\alpha_{\mathcal{R}}}}{\sum_{m=1}^{N} m^{-\alpha_{\mathcal{R}}} + \frac{(N+1)^{1-\alpha_{\mathcal{R}}}}{\alpha_{\mathcal{R}} - 1}}. \tag{29}$$

The discrepancy $E_B = \mathrm{abs}\left(B - \hat{B}(K)\right)/B$, can be used to choose an appropriate value of K as described above, and in more detail in [20].

More complicated cases may be treated at the cost of replacing the Riemann zeta-function used above by a non-standard function expressed as an infinite sum, whose estimation may be slightly more costly and whose error more difficult to control. However, in [20] it was shown that, in some cases, the additional difficulty does not prevent effective computation.

[1] We note that we start with a large K and store terms and use other tricks to avoid redundant computations.

VI. Numerical examples

In [8] and [20], the application of the algorithm to the M/G/1 queue was demonstrated for the standard and regularly varying cases, respectively. A number of points have been demonstrated, and we will not seek to duplicate the results here, as the algorithm is identical in all essential details. For the standard case,

- The algorithm could provide accurate results with relatively little computation.
- The algorithm could be used to choose the best value of K.

And specifically for the regularly varying case,

- The algorithm is stable, except in the so called "heavy queueing regime" which is of little interest in practice.
- The algorithm has important application to cases where the occupancy distribution only slowly converges to the asymptotic result.
- The algorithm can be applied to cases with non-trivial slowly varying functions, $L(x)$.
- There is a marked insensitivity in the occupancy distribution to the form of the slowly-varying function $L(x)$.

Rather than replicate these results for the priority queue we simply present examples that illustrate the accuracy of this method.

We present first an example with two customer classes. We examine, in Figure 1, the four possible cases of this system: namely where both classes, only class 1, only class 2, or neither of the classes contain customers with regularly varying service times. The figure shows the computed occupancy distribution of both classes for these four cases (solid lines), and verifies the results with simulation (shown as crosses).

It should be noted that simulation of queues with regularly varying properties must be approached with great care [7]), and that such simulations converge *very* slowly, for instance see [11] and [16]. Typically the simulation results below are based on 10 million departures, and required about 1 hour to converge sufficiently, whereas the algorithm described in this paper took less than 10 seconds.

We have used the version of the Pareto distribution given by

$$1 - F_{\tilde{x}_j}(x) = \left(\frac{\beta_j}{x + \beta_j}\right)^{\alpha_j},$$

to provide a simply regularly varying distribution, and the distributions for classes which are not regularly varying is an Erlang-N distribution given by density

$$p_{\tilde{x}_j}(x) = \frac{(x/\mu_j)^{N-1} \exp(-x/\mu_j)}{\mu_j (N-1)!}.$$

Figure 2 shows two examples with more than two classes. Parameters are noted in the figures, where service times in these figures are either Pareto or exponential as indicated by the appropriate parameters, and the times taken for simulation and calculation are shown in the captions. The simulaton and calculated results show excellent agreement.

