

# Geometric Aspects of Online Packet Buffering: An Optimal Randomized Algorithm for Two Buffers<sup>\*</sup>

Marcin Bienkowski<sup>1</sup> and Aleksander Mądry<sup>1,2</sup>

<sup>1</sup> Institute of Computer Science, University of Wrocław, Poland

<sup>2</sup> CSAIL, MIT, Cambridge, MA, USA

**Abstract.** We study packet buffering, a basic problem occurring in network switches. We construct an optimal 16/13-competitive randomized online algorithm PB for the case of two input buffers of arbitrary sizes. Our proof is based on geometrical transformations which allow to identify the set of sequences incurring extremal competitive ratios. Later we may analyze the performance of PB on these sequences only.

**Key words:** online algorithms, network problems, packet buffering.

## 1 Introduction

Nowadays, the performance of network backbones depends on the speed, with which network devices can switch data packets arriving at the input ports to the appropriate output ports. Since the traffic is usually bursty, the rate of arriving packets might be much higher than the rate with which the device can transmit them, and in result packets might get lost. This motivates the use of buffers attached to the input ports; these buffers can accumulate incoming packets and store them for later transmission. The capacity of buffers — although usually large — is limited, which makes buffer management techniques crucial for minimizing the data loss.

We study a basic problem in this context. We consider a network device which has  $m$  input ports and one output port. Each input port has an attached buffer which can store up to  $B$  packets; we assume that all packets are of unit size. Time is slotted into time steps. At any time step, any number of packets may arrive at the input ports and they are appended to the appropriate buffers. If a buffer cannot accommodate all the packets, the excess is lost. At any time step, the device can transmit one packet from one buffer; the buffer managing algorithm has to choose which buffer to send from. The scenario described above is typical for input-queued switches or routers. Additionally, this model, in which packets are equally important, is typical for current IP networks.

In our setting no information about the future is available to the algorithm. In particular, we make no probabilistic assumptions about the input. For analyzing the efficiency of our algorithms we use competitive analysis [13], and — on any input sequence — compare the throughput (the number of packets transmitted) of our algorithm and the optimal offline schedule. For any algorithm  $A$  and any sequence of packets arrival  $\tau$ , we denote the throughput of  $A$  on  $\tau$  by  $T_A(\tau)$ . We call a deterministic algorithm ALG  $c$ -competitive if for all sequences  $\tau$ , it holds that  $c \cdot T_{\text{ALG}}(\tau) \geq T_{\text{OPT}}(\tau)$ , where OPT denotes the optimal *offline* algorithm. Number  $c$  is called a *competitive ratio* of the algorithm ALG. If ALG is a randomized algorithm, then in the definition above we replace  $T_{\text{ALG}}(\tau)$  with its expected value.

---

<sup>\*</sup> Full version. Extended abstract appeared in the proceedings of LATIN 2008. Research supported by MNiSW grant number N206 001 31/0436, 2006–2008, MNiSW grant number N N206 1723 33, 2007–2010, and by an Akamai Presidential Fellowship.

**Previous Results.** There are several results for the basic model described above. As the optimal competitive ratios can differ depending on the values  $B$  and  $m$ , the results address particular classes of these values.

First, we consider deterministic algorithms. The general upper bound holding for all values of  $B$  and  $m$  was given by Azar and Richter [2]. They proved that any deterministic *work-conserving* (i.e. serving a non-empty queue) algorithm is 2-competitive. They showed that for  $B = 1$  no deterministic strategy can be better than  $(2 - \frac{1}{m})$ -competitive and presented a lower bound on the competitive ratio of  $1.366 - \Theta(\frac{1}{m})$  which holds for any fixed  $B$ . Albers and Schmidt [1] improved that bound showing that for any fixed  $B$  and for large  $m$  the lower bound can be arbitrarily close to  $e/(e - 1) \approx 1.582$ . They also showed an algorithm *Semi-Greedy* which is 1.944-competitive for  $B \geq 2$  [12]. For  $B = 2$  this algorithm is optimal and 1.857-competitive; for  $B \rightarrow \infty$  the algorithm is 1.889-competitive. For  $m = 2$  and  $B \rightarrow \infty$ , Schmidt [11] demonstrated a lower bound of  $16/13 = 1.231$  and proved that a greedy algorithm achieves a ratio of  $9/7 \approx 1.286$ .

Randomized algorithms were also considered: Schmidt [11] showed a  $3/2$ -competitive *Random Permutation* algorithm; the competitive ratio holds for any values of  $B$  and  $m$ . For the lower bound, define  $h(n, k)$  as

$$h(n, k) = \frac{k + n}{k + 1 + \frac{(n-1)^{k+1}}{n^k}} \quad \text{and} \quad h(n) = \min_{k \in \mathbb{N}} h(n, k) . \quad (1)$$

A lower bound claimed by Albers and Schmidt in [1], whose proof can be found in [12], states that for any value of  $B$  the competitive ratio of any randomized algorithm is at least  $h(m)$ . This value is equal to  $16/13$  for  $m = 2$  and approaches  $1.466$  for  $m \rightarrow \infty$ .

**Our Contribution and Paper Outline.** In this paper, we present a randomized online algorithm which — for the case of  $m = 2$  buffers with arbitrary buffer size  $B$  — achieves the optimal competitive ratio of  $16/13 \approx 1.231$ .

Most papers on packet buffering concentrate around developing a smart algorithm and then comparing its behavior to the optimal one. Hence, the optimal algorithm is considered only in the analysis. We employ a different approach. In each step, we trace the set of possible states of the algorithm, which would so far serve the sequence in optimal manner. Then, by keeping the state of our online algorithm as close to the center of this set as possible, we ensure that it performs well compared to the optimal solution. This technique bears some similarities to the well-known *work-function* technique, used for constructing many optimal or almost optimal online algorithms (for example, for  $k$ -server [7] or page migration [3]).

Initially, we construct and analyze a deterministic algorithm PBF in a setting that allows PBF to have fractional number of packets in its buffers. We note that the lower bound on the competitive ratio of  $h(2) = 16/13$  holds also for such model. The proof of PBF optimality consists of two parts. In Sect. 3.2, we show how the hardest sequences for PBF look like. We show that these input sequences (we call them *regular*) have very special structure, which we exploit to bound the competitiveness of PBF. We prove this by developing a geometric view on the packet buffering problem; such approach turns out to be surprisingly successful. Finally, using a potential function-like argument, in Sect. 3.3, we show that the performance ratio of PBF on any regular sequence is at most  $16/13$ . We note that the idea of reductions of arbitrary sequences to the most difficult ones can be found in the previous papers, for example in [11].

As we mentioned above, PBF is a deterministic algorithm which is optimal in an extended, fractional model. We note that the most straightforward translation of this solution into a randomized non-fractional one does not work. Instead, using techniques similar to randomized

rounding [10], we construct a two-dimensional rounding technique, which yields an optimal algorithm PB.

For clarity, the proofs of technical lemmas were moved to the appendix.

**Related Work.** One of the most straightforward generalizations of the considered simple scenario is the model in which packets have values and the objective is to maximize the total value of packets sent. Although yet not commonly seen in practice, these *Differentiated Services* allow Internet Service Providers to assign different levels of *Quality of Service* to different data streams.

There are several results concerning the case of maintaining a single buffer, where packets have to be transmitted in FIFO order and where preemption (eviction of packets already in buffer) is allowed. Currently, the best deterministic *preemptive greedy* algorithm due to Englert and Westermann [4] achieves a competitive ratio of 1.732 and the best known lower bound for this problem, 1.419, is due to Kesselman et al [6]. There has also been work on a so-called *bounded-delay model*, in which no FIFO order is enforced but packets have deadlines (see, e.g., [5]). Azar and Richter [2] showed how to cope with multiple queues, presenting a general technique of transforming algorithms for single queue into multiple queue algorithms, losing factor 2 in the competitive ratio.

The packet buffering problems were also considered under some probabilistic assumptions on the input sequence (see e.g. [9]). However, there is an observed evidence that the nature of data traffic in networks is chaotic [8] and does not follow standard patterns like Poisson arrival model.

## 2 Preliminaries

First, let us formally define the input sequence. We transform a description of packets arrivals  $\tau$  into a sequence of requests  $\sigma$  with more convenient form. For each time step in which there are no new incoming packets, we append a request IDLE to sequence  $\sigma$ . For a step in which there are new packets at input ports, say  $x_0$  packets at buffer 0 and  $x_1$  packets at buffer 1, we append  $\text{ADD}(0)^{x_1} \text{ADD}(1)^{x_2} \text{IDLE}$  to  $\sigma$ .

By  $\sigma_t$  we understand the  $t$ -th element of  $\sigma$  and by  $\sigma|_a^b$  the contiguous subsequence of  $\sigma$  starting at position (step)  $a$  and ending at  $b$ . Additionally we define  $\sigma^t$  as the first  $t$  elements of  $\sigma$ , i.e.  $\sigma|_1^t$ ; in particular,  $\sigma^0$  denotes an empty sequence  $\epsilon$ . We say that the request  $\sigma_t$  is processed in *step*  $t$ . For any two sequences  $\sigma$  and  $\sigma'$  we denote their concatenation by  $\sigma\sigma'$ .

**Semantics of Request Sequence.** For any algorithm ALG, the *state* of its buffers at the end of a given step can be described by a pair  $\mathbf{x}^{\text{ALG}} \in \{0, \dots, B\} \times \{0, \dots, B\}$ , where the coordinates denote the number of packets in the respective buffers. For any request (sub)sequence  $\sigma$ , we use  $\mathbf{x}^{\text{ALG}}(\sigma)$  to denote the state of ALG after it starts with empty buffers and serves the sequence  $\sigma$ . In particular, by  $\mathbf{x}^{\text{ALG}}(\sigma^t)$  we mean the state of ALG after processing the first  $t$  steps of  $\sigma$ , and thus  $\mathbf{x}^{\text{ALG}}(\sigma^0) = (0, 0)$ .

The semantics of the requests from  $\sigma$  sequence is straightforward. Fix any step  $t$ . For an IDLE request, ALG may choose a non-empty buffer  $i$  and transmit one packet from it. Although the algorithm may also choose not to transmit a packet, any such algorithm can be transformed to a *work-conserving* one, which sends a packet whenever possible, and the competitiveness of the obtained algorithm is not worse. The way of choosing the buffer for transmission is called *pivoting rule*. Note that this rule is the only factor that determines the behavior of the algorithm.

If  $\sigma_t$  is an  $\text{ADD}(i)$  request, a new packet is *added* to the buffer  $i$ . If the number of packets at the  $i$ -th buffer is already  $B$ , then the packet is immediately lost. Let  $\ell^{\text{ALG}}(\sigma^t)$  be the number of packets that are actually added to the buffer in step  $t$ ,  $\ell^{\text{ALG}}(\sigma^t) \in \{0, 1\}$ . We extend this definition to  $\text{IDLE}$  steps, setting  $\ell^{\text{ALG}}(\sigma^t) = 0$  for them.

At the end of the input sequence, the algorithm empties all the buffers; we may assume that they are all transmitted in one batch.

