



HAL
open science

Randomized Two-Valued Bounded Delay Online Buffer Management

Christoph Dürr, Shahin Kamali

► **To cite this version:**

Christoph Dürr, Shahin Kamali. Randomized Two-Valued Bounded Delay Online Buffer Management. 2020. hal-02558500v1

HAL Id: hal-02558500

<https://hal.science/hal-02558500v1>

Preprint submitted on 29 Apr 2020 (v1), last revised 18 Nov 2020 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Randomized Two-Valued Bounded Delay Online Buffer Management

Christoph Dürr* Shahin Kamali†

April 29, 2020

Abstract

In the bounded delay buffer management problem unit size packets arrive online to be sent over a network link. The objective is to maximize the total weight of packets sent before their deadline. In this paper we are interested in the two-valued variant of the problem, where every packet has either low (1) or high priority weight ($\alpha > 1$). We show that its randomized competitive ratio against an oblivious adversary is $1 + (\alpha - 1)/(\alpha^2 + \alpha)$.

Keywords: competitive ratio, oblivious adversary, buffer management, maximum throughput scheduling.

1 Introduction

Online Buffer Management is a scheduling problem which arises in network routers. Packets arrive online to be sent on a specific link, the goal is to maximize the total weight of sent packets under various constraints. Different models have been studied in the past, the FIFO model, where the router stores pending packets in a limited capacity buffer, and the bounded delay model where packets can stay only limited time in the buffer.

We are particularly interested in the later model, which can be formalized as the following scheduling problem. Jobs of unit length arrive online on a single machine. Job j arrives at a release time $r_j \in \mathbb{N}$, has a deadline $d_j \in \mathbb{N}$ and a weight $w_j \in \mathbb{R}^+$. At every time point $t \in \mathbb{N}$, there is a set of pending jobs, which is updated by jobs released at time t and by jobs expiring at time t . The algorithm can choose to schedule one of the pending jobs, at every time point, and the goal is to maximize the total weight of the scheduled jobs. The usual restriction on the online algorithm is that these decisions need to be made without knowledge of future arriving jobs. Its performance is measured by the competitive ratio, which compares the objective value reached by the algorithm with the objective value of the optimal schedule, maximizing the ratio over all possible instances. The deterministic competitive ratio of the problem is defined as the best competitive ratio among all deterministic online algorithms. As randomization usually helps in these online settings, we also consider the randomized competitive ratio against an oblivious adversary. The term *adversary*

*Sorbonne Université, CNRS, Laboratoire d'informatique de Paris 6, LIP6, Paris, France. Corresponding author: christoph.durr@lip6.fr. <https://orcid.org/0000-0001-8103-5333>

†Department of Computer Science, University of Manitoba, Winnipeg, Canada. shahin.kamali@umanitoba.ca. <https://orcid.org/0000-0003-1404-2212>

comes from the game theoretical setting of an online problem, played between an online algorithm which needs to make decisions at every time point and a malicious adversary which generates the instance over time. Other adversarial models exist but are not considered in this paper.

2 Related work

This problem has been introduced in 2001, together with a lower bound of $\varphi \approx 1.618$ on the deterministic competitive ratio [6] and [8] and an upper bound of 2, achieved by the MAX-PACKET algorithm which schedules always the heaviest pending job. This gap between lower and upper bound generated a series of studies of variants of this problem, mostly with restrictions on the deadlines of the job, such as agreeable deadlines (jobs can be ordered by release times and deadlines at the same time) or s -uniform or s -bounded instances ($d_j - r_j$ is s or at most s) for some parameter s . Eventually in SODA 2019 the gap was closed, with a sophisticated algorithm and analysis [9]. The randomized competitive ratio in the oblivious adversary model of this problem is still open, and known to be between 1.582 [1] and 1.25 [4].

The paper [3] considered a variant of the problem, where every job has not unit processing time, but a processing time $p \in \mathbb{N}$, all jobs are tight in the sense $d_j = r_j + p$, and the algorithm can preempt job executions. In this setting the authors give a barely random algorithm with competitive ratio 2, and showed that this optimal for barely random algorithms. By scaling this is equivalent to the problem with unit processing time jobs by fractional release times and deadlines. In our problem these times are restricted to integers, and in this context preemption is of no help to the algorithm.

