

A modified HOL priority scheduling discipline: performance analysis

Tom Maertens*, Joris Walraevens, Herwig Bruneel
Ghent University – UGent
Department of Telecommunications and Information Processing
SMACS Research Group
Sint-Pietersnieuwstraat 41, B-9000 Gent, Belgium
Phone: +32-9-2648901
Fax: +32-9-2644295
E-mail: {tmaerten,jw,hb}@telin.UGent.be

Abstract

In this paper, we introduce and analyze a modified HOL (head-of-the-line) priority scheduling discipline. The modification is incorporated to cope with the so-called starvation problem of regular HOL priority queues. We consider a discrete-time single-server queueing system with two priority queues of infinite capacity and with the introduced priority scheme. We show that the use of probability generating functions is suitable for analyzing the system contents and the packet delay. Some performance measures (such as means and variances) of these stochastic quantities will be derived. Furthermore, approximate expressions of the tail probabilities are obtained from the probability generating functions, by means of the dominant-singularity method. These expressions, together with their characteristics, constitute one of the main contributions of this paper. Finally, the impact and significance of the m-HOL (modified HOL) priority scheduling on these performance measures is illustrated by some numerical examples.

keywords queueing, discrete-time, dynamic priority scheduling, performance analysis

1 Introduction

Network traffic can often be divided into multiple classes of traffic, each having a different QoS (Quality of Service) standard. For delay-sensitive traffic, it is important that mean delay and delay jitter are not too large, while for delay-insensitive traffic, the loss ratio is the restrictive quantity. To support and differentiate several classes of traffic in a network, priority schemes are used. In the Head-Of-the-Line (HOL) priority scheduling discipline for instance, priority is always given to the delay-sensitive traffic, i.e., the delay-sensitive (high-priority) traffic is always scheduled for service before the delay-insensitive (low-priority) traffic. An overview of some basic priority queueing models in continuous time can be found in [14], [15] and [23] (and references therein). Discrete-time priority queues are analyzed in e.g. [17] and [22].

Priority scheduling does indeed provide low delays for the high-priority traffic (see e.g. [4], [25]), but if a large portion of the network traffic consists of high-priority traffic, the performance for the low-priority traffic can be severely degraded. Specifically, HOL priority scheduling can cause excessive delays for the low-priority traffic, especially if the network is highly loaded (which is a major problem if this type of traffic is not entirely delay-insensitive, but has also some delay requirement). This problem is also known as the starvation problem. In order

*Corresponding author

to find a solution for this drawback, several dynamic priority schemes have been proposed in the literature. These schemes are mostly obtained by alternately serving high-priority traffic and low-priority traffic, depending on a certain threshold (Queue-Length-Threshold or Minimum-Laxity-Threshold, see e.g. [10], [13]), or by allowing priority jumps (Head-Of-the-Line with Priority Jumps, see e.g. [4], [18], [19]). In the latter type, when high- and low-priority packets arrive in a high- and low-priority queue respectively, packets of the low-priority queue can jump to the high-priority queue. To decide if and when low-priority packets jump, many criteria can be used: a maximum queuing delay in the low-priority queue, a queue-length-threshold of the low-priority queue, a random jumping probability per time unit, . . . Further, the jumping process is also characterized by the number of packets that jump at the same time and by the specific moments that these packets jump (at the beginning of a time unit, at the end of a time unit, . . .).

In this paper, we analyze a queue where the packet at the HOL-position of the low-priority queue can jump to the high-priority queue. This possible jump depends on the contents of the high-priority queue, i.e., when this queue is non-empty, the packet jumps. When the high-priority queue is empty on the other hand, the low-priority packet at the HOL-position is immediately served by the server. So the difference with the original HOL priority scheduling discipline (see e.g. [25]) appears when the high-priority queue is not empty. As a result, we call this dynamic priority scheme the m-HOL (modified HOL) priority scheduling discipline. This scheduling discipline thus tries to hold the advantage of the HOL priority discipline (i.e., keeping the delay of high-priority packets small), while trying to solve (or at least alleviate) the starvation problem, since the low-priority packets finally jump to the high-priority queue, even when the high-priority load is large.

For assessing the performance of queues with this m-HOL scheduling discipline, we use an analysis based on probability generating functions. In this way, we obtain the probability generating functions of the contents in the high- and low-priority queue and the probability generating functions of the delay of high- and low-priority packets. From these probability generating functions, we can then (easily) calculate expressions for some interesting performance measures, such as mean values, variances and approximate tail probabilities of the studied stochastic quantities. These closed-form expressions require little computational effort and are well-suited for evaluating the impact of the various system parameters on the overall performance.

The contribution of this paper concerns the model that is considered, as well as the combination of the solution technique that we have used and the specific results that are (efficiently) derived from the obtained solution. First, we show that the performance of low-priority traffic in a network can be controlled by letting low-priority traffic *flow* into the high-priority traffic (rather than starve in the low-priority queue). This is easily obtained by slightly modifying the original HOL priority scheduling discipline. Secondly, this paper demonstrates that an analysis based on probability generating functions is suitable for analyzing queues with a dynamic priority scheduling discipline. Once the probability generating functions of the buffer contents and packet delays are derived, expressions for the mean values (and for higher moments) of the analyzed stochastic quantities are easy to obtain. Finally, determining expressions of the tail probabilities of the stochastic quantities from their probability generating functions, is one of the main contributions of this paper. Tail probabilities typically represent the 'exceptional' situations in a queueing

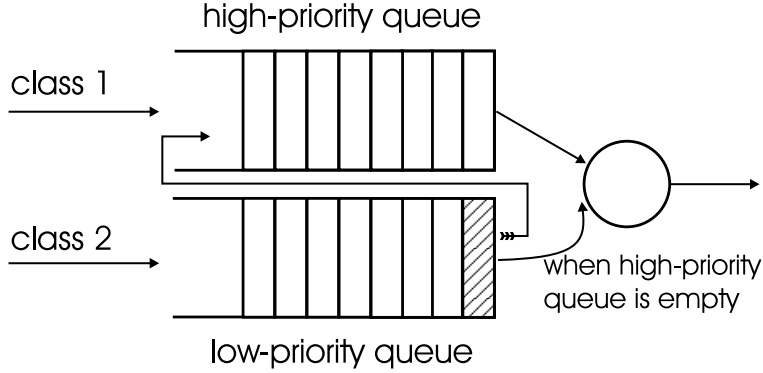


Figure 1: m-HOL priority queueing system

systems and the tail distribution can thus often be used to impose statistical bounds on the guaranteed QoS. Moreover, we will show that the obtained distributions are not necessarily purely geometrical, i.e., besides a dominant geometric factor, distributions may be characterized by an additional (weaker) algebraic factor, which leads to non-geometric tail behavior.

The outline of the paper is as follows. In the following section, we describe the mathematical model. In sections 3 and 4, we analyze the steady-state system contents and packet delay of both classes respectively. We calculate the moments and study the tail behavior of the system contents and packet delay in sections 5 and 6 respectively. We present and discuss some numerical results in section 7 and finally, some conclusions are formulated in section 8.

2 Mathematical model

We consider a discrete-time queueing system, i.e., time is assumed to be slotted. The system consists of one server and two queues of infinite capacity (see Figure 1). Further, the system is characterized by three processes: the arrival process, the service process and the jumping process.

We first describe the arrival process. Two types of traffic arrive in the system, namely packets of class 1 and packets of class 2, which arrive in the first and the second queue respectively. The number of arrivals of class j during slot k is denoted by $a_{j,k}$ ($j = 1, 2$) and the $a_{1,k}$'s and $a_{2,k}$'s are independent and identically distributed (i.i.d.) variables. However, in one slot, the number of arrivals of both classes can be correlated. This dependence is characterized by their joint probability generating function (pgf) $A(z_1, z_2) \triangleq E[z_1^{a_{1,k}} z_2^{a_{2,k}}]$. The total number of arriving packets during slot k is denoted by $a_{T,k} \triangleq a_{1,k} + a_{2,k}$ and its pgf is defined as $A_T(z) \triangleq E[z^{a_{T,k}}] = A(z, z)$. Further, we define the marginal pgf's of the number of class 1 and class 2 arrivals during a slot as $A_1(z) \triangleq E[z^{a_{1,k}}] = A(1, z)$ and $A_2(z) \triangleq E[z^{a_{2,k}}] = A(z, 1)$ respectively. From these pgf's, we can calculate the arrival rate of class j : $\lambda_j \triangleq E[a_{j,k}] = A'_j(1)$. The total arrival rate is the sum of the arrival rates of both classes: $\lambda_T = A'_T(1) = \lambda_1 + \lambda_2$.

