Characterization of Peak-Rate-limited Bandwidth-Efficient Discriminatory Processor Sharing

Pál L. Pályi[1], Attila Kőrösi[1], Balázs Székely[*2], József Bíró[*1], and Sándor Rácz[3]

[1] Department of Telecommunications and Media Informatics, Budapest University of Technology and Economics, H-1529 P.O. box 91, e-mail: palyi@tmit.bme.hu

[2] Institute of Mathematics, Budapest University of Technology and Economics

[3] Traffic Analysis and Network Performance Laboratory, Ericsson Research, Ericsson Hungary

Abstract: In this paper we characterize the state space of the discriminatory processor sharing service discipline with peak-rate limitations of the flows. We analyze a bandwidth-efficient rate sharing model, in which the unused capacity of the server by peak-rate limited flows is re-distributed among the non-limited flows. An efficient algorithmic approach is presented to determine which classes are subject to peak-rate limitations and based on this the bandwidth shares of flows of classes in a given state of this system.

Keywords: Discriminatory Processor Sharing; bandwidth-efficient; Peak rate

1 Introduction

Modern mobile telecommunications networks and high speed data packet services need the elaboration of new resource sharing, congestion avoidance and dimensioning methods, in order to ensure the appropriate Quality of Service (QoS) and Service Differentiation. For link dimensioning [1] [2] purposes bandwidth sharing models are needed, especially for the elastic-type (compressible) traffic flows. Such models describing the flow-level performance of elastic flows have been widely studied in the literature [3, 4, 5, 6]. In this paper, processor sharing-like models are considered, in which the service capacity (the

* The research of Székely and Bíró was partially supported by NKTH-OTKA Foundation research grant #77778 and #77802, respectively.

bandwidth) of the server (the link) is shared among the jobs (flows) according to some sharing principles.

Probably the very first and simplest (egalitarian) processor sharing model was presented by Kleinrock in [7], mainly motivated by the modeling of time-shared computer systems. In [8] a single-server processor-sharing system with several classes was analyzed. Classes are distinguished based on weights, and jobs in the classes have no limitations on their possible service rates (there are no peak-rate limitations), except the server capacity itself. The scheduling strategy considered divides the total capacity in unequal fractions among the different flows according to the corresponding weights; hence such models are called as discriminatory processor sharing (DPS). The paper provides solutions for the conditional expected response time (conditional average waiting time) of a class−*k* job with a given service time requirement (with a given size of the class−*k* flow) as well as for the unconditional response times.

In [9] the authors use the results from [8] and prove that – assuming the system is stable – for each class the expected unconditional response time is finite and that the expected conditional response time has an asymptote.

In [10], for multi-class egalitarian processor sharing queues, the authors show that the marginal queue length distribution for each class is equal to the queue length distribution of an equivalent single class processor sharing model with a random number of permanent customers. Similarly, the mean sojourn time (conditioned on the initial service requirement) for each class can be obtained by conditioning on the number of permanent customers.

Peak-rate limitations have been introduced and analyzed in a single class processor sharing system first in [11] (called M/G/R Processor Sharing Model) and several improved versions have been studied and proposed for dimensioning IP access networks, e.g. in [12] and [13]. In [3] multi-rate (peak rate limited) loss models for elastic traffic are evaluated. A structural characterization of reversibility is developed and used to build a non-egalitarian processor-sharing queueing discipline that admits a product-form solution. However, in this model a special type of Discriminatory Processor Sharing is considered, where corresponding peak rates and weights are in proportion to each other.