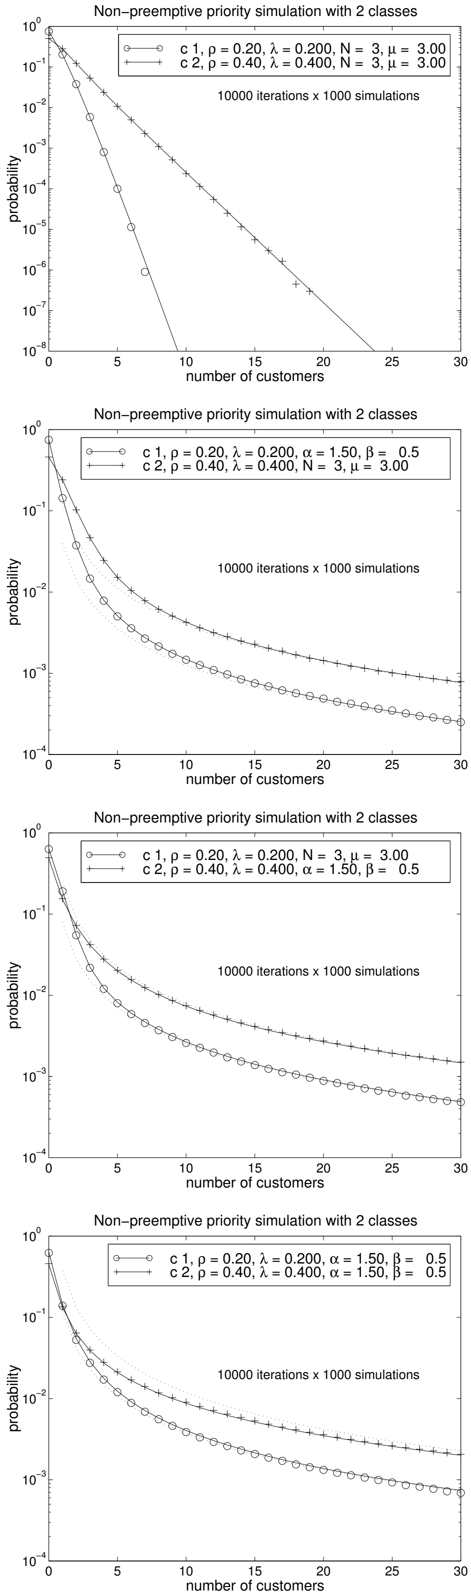

Fig. 1. An example showing the application of the algorithm to the case with a regularly varying tail. The parameters of the system ($\rho_j, \alpha_j, \beta_j, \mu_j$ and λ_j) as well as the simulation parameters (10,000 simulated departures in each of 1000 independent simulations) are indicated in the figure.

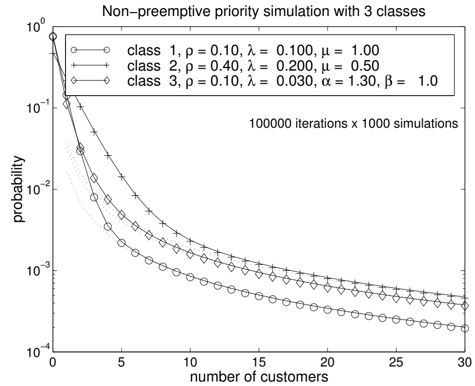

(a) A second example with 3 classes. Computations took ~ 12 seconds while the simulations took nearly 9 hours.

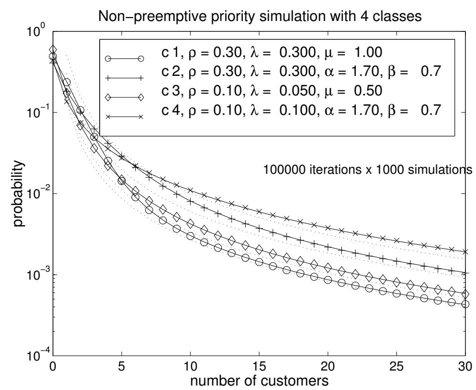

(b) A third example with 4 classes. Computations took $\sim$ 36 seconds while the simulations took nearly 10 hours.

Fig. 2. Examples with three and four class systems.

VII. Conclusion

This paper has presented two main results: an algorithm for calculating the complete occupancy distribution of a HOL priority queue and the asymptotic occupancy distribution of the queue when the service times are regularly varying. Not surprisingly, this asymptotic result very closely parallels that for the standard M/G/1 queue. A consequence is that if any priority class contains customers with regularly varying service times then *all* class feel the effect; their occupancy distributions are all regularly varying!

The algorithm is in essence a method for inverting a PGF. Note that the algorithm, as demonstrated in a number of numerical examples, does not require a closed form expression for the PGF, but rather an iterative procedure is used to evaluate the PGF.

The methods presented here have been applied with success to various systems, including one having eight priority classes where the service times were represented by discrete random variables whose distributions were obtained from measured data. Thus there is good reason to think that this procedure might be usable in more general settings, for instance other queueing problems, or even problems outside of queueing.

References

[1] S. Asmussen and C. Kluppelberg. Stationary M/G/1 excursions in the presence of heavy-tails. *Journal of Applied Probability*, 34:208–212, 1997.