In particular, we consider a variant of the bounded delay buffer management problem, where every job can have either low or high weight, that is $w_j \in \{1, \alpha\}$ for some parameter $\alpha > 1$. The 2-valued variant has been studied in [5] who showed the deterministic competitive ratio is $\min\{1 + 1/\alpha, 1 + (\alpha - 1)/(\alpha + 1)\}$. It is motivated by the fact that some applications, as for example online games, could send packets with two kind of informations: Major updates in the game configuration (such as the arrival or departure of participants), and minor updates (such as not critical moves).

2.1 Contributions

The lower bound of the deterministic competitive ratio for the 2-valued variant can be easily transformed into a randomized lower bound. In addition this transformation suggests a randomized online algorithm, which we show to have optimal competitive ratio $1 + (\alpha - 1)/(\alpha^2 + \alpha)$.

3 Preliminaries

Formally an instance of the bounded delay buffer management problem consists of a finite sequence of jobs. Each job j has a unit processing time, a release time $r_j \in \mathbb{N}$, a deadline $d_j \in \mathbb{N}$ and a priority weight $w_j \in \mathbb{R}_+$. The time line consists of time points (also called time slots) $t \in \mathbb{N}$. During each time point, at most one job can be executed. A job is pending at time t if $r_j \leq t < d_j$ and if it has not already been scheduled. An online algorithm A can schedule at each time point $t \in \mathbb{N}$ at most one pending job, and has to make this decision without knowledge of futur released jobs. The objective value of a schedule is the total priority weight of all scheduled jobs. This value is

compared with the objective value of an optimal schedule, and the resulting value is called the competitive ratio. The optimal schedule can be computed offline by solving a maximum weight bipartite matching problem, matching jobs to time points. The competitive ratio of algorithm A is the supremum of its competitive ratio over all instances. The competitive ratio of the problem is the infimum of the competitive ratio over all online algorithms. Usually the online computation paradigm is seen as a game played between the algorithm (making decisions which job to schedule) and the adversary (generating jobs).

If A is a randomized algorithm, then the performance of A is the expected objective value of the produced schedule. Different adversarial models have been studied in the literature. In this paper we focus on the oblivious adversary, where the input is generated without knowledge of the random choices made by the algorithm.

Online algorithms for this problem, heavily rely on the notion of a *provisional schedule*, which was introduced in [5]. Recall, that for an algorithm A and a point in time t , we define the set of *pending* jobs as the set of jobs released prior or at time t , which have their deadline strictly after t , and which have not yet been scheduled by A . The provisional schedule S is a maximum weight subset of pending jobs, which can be scheduled from time t on. This set S can be computed greedily, by starting with $S = \emptyset$, and processing first all pending heavy jobs, then all pending light jobs, in arbitrary order within each category. Each considered job j is then added to S if and only if the number of jobs $i \in S$ with $d_i \leq d_j$ is strictly less than $d_j - t$. This condition ensures that there is a provisional feasible schedule for S .

There are two natural deterministic algorithms that are defined in [5]. The first algorithm, called MAX-PACKET, schedules at every time step, a job of the provisional schedule of maximum weight, breaking ties by the deadlines of job (earliest deadline has the highest priority). The second algorithm, called FIRST-PACKET schedules at every time step, a job of the provisional schedule of minimum deadline.

The deterministic lower bound consists of an adversary which releases at time 0 two jobs, see Figure 1. A job of weight 1 and deadline 1 and a job of weight α and deadline 2. If the algorithm schedules the heavy job, then no more jobs are released, and the competitive ratio is $\frac{1+\alpha}{\alpha}$. If the algorithm schedules the urgent light job, then at time 1 the adversary releases another job of weight α and deadline 2, leading to the ratio $\frac{2\alpha}{1+\alpha}$. This shows a lower bound of $\min\{\frac{1+\alpha}{\alpha}, \frac{2\alpha}{1+\alpha}\}$ on the deterministic competitive ratio. It also suggests a simple deterministic algorithm. Let $\alpha^* = 1 + \sqrt{2}$ be the solution to the equation $\frac{1+\alpha}{\alpha} = \frac{2\alpha}{1+\alpha}$. If $\alpha \leq \alpha^*$ run FIRST-PACKET, else run MAX-PACKET. In [5] this strategy has shown to be optimal.





Figure 1: The deterministic lower bound from [5].

4 Randomized competitive ratio

The previous lower bound can be extended into a lower bound against an oblivious randomized adversary.