Secondly, the service process is determined by the number of servers, which equals one in this system, and by the service time of a packet, which is also equal to one. Newly arriving packets can enter service at the beginning

of the slot following their arrival slot, at the earliest. Packets in queue 1 have a higher priority than those in queue 2 (we will call queue 1 the high-priority queue and queue 2 the low-priority queue). Thus, when there are packets in the high-priority queue, they have service priority and only when this queue is empty, packets in the low-priority queue can be served. Within a queue, the service discipline is FIFO.

Finally, the system is influenced by a jumping process: in each slot, the packet at the HOL position of the low-priority queue (if one) jumps to the high-priority queue, only if the high-priority queue is not empty (because of the service process, the packet is immediately served if the high-priority queue is empty). In this way, it is avoided that the low-priority traffic starves in the low-priority queue. The jump occurs at the beginning of the slot. This means that the jumping packet is scheduled after the contents of the high-priority queue at the beginning of that slot, but before the packets of class 1 which arrive during the same slot.

3 System contents

In this section, we derive the joint pgf of the system contents of both queues at the beginning of a random slot in steady-state. We assume that the packet in service (if any) is part of the queue that is serviced in that slot. We denote the system contents of the high-priority queue and the low-priority queue at the beginning of slot k by $u_{1,k}$ and $u_{2,k}$ respectively, and the total system contents at the beginning of slot k by $u_{T,k}$. The joint pgf of $u_{1,k}$ and $u_{2,k}$ is then denoted by $U_k(z_1, z_2) \triangleq \mathbb{E}[z_1^{u_{1,k}} z_2^{u_{2,k}}]$. The system contents of both queues evolves in time according to the following system equations:

$$u_{1,k+1} = \begin{cases} [u_{1,k} - 1]^+ + a_{1,k} & \text{if } u_{2,k} = 0 \\ u_{1,k} + a_{1,k} & \text{if } u_{2,k} > 0 \end{cases}, \quad (1)$$

$$u_{2,k+1} = [u_{2,k} - 1]^+ + a_{2,k}, \quad (2)$$

where $[\dots]^+$ denotes the maximum of the argument and zero. When the low-priority queue is empty at the beginning of slot k , no jump occurs. The system equation of the high-priority queue is in that case the normal "FIFO" one. On the other hand, when the low-priority queue is not empty at the beginning of slot k , a low-priority packet jumps to the high-priority queue, if the latter one is non-empty, or is served, if the high-priority queue is empty (at the beginning of slot k). This leads to the second system equation of the high-priority queue. The behavior of the low-priority queue (equation (2)) can then be summarized as follows: a packet leaves in each slot (if any), i.e., either a packet jumps to the high-priority queue, or it is served immediately. Introducing pgf's in the system equations (1) and (2), establishes the following relation between $U_{k+1}(z_1, z_2)$ and $U_k(z_1, z_2)$:

$$U_{k+1}(z_1, z_2) = A(z_1, z_2) \frac{z_1 U_k(z_1, z_2) + (z_2 - z_1) U_k(z_1, 0) + z_2 (z_1 - 1) U_k(0, 0)}{z_1 z_2}. \quad (3)$$

Since we are interested in the *steady-state* distribution of the system contents, we define

$$U(z_1, z_2) \triangleq \lim_{k \rightarrow \infty} U_k(z_1, z_2) = \lim_{k \rightarrow \infty} U_{k+1}(z_1, z_2). \quad (4)$$

Applying this limit in equation (3) and solving the resulting equation for $U(z_1, z_2)$, we find

$$U(z_1, z_2) = A(z_1, z_2) \frac{(z_2 - z_1)U(z_1, 0) + z_1(z_2 - 1)U(0, 0)}{z_1(z_2 - A(z_1, z_2))}. \quad (5)$$

In the right hand side of equation (5), there are two quantities yet to be determined, namely the function $U(z_1, 0)$ and the constant $U(0, 0)$. First, we derive an expression for the function $U(z_1, 0)$. By applying Rouché's theorem, it can be proved (see e.g. [17] and [24]) that for a given value of z_1 inside the unit circle ($|z_1| < 1$), the equation $z_2 = A(z_1, z_2)$ has one solution in the unit circle for z_2 , which will be denoted by $Y(z_1)$ in the remainder, and which is implicitly defined by $Y(z) \triangleq A(z, Y(z))$. Since $Y(z_1)$ is a zero of the denominator of the right hand side of (5) and since a pgf remains finite in the unit circle, $Y(z_1)$ must also be a zero of the numerator. We thus find

$$U(z_1, 0) = U(0, 0) \frac{Y(z_1)(z_1 - 1)}{z_1 - Y(z_1)}. \quad (6)$$

Substituting expression (6) in (5) yields

$$U(z_1, z_2) = U(0, 0) \frac{A(z_1, z_2)(z_1 - 1)}{z_1 - Y(z_1)} \frac{z_2 - Y(z_1)}{z_2 - A(z_1, z_2)}. \quad (7)$$

It is remarkable that expression (7) is identical to the joint pgf of u_1 and u_2 with class 2 packets (!) having a regular HOL priority over class 1 packets (see e.g. [25]). Indeed, $u_2 (> 0)$ is in each slot decreased by one (new arrivals not included), in both the m-HOL priority system as well as the regular HOL priority system. $u_1 (> 0)$ only decreases by one when $u_2 = 0$ at the beginning of the slot. On the contrary, when $u_2 > 0$, the jump of a class 2 packet *catches* the departure of a class 1 packet in the m-HOL priority system, while in the regular HOL priority system no class 1 packet leaves (since class 2 packets have priority).

Next, we determine the constant $U(0, 0)$ from equation (7) by substituting z_1 and z_2 by 1, by applying the normalization condition $U(1, 1) = 1$ and by using de l'Hopital's rule. The result is the probability of having an empty system: $U(0, 0) = 1 - \lambda_T$. From this joint pgf, we can easily obtain an expression for the pgf $U_T(z)$ describing the total system contents:

$$\begin{aligned} U_T(z) &\triangleq \lim_{k \rightarrow \infty} \mathbb{E}[z^{u_T, k}] = U(z, z) \\ &= (1 - \lambda_T) \frac{A_T(z)(z - 1)}{z - A_T(z)}. \end{aligned} \quad (8)$$

This expression is identical to the pgf of the system contents of a queue with a FCFS-discipline and with one class (with an arrival process characterized by $A_T(z)$). This is expected, because for the total system contents, it does not matter in which order the packets are being served (as long as the scheduling discipline is work-conserving). Finally, we can also calculate the marginal pgf's $U_1(z)$ and $U_2(z)$ of the system contents of the high-priority queue and the low-priority queue respectively:

$$U_1(z) \triangleq \lim_{k \rightarrow \infty} \mathbb{E}[z^{u_1, k}] = U(z, 1)$$

$$= (1 - \lambda_T) \frac{A_1(z)(z-1)}{z - Y(z)} \frac{1 - Y(z)}{1 - A_1(z)}, \quad (9)$$

$$\begin{aligned} U_2(z) &\triangleq \lim_{k \rightarrow \infty} E[z^{u_2, k}] = U(1, z) \\ &= (1 - \lambda_2) \frac{A_2(z)(z-1)}{z - A_2(z)}. \end{aligned} \quad (10)$$

As we have noticed before, u_1 and u_2 mutually behave as if class 2 packets have regular HOL priority over class 1 packets. Therefore, expression (9) is similar to the pgf of the system contents of the low-priority queue in a system with a regular HOL-discipline (see e.g. [25]): u_1 is only decreased by one when $u_2 = 0$. Secondly, it is not a surprise that expression (10) is identical to the pgf of the system contents of a queue with a FCFS-discipline and with one class (with arrivals determined by $A_2(z)$). Indeed, one packet (if any) always leaves the low-priority queue: either it is immediately served, or it jumps to the high-priority queue.

4 Packet delay

The packet delay is defined as the total amount of time a packet spends in the system, i.e., the number of slots between the end of the packet's arrival slot and the end of its departure slot. In the current section, we will derive expressions for the pgf's of the packet delay of both classes.

In general, the amount of time a (class 1 or class 2) packet spends in the system can be written as

$$d = [u - 1]^+ + f + p + s. \quad (11)$$

u denotes the number of packets that were present in the system at the beginning of the tagged packet's arrival slot and which have to be served before the tagged packet. f describes the number of packets that arrived during the same slot as the tagged packet but which have to be served before it and s denotes the transmission time of the tagged packet itself (which equals 1 in the analyzed system). The quantity p represents the number of packets of the other class that arrive during slots following the tagged packet's arrival slot, but which have to be served before the tagged one. Specifically, this quantity is caused by the priority scheduling.