[2] N. H. Bingham, C. M. Goldie, and J. L. Teugels. *Regular Variation*. Cambridge University Press, Cambridge England, 1987.

[3] F. Brichet, J. Roberts, A. Simonian, and D. Veitch. Heavy traffic analysis of a storage model with long range dependent On/Off sources. *Queueing Systems*, 23:197–225, 1996.

[4] J. W. Cohen. Some results on regular variation for the distributions in queueing and fluctuation theory. *Journal of Applied Probability*, 10:343–353, 1973.

[5] R. B. Cooper. *Introduction to Queueing Theory*. The Macmillian Company, 1972.

[6] M. E. Crovella and A. Bestavros. Self-similarity in World Wide Web traffic: Evidence and possible causes. *IEEE/ACM Transactions on Networking*, 5(6), December 1997.

[7] M. E. Crovella and L. Lipsky. Long-lasting transient conditions in simulations with heavy-tailed workloads. In S. Andadottir, K. J. Healy, D. H. Withers, and B. L. Nelson, editors, *Proceedings of the 1997 Winter Simulation Conference*, 1997.

[8] J. N. Daigle. Queue length distributions from probability generating functions via discrete Fourier transforms. *Operations Research Letters*, 8:229–236, August 1989.

[9] J. N. Daigle. *Queueing Theory for Telecommunications*. Addison-Wesley, 1992.

[10] D. E. Duffy, A. A. McIntosh, M. Rosenstein, and W. Willinger. Statistical analysis of CCSN/SS7 traffic data from working CCS subnetworks. *IEEE Journal on Selected Areas in Communications*, 12(3), April 1994.

[11] A. Erramilli, J. Gordon, and W. Willinger. Applications of fractals in engineering for realistic traffic processes. In Jaques Labetoulle and James W. Roberts, editors, *Proceedings of the 14th International Teletraffic Congress - ITC 14*, volume 1a, pages 35–44. Elsevier, Amsterdam, 1994.

[12] W. Feller. *An Introducttion to Probability Theory and its Applications*, volume II. John Wiley and Sons, Brisbane, second edition, 1966.

[13] S. W. Fuhrmann and R. B. Cooper. Stochastic decomposition in the M/G/1 queue with generalized vacations. *Operations Research*, 33(5):1117–1129, 1985.

[14] I. S. Gradshteyn and I. M. Ryzhik. *Table of integrals, series and products*. Academic Press, corrected and enlarged edition, 1980.

[15] Gordon Irlam. Unix file size survey - 1993. http://www.base.com/gordoni/ufs93.html, June 1998. Online document.

[16] P. R. Jelenkovic and A. A. Lazar. Multiplexing on-off sources with subexponential on periods: Part II. In V. Ramaswami and P. E. Wirth, editors, *Proceedings of the 15th International Teletraffic Congress - ITC 15*, **2b**, pp. 965–974, Washinton, D. C., USA, June 1997. Elsevier, Amsterdam.

[17] L. Kleinrock. *Queueing Systems*, volume II: Computer Applications. John Wiley and Sons, 1975.

[18] W. E. Leland, M. S. Taqqu, W. Willinger, and D. V. Wilson. On the self-similar nature of Ethernet traffic (extended version). *IEEE/ACM Transactions on Networking*, 2(1):1–15, Feb 1994.

[19] M. Roughan, D. Veitch, and M. P. Rumsewicz. Numerical inversion of probability generating functions of power-law tail queues. Technical Report 0040-SERC, Software Engineering Research Centre, Level 2, 723 Swanston St, Carlton Vic, 3053, Australia, 1997.

[20] M. Roughan, D. Veitch, and M. P. Rumsewicz. Computing queue-length distributions for power-law queues. In *IEEE INFOCOM'98: The Conference on Computer Communications*, San Francisco, USA, March 1998.

[21] J. Abate and W. Whitt. Numerical inversion of probability generating functions *Operations Research Letters*, **12**, pp. 245-251, 1992.

[22] J. Abate and W. Whitt. The Fourier-series method for inverting transforms of probability distributions *Queueing Systems*, **10**, pp. 5-88, 1992.