Proposition 1 *No randomized online algorithm can achieve a competitive ratio better than $\frac{\alpha^2+2\alpha-1}{\alpha^2+\alpha}$.*

Proof We use a game between adversary and the algorithm to derive the result. In the previous lower bound the adversary had two possibilities: either release only two jobs at time 0 (instance 1) or in addition release a job at time 1 (instance 2). The algorithm at time 0 had also two choices, execute the urgent job (FIRST-PACKET) or the heavy job (MAX-PACKET). The competitive ratios resulting from each choice are as follows:

	instance 1	instance 2
MAX-PACKET	$\frac{1+\alpha}{\alpha}$	$\frac{2\alpha}{2\alpha}$
FIRST-PACKET	$\frac{1+\alpha}{1+\alpha}$	$\frac{2\alpha}{1+\alpha}$

Von Neumann's Minimax Theorem tells us that for this game the adversary should play instance 1 with probability q and instance 2 with probability $1 - q$, while the algorithm should play MAX-PACKET with probability p and FIRST-PACKET with probability $1 - p$, for

$$\begin{aligned} q &= \frac{\alpha - 1}{\alpha}, \\ p &= \frac{\alpha^2 - 1}{\alpha^2 + 2\alpha - 1}. \end{aligned} \tag{1}$$

By verifying the expressions, indeed one can show that for the choice of q , MAX-PACKET obtains an expected gain of

$$\frac{\alpha - 1}{\alpha}\alpha + \frac{1}{\alpha}2\alpha = 1 + \alpha$$

which equals the expected gain of FIRST-PACKET, namely

$$\frac{\alpha - 1}{\alpha}(1 + \alpha) + \frac{1}{\alpha}(1 + \alpha) = 1 + \alpha.$$

In addition one can show that for the choice of p , the algorithm has in both instances the competitive ratio

$$\frac{\alpha^2 + 2\alpha - 1}{\alpha^2 + \alpha}.$$

This shows the claimed lower bound. □

Note that the above lower bound is maximized at $\alpha = 1 + \sqrt{2}$ and gives the value $4 + 2\sqrt{2} \approx 1.172$.

4.1 Upper bound

The lower bound construction suggests a simple barely random algorithm. At time 0 it flips a biased coin and based on its outcome decides between the two already mentioned deterministic algorithms. In other words BARELY-RANDOM will run MAX-PACKET with probability p and FIRST-PACKET with probability $1 - p$, for p being the value defined in (1). This surprisingly simple algorithm has optimal competitive ratio.

Theorem 1 *Algorithm BARELY-RANDOM has a competitive ratio of $\frac{\alpha^2+2\alpha-1}{\alpha^2+\alpha}$ against an oblivious adversary.*




Figure 2: Competitive ratios as function of α . The deterministic competitive ratio is the minimum of the competitive ratios of the algorithms FIRST-PACKET and MAX-PACKET.

Proof Fix an instance, and denote by OPT an optimal schedule. We start showing some structural properties of the three schedules OPT, MAX-PACKET and FIRST-PACKET. For any set of unit length jobs S , denote by EDF(S) the schedule obtained by scheduling from left to right at every time point a most urgent pending job. It is known that EDF schedules all jobs as earliest as possible in the sense, that for every time t , it maximizes the number of jobs which can be scheduled among S before time t . In particular it maximizes the number of jobs scheduled among S . In addition for every job j , such that EDF($S \cup \{j\}$) contains one more job than EDF(S), we have that the support of EDF($S \cup \{j\}$) — set of active time slots — is a superset of the support of EDF(S). Moreover every job $i \in S$ is scheduled in EDF($S \cup \{j\}$) at the same time slot or one time slot later than in EDF(S).

These facts imply that MAX-PACKET maximizes the number of heavy jobs, and that its support is included in the support of OPT. Moreover without loss of generality OPT schedules the same heavy jobs, and for every time-slot where MAX-PACKET is idle but not OPT, the later executes a heavy job. The same observation holds for time slots where MAX-PACKET is idle but not FIRST-PACKET. And finally we observe that the schedules OPT and FIRST-PACKET have the same support.

For the remainder of the proof we use the following counters:

- w is the number of α -jobs common to all 3 schedules
- x is the number of α -jobs which are in OPT but not in FIRST-PACKET.
- y is the number of 1-jobs in MAX-PACKET.
- z is the number of idle time slots where MAX-PACKET is idle but not Opt.