Let  $\mathcal{H} = [0, B] \times [0, B]$ . We may view  $\mathbf{x}^{\text{ALG}}$ , the state of  $\text{ALG}$ , as a point from  $\mathcal{H} \cap \mathbb{N}^2$ . Then, the  $\text{IDLE}$  and  $\text{ADD}$  operations described above have obvious geometric interpretation. For any state  $\mathbf{x}$ ,  $x_0$  and  $x_1$  denote the number of packets in the respective buffers, and  $\|\mathbf{x}\| = x_0 + x_1$ .

**Fractional Model.** It is now straightforward to generalize the above description to a *fractional model*, where the state of  $\text{ALG}$  can be any point from  $\mathcal{H}$  (also the one with fractional coordinates). In particular, the algorithm serving  $\text{IDLE}$  request may choose to transmit fractional parts of packets from different buffers, with the only requirement that the total mass of transmitted packets is at most 1. More formally, during an  $\text{IDLE}$  request in step  $t$  of  $\sigma$ ,  $\text{ALG}$  chooses a vector  $\Delta(t) = (\delta_0, \delta_1) \in [0, 1] \times [0, 1]$  such that  $\delta_0 + \delta_1 \leq 1$  and  $x_i^{\text{ALG}}(\sigma^{t-1}) \geq \delta_i(t)$  for  $i \in \{0, 1\}$ . Then the effect of this request is reflected by a new state  $\mathbf{x}^{\text{ALG}}(\sigma^t) = \mathbf{x}^{\text{ALG}}(\sigma^{t-1}) - \Delta(t)$  of the algorithm. Note that the definition of  $\ell^{\text{ALG}}(\sigma^t)$  can be extended to the fractional model in a straightforward manner.

Although we allow online algorithms to use fractional parts of the packets, to simplify our analysis we compare them to the optimal offline algorithm which still works in the standard model.<sup>3</sup> We note that the lower bound of  $16/13$  holds in the fractional model, as well.

**Competitiveness.** For any sequence  $\sigma$  and any deterministic (not necessarily online) algorithm  $\text{ALG}$ , we define a function  $\text{loss}_{\text{ALG}}(\sigma)$  as the number of packets lost by the strategy  $\text{ALG}$  on  $\sigma$ , under the condition that  $\text{ALG}$  starts with empty buffers. For any algorithm  $\text{ALG}$  and two sequences  $\sigma$  and  $\tau$  we define  $\Delta_\sigma \text{loss}_{\text{ALG}}(\tau) = \text{loss}_{\text{ALG}}(\sigma\tau) - \text{loss}_{\text{ALG}}(\sigma)$ .

Obviously, the losses can occur only due to  $\text{ADD}$  requests which overflow some buffer. Therefore, if  $\sigma$  is the whole sequence,  $\text{loss}_{\text{ALG}}(\sigma) = \sum_{t:\sigma_t=\text{ADD}}(1 - \ell^{\text{ALG}}(\sigma^t))$ . Let  $S(\sigma)$  denote the total number of packets added in  $\sigma$ , i.e. the number of  $\text{ADD}$  requests. The throughput of  $\text{ALG}$  (the number of transmitted packets), denoted  $T_{\text{ALG}}(\sigma)$ , is then equal to  $S(\sigma) - \text{loss}_{\text{ALG}}(\sigma)$ .

Consider a sequence  $\sigma$ . Let  $\text{OPT}$  be an optimal (offline) algorithm for the packet buffering problem, i.e. the one which minimizes the number of packets lost. We define the *performance ratio* of (an online) algorithm  $\text{ALG}$  on  $\sigma$  as

$$\mathcal{R}_{\text{ALG}}(\sigma) = \frac{T_{\text{OPT}}(\sigma)}{T_{\text{ALG}}(\sigma)} = \frac{S(\sigma) - \text{loss}_{\text{OPT}}(\sigma)}{S(\sigma) - \text{loss}_{\text{ALG}}(\sigma)}. \quad (2)$$

If  $\Sigma$  is any set of sequences, then  $\mathcal{R}_{\text{ALG}}(\Sigma) = \sup_{\sigma \in \Sigma} \{\mathcal{R}_{\text{ALG}}(\sigma)\}$ . Let  $\mathcal{Q}$  be the set of all possible sequences; then the competitive ratio can be defined as  $\mathcal{R}_{\text{ALG}} = \mathcal{R}_{\text{ALG}}(\mathcal{Q})$ . If  $\mathcal{R}_{\text{ALG}} \leq \alpha$ , then we call  $\text{ALG}$   $\alpha$ -*competitive*.

**OPT State Space** As mentioned in the introduction, we would like to rely the behavior of our algorithm on tracing the state of some optimal off-line algorithm buffers in each step. Obviously, the complete knowledge about this state is not available to an on-line algorithm. Instead, we will focus on extrapolating a set of possible  $\text{OPT}$  states which can be inferred from the already

<sup>3</sup> Although it is not needed for our reasoning, it appears that this restriction does not constrain the power of  $\text{OPT}$ .

seen prefix of  $\sigma$ . In each step we trace a certain set  $\mathcal{I}(\sigma^t)$  of possible states whose relation to an optimal solution is presented in [Lemma 2](#).

We define set  $\mathcal{I}$  inductively as  $\mathcal{I}(\sigma^0) = \{(0, 0)\}$ , and

$$\mathcal{I}_{\text{pre}}(\sigma^t) = \begin{cases} (\mathcal{I}(\sigma^{t-1}) - (1, 0)) \cup (\mathcal{I}(\sigma^{t-1}) - (0, 1)) & \text{if } \sigma_t = \text{IDLE} \\ \mathcal{I}(\sigma^{t-1}) + (1, 0) & \text{if } \sigma_t = \text{ADD}(0) \\ \mathcal{I}(\sigma^{t-1}) + (0, 1) & \text{if } \sigma_t = \text{ADD}(1) \end{cases} , \quad (3)$$

$$\mathcal{I}(\sigma^t) = \begin{cases} \mathcal{I}_{\text{pre}}(\sigma^t) \cap \mathcal{H} & \text{if } \mathcal{I}_{\text{pre}}(\sigma^t) \cap \mathcal{H} \neq \emptyset \\ \mathcal{I}(\sigma^{t-1}) & \text{otherwise} . \end{cases} \quad (4)$$

An intuition behind the set  $\mathcal{I}(\sigma^t)$  is that it contains states of all algorithms which try to greedily reduce their loss, i.e. postpone losing packets. A straightforward induction shows that the number of packets in each state from the set  $\mathcal{I}$  is the same (we denote this number by  $\|\mathcal{I}\|$ ) and  $\mathcal{I}$  consists of non-fractional states contained in an anti-diagonal interval (see [Fig. 1a](#)). We call an ADD request *hit* if it reduces the number of elements in  $\mathcal{I}$ .

We say that a state  $(k_1, k_2)$  is *majorized* by  $(k'_1, k'_2)$  (we write  $(k_1, k_2) \leq (k'_1, k'_2)$ ) if  $k_1 \leq k'_1$  and  $k_2 \leq k'_2$ . We define *shadow of  $\mathcal{I}$* , denoted  $\mathcal{SH}(\mathcal{I})$ , as the set of all the states which are majorized by some state from  $\mathcal{I}$ . The following technical lemmas show the relation between the set  $\mathcal{I}$ , optimal solutions and other work-conserving algorithms.

**Lemma 1.** *For any input  $\sigma$  and any work-conserving algorithm  $\text{ALG}$ ,  $\mathbf{x}^{\text{ALG}}(\sigma) \in \mathcal{SH}(\mathcal{I}(\sigma))$ .*

**Lemma 2.** *There exists an algorithm  $A$ , such that  $\mathbf{x}^A(\sigma^t) \in \mathcal{I}(\sigma^t)$  for any step  $t$ . Every algorithm with such property is optimal and loses a packet in step  $t$  on ADD request if and only if  $\mathcal{I}(\sigma^t) = \mathcal{I}(\sigma^{t-1})$ .*

The main implication of [Lemma 2](#) is that we get a convenient description of the loss of an optimal algorithm. Namely, we can compute  $\text{loss}_{\text{OPT}}(\sigma)$  by counting all ADD requests, for which  $\mathcal{I}(\sigma^t) = \mathcal{I}(\sigma^{t-1})$ . Note that such equality may occur only when  $\mathcal{I}(\sigma^{t-1})$  is a singleton set at one of the upper boundaries.

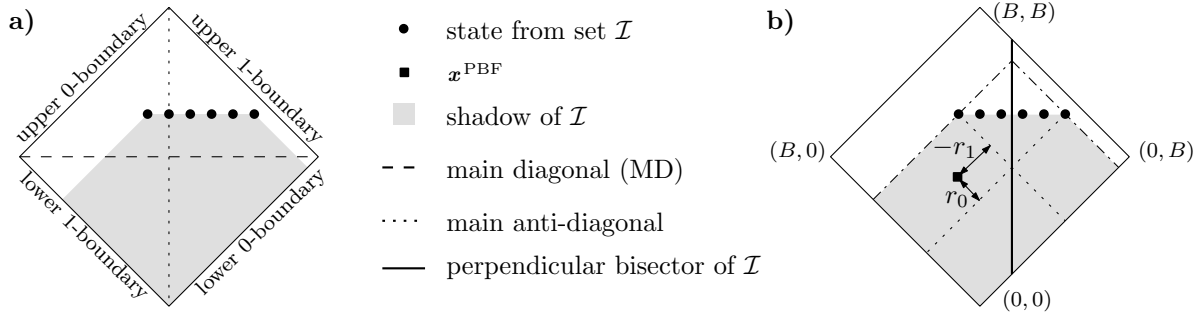
Of course, we would like to achieve the performance of  $A$ , but the main difficulty in constructing an *online* algorithm mimicking such  $A$  is that usually knowing only  $\sigma^t$ , we can neither predict the exact state of  $A$  in step  $t+1$ , nor the future shape of  $\mathcal{I}$  sets. However, as we already said, on the basis of computed  $\mathcal{I}(\sigma^t)$ , we can in some way extrapolate  $\mathcal{I}(\sigma^{t+1})$  and the behavior of an optimal algorithm  $A$ .

### 3 Algorithm PBF

In this section, we present an algorithm PBF, which is optimal in the fractional model. Let

$$r_i(\sigma) = x_i^{\text{PBF}}(\sigma) - \min_{z \in \mathcal{I}(\sigma)} z_i \quad \text{for } i \in \{0, 1\} . \quad (5)$$

Assume that the adversary decides to issue a maximum number of  $\text{ADD}(i)$  requests without incurring a loss to OPT. Then  $r_i$  would be the number of packets lost by PBF. If  $r_i \leq 0$ , PBF cannot lose packets in this way (see  $r_1$  in [Fig. 1b](#)). We note that  $r_i$  can be efficiently computed by an online algorithm (as  $\mathcal{I}$  can be described by a few parameters). Let  $\text{bal}(\sigma) = r_0(\sigma) - r_1(\sigma)$  be called *balance*.



**Fig. 1.** Illustration of the set  $\mathcal{I}$  and PBF parameters;  $r_0 > 0$ ,  $r_1 < 0$

When PBF encounters IDLE request in step  $t$  of  $\sigma$ , it computes a new shape of the set  $\mathcal{I}(\sigma^t)$  first. Then it transmits a total mass of 1 packet, so that the resulting value of  $|\text{bal}(\sigma^t)|$  is as small as possible. This rule can be interpreted geometrically as choosing a new state  $\mathbf{x}^{\text{ALG}}(\sigma^t)$  as close to the perpendicular bisector (hence the abbreviation PB; F stands for fractional model) of the set  $\mathcal{I}$  as possible.