Let us first analyze the delay of a tagged class 1 packet. Slot k is hereby assumed to be the arrival slot of the tagged packet. The quantity u , as described above, here equals the system contents of the high-priority queue at the beginning of slot k ($u_{1,k}$), possibly augmented with the class 2 packet jumping at the beginning of the slot. Indeed, the delay is also influenced by the jumping process: since a possible jump of the packet at the HOL position of the low-priority queue to the high-priority queue, takes place at the beginning of the slot, the newly arriving packets of class 1 are queued after the packet that jumps in the same slot. As a consequence, the packet delay of a tagged class 1 packet not only depends on the system contents of the high-priority queue at the beginning of its arrival slot, but also depends on the system contents of the low-priority queue at the beginning of that slot ($u_{2,k}$). Further, the quantity f is equal to the number of class 1 packets that arrived during slot k , but which have to be served before it ($f_{1,k}$). Finally, p equals 0 in this case, since class 2 packets arriving after the tagged class 1 packet will never be served before that packet (because of the priority scheduling). Thus summarized, the amount of time a

tagged class 1 packet spends in the system is given by the following equation:

$$d_1 = \begin{cases} f_{1,k} + 1 & \text{if } u_{1,k} = 0 \\ u_{1,k} + f_{1,k} & \text{if } u_{1,k} > 0, u_{2,k} = 0 \\ u_{1,k} + f_{1,k} + 1 & \text{if } u_{1,k} > 0, u_{2,k} > 0 \end{cases} \quad (12)$$

The first equation can be explained as follows: if the high-priority queue is empty at the beginning of slot k , the packet at the HOL position of the low-priority queue (if this queue is not empty of course) is immediately served and does not jump (u equals zero). In that case, the delay of the tagged class 1 packet does not depend on the system contents of the low-priority queue. On the other hand, if there is at least one packet present in the high-priority queue, the delay of the tagged class 1 packet also depends on the system contents of the low-priority queue (the second and third equation respectively). The amount of time a class 1 packet spends in the system in the case that the low-priority queue is empty (no jump, i.e., u equals $u_{1,k}$) is one less in comparison with the amount of time in the case that the low-priority queue is not empty (a jump occurs and u thus equals $u_{1,k} + 1$). Introducing pgf's in these equations produces the following expression:

$$D_1(z) = F_1(z) \left((U(z, 0) - U(0, 0))(1 - z) + zU_1(z) \right). \quad (13)$$

The pgf $F_1(z) \triangleq E[z^{f_{1,k}}]$ can be calculated taking into account that an arbitrary tagged packet is more likely to arrive in a larger bulk (see e.g. [7]), yielding

$$F_1(z) = \frac{A_1(z) - 1}{\lambda_1(z - 1)}. \quad (14)$$

Using the latter equation, substituting $U(z, 0)$ by expression (6), $U_1(z)$ by expression (9) and $U(0, 0)$ by $1 - \lambda_T$, we find

$$D_1(z) = \frac{1 - \lambda_T}{\lambda_1} \frac{z(Y(z) - 1)}{z - Y(z)}. \quad (15)$$

Secondly, we analyze the packet delay of a tagged class 2 packet. The amount of time a tagged class 2 packet spends in the system equals

$$d_2 = \left([u_{1,k} + u_{2,k} - 1]^+ + \hat{f}_{1,k} + \hat{f}_{2,k} \right) + \sum_{i=1}^{[u_{2,k}-1]^+ + \hat{f}_{2,k}} a_{1,k+i} + 1, \quad (16)$$

where slot k is assumed to be the arrival slot of the tagged packet. Since a class 2 packet arrives in the low-priority queue, it is obvious that u is equal to the total system contents at the beginning of slot k ($u_{1,k} + u_{2,k}$). f is equal to the sum of the number of high- and low-priority packets that arrived during the same slot as the tagged packet and which are served before this packet ($\hat{f}_{1,k} + \hat{f}_{2,k}$). This explains the first term of the equation. The second term, in which the $a_{1,k+i}$'s denote the number of class 1 arrivals during the slots $k + i$, represents the quantity p . New class 1 packets can indeed enter the system before the tagged packet reaches the HOL position of the low-priority queue,

and because of the priority scheduling, these class 1 packets are served before the tagged packet. The sum is then explained as follows: for each packet that is stored in front of the tagged packet in the low-priority queue at the end of the tagged packet's arrival slot ($[u_{2,k} - 1]^+ + \hat{f}_{2,k}$), we have to take into account $a_{1,k+i}$ class 1 packets. Once the packet reaches the HOL position of the low-priority queue, that packet is either immediately served, or is stored in front of the new class 1 packets which arrive during that same slot, i.e., the tagged packet is then not influenced anymore by newly arriving class 1 packets. Finally, we must also take into account the packet's departure slot, which yields the third term of (16). Introducing pgf's in equation (16) outputs

$$D_2(z) = \frac{F(z, zA_1(z))}{A_1(z)} \left(U(z, 0)(A_1(z) - 1) + U(0, 0)A_1(z)(z - 1) + U(z, zA_1(z)) \right), \quad (17)$$

with $F(z_1, z_2) \triangleq E[z_1^{\hat{f}_{1,k}} z_2^{\hat{f}_{2,k}}]$. Taking into account that an arbitrary tagged packet is more likely to arrive in a larger bulk (e.g. [7]) and that $\hat{f}_{1,k}$, because of the priority scheduling, equals the number of all class 1 packets that arrive during slot k (i.e., $\hat{f}_{1,k} = a_{1,k}$), we obtain the joint pgf

$$F(z_1, z_2) \triangleq E[z_1^{\hat{f}_{1,k}} z_2^{\hat{f}_{2,k}}] = \frac{A(z_1, z_2) - A_1(z_1)}{\lambda_1(z_2 - 1)} \quad (18)$$

Using this equation and equation (7) and substituting $U(z, 0)$ by equation (6) and $U(0, 0)$ by $1 - \lambda_T$, we finally derive

$$D_2(z) = \frac{1 - \lambda_T}{\lambda_2} \frac{z(z - 1)(A(z, zA_1(z)) - A_1(z))(zA_1(z) - Y(z))}{(z - Y(z))(zA_1(z) - 1)(zA_1(z) - A(z, zA_1(z)))}. \quad (19)$$

5 Calculation of the moments

The function $Y(z)$, defined in section 3, can only be explicitly found in case of some simple arrival processes. However, its derivatives for $z = 1$, necessary to calculate the moments of the system contents and the packet delay, can be calculated in closed-form. For example, $Y'(1)$ is given by

$$Y'(1) = \frac{\lambda_1}{1 - \lambda_2}. \quad (20)$$

Now we can calculate the mean values of the studied stochastic variables. To make the expressions more readable, we define λ_{ij} and λ_{TT} as

$$\lambda_{ij} \triangleq \left. \frac{\partial^2 A(z_1, z_2)}{\partial z_i \partial z_j} \right|_{z_1=z_2=1} \quad \text{with } i, j = 1, 2 \quad \text{and} \quad \lambda_{TT} \triangleq \left. \frac{\partial^2 A_T(z)}{\partial z^2} \right|_{z=1}$$

respectively. By taking the first derivative of the respective pgf's for $z = 1$, we get expressions for the mean values of the total system contents and of the system contents of both queues:

$$E[u_T] = \lambda_T + \frac{\lambda_{TT}}{2(1 - \lambda_T)}, \quad (21)$$

$$E[u_1] = \lambda_1 + \frac{\lambda_{11} + 2\lambda_{12}}{2(1 - \lambda_T)} + \frac{\lambda_1 \lambda_{22}}{2(1 - \lambda_T)(1 - \lambda_2)}, \quad (22)$$

$$E[u_2] = \lambda_2 + \frac{\lambda_{22}}{2(1 - \lambda_2)}. \quad (23)$$

It is easily verified that these three equations satisfy $E[u_T] = E[u_1] + E[u_2]$. The expressions for the mean values of the packet delay are respectively

$$E[d_1] = \frac{\lambda_T \lambda_2}{(1 - \lambda_T)(1 - \lambda_2)} - \frac{\lambda_{12}}{\lambda_1} + \frac{2\lambda_T - \lambda_{22}(1 + \lambda_1 - \lambda_2)}{2\lambda_1(1 - \lambda_2)} + \frac{\lambda_{TT}(1 - \lambda_2) - 2\lambda_2}{2\lambda_1(1 - \lambda_T)}, \quad (24)$$

for a class 1 packet,

$$E[d_2] = 1 + \frac{\lambda_{11}}{2(1 - \lambda_T)} + \frac{\lambda_{12}(1 - \lambda_1)}{\lambda_2(1 - \lambda_T)} + \frac{\lambda_{22}((1 + \lambda_2) - (\lambda_T + \lambda_2)^2)}{2\lambda_2(1 - \lambda_2)(1 - \lambda_T)}, \quad (25)$$

for a class 2 packet and

$$E[d] = \frac{\lambda_1}{\lambda_T} E[d_1] + \frac{\lambda_2}{\lambda_T} E[d_2] \quad (26)$$

for an arbitrary packet. Notice that $E[u_j] = \lambda_j E[d_j]$ ($j = 1, 2$) does *not* hold, as one would at first expect, according to Little's law. The reason for this is that in the calculation of the system contents, packets of the second queue jump to the first one and from that moment on, they are treated as part of the system contents of queue 1. This is of course not the case in the calculation of the packet delay. So basically, Little's law does not hold with respect to each queue separately, because the system contents is analyzed on a "queue"-basis, while the packet delay is analyzed on a "packet"-basis. However, Little's law does hold for the complete system. Indeed, it is easily verified that $E[u_T] = \lambda_T E[d]$.

