On the Statistical Structure of Losses Caused by the Buffer Overflow

ANDRZEJ CHYDZINSKI*
Institute of Computer Sciences,
Silesian University of Technology
Akademicka 16, 44-100 Gliwice
POLAND

Abstract: – The paper gives the answer to the question, whether calculating the cell loss ratio gives us information about the statistical structure of losses caused by the buffer overflow. In other words, can two queueing systems having exactly the same cell loss ratio have quite different consecutive cell loss probabilities? The answer to this question is studied for a variety of queueing systems with different buffer capacities, traffic intensities and the distributions of the service time.

Key-Words: – queueing systems, buffer overflow, consecutive cell losses, remaining service time

1 Problem formulation

If we want to estimate the performance of some queueing system with a view towards packet telecommunication, we usually think first of calculating the cell loss ratio. This coefficient is equal to the percentage of cells (packets) lost in the system due to the buffer overflow in a long time interval. Obviously, we want this ratio to be as small as possible, but it almost never (excluding some special cases) can be zero in real systems, in which the capacity of the buffer is finite.

The natural question that arises regarding the cell loss ratio is whether this coefficient tells us something about the "statistical structure" of losses caused by the buffer overflow or not. By the statistical structure of losses we mean just the distribution of the number of cells lost during one buffer overflow period. This distribution may be especially important in these applications, in which we cannot afford losing groups of consecutive arrivals. The question presented above may be formulated also in the following manner: Can two queueing systems with the same cell loss ratio have significantly different distributions of consecutive losses?

Apparently, it is possible. Moreover, the difference can be really great. In Example 4, Section 3 the probability of losing 10 cells in a row is almost 10^7 times greater in the second system than in the first one, although they have the same loss ratio.

What is even more interesting, this effect can be observed even if these two systems have a common buffer capacity, traffic intensity (ρ) or the distribution of the service time.

To prove all of this, four examples of pairs of queues with detailed calculations are presented in Section 3. Carrying out these calculations was possible thanks to recent results devoted to duration of the buffer overflow period [3]. Clearly, the distribution of the length of the buffer overflow period is the basis for the calculation of consecutive cell loss probabilities (see formula (1)).

The buffer overflow period (or, equivalently, the remaining service time) may be defined in the following manner (see, for instance [3]).

Let $X(t)$ denote the queue length at the moment, t of a single-server queue with a buffer of some finite size b. Let the initial queue size be

*This material is based upon work supported by the Polish Ministry of Scientific Research and Information Technology under Grant No. 3 T11C 014 26.

$X(0) = n$ and let $\tau^+(n, b)$ denotes the first moment in which the buffer is overflowed: $\tau^+(n, b) = \inf\{t > 0 : X(t) = b\}$. If $\zeta(n, b)$ stands for the first departure moment after $\tau^+(n, b)$ then the *buffer overflow period* is defined to be $\beta(n, b) = \zeta(n, b) - \tau^+(n, b)$.

In this paper we deal only with the case $n = b - 1$. This is due to the fact, that in queueing systems with the Poisson input stream only the distribution of the first buffer overflow period depends on initial queue length n. For every other the initial queue length is $b-1$. Thus $\beta(b-1, b)$ will be called *subsequent buffer overflow period* and the probability density function of $\beta(b-1, b)$ will be denoted by $h(z)$. The following model of the queue is considered: cells (packets) arrive according to the Poisson process with intensity ν. The service time is distributed according to a distribution function $F(x)$. The capacity of the buffer is finite and equal to b (including service position). In Kendall's notation such a model is called $M/G/1/b$.

Results connected with the distribution of the remaining service time are presented in articles [1, 2, 3, 4, 5, 6, 7]. Chronologically, the equilibrium distributions of past and remaining service times upon reaching a given level in an $M/G/1$ queue were shown in [1]. The properties of mean remaining service time in a $G/G/1$ queue were investigated in [6]. In [7] the formula for the mean remaining service time for a queue with a constant service rate and the Poisson input stream was obtained. A limiting formula (as $b \to \infty$) for the remaining service time in the $M/G/1/b$ model was given in [2]. An explicit form of the distribution of the remaining service time in the $M/G/1/b$ system was presented in [3]. Finally, the properties of the remaining service time in a batch arrival queue were studied in [4, 5].