### 3.1 Outline of the Proof

In the remaining part [Sect. 3](#), we prove the following theorem.

**Theorem 1.** *For any buffer size  $B$ , PBF on two buffers is 16/13-competitive.*

We prove it inductively on  $B$ . Obviously, for  $B = 0$ , the theorem trivially holds (with competitiveness 1). Below we present the roadmap of the proof. As we mentioned in the introduction, the proof is divided into two parts. In the first one, we narrow down instances on which PBF has high competitive ratio. In the second part, we restrict our analysis only to these instances.

For any integers  $x$  and  $y$ , by a *block*  $\mathcal{B}(x, y)$  we denote a subsequence  $\text{IDLE}^x \text{ADD}(0)^y$ . We also denote the sequence  $\text{ADD}(0)^B \text{ADD}(1)^B$  by  $\mathcal{A}$ . By *main diagonal* (MD) we mean the diagonal of the square, which contains all fractional states  $\mathbf{z}$  such that  $\|\mathbf{z}\| = B$ . We say that  $\mathcal{I}$  is *above*, *at*, or *below* MD if  $\|\mathcal{I}\|$  is, respectively, greater, equal, or less than  $B$ . We introduce the following classes of sequences.

**Definition 1.** *We denote the set of all sequences by  $\mathcal{Q}$ . We also define the following sets of sequences.*

- $\mathcal{Q}_N$ : non-trivial.  
 $\sigma \in \mathcal{Q}_N$  if it ends with ADD incurring loss to PBF and  $\mathbf{x}^{\text{PBF}}(\sigma^t) \neq (0, 0)$  for  $t > 0$ .
- $\mathcal{Q}_P$ : proper.  
 $\sigma \in \mathcal{Q}_P$  if it is non-trivial, starts with  $\mathcal{A}$ , and  $\mathcal{I}(\sigma^t)$  is above MD for all  $t \geq |\mathcal{A}|$ .
- $\mathcal{Q}_U$ : uniform.  
 $\sigma \in \mathcal{Q}_U$  if  $\sigma$  is proper and after initial  $\mathcal{A}$ , consists only of ADD(0) and IDLE requests.
- $\mathcal{Q}_R$ : regular.  
 $\sigma \in \mathcal{Q}_R$  if  $\sigma$  is uniform and has the form  $\mathcal{A}\mathcal{B}(x_1, y_1)\mathcal{B}(x_2, y_2) \dots \mathcal{B}(x_n, y_n)$ , where after each block  $\mathcal{B}(x_i, y_i)$ ,  $\mathbf{x}_0^{\text{PBF}} = B$ .

The course of the proof is to show each of consecutive relations below.

$$\mathcal{R}_{\text{PBF}} \leq \mathcal{R}_{\text{PBF}}(\mathcal{Q}_N) \leq \mathcal{R}_{\text{PBF}}(\mathcal{Q}_P) = \mathcal{R}_{\text{GR}}(\mathcal{Q}_P) \leq \mathcal{R}_{\text{GR}}(\mathcal{Q}_U) \leq \mathcal{R}_{\text{GR}}(\mathcal{Q}_R) \quad (6)$$



$\mathcal{R}_{\text{GR}}$  denotes the performance ratio of a natural GREEDY algorithm (see [11]), which is defined formally later. We show these inequalities in Sect. 3.2. Then, in Sect. 3.3, we prove that  $\mathcal{R}_{\text{GR}}(\mathcal{Q}_R) \leq 16/13$ . On the other hand, on the sequence  $\nu(B) = \mathcal{A}\mathcal{B}(B, B)\mathcal{B}(B, B)$ . PBF transmits  $\frac{13}{4} \cdot B$  packets, whereas OPT transmits  $4 \cdot B$ . Hence, for  $B \geq 1$ , it holds that  $\mathcal{R}_{\text{PBF}} \geq \mathcal{R}_{\text{PBF}}(\nu(B)) = \frac{16}{13}$ , and therefore all inequalities above can be replaced by equalities. We note that by [1,12], the competitive ratio of any online algorithm ALG is at least  $h(2) = \frac{16}{13}$ , and therefore PBF is optimal.

### 3.2 Worst-case Sequences

In this section, we consecutively prove all inequalities of (6) but the last one. In general, in order to show that  $\mathcal{R}_{\text{ALG}}(\mathcal{Q}_1) \leq \mathcal{R}_{\text{ALG}}(\mathcal{Q}_2)$ , we show that for any  $\sigma \in \mathcal{Q}_1$ , we can transform it to obtain a sequence  $\hat{\sigma} \in \mathcal{Q}_2$ , such that the performance ratio of ALG does not decrease. Intuitively,  $\hat{\sigma}$  is more difficult for ALG than  $\sigma$ .

**Changes in Balance.** We start from several simple definitions and observations. We define  $\text{len}(\mathcal{I})$  as the length of the smallest interval containing set  $\mathcal{I}$ . This amount is equal to the number of  $\mathcal{I}$  elements minus 1.

We conceptually divide each step into two parts. In the first one, the adversary issues a request and as a result the set  $\mathcal{I}$  is changed. In case of an ADD request,  $\mathbf{x}^{\text{PBF}}$  is changed as well. In the second part, which is present only for IDLE requests, PBF transmits some packets.

As a result of an ADD( $i$ ) request,  $r_i$  may decrease by one. This can happen only if just before this request the set  $\mathcal{I}$  touches (i.e. has non-empty intersection with) the upper  $i$ -boundary of  $\mathcal{H}$ , where *upper  $i$ -boundary* is  $\mathcal{H} \cap \{(k_0, k_1) : k_i = B\}$ . We call such an ADD( $i$ ) request a *hit*; such an ADD reduces the value of  $\text{len}(\mathcal{I})$  by one.

If  $\mathcal{I}$  is above MD, then in a step with IDLE request, both  $r_i$  increase by 1, and therefore the balance remains unchanged. On the other hand, if  $\mathcal{I}$  is at or below MD and it touches the lower  $i$ -boundary of  $\mathcal{H}$ , the corresponding  $r_i$  remains unchanged. If  $\mathcal{I}$  touches only one boundary, the balance may therefore change by one. In total, upon an IDLE request, in the first part of a step  $\text{len}(\mathcal{I})$  increases by 1 minus the number of lower boundaries  $\mathcal{I}$  touches and in the second part PBF changes its state according to its pivoting rule. These observations lead to the following technical lemmas.

**Lemma 3.** *For any sequence  $\sigma$ , if  $\mathcal{I}(\sigma)$  is below or at the main diagonal, then  $\text{bal}(\sigma^t) = 0$ .*

**Lemma 4.** *For any sequence  $\sigma$ , there exists a non-trivial sequence  $\hat{\sigma}$ , such that  $\mathcal{R}_{\text{PBF}}(\sigma) \leq \mathcal{R}_{\text{PBF}}(\hat{\sigma})$ .*

**Proper Sequences.** Consider a non-trivial sequence  $\sigma$ . By the definition, at the end of  $\sigma$ , set  $\mathcal{I}$  touches an upper boundary (in the following informal description, we assume that it touches both boundaries). However, if we look from the adversary point of view, it is not obvious when this should happen for the first time. It is also not clear that keeping the set  $\mathcal{I}$  below the main diagonal is not preferable — even if this constraints the growth of the set  $\mathcal{I}$ , the resulting constraints on PBF’s behavior might be beneficial for the adversary.

We address the issues above by showing that proper sequences incur the worst performance ratio of PBF. Namely, we prove that for any non-trivial sequence  $\sigma$ , there exists another sequence  $\hat{\sigma}$  with not smaller performance ratio, such that  $\hat{\sigma}$  starts with filling the buffers with packets and later it keeps the set  $\mathcal{I}$  all the time above the diagonal.

We construct  $\hat{\sigma}$  on the basis of  $\sigma$ , so that after initial filling the buffers,  $\hat{\sigma}$  contains almost the same steps as  $\sigma$ . We can imagine that we have two instances of PBF running “in parallel” on  $\sigma$  and on  $\hat{\sigma}$ . We show that it is possible to maintain an invariant that the set  $\mathcal{I}$  and the point  $\mathbf{x}^{\text{PBF}}$  for  $\hat{\sigma}$  are equal to  $\mathcal{I}$  and  $\mathbf{x}^{\text{PBF}}$  for  $\sigma$  translated by some vector. This invariant allows us to prove that the performance of PBF can only worsen by replacing  $\sigma$  with  $\hat{\sigma}$ .

How do we create  $\hat{\sigma}$ ? If the request of  $\sigma$  does not change  $\text{len}(\mathcal{I})$  and the spatial relation between  $\mathcal{I}$  and  $\mathbf{x}^{\text{PBF}}$ , we do not append anything to  $\hat{\sigma}$ . Otherwise, if we have an ADD request in  $\sigma$ , which is a hit, then this ADD appended to  $\hat{\sigma}$  is a hit as well (as  $\mathcal{I}(\hat{\sigma})$  touches both boundaries). The only problem arises when we have an IDLE request in  $\sigma$  occurring when  $\mathcal{I}$  touches both lower boundaries, as in this case  $\text{len}(\mathcal{I})$  decreases. To simulate such change also on the instance  $\hat{\sigma}$ , we introduce an additional request LIFT. We justify this enhancement later in this section, by showing that the adversary does not need LIFT to impose the worst competitive ratio.

**Making Sequences Proper.** For the formal proof, we have to extend the notion of proper sequences. Let  $\mathcal{H}(k_0, k_1)$  be the largest square contained entirely in  $\mathcal{H}$ , with an upper corner at  $(k_0, k_1)$ , i.e., the edge length of such square is  $\min\{k_0, k_1\}$ . Note that  $\mathcal{H}(B, B) = \mathcal{H}$ . We say that a sequence  $\sigma$  is  $(k_0, k_1)$ -proper (or just *proper* if  $k_0 = k_1 = B$ ) if it begins with  $\text{ADD}(0)^{k_0} \text{ADD}(1)^{k_1}$  and  $\mathcal{I}(\sigma^t)$  is always contained in the triangle corresponding to the half of  $\mathcal{H}(k_0, k_1)$  above its main diagonal. By a straightforward induction, we get the following fact.

**Fact 2** *If  $\sigma$  is  $(k_0, k_1)$ -proper, then  $\mathcal{I}(\sigma^t)$  has always maximal possible length, i.e., it contains all states of  $\mathcal{H}(k_0, k_1)$  with  $\|\mathcal{I}(\sigma^t)\|$  packets.*

Let  $b_i(\sigma^t) = \min_{z \in \mathcal{I}(\sigma^t)} B - z_i$  be the distance between  $\mathcal{I}$  and the upper  $i$ -boundary in step  $t$ . Let  $b_i^*(\sigma) = \min_{1 \leq t \leq |\sigma|} b_i(\sigma^t)$ . We note that both  $\|\mathbf{b}\|$  and  $\|\mathbf{b}^*\|$  start from  $2 \cdot B$ . Moreover,  $b_i^*$  never increases during runtime.