Expressions of higher and cross-moments can also be obtained in the same way as for the mean values (i.e., by taking higher order derivatives of the respective pgf's), but are omitted because of their size. Variances are however illustrated in figures in the next section.

6 Calculation of the tail probabilities

6.1 Preliminaries

Another important performance characteristic, besides the moments, is the (tail) distribution of the studied stochastic variables. The tail probabilities, i.e., the probability mass function (pmf) for larger values, typically represent the 'exceptional' situations in a queueing system. E.g. the probability that the delay is larger than a given value N or the packet loss are examples of interesting performance measures for which the calculation of the tail probability is usually sufficient. Consequently, the tail distribution is often used to impose statistical bounds on the guaranteed QoS for both classes.

Exact theoretical solutions for this inversion problem make use of the probability generating property of pgf's

and of residue theory. However, since these solution methods need a lot of derivations, they are often quite unpractical. Therefore, we will use an *approximate* solution technique, which is known to be quite popular: the dominant-singularity approximation. In e.g. [9], it has been shown that the pmf $x(n)$ of a discrete variable X is - for high n - dominated by the contribution of the singularity with the smallest absolute value, of the corresponding pgf $X(z)$. Because of a property of pgf's, this *dominant* singularity is necessarily positive real and larger than 1. Further in this section, expressions of the tail probabilities of the total system contents, of the system contents of the high- and low-priority queue, and of the packet delay of class 1 and class 2 packets are derived using this dominant-singularity approximation (and Darboux's theorem, which is explained in Appendix A).

We assume in the remainder that the pgf's of $A_T(z)$, $A_1(z)$ and $A_2(z)$ and their derivatives go to infinity for z equal to their radii of convergence or for $z \rightarrow \infty$. This includes all 'usual' arrival processes, and is thus not a particular, restrictive assumption. Note however that these assumptions mean that the analytic regions of these functions include the unit circle (i.e., $|z| \in (1, 1 + \varepsilon)$, $\varepsilon > 0$). Therefore, heavy tailed arrival processes are not included.

6.2 Total system contents

The tail behavior of $U_T(z)$ has already been investigated in e.g. [25]. The dominant singularity s_T of $U_T(z)$ is a zero of $z - A_T(z)$ and since the first derivative of $U_T(z)$ stays finite for $z = s_T$, this singularity is a pole with multiplicity 1. Hence, we can approximate $U_T(z)$ by $\frac{K_T}{s_T - z}$, in the neighbourhood of its pole. K_T can be obtained by calculating $\lim_{z \rightarrow s_T} U_T(z)(s_T - z)$, with $U_T(z)$ given by (8):

$$K_T = (1 - \lambda_T) \frac{s_T(s_T - 1)}{A_T'(s_T) - 1}. \quad (27)$$

Using Darboux's theorem, we then find

$$\begin{aligned} u_T(n) &\triangleq \text{Prob}[u_T = n] \\ &\sim K_T s_T^{-n-1}. \end{aligned} \quad (28)$$

This constitutes a typical *geometric* (exponential) behavior (for convenience, we will call this *behavior A* from here on).

6.3 System contents of the low-priority queue

Since the pgf of the system contents of the low-priority queue (see expression (10)) is similar to the one of the total system contents, the system contents of the low-priority queue has an identical tail behavior (A) in the neighbourhood of its dominant singularity s_2 , i.e., $U_2(z) \approx \frac{K_2}{s_2 - z}$ (s_2 is a zero of $z - A_2(z)$ and lies on the positive real axis). Obviously, K_2 can be obtained in a similar way as K_T . Using Darboux's theorem yields

$$u_2(n) \triangleq \text{Prob}[u_2 = n]$$

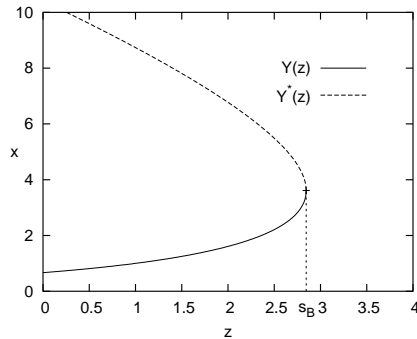


Figure 2: Solutions of $x - A(z, x) = 0$

$$\sim (1 - \lambda_2) \frac{(s_2 - 1)s_2^{-n}}{A_2'(s_2) - 1}. \quad (29)$$

6.4 The function $Y(z)$

The derivation of the tail behavior of the system contents of the high-priority queue, and of the packet delay of class 1 and class 2 packets, is not so straightforward, since it is not a priori clear from expressions (9), (15) and (19) which singularity of $U_1(z)$, $D_1(z)$ and $D_2(z)$ is dominant. This is due to the occurrence of the function $Y(z)$ in these expressions. This function, which is only implicitly defined, will be examined first (on the positive real axis).

As z increases along the positive real axis, a branch point s_B is encountered in $Y(z)$ where $Y'(z) \rightarrow \infty$ (see e.g. [5], [17] and [25] for similar cases). As a consequence, s_B is the solution of

$$\begin{cases} Y(s_B) = A(s_B, Y(s_B)) \\ Y'(s_B) \rightarrow \infty \end{cases} \Rightarrow \begin{cases} Y(s_B) - A(s_B, Y(s_B)) = 0 \\ A^{(2)}(s_B, Y(s_B)) = 1 \end{cases}, \quad (30)$$

with $A^{(2)}(z_1, z_2) \triangleq \frac{\partial A(z_1, z_2)}{\partial z_2}$. For values beyond s_B , $Y(z)$ is no longer properly defined. However, a second real and positive solution $Y^*(z)$ of the functional equation $x - A(z, x) = 0$, which is real and positive for z real, exists and it decreases as z increases (see Figure 2). Both solutions coincide for $z = s_B$. Note that $Y'(s_B)$ is infinite but that $Y(s_B)$ remains finite. Applying the results of [12], one can show that in the neighbourhood of s_B , $Y(z)$ is approximately given by

$$Y(z) \approx Y(s_B) - K_Y (s_B - z)^{1/2}, \quad (31)$$

with

$$K_Y = \sqrt{\frac{2A^{(1)}(s_B, Y(s_B))}{A^{(22)}(s_B, Y(s_B))}}, \quad (32)$$

where $A^{(1)}(z_1, z_2) \triangleq \frac{\partial A(z_1, z_2)}{\partial z_1}$ and $A^{(22)}(z_1, z_2) \triangleq \frac{\partial^2 A(z_1, z_2)}{\partial z_2^2}$ and which is found by substituting $z = s_B$ in expression (31). Since $Y(z)$ appears in the expressions of $U_1(z)$, $D_1(z)$ and $D_2(z)$, s_B is also a singularity of these

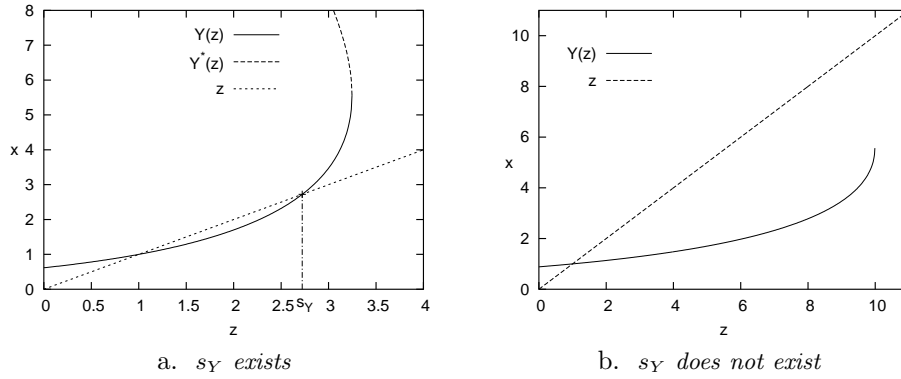


Figure 3: The singularity s_Y

pgf's, and may thus play a role in the tail behavior of these stochastic variables.