1.1 Discriminatory Processor Sharing

Discriminatory Processor Sharing (DPS) [14] is an important generalization of the (multi-class) egalitarian processor sharing discipline. In DPS, to each traffic class a weight is assigned; the weight of class-*i* is denoted by g_i. The bandwidth shares of flows are proportional to these weights. Flows from the same class always get the same bandwidth share. More formally, two requirements can be identified on the capacity shares in DPS: requirement-A: $c_i/c_j=g_i/g_j$ and requirement-B:

$\sum_{i=1}^{K} N_i c_i = C$ (c_i's are the bandwidth shares, and C is the server capacity), which are uniquely fulfilled[1] by

$$c_i = \frac{g_i C}{\sum_{j=1}^{K} g_j N_j}, \; i \in \{1,\ldots,K\},$$

where N_i is the number of class−i users in the system. This bandwidth share is also the solution of the following optimization problem:

$$\max_{\underline{x}} \sum_{i=1}^{K} N_i g_i \log x_i \quad \text{s.t.} \sum_{i=1}^{K} N_i x_i = C. \tag{1}$$

In this paper, we characterize the state space and the bandwidth sharing scheme of the peak-rate limited DPS with bandwidth-efficient rate sharing. The peak rate limitation means that each traffic class has its own maximal rate that is denoted by b_i for class-i. If there is enough capacity then the flows receive their peak bandwidths. When there is not enough capacity for all ongoing flows to get their peak rates, that is, $\sum_{i=1}^{K} N_i b_i > C$, then some flows or all flows will be "compressed" in the sense of their reduced service rates. This is the elastic "regime" of the model. On bandwidth-efficient rate sharing we mean that requirement-B should be fulfilled in the elastic regime of the model; that is, all bandwidth left by the uncompressed flows is to be redistributed among the compressed flows. This type of rate sharing is also referred to as Pareto-efficient in the literature [14].

The paper is organized as follows. In the next section we show that the bandwidth redistribution leads to a simpler and well interpretable bandwidth share calculation in the case of compressed flows. In Section 3 we present that there is a strict order of compression and we give a method for determining the set of compressed classes and the bandwidth shares.

1 Work-conserving property, i.e., either all flows get all the bandwidth they required or the system is serving on its full capacity.

2 Bandwidth Share Calculations in Peak-Rate Limited DPS

The non bandwidth-efficient processor sharing has been widely studied in the literature, but it does not prove to be a realistic model for real systems. In a non bandwidth-efficient case the total available capacity may not be used because residual capacity left by peak-rate limited flows is not (fully) redistributed among non peak-rate limited flows. In the models analyzed in [14, 15] there is no redistribution at all of unused capacity, hence bandwidth share of flow-*i* can be simply calculated in the following way:

$$c_i = \min\left(\frac{g_i C}{\sum_{j=1}^{K} g_j N_j}, b_i\right). \tag{2}$$

In [14] only the non bandwidth-efficient case is discussed and the bandwidth-efficient case is considered to be harder to analyze.

According to bandwidth-efficient rate sharing, unused capacity is redistributed among flows in proportion to their weights, so the calculation of the bandwidth shares and determining the set of compressed classes are somewhat more complicated. For a while, let us assume that the set of compressed ($\mathfrak{Z} : \{\forall i, c_i < b_i\}$) and uncompressed ($\mathfrak{U}$) classes are known in a given state $\underline{N} = (N_i,\ i \in \{1,\ldots,K\})$. In this case, $c_i = b_i,\ i \in \mathfrak{U}$. Since these flows cannot utilize their bandwidth shares, they leave

$$\sum_{i \in \mathfrak{U}} \left(\frac{g_i N_i}{\sum_{j=1}^{K} g_j N_j} C - N_i b_i\right)$$

Capacity, which is re-distributed among compressed flows. If $j \in \mathfrak{Z}$, the original bandwidth share is increased due to the redistribution. The redistribution should be proportional to the weights g_i, in order to keep a similar requirement to requirement-A. Between two compressed classes, $c_i / c_j = g_i / g_j,\ i, j \in \mathfrak{Z}$, and between a compressed and an uncompressed class, $c_i > c_k \frac{g_i}{g_k}, \forall i \in \mathfrak{Z}, \forall k \in \mathfrak{A}$.[2] This results

$$c_i = \frac{g_i}{\sum_{j \in \mathfrak{A} \bigcup \mathfrak{Z}} g_j N_j} C + \frac{g_i}{\sum_{k \in \mathfrak{Z}} g_k N_k} \sum_{l \in \mathfrak{A}} \left(\frac{g_l N_l}{\sum_{j=1}^{K} g_j N_j} C - N_l b_l \right), \ i \in \mathfrak{Z}. \tag{3}$$

Due to our assumption, constraint $c_i \leq b_i$ is fulfilled for $i \in \mathfrak{Z}$. This formula shows that identifying the service rate of classes and the set of compressed classes are more complicated in the bandwidth-efficient approach.