We also introduce a new request LIFT, which adds a half of a packet to both buffers of PBF. When it is issued in step  $t$ , it changes set  $\mathcal{I}$  in a way opposite to the effect an IDLE request would have. Namely, we extend definition (3) by the following case.

$$\mathcal{I}_{\text{pre}}(\sigma^t) = (\mathcal{I}(\sigma^{t-1}) + (1, 0)) \cap (\mathcal{I}(\sigma^{t-1}) + (0, 1)) \cap \mathcal{H}(b_0(\sigma^{t-1}), b_1(\sigma^{t-1})) \quad \text{if } \sigma_t = \text{LIFT} \quad (7)$$

It appears that this artificial request helps us to normalize the behavior of PBF. Our goal is to show a transformation from an arbitrary sequence without LIFTS to a proper sequence possibly containing LIFTS. The most important property of this transformation is that in each step  $t$  it preserves the spatial relation between  $\mathcal{I}(\sigma^t)$  and  $\mathbf{x}^{\text{PBF}}(\sigma^t)$ .


Let  $\mathcal{P}_\sigma(\sigma)$  be constructed in the following manner. We define  $\Delta_\sigma \text{len}(a) \equiv \text{len}(\mathcal{I}(\sigma a)) - \text{len}(\mathcal{I}(\sigma))$ . Let  $\mathcal{P}_\sigma(\epsilon) = \text{ADD}(0)^{B-b_0^*(\sigma)} \text{ADD}(1)^{B-b_1^*(\sigma)}$ . Let

$$\mathcal{P}_\sigma(\sigma^{t+1}) = \begin{cases} \mathcal{P}_\sigma(\sigma^t) \text{ IDLE} & \text{if } \sigma_{t+1} = \text{IDLE and } \Delta_{\sigma^t} \text{len}(\sigma_{t+1}) = 1, \\ \mathcal{P}_\sigma(\sigma^t) \text{ LIFT} & \text{if } \sigma_{t+1} = \text{IDLE and } \Delta_{\sigma^t} \text{len}(\sigma_{t+1}) = -1, \\ \mathcal{P}_\sigma(\sigma^t) \text{ ADD}(i) & \text{if } \sigma_{t+1} = \text{ADD}(i) \text{ and this ADD}(i) \text{ is a hit,} \\ \mathcal{P}_\sigma(\sigma^t) & \text{otherwise.} \end{cases} \quad (8)$$

Although this is not necessary in our reasoning, we observe that if  $\sigma$  is proper, then  $\mathcal{P}_\sigma(\sigma) = \sigma$ .

**Lemma 5.** *For any  $\sigma$  and corresponding  $\mathcal{P}_\sigma(\sigma)$ , the following invariants hold for all  $0 \leq t \leq |\sigma|$ :*






**Fig. 2.** Transformation  $\mathcal{P}$

- (i)  $\text{len}(\mathcal{P}_\sigma(\sigma^t)) = \text{len}(\sigma^t)$ ;
- (ii)  $r_i(\mathcal{P}_\sigma(\sigma^t)) = r_i(\sigma^t)$  for  $i \in \{0, 1\}$ ;
- (iii)  $S(\sigma^t) - S(\mathcal{P}_\sigma(\sigma^t)) + \|\mathbf{b}^*(\sigma)\| + \min_i \{b_i(\sigma^t) - b_i^*(\sigma^t)\} \geq 0$ ;
- (iv)  $\text{loss}_{\text{OPT}}(\mathcal{P}_\sigma(\sigma^t)) = \text{loss}_{\text{OPT}}(\sigma^t)$  and  $\text{loss}_{\text{PBF}}(\mathcal{P}_\sigma(\sigma^t)) = \text{loss}_{\text{PBF}}(\sigma^t)$ ;
- (v)  $\mathcal{P}_\sigma(\sigma^t)$  is  $(B - b_0^*(\sigma), B - b_1^*(\sigma))$ -proper.

*Proof.* We prove the invariants inductively. At the very beginning, for  $t = 0$ , invariants (i), (ii), (iv) and (v) trivially hold. For invariant (iii), we note that  $S(\epsilon) = 0$ ,  $S(\mathcal{P}_\sigma(\epsilon)) = \|\mathbf{b}^*(\sigma)\|$ , and  $\mathbf{b}^*(\sigma^0) = \mathbf{b}(\sigma^0)$ .

Assume that invariants hold for some  $t < |\sigma|$ ; we show that they hold also for  $t + 1$ . The cases are depicted on Fig. 3. By MD' we mean the main diagonal of  $\mathcal{H}(B - b_0^*(\sigma), B - b_1^*(\sigma))$ .

1.  $\sigma_{t+1} = \text{IDLE}$ . Invariant (iv) holds trivially as IDLE request cannot generate loss. Since  $\sigma$  is non-trivial, it holds that  $\|\mathbf{x}^{\text{PBF}}(\sigma^t)\| \geq 1$ . We prove invariants (i), (ii), (iii) and (v) for the following three subcases, depending on the value of  $\Delta_{\sigma^t} \text{len}(\sigma_{t+1})$ .
  - (a) If  $\Delta_{\sigma^t} \text{len}(\sigma_{t+1}) = 1$ , then an IDLE request is appended to  $\mathcal{P}_\sigma(\sigma^t)$ . We note that in this case  $\text{len}(\sigma^t)$  has to be shorter than the length of MD'. By invariant (i),  $\text{len}(\mathcal{P}_\sigma(\sigma^t)) = \text{len}(\sigma^t)$ , and thus Fact 2 implies that  $\mathcal{I}(\mathcal{P}_\sigma(\sigma^t))$  lies strictly above MD'. Therefore,  $\Delta_{\mathcal{P}_\sigma(\sigma^t)} \text{len}(\text{IDLE}) = 1$ , which proves invariants (i) and (v). Moreover, both values of  $r_i$  for  $\sigma^t$  and  $\mathcal{P}_\sigma(\sigma^t)$  increase by 1, which yields invariant (ii). Invariant (iii) is fulfilled trivially, as none of the terms on its left hand side changes.
  - (b) If  $\Delta_{\sigma^t} \text{len}(\sigma_{t+1}) = 0$ , nothing is appended to  $\mathcal{P}_\sigma(\sigma^t)$ , and thus  $\text{len}(\mathcal{P}_\sigma(\sigma^t))$  does not change. This proves invariants (i) and (v). The considered case may occur only if  $\mathcal{I}(\sigma^t)$  is at or below the diagonal of  $\mathcal{H}(B - b_0^*(\sigma), B - b_1^*(\sigma))$ , which, by Lemma 3, implies that  $\text{bal}(\sigma^t) = \text{bal}(\sigma^{t+1}) = 0$ . As  $\|\mathbf{x}^{\text{PBF}}(\sigma^t)\| \geq 1$ , PBF is able to transmit a packet, and it can compensate the increase of one of  $r_i$  caused by the shift of  $\mathcal{I}$ . This proves invariant (ii). For proving invariant (iii), we note that only  $b_i$  can change, but they may only increase.
  - (c) If  $\Delta_{\sigma^t} \text{len}(\sigma_{t+1}) = -1$ , then a LIFT is appended to  $\mathcal{P}_\sigma(\sigma^t)$ . Since  $\mathcal{P}_\sigma(\sigma^t)$  is  $(B - b_0^*(\sigma), B - b_1^*(\sigma))$ -proper, by Fact 2, the corresponding set  $\mathcal{I}(\mathcal{P}_\sigma(\sigma^t))$  has maximal possible length and therefore  $\Delta_{\mathcal{P}_\sigma(\sigma^t)} \text{len}(\text{LIFT}) = -1$ , which shows invariant (i) and (v). The proof of invariant (ii) is the same as in the previous case, because  $\text{bal}(\mathcal{P}_\sigma(\sigma^t) \text{ LIFT}) = \text{bal}(\mathcal{P}_\sigma(\sigma^t))$ . For proving invariant (iii), we note that the increase of  $S(\mathcal{P}_\sigma(\sigma^t))$  is compensated by the increase of  $b_0$  and  $b_1$ .
2.  $\sigma_{t+1} = \text{ADD}(j)$  request. We consider two subcases, depending on whether this request was a hit.



**Fig. 3.** Illustration for the case analysis of [Lemma 5](#).

- (a) If  $\text{ADD}(j)$  is not a hit, then  $\mathcal{P}_\sigma(\sigma^{t+1}) = \mathcal{P}_\sigma(\sigma^t)$ , and thus the increase of  $\text{len}(\mathcal{I}(\mathcal{P}(\sigma)))$  and  $\text{len}(\mathcal{I}(\sigma))$  is equal to 0. In this case,  $r_i$  values remain unchanged both for sequence  $\sigma^t$  and  $\mathcal{P}_\sigma(\sigma^t)$ . Additionally, there is no loss incurred neither on OPT nor on PBF since by [Lemma 1](#) only a hit can incur loss to PBF. This proves invariants (i), (ii), (iv), and (v). For invariant (iii), we note that  $S(\sigma^t) - S(\mathcal{P}_\sigma(\sigma^t))$  increases by 1 and  $b_j$  decreases by 1. If  $b_j^*(\sigma^t)$  does not change, then the invariant holds. If  $b_j^*$  decreases, it may happen only if  $b_j^*(\sigma^t) = b_j(\sigma^t)$ . In this case, the last term on the left side of the inequality is equal to zero, and the invariant holds as well.
- (b) If  $\text{ADD}(j)$  is a hit, then  $\text{ADD}(j)$  is appended to  $\mathcal{P}_\sigma(\sigma^t)$  sequence. Obviously, invariant (v) is preserved. By [Fact 2](#),  $\text{ADD}(j)$  is also a hit in  $\mathcal{P}_\sigma(\sigma)$ , i.e., both  $\mathcal{I}(\sigma^t)$  and  $\mathcal{I}(\mathcal{P}(\sigma^t))$  touch the upper  $j$ -boundary. By invariant (i), we have that  $\mathcal{I}(\mathcal{P}_\sigma(\sigma^t)) = \mathcal{I}(\sigma^t) + (\delta_0, \delta_1)$ , where  $\delta_{1-j} \geq 0$  and  $\delta_j = 0$ . Furthermore, this observation combined with invariant (ii) for step  $t$  implies  $\mathbf{x}^{\text{PBF}}(\mathcal{P}_\sigma(\sigma^t)) = \mathbf{x}^{\text{PBF}}(\sigma^t) + (\delta_0, \delta_1)$ . Because the shift caused by  $\text{ADD}(j)$  is perpendicular to the vector  $(\delta_0, \delta_1)$ ,  $\text{loss}_{\text{OPT}}$  and  $\text{loss}_{\text{PBF}}$  incurred in this step depend only on the length  $\text{len}(\mathcal{I})$  and  $r_i$  values. We conclude that invariant (ii) and (iv) are preserved. Finally, since  $\text{ADD}(j)$  is a hit,  $\mathbf{b}^*$  and  $\mathbf{b}$  do not change and both  $S(\cdot)$  increase by one, which proves invariant (iii).

This finishes the proof of all invariants.  $\square$

We use the lemma above to show that in the analysis of PBF, we may restrict our consideration to proper sequences only.