6.5 System contents of the high-priority queue

A second potential singularity s_Y - besides s_B - of $U_1(z)$ on the positive real axis is given by the (dominant) solution of $z - Y(z) = 0$ (see Figure 3a.). Note that since $Y(z)$ appears in this equation and since $Y(z)$ does not exist for $z > s_B$, this possible solution s_Y has to be smaller than s_B . However, the equation $z - Y(z) = 0$ does not always have a solution, in which case s_Y does not exist (see Figure 3b.). If s_Y does exist, then it is easily proven that $s_Y = s_T$.

The tail behavior of the system contents of the high-priority queue is thus characterized by s_Y or s_B , depending on which is the dominant singularity. Three cases may occur: s_Y exists and $s_Y = s_T < s_B$, s_Y exists and $s_Y = s_T = s_B$ or s_Y does not exist. In the first case, the singularity $s_Y (= s_T)$ is dominant. This singularity is a pole with multiplicity 1 (since the first derivative of $U_1(z)$ stays finite for $z = s_Y$). Consequently, $U_1(z) \approx \frac{K_1^{(1)}}{s_Y - z}$, for $z \rightarrow s_Y$. $K_1^{(1)}$ can be obtained in a similar way as K_T (see subsection 6.2):

$$K_1^{(1)} = (1 - \lambda_T) \frac{A_1(s_Y)(s_Y - 1)^2}{Y'(s_Y) - 1}. \quad (33)$$

In the second case, furthermore called a *boundary* case, s_Y and s_B coincide. We first study the behavior of $U_1(z)$ in the neighbourhood of s_B . Using expression (31) in (9), yields

$$U_1(z) \approx (1 - \lambda_T) \frac{A_1(z)(z - 1) \left(s_B - K_Y(s_B - z)^{1/2} - 1 \right)}{(s_B - z)^{1/2} \left((s_B - z)^{1/2} + K_Y \right) (A_1(z) - 1)}, \quad (34)$$

where we have also used the fact that $Y(s_B) = s_B$ (since $s_Y = s_B$ and $s_Y = Y(s_Y)$). This leads to $U_1(z) \approx \frac{K_1^{(2)}}{(s_B - z)^{1/2}}$ in the neighbourhood of its dominant singularity, with

$$K_1^{(2)} = (1 - \lambda_T) \frac{A_1(s_B)(s_B - 1)^2}{K_Y(A_1(s_B) - 1)}. \quad (35)$$

In the third case, i.e., when s_Y does not exist, the branch point s_B is dominant. By substituting expression (31)

in (9), the behavior of $U_1(z)$ in the neighbourhood of s_B is investigated:

$$U_1(z) \approx (1 - \lambda_T) \frac{\left\{ \begin{array}{l} A_1(z)(z-1) \left(Y(s_B) - K_Y(s_B - z)^{1/2} - 1 \right) \\ \times \left(z - Y(s_B) - K_Y(s_B - z)^{1/2} \right) \end{array} \right\}}{\left((z - Y(s_B))^2 - K_Y^2(s_B - z) \right) (A_1(z) - 1)}. \quad (36)$$

This expression leads to $U_1(z) \approx U_1(s_B) - K_1^{(3)}(s_B - z)^{1/2}$ in the neighbourhood of s_B , with

$$K_1^{(3)} = (1 - \lambda_T) \frac{K_Y A_1(s_B) (s_B - 1)^2}{(A_1(s_B) - 1) (s_B - Y(s_B))^2}. \quad (37)$$

We have now expressions for the behavior of $U_1(z)$ in its dominant singularity for the three possible cases. By furthermore using these expressions, and Darboux's theorem, we find the tail probabilities for these three cases:

$$u_1(n) \triangleq \text{Prob}[u_1 = n] \sim \begin{cases} K_1^{(1)} s_Y^{-n-1} & \text{if } s_Y = s_T < s_B \\ \frac{K_1^{(2)} n^{-1/2} s_B^{-n}}{\sqrt{\pi s_B}} & \text{if } s_Y = s_T = s_B \\ \frac{K_1^{(3)} n^{-3/2} s_B^{-n}}{2\sqrt{\pi/s_B}} & \text{if } s_Y \text{ does not exist} \end{cases}, \quad (38)$$

with the constants $K_1^{(i)}$ ($i = 1, 2, 3$) given by (33), (35) and (37) respectively. The first expression constitutes a typical geometric tail behavior (i.e., behavior A), while the two others are of a *non-geometric* nature. We call the two latter ones *behavior B* and *behavior C* respectively.

6.6 Packet delay of class 1

The tail behavior of $D_1(z)$ is characterized by the same singularities that characterize the tail behavior of $U_1(z)$, i.e., s_B and s_Y . Consequently, the tail behavior of $D_1(z)$ can also be divided into three types. In the same way as $U_1(z)$, $D_1(z)$ can be approximated in the neighbourhood of its dominant singularity by

$$D_1(z) \approx \begin{cases} \frac{\hat{K}_1^{(1)}}{s_Y - z} & \text{if } s_Y = s_T < s_B \\ \frac{\hat{K}_1^{(2)}}{(s_B - z)^{1/2}} & \text{if } s_Y = s_T = s_B \\ D_1(s_B) - \hat{K}_1^{(3)}(s_B - z)^{1/2} & \text{if } s_Y \text{ does not exist} \end{cases}, \quad (39)$$

where the constants $\hat{K}_1^{(i)}$ ($i = 1, 2, 3$) can be found similarly as the constants $K_1^{(i)}$ ($i = 1, 2, 3$), i.e., by investigating the behavior of $D_1(z)$ in the neighbourhood of its dominant singularity. Using Darboux's theorem, we find the tail probabilities for the three possible cases:

$$d_1(n) \triangleq \text{Prob}[d_1 = n]$$

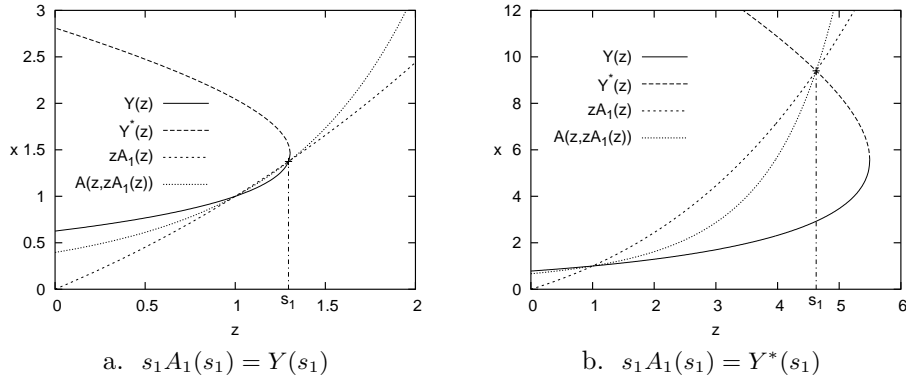


Figure 4: The singularity s_1

$$\sim \begin{cases} \frac{(1 - \lambda_T)}{\lambda_1} \frac{(s_Y - 1)}{(Y'(s_Y) - 1)} s_T^{-n} & \text{if } s_Y = s_T < s_B \\ \frac{(1 - \lambda_T)}{\lambda_1 K_Y} \frac{s_B (s_B - 1)}{\sqrt{\pi s_B}} n^{-1/2} s_B^{-n} & \text{if } s_Y = s_T = s_B \\ \frac{(1 - \lambda_T) K_Y}{\lambda_1} \frac{s_B (s_B - 1)}{2\sqrt{\pi/s_B} (s_B - Y(s_B))^2} n^{-3/2} s_B^{-n} & \text{if } s_Y \text{ does not exist} \end{cases} . \quad (40)$$

6.7 Packet delay of class 2

Finally, we investigate the tail behavior of $D_2(z)$. The tail behavior of $D_2(z)$ may again be influenced by the branch point s_B . Furthermore, two other singularities play a role, namely the dominant positive real (> 1) zeros s_Y of $z - Y(z)$ and s_1 of $zA_1(z) - A(z, zA_1(z))$. The potential singularity s_Y appears also in the tail behavior of $U_1(z)$ and $D_1(z)$, so we refer to the corresponding subsections for more details on s_Y . In this subsection, we first focus on s_1 and we will show that s_1 is not always a singularity.