For an illustration of the differences between the bandwidth-efficient and the non bandwidth-efficient approaches see the Appendix. In the following, we consider the bandwidth-efficient method.

For implementing a calculation of bandwidth-efficient rate shares based on (3), we first show a simpler form of that equation, and then using this simpler form, we present a method for determining $\mathfrak{Z}$ and $\mathfrak{A}$.

Lemma 1 *The service rate of the compressed classes' users formulated in* (3) *can be re-written as*

$$c_i = \frac{g_i}{\sum_{j \in \mathfrak{Z}} g_j N_j} \left(C - \sum_{k \in \mathfrak{A}} N_k b_k \right), \ i \in \mathfrak{Z}. \tag{4}$$

The proof of this lemma is based on taking the right-hand side of (3) over a common denominator and performing a simplification, which eventually results in the right-hand side of (4). The immediate consequence is that c_i, $i \in \mathfrak{Z}$ can be considered as the bandwidth allocation of a reduced Discriminatory Processor Sharing system with capacity $(C - \sum_{k \in \mathfrak{A}} N_k b_k)$ and traffic classes $\mathfrak{Z}$ in state $\underline{N}$.

Note that $\mathfrak{Z}$ is unique for a given $\underline{N}$. We can distinguish between two cases. If $\sum_{i=1}^{K} N_i b_i \leq C$, $\mathfrak{Z}$ is empty so solution $c_i = b_i$ evidently fulfills constraint $\sum_{i=1}^{K} N_i c_i \leq C$. In the second case, $\sum_{i=1}^{K} N_i b_i > C$ and hence

2 This requirement is needed to ensure that $\mathfrak{Z}$ is unique for given $\underline{N}$.

$\sum_{i=1}^{K} N_i c_i = C$. From this, it follows that there exists a class i^* for which $c_{i^*} < b_{i^*}$, that is, there is at least one compressed class (i^*). If class j is also compressed, $\frac{g_{i^*}}{g_j} = \frac{c_{i^*}}{c_j}$ holds by definition. Thus, $c_j = \min\left(b_j, \frac{g_j}{g_{i^*}} c_{i^*}\right)$, which means that $c_j, \forall j$ can be calculated from the bandwidth share c_{i^*} of one compressed class i^*. As a consequence,

$$\sum_{j=1}^{K} N_j \cdot \min\left(b_j, \frac{g_j}{g_{i^*}} c_{i^*}\right) = C. \tag{5}$$

The left-hand side of (5) is monotonously increasing function of c_{i^*}, so while $\sum_{i=1}^{K} N_i b_i > C$, there is one solution for c_{i^*}. Therefore, there is one solution for each c_j.

Preliminary numerical calculations lead us to conjecture that the bandwidth allocation from (4) is a global solution of the following optimization problem, which differs from (1) in the constraint $x_i \in [0, b_i]$:

$$\max_{\underline{x}} \sum_{i=1}^{K} N_i g_i \log x_i \quad \text{s.t.} \quad \sum_{i=1}^{K} N_i x_i \le C \ \ \& \ \ \forall i \ x_i \in [0, b_i].$$

In the following section we present an algorithmic approach to determine the set of compressed classes $\mathfrak{Z}$.

3 Determining the Compression of Classes

In this section, we propose a method for determining the set of compressed classes $\mathfrak{Z}$ and also the bandwidth shares of flows from each class.

Let C denote the considered capacity. We distinguish among three disjunct cases considering the compression of classes; namely all classes are uncompressed, all classes are compressed, and there are compressed classes but at least one class is uncompressed.

Let $\mathfrak{S}$ be a subset of $\{1,\ldots,K\}$. In the following three steps below, $\mathfrak{S}$ will be adjusted. Initially, let $\mathfrak{S} = \{1,\ldots,K\}$.