**Lemma 6.** *For any non-trivial sequence  $\sigma$ , there exists a proper sequence  $\hat{\sigma}$  (possibly containing LIFT requests) such that  $\mathcal{R}_{\text{PBF}}(\sigma) \leq \mathcal{R}_{\text{PBF}}(\hat{\sigma})$ .*

*Proof.* As  $\sigma$  is non-trivial, by [Lemma 1](#), it ends with a hit. W.l.o.g., we may assume that it ends with  $\text{ADD}(0)$ . Therefore,  $b_0^*(\sigma) = b_0(\sigma) = 0$ , and by the technical lemma above,  $\mathcal{P}_\sigma(\sigma)$  is  $(B, B - b_1^*(\sigma))$ -proper. If additionally  $b_1^*(\sigma) = 0$ , then by invariant (iii) of [Lemma 5](#),  $S(\sigma) \geq S(\mathcal{P}_\sigma(\sigma))$ . This combined with invariant (iv) of the same lemma allows to take  $\hat{\sigma} = \mathcal{P}_\sigma(\sigma)$  and obtain  $\mathcal{R}_{\text{PBF}}(\sigma) \leq \mathcal{R}_{\text{PBF}}(\hat{\sigma})$ .

The problem arises if  $b_1^*(\sigma) > 0$ . Consider  $\mathcal{P}'_\sigma(\sigma)$  equal to  $\mathcal{P}_\sigma(\sigma)$  but with  $b_1^*(\sigma)$  of ADD(0) requests removed from its beginning. By Lemma 5,  $\mathcal{P}'_\sigma(\sigma)$  is a proper sequence in square  $[0, B - b_1^*(\sigma)] \times [0, B - b_1^*(\sigma)]$ . The only difference between  $\mathcal{P}'_\sigma(\sigma)$  and  $\mathcal{P}_\sigma(\sigma)$  is that  $\mathcal{P}'_\sigma(\sigma)$  uses  $b_1^*(\sigma)$  packets less and thus, by invariant (iii) of Lemma 5,  $S(\sigma) \geq S(\mathcal{P}'_\sigma(\sigma))$ . In effect,  $\mathcal{R}_{\text{PBF}}(\sigma) \leq \mathcal{R}_{\text{PBF}}(\mathcal{P}'_\sigma(\sigma))$ . We note that the latter performance ratio holds if it is computed for smaller buffer sizes, namely for  $B' = B - b_1^*(\sigma)$ . However, by the inductive assumption, PBF is 16/13-competitive for  $B'$ , and hence  $\mathcal{R}_{\text{PBF}}(\sigma) \leq 16/13$ . At this point we may stop our analysis, because this is what we want to prove in Thm. 1. On the other hand, we may prove the lemma itself by choosing  $\hat{\sigma} = \nu(B)$  or any other proper sequence with the performance ratio equal to 16/13.  $\square$

**Uniform Sequences.** An important observation at this stage is that when we constrain our consideration only to proper sequences  $\sigma$  (possibly containing LIFTS), PBF behaves like a GREEDY algorithm. The set  $\mathcal{I}$  always touches both upper boundaries of  $\mathcal{H}$  (and in fact, the whole set  $\mathcal{I}$  is therefore defined by a single variable  $\|\mathcal{I}\|$ ). Thus, its perpendicular bisector always coincides with the main anti-diagonal of  $\mathcal{H}$  and PBF just tries to move as close to it as possible, i.e., transmits packets to minimize the maximal level of packets in its buffers. This is exactly the pivoting rule of GREEDY. Note that it does not mean that PBF is just a GREEDY algorithm, which is known to be not optimal. It only means that it behaves in such manner on special kind of sequences (the proper ones). The worst-case behavior for the GREEDY algorithm is incurred by a sequence which do not start from full buffers (see [11]). Now, we want to further simplify the structure of the worst-case instances.

**Lemma 7.** *For any proper sequence  $\sigma$ , which may contain LIFT requests, there exists a uniform sequence  $\hat{\sigma}$ , such that  $\mathcal{R}_{\text{GR}}(\sigma) \leq \mathcal{R}_{\text{GR}}(\hat{\sigma})$ .*

*Proof.* We show a series of transformations of  $\sigma$ , which eventually lead to a uniform sequence  $\hat{\sigma}$ . If  $\sigma$  is not uniform, it contains some LIFT and/or ADD(1) requests; let  $\sigma_{k+1}$  be the last of them. We show that if we replace  $\sigma_{k+1}$  by ADD(0) and possibly exchange the roles of ADD(0) and ADD(1) in  $\sigma^k$ , then the resulting sequence, called  $\tilde{\sigma}$ , incurs non-smaller loss on GREEDY. As the number of injected packets is the same for  $\sigma$  and  $\tilde{\sigma}$  and on proper sequences the set  $\mathcal{I}$  behaves identically for any injection request (ADD(0), ADD(1) or LIFT), we obtain  $\mathcal{R}_{\text{GR}}(\tilde{\sigma}) \geq \mathcal{R}_{\text{GR}}(\sigma)$ . After repeating the above operation at most  $|\sigma|$  times, we end up with a desired uniform sequence  $\hat{\sigma}$ . (This sequence is equal to  $\sigma$  with all ADD(1) and LIFT requests replaced by ADD(0).)

In the following, for succinctness, we omit GR subscripts and superscripts. If  $x_0(\sigma^k) \geq x_1(\sigma^k)$ , then we set  $\tilde{\sigma}^k = \sigma^k$ , otherwise we set  $\tilde{\sigma}^k$  to be equal to  $\sigma^k$  with the roles of ADD(0) and ADD(1) exchanged. This way, it is guaranteed that  $x_0(\tilde{\sigma}^k) \geq x_1(\tilde{\sigma}^k)$  and  $\|\mathbf{x}(\tilde{\sigma}^k)\| = \|\mathbf{x}(\sigma^k)\|$ . As already stated above,  $\tilde{\sigma}_{k+1} = \text{ADD}(0)$ . Moreover,  $\tilde{\sigma}$  has the same length as  $\sigma$  and  $\tilde{\sigma}|_{k+2}^{|\sigma|} = \sigma|_{k+2}^{|\sigma|}$ .

Recall that  $x_0(\tilde{\sigma}^k) \geq x_1(\tilde{\sigma}^k)$ . Furthermore, all the requests within  $\tilde{\sigma}|_{k+1}^{|\sigma|}$  are either ADD(0) that increase only  $x_0$  coordinate or IDLE requests that (for GREEDY) cannot change the ordering of  $x_0$  and  $x_1$ . Therefore,

$$x_0(\tilde{\sigma}^t) \geq x_1(\tilde{\sigma}^t) \tag{9}$$

for any  $t \in \{k, \dots, |\sigma|\}$ . Now, we show inductively show that the following invariants hold for any  $t \in \{k, \dots, |\sigma|\}$ .

- (i)  $\|\mathbf{x}(\sigma^t)\| \geq \|\mathbf{x}(\tilde{\sigma}^t)\|$
- (ii)  $x_0(\tilde{\sigma}^t) - x_1(\tilde{\sigma}^t) \geq x_0(\sigma^t) - x_1(\sigma^t)$

The induction basis holds trivially by the choice of  $\tilde{\sigma}^k$ . We assume that these invariants hold for some  $t < |\sigma|$  and we consider two cases.

1.  $\tilde{\sigma}_{t+1} = \text{IDLE}$ . In this case, invariant (i) is trivially preserved. For showing invariant (ii), we define  $\delta = x_0(\sigma^t) - x_1(\sigma^t)$ . If  $\delta \geq 1$ , then both instances of GREEDY transmit a packet from buffer 0 and thus invariant (ii) is preserved. If  $0 \leq \delta < 1$ , then  $x_0(\sigma^{t+1}) - x_1(\sigma^{t+1}) = 0$ . If  $\delta < 0$ , then again we use the fact that GREEDY does not change the relation between  $x_0$  and  $x_1$ , and therefore  $x_0(\sigma^{t+1}) - x_1(\sigma^{t+1}) \leq 0$ . But (9) implies that  $x_0(\tilde{\sigma}^{t+1}) - x_1(\tilde{\sigma}^{t+1}) \geq 0$ , which settles invariant (ii).
2.  $\tilde{\sigma}_{t+1} = \text{ADD}(0)$ . By adding invariants (ii) and (i) for step  $t$  and reorganizing terms, we obtain  $x_1(\sigma^t) \geq x_1(\tilde{\sigma}^t)$ . As the buffer 1 is left untouched in step  $t + 1$ ,  $x_1(\sigma^{t+1}) \geq x_1(\tilde{\sigma}^{t+1})$ . For showing invariant (i), we observe that if  $x_0(\sigma^t) \leq B - 1$ , then the LHS of invariant (i) increases by 1 and the RHS increases by at most 1. Otherwise,  $\|\mathbf{x}(\sigma^{t+1})\| = B + x_1(\sigma^{t+1}) \geq B + x_1(\tilde{\sigma}^{t+1}) \geq \|\mathbf{x}(\tilde{\sigma}^{t+1})\|$ . For showing invariant (ii), we observe that if  $x_0(\tilde{\sigma}^t) \leq B - 1$ , then the LHS of invariant (ii) increases by 1 and the RHS increases by at most 1. Otherwise,  $x_0(\tilde{\sigma}^{t+1}) - x_1(\tilde{\sigma}^{t+1}) = B - x_1(\tilde{\sigma}^{t+1}) \geq B - x_1(\sigma^{t+1}) \geq x_0(\sigma^{t+1}) - x_1(\sigma^{t+1})$ .

Finally, we may compare the losses of GREEDY on  $\sigma$  and  $\tilde{\sigma}$ . During any packet injection, the loss of an algorithm plus the increase of  $\|\mathbf{x}\|$  is equal to 1. Therefore, invariant (i) implies that  $\text{loss}(\tilde{\sigma}) \geq \text{loss}(\sigma)$ .  $\square$

**Regular Sequences.** We cannot analyze uniform sequences easily, because IDLE and ADD(0) requests can be arbitrarily mixed. In order to alleviate this problem, we show how to change a uniform sequence into a regular one.

**Lemma 8.** *For any uniform sequence  $\sigma$ , there exists a regular sequence  $\hat{\sigma}$ , such that  $\mathcal{R}_{\text{GR}}(\sigma) \leq \mathcal{R}_{\text{GR}}(\hat{\sigma})$ .*

*Proof.* We denote a subsequence ADD(0) IDLE by  $\mathcal{F}$ . We process  $\sigma$  from the beginning to the end, looking for  $\mathcal{F}$ . We show that if  $\sigma$  contains  $\mathcal{F}$  and  $x_0^{\text{GR}} \leq B - 1$  before processing this  $\mathcal{F}$ , then we show that such  $\mathcal{F}$  can be removed from  $\sigma$  without decreasing the performance ratio. Moreover, after the removal the sequence remains proper (and thus also uniform). By applying this removal inductively to any occurrence of  $\mathcal{F}$  in  $\sigma$ , we eventually get a regular  $\hat{\sigma}$ .