For convenience we use upper case letters for the corresponding sets. In particular X, Y, W are job sets, while Z is a set of time slots.

We claim that $x + z \leq w$. For this we construct a surjective mapping from $X \cup Z$ to W .

Consider a job k scheduled at time t in OPT which belongs to X . Since k did not make it into the plan of FIRST-PACKET, during all its lifespan $[r_k, d_k)$, FIRST-PACKET executes only α jobs. So we map k to the α -job j scheduled by FIRST-PACKET at the same time t .

Consider a time $t \in Z$, i.e. MAX-PACKET is idle, but OPT executes some job. We observed previously that both OPT and FIRST-PACKET execute an α -job at time t . Let j be the α -job scheduled by FIRST-PACKET at time t , and let k be the job scheduled by OPT at time t . If $k \notin X$, then $k \in W$ and we map the time t to job j .

Otherwise if $k \in X$, then we follow a *pointer chasing* construction as follows, see Figure 3 for an illustration. Job k was scheduled by MAX-PACKET at some earlier time step $t' < t$. Let j' be the job scheduled by FIRST-PACKET at time t' and k' be the job scheduled by OPT at time t' . We know that j' is an α -job, because this is the case for all jobs scheduled by FIRST-PACKET during the lifespan of k . We repeat the same argument, namely, if $k' \notin X$, then we map t to j' otherwise we recurse on a earlier time step, where MAX-PACKET schedules k' .

The mapping is surjective, because every job in $X \cup Z$ is executed in a different time slot in OPT. Either it is scheduled to a job in FIRST-PACKET scheduled at the same time or to job in another time slot, find is found through the previously mentioned pointer chasing procedure. The construction of this procedure ensures that different initial time slots, lead to different final time slots.

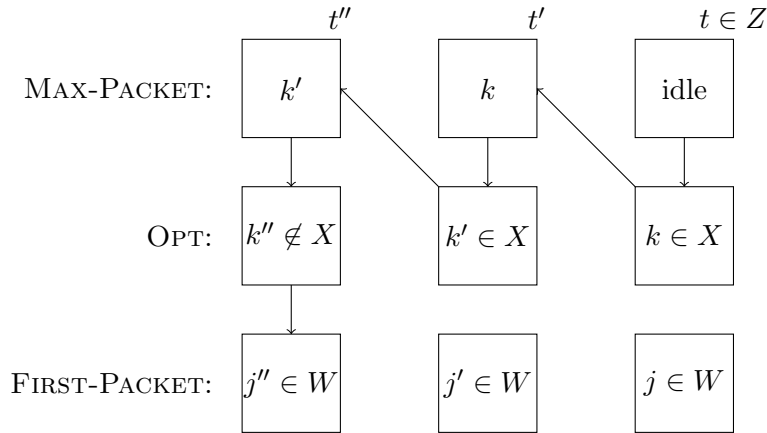


Figure 3: Mapping from $X \cup Z$ to W .

It remains to show that from the inequality $x + z \leq w$ we can derive the claimed competitive ratio $R = \frac{\alpha^2 + 2\alpha - 1}{\alpha^2 + \alpha}$. Recall the probability $p = \frac{\alpha^2 - 1}{\alpha^2 + 2\alpha - 1}$, $1 - p = \frac{2\alpha}{\alpha^2 + 2\alpha - 1}$.

Using the counters w, x, y, z we can express the profits of the schedules as follows

$$\begin{aligned} \text{MAX-PACKET} &= w\alpha + x\alpha + y \\ \text{OPT} &= w\alpha + x\alpha + y + z \\ \text{FIRST-PACKET} &= w\alpha + x + y + z. \end{aligned}$$

Here for convenience we used the same names for the schedules and for their profits. Note that in expression of the profit of FIRST-PACKET has the same variables as for OPT because their support is the same, however the coefficients for x differ. Also OPT schedules the same heavy jobs as MAX-PACKET, this why they both have the same variables with coefficient α , but OPT has the additional term z , taking into account the additional support of OPT. We show the following inequality through a sequence of simplifications.