Step-1: Check whether all classes are uncompressed. If it is true, then C is enough for peak rates of all flows, i.e.

$$\sum_{i \in \mathfrak{S}} N_i b_i \leq C . \tag{6}$$

In this case every flow gets its peak rate: $c_i = b_i, \ \forall i \in \mathfrak{S}$ and all classes are contained by $\mathfrak{A}$. No further steps are needed.

Step-2: Check whether all classes are compressed. If it is true, the bandwidth share of all flows are less than their peak rates, and there is no unused capacity to be redistributed, i.e.,

$$\frac{g_i}{\sum_{j \in \mathfrak{S}} N_j g_j} C < b_i, \forall i \in \mathfrak{S} . \tag{7}$$

The bandwidth share of flows are

$$c_i = \frac{g_i}{\sum_{j \in \mathfrak{S}} N_j g_j} C, \ \forall i \in \mathfrak{S},$$

so every class gets bandwidth share in proportion to their weights and all classes are contained by $\mathfrak{Z}$. No further steps are needed.

Step-3: In this case there are compressed classes, but at least one class is uncompressed, because none of the conditions in Step-1 (6) and Step-2 (7) is fulfilled. To determine which class is surely uncompressed we use the following equivalence

$$\frac{g_i}{b_i} < \frac{\sum_{j \in \mathfrak{S}} N_j g_j}{C}, \forall i \in \mathfrak{S} \Leftrightarrow \max_{i \in \mathfrak{S}} \frac{g_i}{b_i} < \frac{\sum_{j \in \mathfrak{S}} N_j g_j}{C} \tag{8}$$

The left-hand side of the relation simply comes from (7) by rearrangement. The equivalence above also means that if (7) is not fulfilled, then the right-hand side of (8) is also not fulfilled. Consequently, this surely uncompressed class is i',

$$i' = \arg\max_{i \in \mathfrak{S}} \frac{g_i}{b_i}, \text{ with bandwidth share } c_{i'} = b_{i'}.$$

For the remaining classes we should evaluate a reduced system where the effect of this class is considered according to Lemma-1; that is, the considered capacity is reduced by $N_{i'} b_{i'}$ ($C \leftarrow C - N_{i'} b_{i'}$) and only the remaining classes are considered ($\mathfrak{S} \leftarrow \mathfrak{S} \setminus \{i'\}$). For the reduced system we continue with Step-2.

The above described method can be summarized as the following algorithm (Algorithm 1):

1. $\mathfrak{Z} = \{1, 2, ..., K\}$

2. while $\max_{i \in \mathfrak{Z}} \left\{ \frac{g_i}{b_i} \right\} \geq \frac{\sum_{j \in \mathfrak{Z}} N_j g_j}{C}$ and $\mathfrak{Z} \neq \varnothing$ do

$$i' = \arg\max_{i \in \mathfrak{Z}} \left\{ \frac{g_i}{b_i} \right\}$$

$$\mathfrak{Z} \leftarrow \mathfrak{Z} \setminus \{i'\}$$

$$C \leftarrow C - N_{i'} b_{i'}$$

3. for $i = 1$ to K do

$$\text{if } i \in \mathfrak{Z} \text{ then } c_i = \frac{g_i}{\sum_{i \in \mathfrak{Z}} N_i g_i} C$$

$$\text{else } c_i = b_i .$$

Algorithm 1 Determining the set of compressed classes

An important consequence of the above described method is that the following order of classes:

$$\frac{g_1}{b_1} \leq \frac{g_2}{b_2} \leq \ldots \frac{g_K}{b_K},$$

based on the ratios g_i/b_i, is directly related to the compressed and uncompressed classes in such a way that:

if a class with higher *g*/*b* value is compressed, then all classes with lower *g*/*b* are compressed, and if a class with lower *g*/*b* value is uncompressed, then all classes with higher *g*/*b* are uncompressed. Also, observe that the compression order depends on neither the server capacity nor the number of users. In addition to this, the above described method also provides bandwidth shares of flows in each classes, as formulated in (4).