Assume that  $\sigma = \sigma_{\text{prec}} \mathcal{F} \sigma_{\text{succ}}$ . We look separately at the change of the throughput of OPT and GREEDY. If  $x_0^{\text{GR}}(\sigma_{\text{prec}}) \leq B - 1$ , then obviously the removal of  $\mathcal{F}$  from  $\sigma$  does not change the state of GREEDY and decreases its throughput by 1. The change of  $T_{\text{OPT}}$  is twofold. First, it decreases by 1, since one IDLE was removed. Second, the throughput on  $\sigma_{\text{succ}}$  may only increase, because  $\|\mathcal{I}(\sigma_{\text{prec}})\| \geq \|\mathcal{I}(\sigma_{\text{prec}} \mathcal{F})\|$ . Moreover, in either case,  $\|I\|$  can only increase after the removal, which implies that the new sequence is also uniform.  $\square$

### 3.3 Performance Ratio on Regular Sequences

In this section we prove that the performance ratio of GREEDY on any regular sequence is at most 16/13. By the previous section, this will prove the competitiveness of PBF. We begin with an observation on the behavior of GREEDY on regular sequences.

**Lemma 9.** *Fix any regular sequence  $\sigma = \mathcal{A}\mathcal{B}(x_1, y_1), \mathcal{B}(x_2, y_2) \dots \mathcal{B}(x_n, y_n)$ . Then  $\mathbf{x}^{\text{GR}}(\sigma) = (B, B - \gamma_i)$ , where  $\gamma_0 = 0$  and  $\gamma_i \in [0, B)$  for all  $i \leq n$ . Moreover, we have the following recurrence relation for  $\gamma_i$*

$$\gamma_i = \begin{cases} \gamma_{i-1} & \text{if } x_i \leq \gamma_{i-1} \\ \frac{\gamma_{i-1} + x_i}{2} & \text{if } x_i > \gamma_{i-1} \end{cases}.$$



**Fig. 4.** Function  $s(x)$

*Proof.* We prove the lemma inductively. For  $i = 0$ , GREEDY ends at  $(B, B)$  and the lemma follows trivially.

Assume that the lemma holds for  $i - 1$ . If  $x_i \leq \gamma_{i-1}$ , then during  $\text{IDLE}^{x_i}$  steps the algorithm transmits from the first buffer, ending in the state  $(B - x_i, B - \gamma_i)$ . On the other hand, if  $x_i > \gamma_{i-1}$ , then during the  $\text{IDLE}^{x_i}$  requests GREEDY first transmits  $\gamma_{i-1}$  packets from the buffer 0 and then  $\frac{1}{2} \cdot (x_i - \gamma_{i-1})$  packets from each buffer, ending at the state  $(B - \frac{1}{2}(\gamma_{i-1} + x_i), B - \frac{1}{2}(\gamma_{i-1} + x_i))$ .

Since  $\sigma$  is regular, within the following  $\text{ADD}(0)^{y_i}$  requests, GREEDY fills up buffer 0, ending at state  $(B, B - \gamma_{i-1})$  or  $(B, B - \frac{1}{2}(\gamma_{i-1} + x_i))$ , respectively. As  $\sigma$  is proper,  $x_i \leq B$ , and therefore  $\gamma_i < B$  for all  $i \leq n$ .  $\square$

**Function  $s$  and the Competitive Ratio.** Before we continue with the proof of the competitiveness, we need to introduce a function  $s$  on interval  $[0, 1)$ . We use it to lower-bound the number of packets needed by the adversary to achieve a specific state of the algorithm. We refrain from giving a closed-form formula, as it makes the construction unnecessarily complicated. Fix any  $x \in [0, 1)$ . Let  $i$  be an integer such that  $x \in [1 - 2^{-i}, 1 - 2^{-(i+1)})$ . We define

$$s(x) = i + \frac{x - (1 - 2^{-i})}{2^{-(i+1)}}. \quad (10)$$

It is easy to check that  $s(x)$  is continuous, piecewise linear, and monotonically increasing. Its plot is presented in Fig. 4. In the appendix, we present the proofs of the two following technical lemmas.

**Lemma 10.** For any  $0 \leq a < b \leq 1$ , it holds that  $s(\frac{a+b}{2}) - s(a) \leq b$ .

**Lemma 11.** For any  $a \in [0, 1)$ , it holds that  $s(a) \geq \frac{16}{3} \cdot a - 2$ .


**Lemma 12.** For any regular sequence  $\sigma = \mathcal{AB}(x_1, y_1) \mathcal{B}(x_2, y_2) \dots \mathcal{B}(x_n, y_n)$ , and a corresponding sequence  $\gamma_1, \gamma_2, \dots, \gamma_n$ , it holds that  $\sum_{i=1}^n x_i/B \geq s(\gamma_n/B)$ .

*Proof.* We note, without proof, that if all  $x_i$  were equal to  $B$ , then in the lemma statement we would have an equality. We prove the lemma inductively, i.e., we show that  $\sum_{i=1}^j x_i/B \geq s(\gamma_j/B)$  holds for any  $1 \leq j \leq n$ . For  $j = 0$ , we have  $s(\gamma_0) = s(0) = 0$ .

Assume that the bound holds for  $j - 1$ . We consider two cases. If  $x_j \leq \gamma_{j-1}$ , then by Lemma 9,  $\gamma_j = \gamma_{j-1}$  and the lemma follows. If  $x_j > \gamma_{j-1}$ , then by Lemma 10 and Lemma 9, it holds that

$$x_j/B \geq s\left(\frac{x_j/B + \gamma_{j-1}/B}{2}\right) - s(\gamma_{j-1}/B) = s(\gamma_j/B) - s(\gamma_{j-1}/B).$$

By adding the inequality above to the induction hypothesis, we get  $\sum_{i=1}^j x_i/B \geq s(\gamma_j/B)$ .  $\square$



**Fig. 5.** Distribution of PB possible states

**Lemma 13.** For any regular sequence  $\sigma$ ,  $\mathcal{R}_{\text{GR}}(\sigma) \leq 16/13$ .

*Proof.* Let  $\sigma = \mathcal{A}\mathcal{B}(x_1, y_1) \mathcal{B}(x_2, y_2) \dots \mathcal{B}(x_n, y_n)$ . This determines the sequence  $\gamma_1, \gamma_2, \dots, \gamma_n$ . Since  $\sigma$  is non-trivial, both OPT and GREEDY transmit packets during all  $\sum_{i=1}^n x_i$  IDLE requests of  $\sigma$ . Afterwards,  $\mathbf{x}^{\text{GR}} = (B - \gamma_n, B)$  and OPT has at most  $2 \cdot B$  packets; these packets are transmitted at the end of  $\sigma$ . Therefore, the reciprocal of the performance ratio on  $\sigma$  is

$$\frac{1}{\mathcal{R}(\sigma)} = \frac{(\sum_{i=1}^n x_i) + B + (B - \gamma_n)}{(\sum_{i=1}^n x_i) + B + B} = 1 - \frac{\gamma_n/B}{2 + \sum_{i=1}^n x_i/B} \geq 1 - \frac{\gamma_n/B}{2 + s(\gamma_n/B)} \geq 1 - \frac{3}{16} = \frac{13}{16},$$

where the last two inequalities follow by Lemma 12 and Lemma 11, respectively.  $\square$

We note that (6) (and thus Thm. 1) follows directly from combining Lemmas 4, 6, 7, 8, and 13.

## 4 Randomization

In this section, we describe a randomized algorithm PB, which works in a standard model and whose expected loss on a sequence  $\sigma$  is exactly the same as the loss of PBF on  $\sigma$  in the fractional model.

PB traces the current state of PBF. In case of IDLE requests, PB tries to transmit a packet in such a way, that the expected number of packets in its buffers is equal to the actual number of packets in the buffers of PBF. In the appendix we show an argument why a straightforward randomization of the fractional solution is inappropriate. We define PB algorithm implicitly, i.e., in each step we show what its probability distribution over possible states should be. First, we define this distribution for any possible state of PBF. Later, we show how to infer the actual behavior of the algorithm on the basis of the distribution.

**Definition 2.** Fix any (fractional) state of the buffers  $\mathbf{x} = (k_0 + a, k_1 + b)$ , where  $k_0, k_1 \in \mathbb{N}$  and  $a, b \in [0, 1]$ . Let  $\mu(\mathbf{x})$  be a random variable with the following distribution.

– If  $a + b \leq 1$ , then

$$\mu(\mathbf{x}) = (k_0, k_1) + \begin{cases} (1, 0) & \text{with probability } a \\ (0, 1) & \text{with probability } b \\ (0, 0) & \text{with probability } 1 - a - b \end{cases}$$



– If  $a + b \geq 1$ , then

$$\mu(\mathbf{x}) = (k_0, k_1) + \begin{cases} (1, 0) & \text{with probability } 1 - b \\ (0, 1) & \text{with probability } 1 - a \\ (1, 1) & \text{with probability } a + b - 1 \end{cases}$$

We observe that  $\mathbf{E}[\mu(\mathbf{x})] = \mathbf{x}$ . We note that this definition is purposely made ambiguous. In particular, it can be easily verified that if a number of packets in the first buffer is an integer, then the resulting value of  $\mu$  is the same, no matter whether we pick  $a = 0$  or  $a = 1$ . The same holds for the second buffer. Also, if  $a + b = 1$ , we may choose any of the two rules above for setting  $\mu(\mathbf{x})$  above, and both would yield the same distribution. For more intuitions about function  $\mu$ , see Fig. 5. Each point from the square represents a legal state of an algorithm in the fractional model. Legal states of the algorithm in the standard model are represented by dots (points with integer coordinates). Then  $\mu(\mathbf{x})$  is just a function which assigns probabilities to the vertices of a triangle enclosing  $\mathbf{x}$ . If  $\mathbf{x}$  lies on a triangle edge or at a triangle vertex, then  $\mu(\mathbf{x})$  has non-zero probability on two points (ends of the edge) or at one point (the vertex), respectively.

The following lemma shows that whenever PBF changes its state from  $\mathbf{x}$  to  $\mathbf{x}'$ , it is possible for PB to change the probability distribution from  $\mu(\mathbf{x})$  to  $\mu(\mathbf{x}')$ .

**Lemma 14.** *It is possible to construct a randomized online algorithm PB for the standard model, such that  $\mathbf{x}^{\text{PB}}(\sigma) = \mu(\mathbf{x}^{\text{PBF}}(\sigma))$  for any sequence  $\sigma$ . In effect,  $\text{loss}_{\text{PB}}(\sigma) = \text{loss}_{\text{PBF}}(\sigma)$ .*

**Theorem 3.** *PB is  $\frac{16}{13}$ -competitive for the packet buffering problem on two buffers.*

## 5 Conclusions

Although we presented an optimal algorithm for a specific case of two buffers, its main idea of tracing set  $\mathcal{I}$  and trying to stay close to its center can be generalized; the game is then played in a  $m$ -dimensional cube  $\mathcal{H}$ . The main open question is whether such algorithms can approach the lower bound of  $h(m)$ .