2 Notation

Throughout the article the following notation will be used:

ν – the intensity of the input stream

$F(z)$ – the distribution function of service time, $z > 0$.

b – the capacity of the buffer (including service position)

$\rho = \nu \int_0^\infty x dF(x)$ – the offered load (traffic intensity) of the system

$h(z)$ – the probability density function for the duration of the subsequent buffer overflow period

r_n – the probability of n consecutive losses during the subsequent buffer overflow period

P_b – the cell loss ratio (blocking probability)

$\delta_{i,j}$ - the Kronecker symbol ($\delta_{i,j} = 1$ if $i = j$ and 0 otherwise)

$$I(x > y) = \begin{cases} 1 & \text{if} \quad x > y, \\ 0 & \text{otherwise} \end{cases}$$

3 Problem solution

The basic characteristic we are interested in is r_n – the probability of n consecutive losses during the subsequent buffer overflow period. It is easy to see that

$$r_n = \int_0^\infty \frac{e^{-\nu u}(\nu u)^n}{n!} h(u) du, \tag{1}$$

and now we need an effective way for finding $h(u)$. The formula for the tail of the distribution of the subsequent buffer overflow period was proven in [3]. Thus differentiating (14) in [3] we get:

$$h(z) = \frac{\sum_{k=1}^{b-1} a_k(z)(R_{b-k} - R_{b-1-k})}{\sum_{k=0}^{b-1} q_k (R_{b-k} - R_{b-1-k})}, \tag{2}$$

where

$$a_k(z) = \int_0^\infty \nu \frac{e^{-\nu u}(\nu u)^{k-1}}{(k-1)!} dF_u(z), \quad F_u(z) = F(u+z),$$

$$q_k = \int_0^\infty \frac{e^{-\nu u}(\nu u)^k}{k!} dF(u),$$

$$R_0 = 0,$$

$$R_{k+1} = \frac{1}{q_0}\left(\delta_{0,k} + R_k - \sum_{n=0}^{k} q_{n+1} R_{k-n}\right), \quad k \geq 0.$$

As the basic assumption is that the cell loss ratio (blocking probability) is the same in both queueing systems, we also have to control this characteristic. Its value can be calculated in the following manner (see [8], p. 202, formula (1.18b)):

$$P_b = 1 - \frac{1}{\pi_0 + \rho}, \quad \pi_0 = \left(\sum_{k=0}^{b-1} \pi'_k\right)^{-1},$$

$$\pi'_0 = 1, \quad \pi'_{k+1} = \frac{1}{q_0}\left(\pi'_k - \sum_{j=1}^{k} \pi'_j q_{k-j+1} - q_k\right).$$

Now we are in a position to carry out all necessary calculations. P

3.1 Example 1

In the first example we consider System A with constant service time (= 1) and System B with exponential one. This reduces the necessary calculations, as in the second system $h(z)$ is also exponential. The capacity of the buffer is the same in both systems. Namely, we put:
System A:

$$\nu = 0.2, \quad b = 5,$$
$$F(z) = I(z > 1).$$

System B:

$$\nu = 9.82711 \cdot 10^{-2}, \quad b = 5,$$
$$F(z) = 1 - e^{-z}.$$

The value of the input stream intensity in System B was adjusted for giving the same cell loss ratio in both systems. Its common value is equal to:

$$P_b = 8.2643 \cdot 10^{-6}.$$

Shapes of distributions of the subsequent buffer overflow periods for this input data are presented in Figure 1. Consecutive cell loss probabilities are gathered in the following table:

n	r_n (System A)	r_n (System B)
1	$5.5024 \cdot 10^{-2}$	$8.1471 \cdot 10^{-2}$
2	$2.6790 \cdot 10^{-3}$	$7.2899 \cdot 10^{-3}$
3	$1.0987 \cdot 10^{-4}$	$6.5229 \cdot 10^{-4}$
4	$3.8484 \cdot 10^{-6}$	$5.8365 \cdot 10^{-5}$
5	$1.1684 \cdot 10^{-7}$	$5.2224 \cdot 10^{-6}$
6	$3.1173 \cdot 10^{-9}$	$4.6729 \cdot 10^{-7}$
7	$7.3984 \cdot 10^{-11}$	$4.1812 \cdot 10^{-8}$
8	$1.5787 \cdot 10^{-12}$	$3.7413 \cdot 10^{-9}$
9	$3.0565 \cdot 10^{-14}$	$3.3476 \cdot 10^{-10}$
10	$5.4120 \cdot 10^{-16}$	$2.9954 \cdot 10^{-11}$

We see that losing two cells in row is three times more probable in the second system than in the first one and this difference increases with n. Losing five cells in row is 50 times more probable in the second system than in the first one, etc.