Appendix

I. Proof of Proposition A

This appendix provides a proof of Proposition A which is restated below. A minimum of notation and ideas are repeated in this appendix to facilitate continuity.

Proposition A: *For a particular value of z on the unit circle of the complex plane, the value $\nu = F^*_{\tilde{y}_{HH}}(\lambda_j[1-z])$ can always be determined from the expression $\nu_i = F^*_{\tilde{x}_H}(\lambda_j[1-z]+\lambda_H[1-\nu_{i-1}])$ by iteration on i, starting with $\nu_0 = z$.*

Proof: First, note that the right hand side of the expression $\nu_i = F^*_{\tilde{x}_H}(\lambda_j[1-z]+\lambda_H[1-\nu_{i-1}])$ is the joint probability generating function for the number of class j and class H customers who arrive from independent Poisson processes during a period of time $\tilde{x}_H$. Therefore, if $|z| \le 1$ and $|\nu_{i-1}| \le 1$, then $|\nu_i| \le 1$.

The basis of the proof of Proposition A is Banach's fixed point theorem, which is as follows. If for any two points in the domain of $T(\nu)$, ν_0 and ν_1, there exists some positive real-valued $\alpha < 1$, such that $|T(\nu_0) - T(\nu_1)| \le \alpha|\nu_0 - \nu_1|$, then, $T(\nu)$ is a contraction mapping of ν into itself, $\nu = T(\nu)$ is a unique fixed point, and $\nu = \lim_{i\to\infty}\nu_i$, where $\nu_i = T(\nu_{i-1})$.

With regard to the problem at hand, define

$$\begin{aligned} T(\nu_i) \quad &= F^*_{\tilde{x}_H H}(\lambda_j[1-z]+\lambda_H[1-\nu_i]) \\ &= \int_0^\infty e^{-(\lambda_j[1-z]+\lambda_H[1-\nu_i])x} dF_{\tilde{x}_H}(x). \end{aligned}$$

Define $a_i = \lambda_H[1-\nu_i]$. Then, $\mathrm{Re}[a_i] \ge 0$ and

$$T(\nu_0)-T(\nu_1) = \int_0^\infty e^{-(\lambda_j[1-z])x}\left[e^{-a_0x} - e^{-a_1x}\right] dF_{\tilde{x}_H}(x).$$

Assume first that $\mathrm{Re}[a_0] \le \mathrm{Re}[a_1]$, and define $b = a_1 - a_0$ so that $\mathrm{Re}[b] \ge 0$. Then

$$T(\nu_0)-T(\nu_1) = \int_0^\infty e^{-(\lambda_j[1-z])x}e^{-a_0x}\left[1 - e^{-bx}\right] dF_{\tilde{x}_H}(x).$$

Since, for all $|\nu| \le 1$, $|e^{-a_0x}| \le 1$, $|e^{-(\lambda_j[1-z])x}| \le 1$, it follows that

$$|T(\nu_0) - T(\nu_1)| \le \int_0^\infty \left|1 - e^{-bx}\right| dF_{\tilde{x}_H}. \qquad (A.1)$$

Since $|1-e^{-bx}| = \left[(1-e^{-bx})(1-e^{-bx})^* + 2e^{-b_ix} - 2e^{-b_ix}\right]^{\frac{1}{2}}$, we find after some algebra that

$$|1-e^{-bx}| = \left[\left(1-e^{-b_rx}\right)^2 + 4e^{-b_rx}\sin^2\left(\frac{b_ix}{2}\right)\right]^{\frac{1}{2}}, \quad (A.2)$$

where b_r and b_i denote the real and imaginary parts of b. Now, since $b_r \ge 0$ and $x \ge 0$, $1 - e^{-b_rx} \le b_rx$, and $e^{-b_rx} \le 1$. Also, for any θ, $|\sin\theta| \le |\theta|$. It then follows from (A.2) that

$$\left|1 - e^{-bx}\right| \le \left[(b_rx)^2 + (b_ix)^2\right]^{\frac{1}{2}} = |b|x. \qquad (A.3)$$

Using (A.3) in (A.1) then results in

$$|T(\nu_0) - T(\nu_1)| \le \int_0^\infty |b| x dF_{\tilde{x}_H}(x) = |b|\ E[\tilde{x}_H]. \quad (A.4)$$