## References

1. S. Albers and M. Schmidt. On the performance of greedy algorithms in packet buffering. In *Proc. of the 36th ACM Symp. on Theory of Computing (STOC)*, pages 35–44, 2004.
2. Y. Azar and Y. Richter. Management of multi-queue switches in QoS networks. In *Proc. of the 35th ACM Symp. on Theory of Computing (STOC)*, pages 82–89, 2003.
3. M. Chrobak, L. L. Larmore, N. Reingold, and J. Westbrook. Page migration algorithms using work functions. *Journal of Algorithms*, 24(1):124–157, 1997. Also appeared in *Proc. of the 4th ISAAC*, pages 406–415, 1993.
4. M. Englert and M. Westermann. Lower and upper bounds on FIFO buffer management in QoS switches. In *Proc. of the 14th European Symp. on Algorithms (ESA)*, pages 352–363, 2006.
5. M. Englert and M. Westermann. Considering suppressed packets improves buffer management in QoS switches. In *Proc. of the 18th ACM-SIAM Symp. on Discrete Algorithms (SODA)*, pages 209–218, 2007.
6. A. Kesselman, Y. Mansour, and R. van Stee. Improved competitive guarantees for QoS buffering. *Algorithmica*, 43(1–2):63–80, 2005. Also appeared in *Proc. of the 11th ESA*, pages 361–372, 2003.
7. E. Koutsoupias and C. H. Papadimitriou. On the  $k$ -server conjecture. *Journal of the ACM*, 42(5):971–983, 1995. Also appeared in *Proc. of the 26th STOC*, pages 507–511, 1994.
8. W. E. Leland, M. S. Taqqu, W. Willinger, and D. V. Wilson. On the self-similar nature of ethernet traffic (extended version). *IEEE/ACM Transactions on Networking*, 2(1):1–15, 1994.
9. M. May, J. Bolot, A. Jean-Marie, and C. Diot. Simple performance models of differentiated services schemes for the internet. In *Proc. of the IEEE INFOCOM*, pages 1385–1394, 1999.

10. P. Raghavan and C. D. Tompson. Randomized rounding: a technique for provably good algorithms and algorithmic proofs. *Combinatorica*, 7(4):365–374, 1987.
11. M. Schmidt. Packet buffering: Randomization beats deterministic algorithms. In *Proc. of the 22nd Symp. on Theoretical Aspects of Computer Science (STACS)*, pages 293–304, 2005.
12. M. Schmidt. *Online Packet Buffering*. PhD thesis, Albert-Ludwigs-Universität Freiburg, 2006.
13. D. D. Sleator and R. E. Tarjan. Amortized efficiency of list update and paging rules. *Communications of the ACM*, 28(2):202–208, 1985.

## A Proofs of Technical Lemmas

### A.1 The Set $\mathcal{I}$

*Proof (of Lemma 1).* We prove the lemma inductively. At the very beginning  $\mathbf{x}^{\text{ALG}} = (0, 0) \in \mathcal{I} \subseteq \mathcal{SH}(\mathcal{I})$ . Assume that this lemma holds for sequence  $\sigma^{t-1}$ . It means that there exists a state  $(k_0, k_1) \in \mathcal{I}(\sigma^{t-1})$ , which majorizes  $\mathbf{x}^{\text{ALG}}(\sigma^{t-1})$ .

- If  $\sigma_t = \text{IDLE}$ , then we consider two cases. If  $\mathbf{x}^{\text{ALG}}(\sigma^{t-1}) = (0, 0)$ , then ALG remains in this state and the lemma trivially holds. Otherwise ALG transmits a packet from a non-empty buffer, say from buffer 0, which implies  $k_0 \geq 1$ . But  $(k_0 - 1, k_1) \in \mathcal{I}(\sigma^t)$ , and this state majorizes  $\mathbf{x}^{\text{ALG}}(\sigma^{t-1}) - (1, 0) = \mathbf{x}^{\text{ALG}}(\sigma^t)$ .
- If  $\sigma_t = \text{ADD}(0)$ , then we consider three subcases. Analogous reasoning applies if  $\sigma_t = \text{ADD}(1)$ .
  - $k_0 < B$ . In this case,  $\mathbf{x}^{\text{ALG}}(\sigma^t) \leq \mathbf{x}^{\text{ALG}}(\sigma^{t-1}) + (1, 0) \leq (k_0 + 1, k_1) \in \mathcal{I}(\sigma^t)$ .
  - $k_0 = B$  and  $\mathcal{I}(\sigma^{t-1})$  is a singleton set. In this case  $\mathcal{I}(\sigma^t) = \mathcal{I}(\sigma^{t-1})$ . We get  $x_0^{\text{ALG}}(\sigma^t) \leq B = k_0$  and  $x_1^{\text{ALG}}(\sigma^t) = x_1^{\text{ALG}}(\sigma^{t-1}) \leq k_1$ , and therefore  $\mathbf{x}^{\text{ALG}}(\sigma^t) \leq (k_0, k_1) \in \mathcal{I}(\sigma^t)$ .
  - $k_0 = B$  and  $\mathcal{I}(\sigma^{t-1})$  is not a singleton set. By the construction of  $\mathcal{I}$ ,  $\mathcal{I}(\sigma^{t-1})$  contains also  $(k_0 - 1, k_1 + 1)$ , and therefore  $\mathcal{I}(\sigma^t)$  contains  $(k_0, k_1 + 1)$ . Similarly to the previous subcase,  $x_0^{\text{ALG}}(\sigma^t) \leq B = k_0$ ,  $x_1^{\text{ALG}}(\sigma^t) = x_1^{\text{ALG}}(\sigma^{t-1}) \leq k_1 < k_1 + 1$ , and thus  $\mathbf{x}^{\text{ALG}}(\sigma^t) \leq (k_0, k_1 + 1)$ .

Hence, in either case the lemma holds for  $\sigma^t$ . □

*Proof (of Lemma 2).* In order to prove the first part of the lemma, we show that for each step  $t \geq 1$  and  $\mathbf{x} \in \mathcal{I}(\sigma^t)$ , there exists a state  $\mathbf{x}_{\text{prev}} \in \mathcal{I}(\sigma^{t-1})$ , such that after request  $\sigma_t$  and a proper choice of packet transmitted,  $\mathbf{x}_{\text{prev}}$  becomes  $\mathbf{x}$ . Thus, if we choose any sequence  $\sigma$  and any final state from  $\mathcal{I}(\sigma)$ , we can reconstruct the sequence of states of an algorithm  $A$ , which remains for the whole sequence  $\sigma$  in the corresponding  $\mathcal{I}$  set.

We introduce a notion of  $\ell^*$  which is the same measure for set  $\mathcal{I}$  as  $\ell^{\text{ALG}}$  is for an algorithm ALG, i.e.,  $\ell^*(\sigma^t)$  is the number of packets added during an ADD request in step  $t$  to a state from  $\mathcal{I}$ . For an IDLE request in step  $t$ ,  $\ell^*(\sigma^t) = 0$ . This means that for an ADD(0) request in step  $t$ ,  $\mathcal{I}(\sigma^t) = \mathcal{I}(\sigma^{t-1}) + (\ell^*, 0)$  and for ADD(1),  $\mathcal{I}(\sigma^t) = \mathcal{I}(\sigma^{t-1}) + (0, \ell^*)$ .

If  $\sigma_t = \text{IDLE}$ , then it follows from the definition of  $\mathcal{I}(\sigma^t)$  that either  $\mathbf{x} + (1, 0)$ ,  $\mathbf{x} + (0, 1)$  or  $\mathbf{x}$  itself belongs to  $\mathcal{I}(\sigma^{t-1})$ . Choosing  $\mathbf{x}_{\text{prev}}$  as such element of  $\mathcal{I}(\sigma^{t-1})$  guarantees that  $\mathbf{x}$  can be reached from  $\mathbf{x}_{\text{prev}}$  by transmitting packet from an appropriate buffer or not transmitting at all. On the other hand, if  $\sigma_t = \text{ADD}(0)$ , then we choose  $\mathbf{x}_{\text{prev}} = \mathbf{x} - (\ell^*, 0)$ . By the observation above,  $\mathbf{x}_{\text{prev}} \in \mathcal{I}(\sigma^{t-1})$ . The case of  $\sigma_t = \text{ADD}(1)$  is analogous.

Moreover, we observe that for all the requests, it holds that  $\ell^A \equiv \ell^*$ , which means that  $A$  loses a packet during an ADD request which does not change the set  $\mathcal{I}$ .<sup>4</sup>

For proving the second part of the lemma, we fix any algorithm  $A$ , which is always in the set  $\mathcal{I}$ . Note that  $\|\mathbf{x}^A\| \equiv \|\mathcal{I}\|$ . For proving the optimality of  $A$ , we show that for any algorithm

<sup>4</sup> It may only happen if  $\mathcal{I}$  is a singleton set.

ALG and any input sequence  $\sigma$ ,  $\text{loss}^{\text{ALG}}(\sigma) \geq \text{loss}^A(\sigma)$ . Without loss of generality, we may assume that ALG is work-conserving.

When we look at the number of packets lost by ALG in the consecutive steps, we notice that the fewer packets it has, the harder it is to force it to lose another packets. Thus, one may try a strategy of losing more packets than strictly necessary at some step, and then winning back not only the surplus lost, but also some additional ones. We show that no such strategy is feasible.

We inductively prove that for any input  $\sigma$  and step  $0 \leq t \leq |\sigma|$ , it holds that

$$\|\mathbf{x}^{\text{ALG}}(\sigma^t)\| + \text{loss}_{\text{ALG}}(\sigma^t) \geq \|\mathbf{x}^A(\sigma^t)\| + \text{loss}_A(\sigma^t) . \quad (11)$$

Obviously, it holds for  $t = 0$ . For an ADD requests, both sides of inequality increase by 1, i.e., either the ADD request incurs a loss to an algorithm or shifts its position in appropriate direction. For an IDLE request, there is no loss generated. By Lemma 1,  $\|\mathbf{x}^{\text{ALG}}\| \leq \|\mathbf{x}^A\|$ , and therefore if ALG transmits a packet,  $A$  does this as well. Hence, (11) holds.

Finally,  $\|\mathbf{x}^{\text{ALG}}\| \leq \|\mathbf{x}^A\|$  together with (11) implies  $\text{loss}_{\text{ALG}}(\sigma) \geq \text{loss}_A(\sigma)$ , which finishes the proof.  $\square$

## A.2 Non-trivial Sequences

*Proof (of Lemma 3).* It is sufficient to show that the balance is at most  $\max\{\|\mathcal{I}\| - B, 0\}$ . Initially, this condition is trivially fulfilled. Assume that this condition holds at the end of step  $t$ .

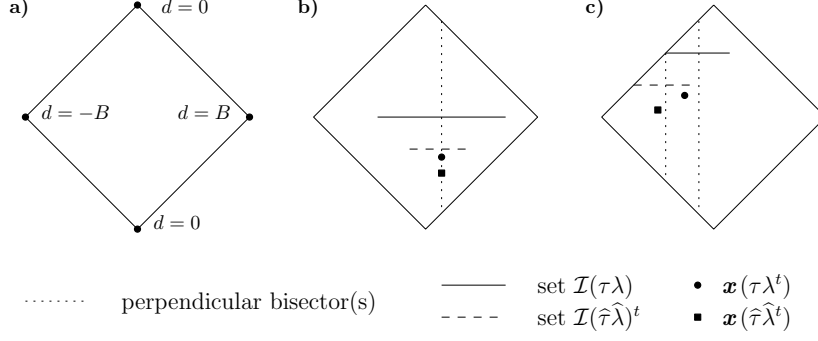
If  $\sigma_{t+1} = \text{ADD}(i)$ , then the balance changes by 1 only if this  $\text{ADD}(i)$  is a hit, in which case  $\mathcal{I}$  has to be above or at MD. As a result,  $\max\{\|\mathcal{I}\| - B, 0\}$  increases as well.