Conclusion

This paper is concerned with discriminatory processor sharing model with peak-rate limitations and a bandwidth-efficient rate sharing. We have shown that at a given state space (as a snapshot) the peak-rate limited DPS system includes a reduced-capacity DPS system over the compressed classes. Based on this, we have given a method to determine the set of compressed and uncompressed classes of flows and their bandwidth shares. It has also been shown that there is a strict order of classes in which they become compressed, and this order coincides with the order of ratios of the corresponding class weights and class peak rates (g_i/b_i).

The significance of the result presented in this paper lies in the fact that these are inevitably the starting points for both the analysis and simulation of Discriminatory Processor Sharing constrained by peak-rate limitations and unused capacity redistribution.

Acknowledgment

The fruitful discussions with Szilveszter Nádas from Ericsson are highly appreciated.

References

[1] R. Jain, “Congestion Control and Traffic Management in Atm Networks: Recent Advances and a Survey,” in *COMPUTER NETWORKS AND ISDN SYSTEMS*, 1995, pp. 1723-1738

[2] S. Nádas, S. Rácz, and P. Pályi, *Handbook of HSDPA/HSUPA Technology*. CRC Press, Taylor & Francis Group, 2010, ch. HSPA Transport Network Layer Congestion Control, pp. 297-330

[3] V. Koukoulidis, “A Characterization of Reversible Markov Processes with Applications to Shared-Resource Environments,” Ph.D. dissertation, Concordia University Montreal, Canada, 1993

[4] S. Rácz, B. P. Gerö, and G. Fodor, “Flow Level Performance Analysis of a Multi-Service System Supporting Elastic and Adaptive Services,” *Performance-Evaluation*, Vol. 49, pp. 451-469, 2002

[5] P. A. V. J. Lassila, P., “Dimensioning of Data Networks: a Flow-Level Perspective,” *European Transactions on Telecommunications*, Vol. 20, No. 6, pp. 549-563, 2008

[6] J. Roberts, “A Survey of Statistical Bandwidth Sharing,” *Computer Networks*, Vol. 45, No. 3, 2004

[7] L. Kleinrock, “Time-shared Systems: A Theoretical Treatment,” *J. of ACM*, Vol. 14, No. 2, pp. 242-261, 1967

[8] G. Fayolle, I. Mitrani, and R. Iasnogorodski, “Sharing a Processor among Many Job Classes,” *J. ACM*, Vol. 27, No. 3, pp. 519-532, 1980

[9] K. Avrachenkov, U. Ayesta, P. Brown, and R. Núńez-Queija, “Discriminatory Processor Sharing Revisited,” in *In: Proc. IEEE Infocom 2005, Miami FL*, 2005, pp. 784-795

[10] S. Cheung, H. van den Berg, and R. Boucherie, “Decomposing the Queue Length Distribution of Processor-Sharing Models into Queue Lengths of Permanent Customer Queues,” *Performance Evaluation*, Vol. 62, No. 1-4, pp. 100-116, 2005

[11] K. Lindberger, "Balancing Quality of Service, Pricing and Utilisation in Multiservice Networks with Stream and Elastic Traffic," in *ITC-16: International Teletraffic Congress*, 1999, pp. 1127-1136

[12] A. Riedl, T. Bauschert, and J. Frings, "A Framework for Multi-Service IP Network Planning," in *International Telecommunication Network Strategy and Planning Symposium (Networks)*, 2002, pp. 183-190

[13] Z. Fan, "Dimensioning Bandwidth for Elastic Traffic," *NETWORKING 2002: Networking Technologies, Services, and Protocols; Performance of Computer and Communication Networks; Mobile and Wireless Communications*, pp. 826-837, 2006

[14] U. Ayesta and M. Mandjes, "Bandwidth-Sharing Networks under a Diffusion Scaling," *Annals Operation Research*, Vol. 170, No. 1, pp. 41-58, 2009

[15] A. Lakshmikantha, R. Srikant, and C. Beck, "Differential Equation Models of Flow-Size-based Priorities in Internet Routers," *International Journal of Systems, Control and Communications*, Vol. 2, No. 1, pp. 170-196, 2010