Since s_1 is a zero of $zA_1(z) - A(z, zA_1(z))$, it is easily seen that $(x, z) = (s_1 A_1(s_1), s_1)$ is a solution of $x - A(z, x) = 0$. As a consequence, s_1 has to be smaller than s_B , since this equation has no solution for $z > s_B$ (see subsection 6.4). Furthermore, we have shown in subsection 6.4 that the equation $x - A(z, x) = 0$ has two possible real solutions - namely $(Y(z), z)$ and $(Y^*(z), z)$ - for $z < s_B$ positive real. Therefore, $s_1 A_1(s_1) = Y(s_1)$ or $s_1 A_1(s_1) = Y^*(s_1)$, depending on the pgf's $zA_1(z)$ and $A(z, zA_1(z))$. Both cases are illustrated in Figure 4, where the functions $Y(z)$, $Y^*(z)$, $zA_1(z)$ and $A(z, zA_1(z))$ are shown. Note that $s_1 A_1(s_1) = Y(s_1)$ if $s_1 A_1(s_1) < Y(s_B)$ and $s_1 A_1(s_1) = Y^*(s_1)$ if $s_1 A_1(s_1) > Y(s_B)$ and that it is also possible that $s_1 A_1(s_1) = Y(s_B)$, in which case the branch point s_B and s_1 coincide. We can now easily verify that in case $s_1 A_1(s_1) = Y(s_1)$, s_1 is also a zero of the numerator of $D_2(z)$ (see expression (19)). As a result, s_1 is not a singularity of $D_2(z)$ in this case. If $s_1 A_1(s_1) = Y^*(s_1)$ on the other hand, s_1 is not a zero of the numerator. In this case, s_1 is a (potential dominant) singularity.

Summarizing, the tail behavior of $D_2(z)$ is characterized by s_B , s_Y or s_1 , depending on which singularity is dominant. As a consequence, 7 cases may occur:

- s_Y is dominant
- s_1 is dominant

- $s_Y = s_1$ is dominant
- s_B is dominant
- $s_Y = s_B$ is dominant
- $s_1 = s_B$ is dominant
- $s_Y = s_1 = s_B$ is dominant

The behavior of $D_2(z)$ for all these cases can be investigated similarly to the behavior of $U_1(z)$, as described in subsection 6.5. Do notice that in the third case (a boundary case), the coinciding poles s_Y and s_1 lead to a dominant singularity with multiplicity 2. $D_2(z)$ can then, in the neighbourhood of its dominant singularity $s_Y = s_1$, be approximated by $\frac{\hat{K}_2^{(3)}}{(s_Y - z)^2}$ (where $\hat{K}_2^{(3)}$ can be calculated by taking $\lim_{z \rightarrow s_T} D_2(z)(s_Y - z)^2$). We denote this type of behavior, which was not encountered before, by *behavior D*.

Using Darboux's theorem (see Appendix A), we eventually find

$$d_2(n) \triangleq \text{Prob}[d_2 = n] \sim \begin{cases} \hat{K}_2^* s_*^{-n-1} & \text{behavior A} \\ \frac{\hat{K}_2^* n^{-1/2} s_*^{-n}}{\sqrt{\pi s_*}} & \text{behavior B} \\ \frac{\hat{K}_2^* n^{-3/2} s_*^{-n}}{2\sqrt{\pi/s_*}} & \text{behavior C} \\ \hat{K}_2^* (n+1) s_*^{-n-2} & \text{behavior D} \end{cases}, \quad (41)$$

with s_* and \hat{K}_2^* general notations for the dominant singularity and the constants which can in each case separately be derived from the expression of $D_2(z)$. Behavior A constitutes a geometric behavior, as we have seen before. The others are of a non-geometric nature.

6.8 Concluding remarks

Note first that the behaviors exhibited in expressions (38), (40) and (41) do not generally occur in the “basic” one-dimensional queueing systems. However, they are quite common for the marginal distributions in various two-dimensional models (see e.g. [6], [16] and [20], and references therein). Secondly, it should be stated that it is also possible to give results more uniform than those in expression (41). Specifically, when e.g. the dominant singularity leading to behavior A is close to (but not coinciding with) the branch point leading to behavior C, then a transition range result should be obtainable by taking into account both singularities. Similarly, the transition between behaviors A and D, where two simple poles are close to each other, can be considered. If the dominant singularity approximations are thus too crude, one can opt for taking into account more singularities. We will show however in the next section, that the obtained approximations are good enough for our purpose, namely illustrating the performance of the m-HOL priority queue.

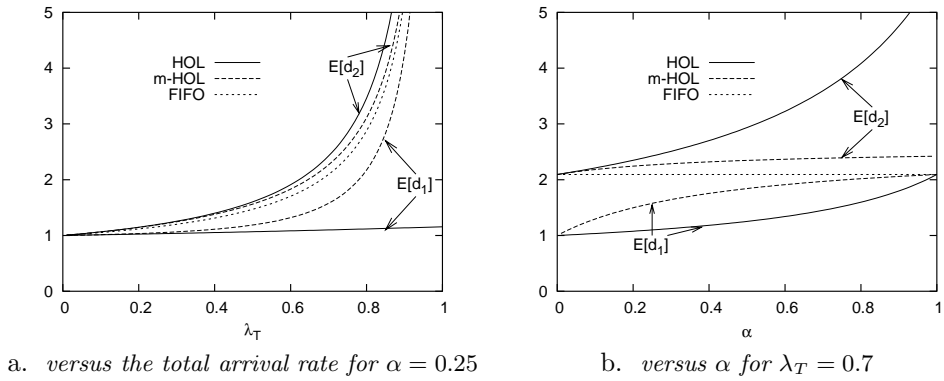


Figure 5: Mean value of packet delays

7 Numerical results

In this section, we present some numerical examples. We assume that the traffic of the two classes arrives according to a two-dimensional binomial process, which is fully characterized by the following joint pgf $A(z_1, z_2) = (1 - \frac{\lambda_1}{N}(1 - z_1) - \frac{\lambda_2}{N}(1 - z_2))^N$ (with $N = 16$ in the figures). The arrival rate of class- j traffic is then given by λ_j , and α is defined as the fraction of (high-priority) class-1 traffic in the overall traffic mix (i.e., $\alpha = \lambda_1/\lambda_T$). This is the arrival process to a queue in an $N \times N$ output-queueing switch with Bernoulli arrivals at its inlets and with uniform routing. Obviously, the numbers of class-1 and class-2 arrivals at an output queue during a slot are correlated: when m class-1 packets arrive during a slot ($0 \leq m \leq N$), the maximum number of class-2 arrivals during the same slot is limited by $N - m$ (because there are only N inlets).

We will now investigate the effect of the m-HOL priority scheduling discipline on some performance measures, focusing on the mean packet delay and the approximate tail probabilities. In the next two figures, we concentrate on the comparison between queues with a m-HOL priority scheduling, a HOL priority scheduling (see e.g. [25]) and a FIFO scheduling. In the latter, the packet delays for class 1 and class 2 packets are equal (i.e., independent of α), and can thus be calculated as if there is only one class of traffic arriving (determined by $A(z, z)$, see [8]). In Figure 5a., the mean value of the packet delay of both classes is shown for $\alpha = 0.25$, as a function of the total load for the m-HOL, HOL and FIFO scheduling respectively. In the case of a FIFO scheduling, the packet delay is the same for class 1 and class 2 packets (independent of α), and can thus be calculated as if there is only one class arriving according to an arrival process with pgf $A(z, z)$ (see e.g. [8]). First, we compare the m-HOL priority scheduling with a FIFO scheduling. The figure shows that the mean packet delay of class 1 is smaller for the m-HOL scheduling than for FIFO scheduling, while for the mean packet delay of class 2, the opposite again holds. This is quite logic: the packets of class 1 have priority over those of class 2. The difference is however limited: the mean delay of class 1 packets reduces moderately compared to FIFO scheduling, while the price to pay, a higher mean delay for class 2 packets, is kept small. Secondly, we compare the m-HOL priority scheduling with the HOL priority scheduling. In Figure 5a., we see that the mean packet delay of class 1 is smaller for HOL priority than for m-HOL priority. For the mean packet delay of class 2, the opposite holds. This is also expected: because of the jumps, packets of class 2 have a negative influence on the mean delay of class 1 packets. Note that the delays in the modified HOL priority

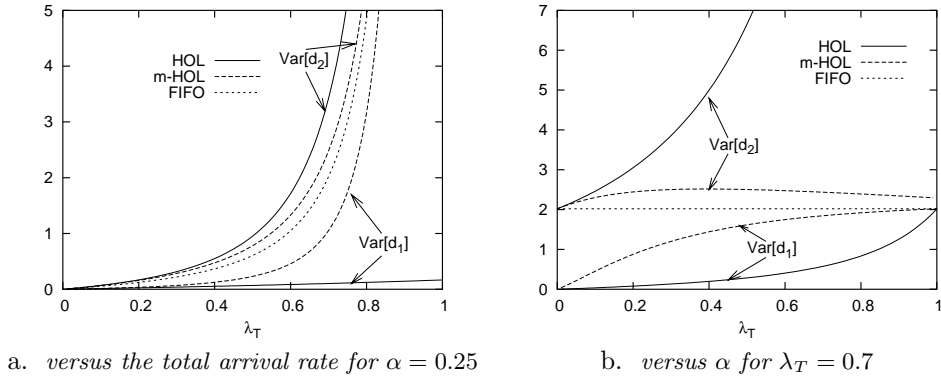


Figure 6: Variance of packet delays

queue can differ considerably from the mean packet delays in the original HOL priority queue, especially for high arrival rates. This is due to the fact that in every slot that both priority queues are not empty, a low-priority packet jumps to the high-priority queue. These jumps increase the delay of the high-priority packets (compared with HOL priority).