But from (A.3), $|b| = |a_0 - a_1| = \lambda_H |\nu_0 - \nu_1|$, so

$$|T(\nu_0) - T(\nu_1)| \le \rho_H |\nu_0 - \nu_1| \quad (A.5)$$

Since $\rho_H < 1$ is required for stability, (A.5) implies that for $|\nu| \le 1$ and $|z| \le 1$, $T(\nu)$ is a contraction. If we now assume $\mathrm{Re}[a_0] > \mathrm{Re}[a_1]$ and define $b = a_0 - a_1$, the steps of the proof are identical, and the proof is complete.

II. Proof of Proposition B

This appendix provides a proof of Proposition B which is restated below.

Proposition B: *Given the HOL priority system described above, $\rho_H + \rho_j < 1$, $\alpha_{\mathcal{R}} = \min_{j=1,\dots,J} \alpha_j \in (1,2)$, $\mathcal{R} = \{j = 1,\dots,J | \alpha_j = \alpha_{\mathcal{R}}\}$, and $\lambda_{\mathcal{R}} L_{\mathcal{R}} = \sum_{i\in\mathcal{R}} \lambda_i L_{\tilde{x}_i}$, then*

$$1 - F^*_{\tilde{w}_j}(s) \sim \frac{\Gamma(2-\alpha_{\mathcal{R}})}{(\alpha_{\mathcal{R}}-1)(1-\rho_j-\rho_H)} \lambda_{\mathcal{R}} L_{\mathcal{R}}\left(\frac{1}{s}\right)\left(\frac{s}{1-\rho_H}\right)^{\alpha_{\mathcal{R}}-1}.$$

Proof: The LST $F^*_{\tilde{w}_j}(s)$ is given in Equation (2). We start by calculating $G^*_H(s) = s + \lambda_H\left[1 - F^*_{y_{HH}}(s)\right]$ which features prominently in (2). We use the asymptotic form of $F^*_{y_{HH}}(s)$ given in (15) and the fact that $\rho_H = \lambda_H/\mu_H$ to get

$$\begin{aligned} G^*_H(s) &= \frac{1}{1-\rho_H} s \\ &- \frac{\Gamma(2-\alpha_H)}{(\alpha_H-1)(1-\rho_H)} \lambda_H L_{\tilde{x}_H}\left(\frac{1}{s}\right)\left(\frac{s}{1-\rho_H}\right)^{\alpha_H} \\ &+ O(s^2)\,. \end{aligned} \quad \text{(B.1)}$$

Now the asymptotic form of $F^*_{\tilde{x}_L}(s)$ is given in (14). Substituting (B.1) we get

$$\begin{aligned} \lambda_L\left[1 - F^*_{\tilde{x}_L}\left(G^*_H(s)\right)\right] &= \frac{\rho_L}{1-\rho_H} s \\ &- \rho_L \frac{\Gamma(2-\alpha_H)}{(\alpha_H-1)(1-\rho_H)} \lambda_H L_{\tilde{x}_H}\left(\frac{1}{s}\right)\left(\frac{s}{1-\rho_H}\right)^{\alpha_H} \\ &- \frac{\Gamma(2-\alpha_L)}{\alpha_L-1} \lambda_L L_{\tilde{x}_L}\left(\frac{1}{s}\right)\left(\frac{s}{1-\rho_H}\right)^{\alpha_L} + O(s^2) \end{aligned} \quad \text{(B.2)}$$

which, leads to the numerator of (2) being given by

$$\begin{aligned} (1-\rho)G^*_H(s) + \lambda_L\left[1 - F^*_{\tilde{x}_L}\left(G^*_H(s)\right)\right] &= \frac{1-\rho+\rho_L}{1-\rho_H} s \\ &- \frac{\Gamma(2-\alpha_H)(1-\rho+\rho_L)}{(\alpha_H-1)(1-\rho_H)} \lambda_H L_{\tilde{x}_H}\left(\frac{1}{s}\right)\left(\frac{s}{1-\rho_H}\right)^{\alpha_H} \\ &- \frac{\Gamma(2-\alpha_L)}{\alpha_L-1} \lambda_L L_{\tilde{x}_L}\left(\frac{1}{s}\right)\left(\frac{s}{1-\rho_H}\right)^{\alpha_L} + O(s^2)\,. \end{aligned} \quad \text{(B.3)}$$