Appendix

Comparison of bandwidth-efficient and non bandwidth-efficient limited approaches

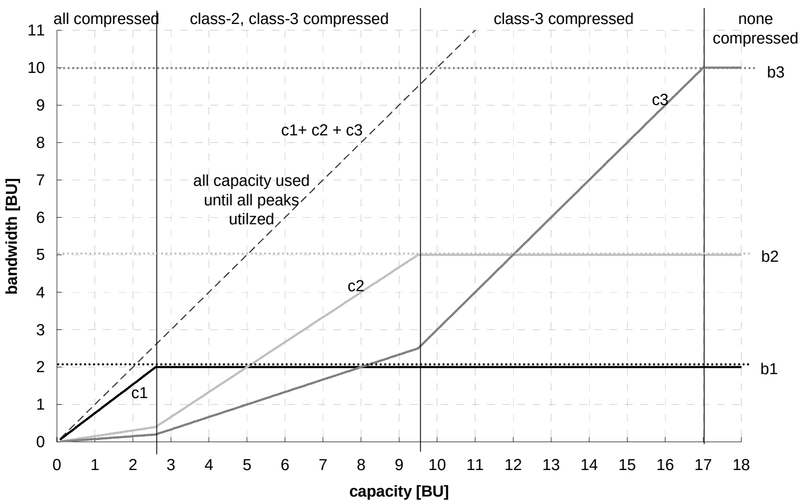

Figure 1

Peak-rate limited bandwidth-efficient DPS

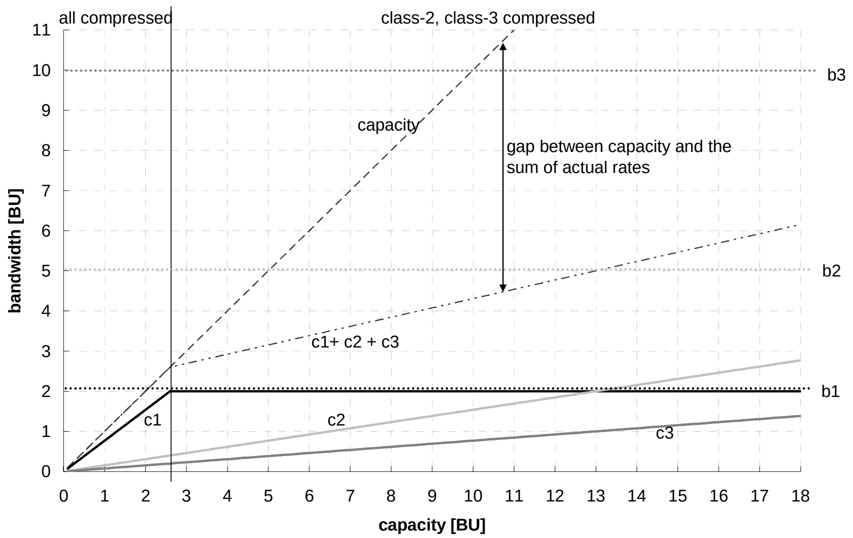

Figure 2

Peak-rate limited *non* bandwidth-efficient DPS

Table 1

Parameter settings

	Class-1	Class-2	Class-3
Peak rate (b_i)	2	5	10
Weight (g_i)	10	2	1
Number of users (N_i)	1	1	1

Figure 1 and **Figure 2** illustrate how the bandwidth-efficient and non bandwidth-efficient approaches are different from each other. Both figures show the same scenario, the only difference is the bandwidth-efficiency. See class parameters in **Table 1**. On the horizontal axis the capacity is shown. It is increased from 0 to 18 Bandwidth Units. On the vertical axis the actual service rate of the given class (c_i) is plotted. The peak rate of each class (b_i) is also shown.

The main difference between the two approaches is that in the case of the non-bandwidth-efficient approach, the total capacity is fully utilized only if all classes are compressed (**Figure 2**, where capacity is less than 2.7), i.e., not the peak rates are the limiting factors in the service rates. Otherwise, in this case the capacity is not utilized because residual capacity left by the peak-rate limited class 1 is not fully redistributed among non peak-rate limited class 2 and 3. In the case of the bandwidth-efficient approach (**Figure 1**), the available capacity is always fully utilized, except when the service rates of all flows are limited by their peak rates.