Figure 5b. shows the mean value of the packet delay of both classes when $\lambda_T = 0.7$ as a function of α , again for the m-HOL priority scheduling, HOL priority scheduling and FIFO scheduling respectively. In case of FIFO scheduling, the packet delay of both classes is independent of α . In case of the m-HOL priority scheduling and the HOL priority scheduling, the higher α (and thus the higher the fraction of class 1 packets in the overall traffic mix), the larger the mean delay of both classes. For these scheduling disciplines, when $\alpha = 0$ or $\alpha = 1$ (i.e., the overall traffic mix only exists of class 2 packets and class 1 packets respectively), the mean packet delay of class 2 and class 1 respectively, equals the mean delay for FIFO scheduling. This is quite obvious: there is only one type of packets in these cases. Note that similar determinations can be derived with respect to the variance of the packet delay, which is shown in Figure 6, as a function of λ_T (for $\alpha = 0.25$) and of α (for $\lambda_T = 0.7$) respectively.

From Figures 5 and 6, we can conclude that the differentiation caused by the m-HOL priority scheduling discipline is not as pronounced as that of the HOL priority scheduling. This m-HOL priority is clearly seen to be a compromise between HOL priority and FIFO scheduling, when the class 2 traffic has some delay requirements and/or when a moderate differentiation between both classes is wanted.

Furthermore, we have shown in section 6 that the tails of the system contents of the high-priority queue and of the packet delay of class 1 and class 2 packets can have several types of behavior, depending on which singularity of the respective pgf's is dominant. In case of the two-dimensional binomial arrival process considered in this section, the curves in Figures 7a. and 7b. show for which combinations of class 1 and class 2 arrival rates (λ_1 and λ_2), $s_Y = Y(s_B)$ (i.e., the singularities s_Y and s_B coincide), $s_1 A_1(s_1) = Y(s_B)$ (i.e., s_1 and s_B coincide) and $s_1 = s_Y$ (i.e., s_1 and s_Y coincide). The curve $s_Y = Y(s_B)$ represents the following: below the curve, s_Y does not exist, while above the curve, it exists. The curve $s_1 A_1(s_1) = Y(s_B)$ can be interpreted in a similar way: above the curve $s_1 A_1(s_1) = Y(s_1) < Y(s_B)$, while below the curve $s_1 A_1(s_1) = Y^*(s_1) > Y(s_B)$. In other words, below this curve, s_1 is a potential dominant singularity, while above the curve, s_1 is not a singularity (see subsection 6.7). Finally,

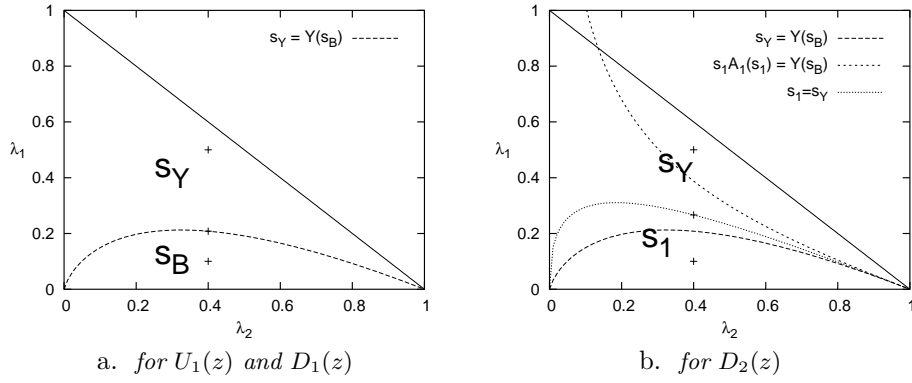


Figure 7: Regions for tail behavior as a function of the arrival rates of both classes

below the curve on which $s_1 = s_Y$, s_1 is smaller than s_Y , while above this curve s_1 is larger than s_Y . Note that in the area above the linear line $\lambda_1 + \lambda_2 = 1$, the total load is larger than 1, resulting in an unstable system. It can easily be seen that these curves split the (λ_2, λ_1) -space in regions in which one particular singularity is dominant, depending on which singularity (of those who exist in a region) has the smallest value. On the curves, singularities coincide. In Figure 7a. (i.e., the tail behavior of $U_1(z)$ and $D_1(z)$), the one curve $s_Y = Y(s_B)$ splits the (λ_2, λ_1) -space in two regions. Since s_Y does not exist below the curve, s_B is dominant in that region. In the region above the curve however, s_Y is dominant, since s_Y , when it exists, is smaller than s_B (see subsection 6.5). The tail behavior of $D_2(z)$ is characterized by three singularities (i.e., s_Y , s_1 and the branch point s_B), and is, as a consequence, split in more than two regions. By observing the values of the existing singularities in the several regions, one can easily determine that only the singularities s_Y and s_1 play a role in the tail behavior of $D_2(z)$ and that the branch point s_B is thus never dominant (see Figure 7b.).

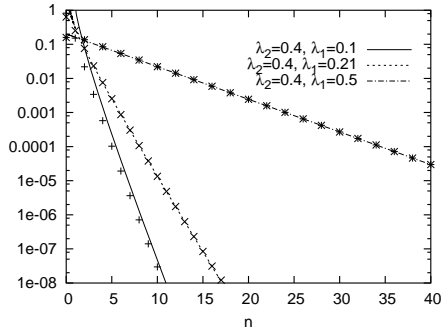


Figure 8: Tail behavior of $U_1(z)$

n	$\lambda_1 = 0.1$ (C)	$\lambda_1 = 0.21$ (B)	$\lambda_1 = 0.5$ (A)
10	0.698659	0.000109	0.000925
25	0.312592	0.000119	0.000003
40	0.203926	0.000098	0.000000
55	0.151814	0.000079	0.000000
70	0.121065	0.000065	0.000000
85	0.100732	0.000056	0.000000
100	0.086274	0.000048	0.000000
115	0.075458	0.000043	0.000000
130	0.067060	0.000038	0.000000

Table 1: Relative error

In the remaining figures, we compare, for the (λ_2, λ_1) -combinations indicated by the marks in Figures 7a. and 7b., the calculated asymptotic (tail) approximations for $\text{Prob}[u_1 = n]$ (38), $\text{Prob}[d_1 = n]$ (40) and $\text{Prob}[d_2 = n]$ (41) with exact numerical values computed via the so-called Cooley-Tukey Fast Fourier transform (FFT) algorithm (see e.g. [11] and [21]), i.e., an efficient algorithm to obtain the discrete Fourier transform (DFT) and its inverse (see e.g. [1] and [2]). Note that we have also performed an extensive number of simulations, but since the simulation results coincide with the exact numerical values, they are omitted in the figures. In the tables next to the figures,

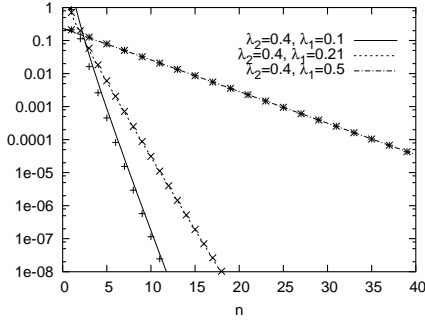


Figure 9: Tail behavior of $D_1(z)$

n	$\lambda_1 = 0.1$ (C)	$\lambda_1 = 0.21$ (B)	$\lambda_1 = 0.5$ (A)
10	0.552651	0.036175	0.001203
25	0.260790	0.014391	0.000003
40	0.172694	0.008981	0.000000
55	0.129488	0.006527	0.000000
70	0.103698	0.005126	0.000000
85	0.086523	0.004221	0.000000
100	0.074252	0.003587	0.000000
115	0.065040	0.003119	0.000000
130	0.057868	0.002758	0.000000