$$\begin{aligned}
R \cdot \text{BARELY-RANDOM} &\geq \text{OPT} \\
R \cdot p \cdot \text{MAX-PACKET} + R \cdot (1 - p) \cdot \text{FIRST-PACKET} &\geq \text{OPT} \\
\frac{\alpha^2 - 1}{\alpha^2 + \alpha}(w\alpha + x\alpha + y) + \frac{2\alpha}{\alpha^2 + \alpha}(w\alpha + x + y + z) &\geq w\alpha + x\alpha + y + z \\
\alpha^2(w\alpha + x\alpha + y) - (w\alpha + x\alpha + y) + 2\alpha(w\alpha + x + y + z) &\geq (\alpha^2 + \alpha)(w\alpha + x\alpha + y + z) \\
(\alpha^2 - \alpha)w + \alpha(x + y + z) &\geq (\alpha^2 - \alpha)(x + z) + \alpha(w + x) + y
\end{aligned}$$

which holds by $x + z \leq w$, $\alpha \geq 1$ and $\alpha^2 - \alpha \geq 0$. □

5 Conclusion

Now that the randomized competitive ratio of the two valued bounded delay online buffer management problem is tackled — against the oblivious adversary — it would be interesting to see how it can be generalized to more values, for example job weights belonging to the set $\{1, \alpha, \alpha^2\}$, or even to the general problem. Current the only known randomized algorithms for the general problem are RMIX and REMIX.

The algorithm RMIX was defined in [2] as follows. At every time point, consider h the heaviest pending job. Denote w_h its weight. Let $x \in [-1, 0]$ be a uniformly chosen real number. Let f be a job with minimal deadline and weight $w_f \geq e^x w_h$. Execute f at this time point. We can observe that, if executed on instance 1 of the lower bound instance from Proposition 1, at time 0 it will choose the heavy job with probability $\min\{1, \ln \alpha\}$. Hence it has a larger tendency to choose the heavy job, than BARELY-RANDOM and is fooled by instance 1.

The algorithm REMIX was defined in [7], and behaves even differently on instance 1. It will execute the heavy job with probability $1 - 1/\alpha$ and the light job with probability $1/\alpha$, but is also not optimal. These comparisons with BARELY-RANDOM are unfair in a sense, because they have been designed for the general weighted case, and mostly for an adaptive adversary. But the comparison indicates room for improvement. The optimal randomized algorithm for the general weighted problem, might combine ideas of all these mentioned algorithms, and in particular borrow from the optimal deterministic algorithm [9].

6 Acknowledgment

This work was partially supported by the French research agency (ANR-19-CE48-0016), the CNRS PEPS project ADVICE as well as by the EPSRC grant EP/S033483/1.

We would like to thank Spyros Angelopoulos for helpful discussions.

References

- [1] Yair Bartal, Francis Y. L. Chin, Marek Chrobak, Stanley P. Y. Fung, Wojciech Jawor, Ron Lavi, Jiří Sgall, and Tomáš Tichý. Online competitive algorithms for maximizing weighted throughput of unit jobs. volume 2996, pages 187–198, 2004.

- [2] Francis Y. L. Chin, Marek Chrobak, Stanley P. Y. Fung, Wojciech Jawor, Jirí Sgall, and Tomáš Tichý. Online competitive algorithms for maximizing weighted throughput of unit jobs. *Journal of Discrete Algorithms*, 4:255–276, 2006.
- [3] Francis Y. L. Chin and Stanley P. Y. Fung. Improved competitiveness algorithms for online scheduling with partial job values. In *9th International Computing and Combinatorics Conference*, pages 425–434, 2003.
- [4] Francis Y. L. Chin and Stanley P. Y. Fung. Online scheduling for partial job values: Does timesharing or randomization help? *Algorithmica*, 37:149–164, 2003.
- [5] Matthias Englert and Matthias Westerman. Considering suppressed packets improves buffer management in QoS switches. pages 209–218, 2007.
- [6] B. Hajek. On the competitiveness of online scheduling of unit-length packets with hard deadlines in slotted time. In *Conference in Information Sciences and Systems*, pages 434–438, 2001.
- [7] Łukasz Jez. A universal randomized packet scheduling algorithm. *Algorithmica*, 67(4):498–515, 2013.
- [8] A. Kesselman, Zvi Lotker, Yishay Mansour, Boaz Patt-Shamir, Baruch Schieber, and Maxim Sviridenko. Buffer overflow management in QoS switches. pages 520–529, 2001.
- [9] Pavel Veselý, Marek Chrobak, Łukasz Jez, and Jirí Sgall. A φ -competitive algorithm for scheduling packets with deadlines. *CoRR*, abs/1807.07177, 2018.