In **Figure 1**, four regions can be distinguished. If capacity is less than 2.7, all classes are compressed. This is the only region, where the bandwidth-efficient and the non bandwidth-efficient approaches give the same service rates, because there is no unused capacity left from peak-rate limited classes. If the capacity is not less than 2.7 but less than 9.5, then class-1 is no longer compressed, i.e., it gets its peak rate (b_1). Class-2 and class-3 are still compressed in proportion of their weights. If the capacity is not less than 9.5 but less than 17 then class-2 receives its peak rate also and is no longer compressed. In this region only class-3 is compressed, but gets all the capacity left from both peak-rate limited classes. If the capacity is not less than 17 then all classes receives their peak rates. In this region, because all classes are limited by their peak rates, further increase of the capacity does not increase the sum of the services rates of the three classes.

In **Figure 2**, only two regions can be distinguished. If capacity is less than 2.7, all classes are compressed. This is the only region, where the bandwidth-efficient and the non bandwidth-efficient approaches give the same service rates. Therefore, this region of **Figure 2** is the same as that of **Figure 1**. If the capacity is not less than 2.7 then class-1 already gets its peak rate (b_1), class-2 and class-3 are still compressed and their service rate is calculated according to the same formula as in region 1. Since capacity left by the peak rate limited class-1 is not redistributed among class-2 and class-3, there is a gap between the capacity and the sum of the actual service rates. The higher the capacity, the larger this gap gets.