3.2 Example 2

In the second example we want the input stream intensity to be common. To get this we have to adjust its value and the buffer sizes. They may be set, for instance:
System A:

$$\nu = 0.5217976, \quad b = 10,$$

$$F(z) = I(z > 1).$$

System B:

$$\nu = 0.5217976, \quad b = 17,$$

$$F(z) = 1 - e^{-z}.$$

The cell loss ratio is then

$$P_b = 7.5362 \cdot 10^{-6}$$

in both systems.

The densities $h(z)$ are presented in Figure 2, while consecutive cell loss probabilities have the following values:

n	r_n (System A)	r_n (System B)
1	$0.1556 \cdot 10^{-1}$	$2.2531 \cdot 10^{-1}$
2	$2.2858 \cdot 10^{-2}$	$7.7256 \cdot 10^{-2}$
3	$2.6918 \cdot 10^{-3}$	$2.6490 \cdot 10^{-2}$
4	$2.6248 \cdot 10^{-4}$	$9.0829 \cdot 10^{-3}$
5	$2.1757 \cdot 10^{-5}$	$3.1144 \cdot 10^{-3}$
6	$1.5651 \cdot 10^{-6}$	$1.0679 \cdot 10^{-3}$
7	$9.9321 \cdot 10^{-8}$	$3.6615 \cdot 10^{-4}$
8	$5.6347 \cdot 10^{-9}$	$1.2555 \cdot 10^{-4}$
9	$2.8888 \cdot 10^{-10}$	$4.3048 \cdot 10^{-5}$
10	$1.3506 \cdot 10^{-11}$	$1.4760 \cdot 10^{-5}$

Again, the difference is significant and increases with n. Pay attention to the fact that the offered load is also common:

$$\rho_A = \rho_B = 0.5217976.$$

3.3 Example 3

The previous result may suggest the question whether it is possible to find an example in which three parameters: ν, b, and ρ are common in both systems. Apparently, a construction of such example is possible.

System A:

$$\nu = 1, \quad b = 30,$$
$$F(z) = I(z > 2).$$

System B:

$$\nu = 1, \quad b = 30,$$
$$F(z) = 1 - e^{-0.5z}.$$

In both systems we have:

$$P_b = 0.5, \quad \rho = 2.$$

The shapes of $h(z)$ are presented in Figure 3. The values o r_n are the following:

n	r_n (System A)	r_n (System B)
1	$3.1138 \cdot 10^{-1}$	$2.2222 \cdot 10^{-1}$
2	$2.0038 \cdot 10^{-1}$	$1.4815 \cdot 10^{-1}$
3	$9.8083 \cdot 10^{-2}$	$9.8765 \cdot 10^{-2}$
4	$3.8683 \cdot 10^{-2}$	$6.5844 \cdot 10^{-2}$
5	$1.2765 \cdot 10^{-2}$	$4.3896 \cdot 10^{-2}$
6	$3.6199 \cdot 10^{-3}$	$2.9264 \cdot 10^{-2}$
7	$8.9971 \cdot 10^{-4}$	$1.9509 \cdot 10^{-2}$
8	$1.9901 \cdot 10^{-4}$	$1.3006 \cdot 10^{-2}$
9	$3.9650 \cdot 10^{-5}$	$8.6708 \cdot 10^{-3}$
10	$7.1863 \cdot 10^{-6}$	$5.7805 \cdot 10^{-3}$

3.4 Example 4

In the final example we will answer the question, if it is possible to construct an example in which service time distributions are equal in both systems. The following settings will solve the problem:

System A:

$$\nu = 0.2, \quad b = 5,$$
$$F(z) = I(z > 1).$$

System B:

$$\nu = 0.9085966, \quad b = 50,$$
$$F(z) = I(z > 1).$$

We have got the same constant service time and the cell loss ratio is in both cases:

$$P_b = 8.2643 \cdot 10^{-6}.$$

Shapes of distributions for this input data are presented in Figure 1. The consecutive cell loss probabilities have the following values:

n	r_n (System A)	r_n (System B)
1	$5.5024 \cdot 10^{-2}$	$2.4840 \cdot 10^{-1}$
2	$2.6790 \cdot 10^{-3}$	$6.7906 \cdot 10^{-2}$
3	$1.0987 \cdot 10^{-4}$	$1.4501 \cdot 10^{-2}$
4	$3.8484 \cdot 10^{-6}$	$2.5294 \cdot 10^{-3}$
5	$1.1684 \cdot 10^{-7}$	$3.7207 \cdot 10^{-4}$
6	$3.1173 \cdot 10^{-9}$	$4.7261 \cdot 10^{-5}$
7	$7.3984 \cdot 10^{-11}$	$5.2787 \cdot 10^{-6}$
8	$1.5787 \cdot 10^{-12}$	$5.2588 \cdot 10^{-7}$
9	$3.0565 \cdot 10^{-14}$	$4.7267 \cdot 10^{-8}$
10	$5.4120 \cdot 10^{-16}$	$3.8694 \cdot 10^{-9}$

We see, that in this example the differences are even greater than in the previous three examples.

4 Conclusions

This article gives the answer to the question, whether calculating the cell loss ratio gives us information about the statistical structure of losses caused by the buffer overflow or not. The answer is negative. This means that two queueing systems having exactly the same cell loss ratios may have quite different consecutive cell loss probabilities.

This effect can be quite strong (Example 4). Pay attention to the fact that all presented examples are rather simple. Introducing more sophisticated distributions of service time (heavy-tailed) would likely intensify the effect.

What is interesting, this effect can be observed even if the buffer capacities, traffic intensities or the distributions of the service time are equal in both systems.

References

[1] S. Asmussen. Equilibrium properties of the M/G/1 queue. Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete, 58:267-281, 1981.

[2] P. T. de Boer, V.F. Nicola and J.C.W. van Ommeren. The Remaining Service Time Upon Reaching a High Level in M/G/1 Queues. Queueing Systems, 39:55-78, 2001.

[3] A. Chydzinski. On the remaining service time upon reaching a target level in M/G/1 queues. Queueing Systems, vol. 47, issue 1/2, pp. 71-80, 2004.

[4] A. Chydzinski. On the Transient Remaining Service Time in a Large Buffer Queue. WSEAS Transactions on Communications. Issue 4, Vol. 3, 2004.

[5] A. Chydzinski. Duration of the buffer overflow period in a batch arrival queue. Submitted for publication in Performance Evaluation, 2004.

[6] D. Fakinos. The expected remaining service time in a single server queue. Operations Research, 30:1014-1018, 1982.

[7] E. Gelenbe, X. Mang and R. Önvural. Diffusion based statistical call admission control in ATM. Performance Evaluation, 27&28:411-436, 1996.

[8] H. Takagi, H. *Queueing analysis, A Foundation of Performance Evaluation, v. 2.* Elsevier Science Publishers, 1993.

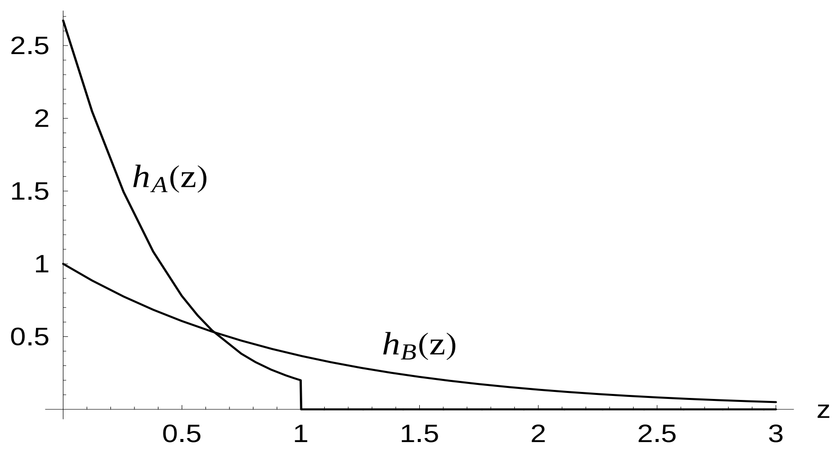

Figure 1: Density functions for the duration of the subsequent buffer overflow period in systems A and B (Example 1).