Table 2: Relative error

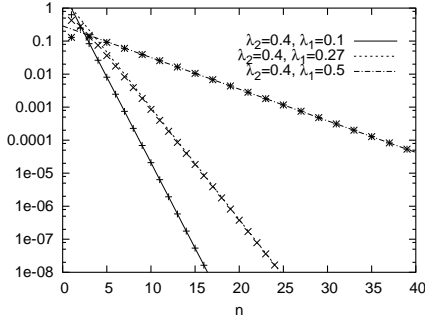


Figure 10: Tail behavior of $D_2(z)$

n	$\lambda_1 = 0.1$ (A)	$\lambda_1 = 0.27$ (D)	$\lambda_1 = 0.5$ (A)
10	0.000928	0.137424	0.005946
25	0.000003	0.036235	0.000032
40	0.000000	0.013792	0.000000
55	0.000000	0.006122	0.000000
70	0.000000	0.002971	0.000000
85	0.000000	0.001529	0.000000
100	0.000000	0.000822	0.000000
115	0.000000	0.000456	0.000000
130	0.000000	0.000259	0.000000

Table 3: Relative error

one can see the relative error between the asymptotic approximations and the exact values, for some n . The accuracy of the exponential (behavior A) asymptotic approximations is excellent, as can be derived from tables 1 to 2. Specifically, one can read that the relative error for all exponential approximations with respect to the exact numerical values is less than 0.6% for $n \geq 10$. The non-exponential approximations, with, e.g. for behavior C, a relative error of 0.6% for $\text{Prob}[d_1 = n]$, $n = 130$ (see table 2), are not as accurate as the exponential ones. This was also noticed in [3], where the large relative error is imputed to the slower rates of convergence in the corresponding expressions. While the error term is exponentially small for behavior A, the rates of convergence for behavior B, C and D are $O(1/\sqrt{n})$, $O(1/n)$ and $O(1/n)$ respectively. As to behaviors B and D, also called the boundary cases, it appears that the asymptotic approximations, compared to exact numerical values, are close to excellent.

8 Conclusions

In this paper, we have analyzed a queueing system with a modified HOL priority scheduling discipline and with two priority classes. The model included possible correlation between the number of arrivals of the two types of traffic during a slot. We have derived the joint generating function of the contents in the high- and low-priority queue and the generating functions of the delay of high- and low-priority traffic. Performance measures (such as means and variances) can be calculated from these pgf's. Furthermore, the tail distributions of system contents and packet delay are studied. We have shown that non-geometric tails can occur for the system contents of the

high-priority queue and for the packet delay of the high- and low-priority traffic. The impact of the dynamic priority scheduling on the performance characteristics is then shown by some numerical examples and we can see that the m-HOL priority scheduling disciplines (m-HOL) does what it is designed for: lowering the (mean) delay of the high-priority packets in comparison with a FIFO scheduling discipline, but, in contrast with the (original) HOL scheduling discipline, taking into consideration the (mean) delay of the low-priority packets. Finally, we have shown that our obtained approximations of the tails are excellent.

Appendix A: Darboux's theorem

Theorem 1 *Suppose $X(z) = \sum_{n=0}^{\infty} x(n)z^n$ with positive real coefficients $x(n)$ is analytic near 0 and has only algebraic singularities α_k on its circle of convergence $|z| = R$, in other words, in a neighbourhood of α_k we have*

$$X(z) \sim \left(1 - \frac{z}{\alpha_k}\right)^{-\omega_k} G_k(z), \quad (42)$$

where $\omega_k \neq 0, -1, -2, \dots$ and $G_k(z)$ denotes a nonzero analytic function near α_k . Let $\omega = \max_k \operatorname{Re}(\omega_k)$ denote the maximum of the real parts of the ω_k . Then we have

$$x(n) = \sum_j \frac{G_j(\alpha_j)}{\Gamma(\omega_j)} n^{\omega_j-1} \alpha_j^{-n} + o(n^{\omega-1} R^{-n}), \quad (43)$$

with the sum taken over all j with $\operatorname{Re}(\omega_j) = \omega$ and $\Gamma(\omega)$ the Gamma-function of ω (with $\Gamma(n) = (n-1)!$ for n discrete).

Acknowledgement

The authors would like to thank the anonymous referees and the editor for their constructive suggestions, which led to a considerable improvement of this paper. They also like to express their thanks to Dr. Koenraad Laevens for his valuable help to the material in section 7.

References

- [1] J. Abate, G. Choudhury, and W. Whitt. An introduction to numerical transform inversion and its application to probability models. In W. Grassman, editor, *Computational Probability*, pages 257–323. Kluwer, Boston, 2000.
- [2] J. Abate and W. Whitt. The fourier-series method for inverting transforms of probability distributions. *Queueing Systems*, 10:5–88, 1992.
- [3] J. Abate and W. Whitt. Asymptotics for M/G/1 low-priority waiting-time tail probabilities. *Queueing Systems*, 25:173–233, 1997.

- [4] J.J. Bae and T. Suda. Survey of traffic control schemes and protocols in ATM networks. *ACM Transactions on Networking*, 2(5):508–519, 1994.
- [5] E.A. Bender. Asymptotic methods in enumeration. *SIAM Review*, 16:485–515, 1974.
- [6] S. Borst, O. Boxma, and M. van Uitert. The asymptotic workload behavior of two coupled queues. *Queueing Systems*, 43:81–102, 2003.
- [7] H. Bruneel and B.G. Kim. *Discrete-time models for communication systems including ATM*. Kluwer Academic Publishers, Boston, 1993.
- [8] H. Bruneel, B. Steyaert, E. Desmet, and G.H. Petit. An analytical technique for the derivation of the delay performance of ATM switches with multiserver output queues. *International Journal of Digital and Analog Communication Systems*, 5:193–201, 1992.
- [9] H. Bruneel, B. Steyaert, E. Desmet, and G.H. Petit. Analytic derivation of tail probabilities for queue lengths and waiting times in ATM multiserver queues. *European Journal of Operational Research*, 76:563–572, 1994.
- [10] D.I. Choi, B.D. Choi, and D.K. Sung. Performance analysis of priority leaky bucket scheme with queue-length-threshold scheduling policy. *IEE Proceedings-Communications*, 145(6):395–401, 1998.
- [11] J. W. Cooley and J. W. Tukey. An algorithm for the machine calculation of complex fourier series. *Mathematics Computing*, 19:297–301, 1965.
- [12] M. Drmota. Systems of functional equations. *Random Structures & Algorithms*, 10(1-2):103–124, 1997.
- [13] P.A. Ganos, M.N. Koukias, and G.K. Kokkinakis. ATM switch with multimedia traffic priority control. *European Transactions on Telecommunications*, 7(6):527–540, 1996.
- [14] N.K. Jaiswal. *Priority queues*. Academic Press, New York, 1968.
- [15] L. Kleinrock. *Queueing systems volume II: computer applications*. John Wiley and Sons, New York, 1976.
- [16] C. Knessl and C. Tier. Two tandem queues with general renewal input II: asymptotic expansions for the diffusion model. *SIAM Journal on Applied Mathematics*, 59:1960–1997, 1997.
- [17] K. Laevens and H. Bruneel. Discrete-time multiserver queues with priorities. *Performance Evaluation*, 33(4):249–275, 1998.
- [18] Y. Lin and J.E. Kobza. Analysis of a delay-dependent priority discipline in an integrated multiclass traffic fast packet switch. *IEEE Transactions on Communications*, 38(5):659–685, 1990.
- [19] T. Maertens, J. Walraevens, and H. Bruneel. Performance analysis of a single-server queue with a HOL-PJ priority scheduling discipline. In *Proceedings of the Second International Working Conference on Performance Modelling and Evaluation of Heterogeneous Networks (HET-NETs '04)*, pages P42/1–P42/10, 2004.
- [20] J. Morrison. Two discrete-time queues in tandem. *IEEE Transactions on Communications*, 27:563–573, 1979.
- [21] W. H. Press, S. A. Teukosky, W. T. Vetterling, and B. P. Flannery. *Numerical Recipes in C – The Art of Scientific Computing*. Cambridge University Press, second edition, 1994.

- [22] M. Sidi. Discrete-time priority queues with partial interference. *IEEE Journal on Selected Areas in Communications*, 5(6):1041–1050, 1987.
- [23] H. Takagi. *Queueing analysis: a foundation of performance evaluation volume 1: vacation and priority systems, part 1*. Elsevier, Amsterdam, 1991.
- [24] J. Walraevens. *Discrete-time queueing models with priorities*. PhD thesis, Ghent University, 2004.
- [25] J. Walraevens, B. Steyaert, and H. Bruneel. Performance analysis of a single-server ATM queue with a priority scheduling. *Computers and Operations Research*, 30(12):1807–1829, 2003.