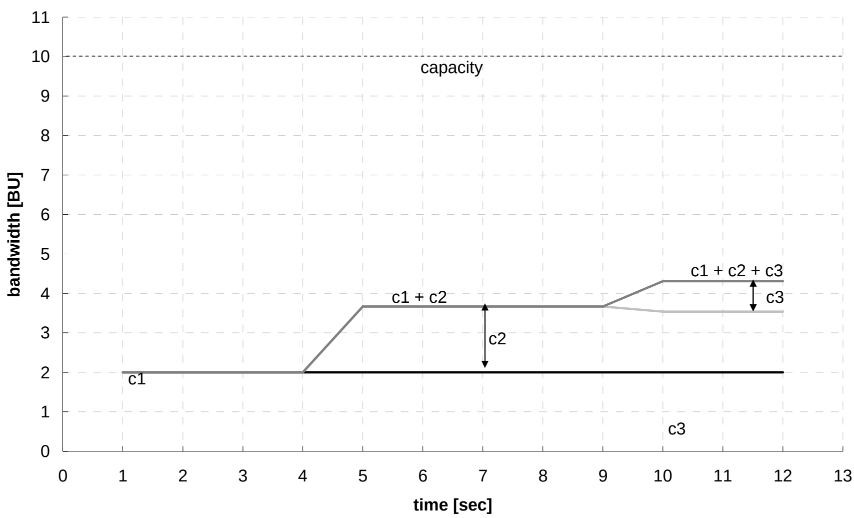

Figure 3

Peak-rate limited *non* bandwidth-efficient DPS, capacity=10 BU

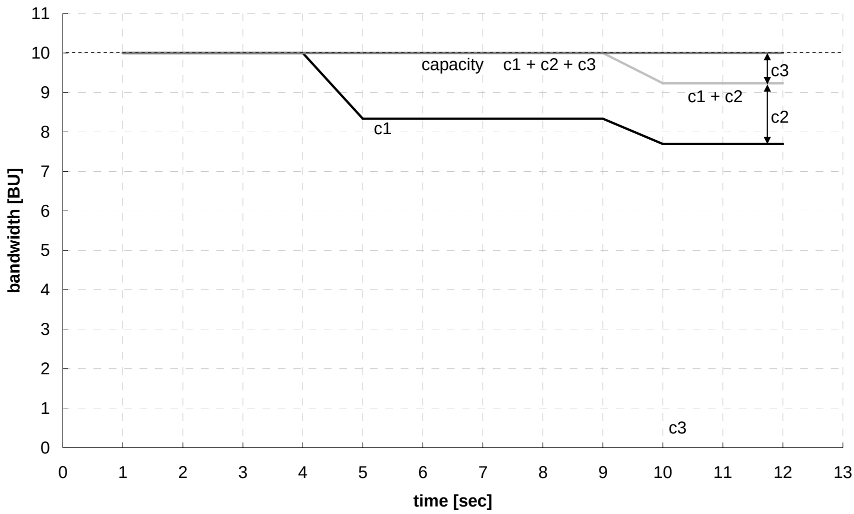

Figure 4
DPS with no peak rates, capacity=10 BU

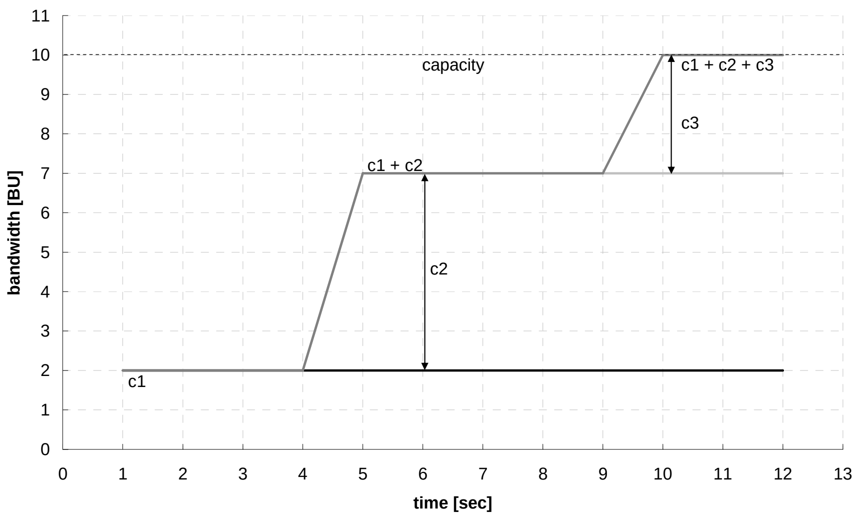

Figure 5
Peak-rate limited bandwidth-efficient DPS, capacity=10 BU

Figure 3 and **Figure 5** give an other illustration of the difference between the bandwidth-efficient and the non-bandwidth efficient approaches. In both figures the same scenario is depicted, apart from the bandwidth-efficiency. See class parameters in **Table 1**. Capacity is now fixed at 10 BU. **Figure 4** also shows the same scenario but no peak rates are used. In these three figures, time is plotted on the horizontal axis. On the vertical axis the actual service rate of the given class (c_i) is shown in a cummalative way. It means that instead of plotting c_1, c_2, c_3 individually, c_1, the sum of c_1 and c_2, and the sum of c_1, c_2, and c_3 is plotted. At 0 sec no flows are in the system. At 1 sec a flow from class-1 arrives, at 5 sec a flow from class-2 arrives, finally at 10 sec, a flow from class-3 arrives.

Figure 3 shows the peak-rate limited non bandwidth-efficient approach. When class-1 arrives at 1 sec, it gets its peak rate. When class-2 arrives at 5 sec, its service rate is calculated according to (2), therefore it cannot utilize its peak rate. The same applies to class-3 when it arrives at 10 sec. The total capacity can still not be utilized.

Figure 4 depicts the same scenario with *no* peak rate limitations. When class-1 arrives at 1 sec, it can use the total capacity. When class-2 arrives at 5 sec the two classes share the capacity in proportion of their weights. When class-3 arrives at 10 sec the capacity is shared among three flows in proportion of their weights. The total capacity is always utilized since no peak rates are limiting the flows service rates.

Figure 5 shows the peak-rate limited bandwidth-efficient approach. When class 1 arrives at 1 sec, it gets its peak rate. When class-2 arrives at 5 sec it also gets its peak rate since the sum of the peak rates (7 BU) is still less than the capacity (10 BU). When class-3 arrives at 10 sec, it gets compressed since it has the smallest weight and otherwise the capacity would be exceed. The total capacity is only utilzed after 10 sec because until this time peak rate is limiting both flows.