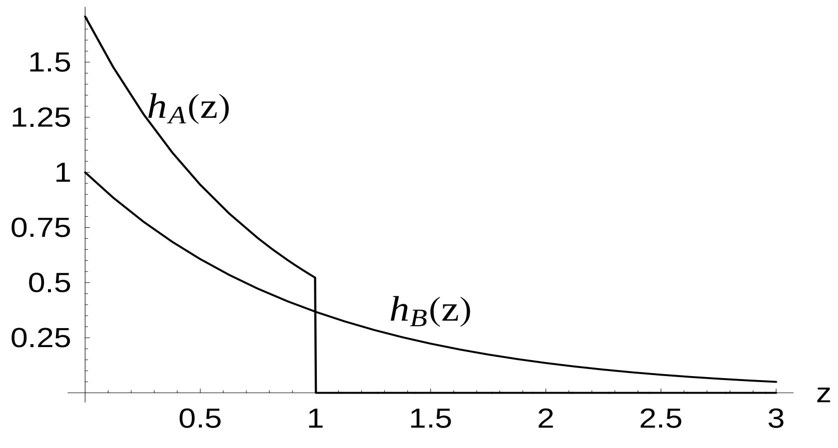

Figure 2: Density functions for the duration of the subsequent buffer overflow period in systems A and B (Example 2).

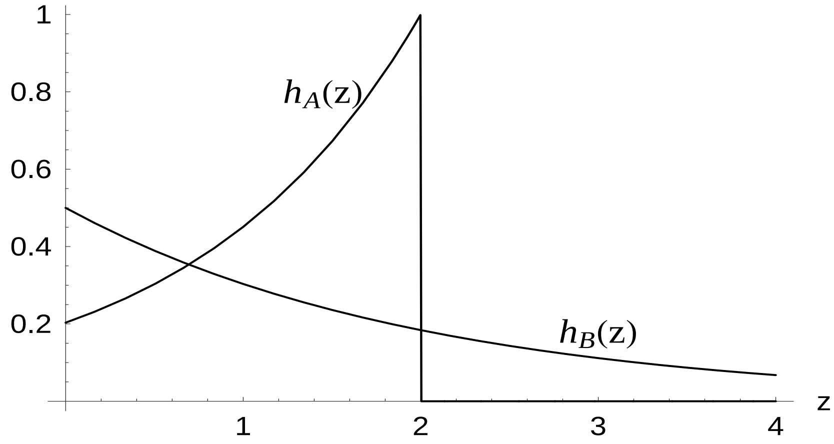

Figure 3: Density functions for the duration of the subsequent buffer overflow period in systems A and B (Example 3).

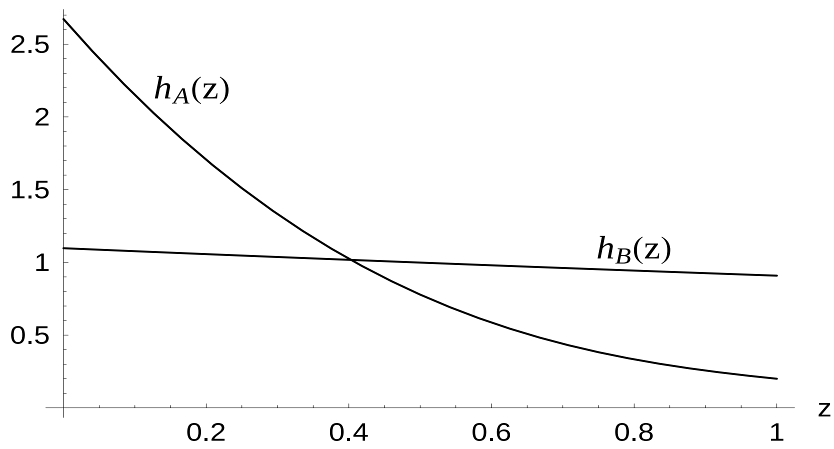

Figure 4: Density functions for the duration of the subsequent buffer overflow period in systems A and B (Example 4).