We obtain the denominator of (2) by noting that the asymptotic form of $F^*_{\tilde{x}_j}(s)$ is the same as that of $F^*_{\tilde{x}_L}(s)$ with L replaced by j throughout. Hence, the denominator of (2) is given by

$$\begin{aligned} s - \lambda_j\left[1 - F^*_{\tilde{x}_j}\left(G^*(s)\right)\right] &= \frac{1-\rho_H-\rho_j}{1-\rho_H} s \\ &+ \rho_j \frac{\Gamma(2-\alpha_H)}{(\alpha_H-1)(1-\rho_H)} \lambda_H L_{\tilde{x}_H}\left(\frac{1}{s}\right)\left(\frac{s}{1-\rho_H}\right)^{\alpha_H} \\ &+ \frac{\Gamma(2-\alpha_j)}{\alpha_j-1} \lambda_j L_{\tilde{x}_j}\left(\frac{1}{s}\right)\left(\frac{s}{1-\rho_H}\right)^{\alpha_j} + O(s^2)\,. \end{aligned} \quad \text{(B.4)}$$

Now, $1 - \rho + \rho_L = 1 - \rho_j - \rho_H$ so we can divide both the numerator and denominator of (2) by $\frac{1-\rho_j-\rho_H}{1-\rho_H}s$. We then exploit the fact that $1/(1+s) \sim 1 - s$ and group terms of $O(s)$ or greater to get

$$\begin{aligned} 1 - F^*_{\tilde{w}_j}(s) &= \frac{\Gamma(2-\alpha_H)}{(\alpha_H-1)(1-\rho_j-\rho_H)} \lambda_H L_{\tilde{x}_H}\left(\frac{1}{s}\right)\left(\frac{s}{1-\rho_H}\right)^{\alpha_H-1} \\ &+ \frac{\Gamma(2-\alpha_L)}{(\alpha_L-1)(1-\rho_j-\rho_H)} \lambda_L L_{\tilde{x}_L}\left(\frac{1}{s}\right)\left(\frac{s}{1-\rho_H}\right)^{\alpha_L-1} \\ &+ \frac{\Gamma(2-\alpha_j)}{(\alpha_j-1)(1-\rho_j-\rho_H)} \lambda_j L_{\tilde{x}_j}\left(\frac{1}{s}\right)\left(\frac{s}{1-\rho_H}\right)^{\alpha_j-1} \\ &+ O(s)\,. \end{aligned} \quad \text{(B.5)}$$

Now, the above result has three different exponents α_L, α_j and α_H. If these are all different then the minimum will dominate and therefore only one of the regularly varying terms above will matter, but if two of the exponents above are equal then we must include two of the terms above. However, this does not necessarily complicate matters, for the following reason. From the definition

$$\lambda_H L_{\tilde{x}_H}(x) = \sum_{i\in H} \lambda_i L_{\tilde{x}_i}(x),$$
$$\lambda_L L_{\tilde{x}_L}(x) = \sum_{i\in L} \lambda_i L_{\tilde{x}_i}(x),$$

where $L_{\tilde{x}_i}$ is assumed to be identically zero unless class i customers have regularly varying service times. We define $\mathcal{R}$ to be the set of the classes i of traffic with α_i equal to the dominant exponent $\alpha_{\mathcal{R}} = \min_{j=1,\dots,J} \alpha_j$. The terms in (B.5) may be expanded as sums, and then the terms with dominant values of α grouped together, and the rest discarded, to give

$$1 - F^*_{\tilde{w}_j}(s) \sim \frac{\Gamma(2-\alpha_{\mathcal{R}})}{(\alpha_{\mathcal{R}}-1)(1-\rho_j-\rho_H)} \lambda_{\mathcal{R}} L_{\mathcal{R}}\left(\frac{1}{s}\right)\left(\frac{s}{1-\rho_H}\right)^{\alpha_{\mathcal{R}}-1}$$

which is the required result.