If  $\sigma_{t+1} = \text{IDLE}$ , then the balance may increase only if  $\mathcal{I}$  is below or at MD and touches exactly one lower boundary. In this case, the PBF decreases the balance back in the second part of the step.  $\square$

*Proof (of Lemma 4).* First, we show that we may cut off our sequence  $\sigma$  after the last ADD request which incurs a loss to PBF. Indeed, if  $\sigma$  does not end with ADD request inflicting loss to PBF, then let  $\sigma' = \sigma^{|\sigma|-1}$ . We have  $S(\sigma') \leq S(\sigma)$ ,  $\text{loss}_{\text{OPT}}(\sigma') \leq \text{loss}_{\text{OPT}}(\sigma)$ , and  $\text{loss}_{\text{PBF}}(\sigma') = \text{loss}_{\text{PBF}}(\sigma)$ . Therefore,  $\mathcal{R}_{\text{PBF}}(\sigma') \geq \mathcal{R}_{\text{PBF}}(\sigma)$ . By proceeding inductively, we obtain a sequence ending with an ADD request which incurs a loss to PBF.

For proving the second part of the non-triviality property, we show that if  $\sigma$  contains a step  $t$ , such that  $\mathbf{x}^{\text{PBF}}(\sigma^t) = (0, 0)$ , then there exists a strictly shorter sequence  $\tilde{\sigma}$ , such that  $\mathcal{R}_{\text{PBF}}(\tilde{\sigma}) \geq \mathcal{R}_{\text{PBF}}(\sigma)$ . If a newly obtained sequence  $\tilde{\sigma}$  is not non-trivial, then we may apply this scheme inductively on  $\tilde{\sigma}$ . After finitely many steps, we get a non-trivial sequence  $\hat{\sigma}$ , such that  $\mathcal{R}_{\text{PBF}}(\hat{\sigma}) \geq \mathcal{R}_{\text{PBF}}(\sigma)$ .

To show this, we take the last step  $t$  for which  $\mathbf{x}^{\text{PBF}}(\sigma^t) = (0, 0)$ . By a straightforward induction, we observe that it may only happen when  $\mathcal{I}(\sigma^t)$  touches both lower boundaries. Let  $L = \|\mathcal{I}(\sigma^t)\|$ . Our goal is now to show that if we insert  $L$  IDLE requests directly after step  $t$ , then we may change the remaining suffix of  $\sigma$ , so that the performance ratio of PBF does not decrease. Sequences  $\tau = \sigma^t$  and  $\lambda = \sigma_{t+1}^{|\sigma|}$  fulfill the requirements of Lemma 15 (see below). By applying this lemma  $L$  times, we obtain a sequence  $\sigma^t \text{IDLE}^L \eta$ , such that  $\mathcal{R}_{\text{PBF}}(\sigma^t \text{IDLE}^L \eta) \geq \mathcal{R}_{\text{PBF}}(\sigma)$  and  $|\eta| = |\lambda|$ . Moreover,  $\mathbf{x}^{\text{PBF}}(\sigma^t \text{IDLE}^L) = (0, 0)$  and  $\mathcal{I}(\sigma^t \text{IDLE}^L) = \{(0, 0)\}$ , which means that after the first  $t + L$  steps the game between PBF and the adversary starts over. In effect,  $\mathcal{R}_{\text{PBF}}(\sigma^t \text{IDLE}^L \eta) \leq \max\{\mathcal{R}_{\text{PBF}}(\sigma^t), \mathcal{R}_{\text{PBF}}(\eta)\}$ . We pick  $\tilde{\sigma}$  to be either  $\sigma^t$  or  $\eta$ , choosing the one with higher performance ratio. Finally we note that for any such choice  $|\tilde{\sigma}| < |\sigma|$ , and thus the lemma follows.  $\square$



**Fig. 6.** Types of steps defined in the proof of [Lemma 15](#)

**Lemma 15.** Fix any sequence  $\tau$ , such that  $\mathbf{x}^{\text{PBF}}(\tau) = (0, 0)$ ,  $1 \leq \|\mathcal{I}(\tau)\| \leq B$ , and  $\mathcal{I}(\tau)$  touches both lower boundaries. Let  $\hat{\tau} = \tau \text{ IDLE}$ . Fix any sequence  $\lambda$ , such that  $\mathbf{x}^{\text{PBF}}(\tau\lambda^t) \neq (0, 0)$  for any  $t = 1, \dots, |\lambda|$ . There exists a sequence  $\hat{\lambda}$ , such that  $|\hat{\lambda}| = |\lambda|$  and  $\mathcal{R}_{\text{PBF}}(\hat{\tau}\hat{\lambda}) \geq \mathcal{R}_{\text{PBF}}(\tau\lambda)$ .

*Proof.* Note that values of  $S$ ,  $\text{loss}_{\text{OPT}}$  and  $\text{loss}_{\text{PBF}}$  are the same on  $\tau$  and on  $\hat{\tau}$ . In the following, we omit subscripts and superscripts at  $\mathbf{x}$ .

We create sequence  $\hat{\lambda}$  iteratively on the basis of  $\lambda$ , thinking that  $\lambda$  is created by an adversary. We refer to  $\lambda$  as the *original* sequence and to  $\hat{\lambda}$  as the *modified* one. To simplify the analysis later, we want to make  $\hat{\lambda}$  quite similar to  $\lambda$ . In particular,  $\hat{\lambda}$  will have the same length as  $\lambda$  and will inject packets in exactly same steps. That said,  $\hat{\lambda}$  may inject packets to different buffers than  $\lambda$ .

Additionally, we employ the following *swapping operation*. As at the beginning  $\mathcal{I}(\tau)$  and  $\mathcal{I}(\hat{\tau})$  touch both lower boundaries and  $\mathbf{x}(\tau) = \mathbf{x}(\hat{\tau}) = (0, 0)$ , the buffers are completely symmetric. This means that at step  $t$  we may replace all  $\text{ADD}(0)$  requests in  $\lambda|_1^t$  by  $\text{ADD}(1)$  and vice versa. Note that this does not restrict the power of the adversary of creating an arbitrary sequence: we may think that upon swapping operation the adversary performs the same replacement in the not yet processed part of  $\lambda$ . This way, we merely swapped the meaning of two buffers in the original sequence.

We define a couple of general invariants.

- (I1)  $\text{loss}_{\text{OPT}}(\tau\lambda^t) = \text{loss}_{\text{OPT}}(\hat{\tau}\hat{\lambda}^t)$ ;  
(I2)  $\|\mathcal{I}(\tau\lambda^t)\| = \|\mathcal{I}(\hat{\tau}\hat{\lambda}^t)\| + 1$ .

Additionally, we divide steps into two different types and define invariants for each type. To this end, we first introduce a few notions. By  $\text{pb}(\mathcal{I})$  we denote the perpendicular bisector of the smallest interval containing set  $\mathcal{I}$ . For any point of  $\mathbf{y} \in \mathcal{H}$ , we define its *horizontal position* as  $d(\mathbf{y}) = y_1 - y_0$ . Thus, the horizontal position of the leftmost point of  $\mathcal{H}$  is  $-B$  and the rightmost one is  $B$ , cf. [Fig. 6a](#). Note that points from the same vertical lines of  $\mathcal{H}$  (in particular from the  $\text{pb}(\mathcal{I})$ ) have the same horizontal positions. We call a difference between the values of  $d$  of two points the *horizontal distance* between these two points.

A T1-type step  $t$  is depicted in [Fig. 6b](#). Its invariants are:

- (P1a)  $\text{len}(\mathcal{I}(\tau\lambda^t)) = \text{len}(\mathcal{I}(\hat{\tau}\hat{\lambda}^t)) + 1$ ;  
(P1b)  $\text{pb}(\mathcal{I}(\tau\lambda^t)) = \text{pb}(\mathcal{I}(\hat{\tau}\hat{\lambda}^t))$ ;  
(P1c)  $\|\mathbf{x}(\hat{\tau}\hat{\lambda}^t)\| \leq \|\mathbf{x}(\tau\lambda^t)\| \leq \|\mathcal{I}(\hat{\tau}\hat{\lambda}^t)\|$ ;  
(P1d)  $\mathbf{x}(\tau\lambda^t)$  and  $\mathbf{x}(\hat{\tau}\hat{\lambda}^t)$  lie on  $\text{pb}(\mathcal{I}(\tau\lambda^t)) = \text{pb}(\mathcal{I}(\hat{\tau}\hat{\lambda}^t))$ .

A T2-type step  $t$  is depicted in [Fig. 6c](#). Its invariants are:

- (P2a)  $\text{len}(\mathcal{I}(\tau\lambda^t)) = \text{len}(\mathcal{I}(\widehat{\tau}\widehat{\lambda}^t))$ ;  
(P2b) either  $\mathcal{I}(\tau\lambda^t) = \mathcal{I}(\widehat{\tau}\widehat{\lambda}^t) + (0, 1)$  and both  $\mathcal{I}(\tau\lambda^t)$  and  $\mathcal{I}(\widehat{\tau}\widehat{\lambda}^t)$  touch the upper 0-boundary (we call  $\mathcal{I}(\tau\lambda^t)$  *right-shifted*),  
or  $\mathcal{I}(\tau\lambda^t) = \mathcal{I}(\widehat{\tau}\widehat{\lambda}^t) + (1, 0)$  and both  $\mathcal{I}(\tau\lambda^t)$  and  $\mathcal{I}(\widehat{\tau}\widehat{\lambda}^t)$  touch the upper 1-boundary (we call  $\mathcal{I}(\tau\lambda^t)$  *left-shifted*);  
(P2c)  $\|\mathbf{x}(\widehat{\tau}\widehat{\lambda}^t)\| \leq \|\mathbf{x}(\tau\lambda^t)\|$ ;  
(P2d)  $d(\mathbf{x}(\widehat{\tau}\widehat{\lambda}^t)) \leq d(\mathbf{x}(\tau\lambda^t))$ ;  
(P2e)  $d(\mathbf{x}(\widehat{\tau}\widehat{\lambda}^t)) \leq d(\text{pb}(\mathcal{I}(\widehat{\tau}\widehat{\lambda}^t)))$ ;  
(P2f)  $d(\mathbf{x}(\tau\lambda^t)) \leq d(\text{pb}(\mathcal{I}(\tau\lambda^t)))$ .

The inequalities in invariants (P2d)–(P2f) hold for a right-shifted  $\mathcal{I}(\tau\lambda^t)$ . For a left-shifted one, they are reversed.

We inductively show that the invariants hold. Clearly, step  $t = 0$  is of type T1: all invariants are trivially preserved as  $\mathbf{x}(\tau) = \mathbf{x}(\widehat{\tau}) = (0, 0)$ . We assume that the invariants hold for step  $t < T$  and we show them for step  $t + 1$ . We refrain for formally showing invariants related solely to sets  $\mathcal{I}$  (i.e., (I1), (I2), (P1a), (P1b), (P2a), (P2b)) as they will follow immediately from our construction. Unless specified otherwise, we choose  $\lambda_{t+1} = \lambda_{t+1}$ ,

*Preserving invariants when  $t$  is of type T1.*

1.  $\lambda_{t+1} = \text{ADD}$  and this ADD is not a hit in sequence  $\tau\lambda^{t+1}$ . In this case, this ADD is also not a hit in sequence  $\widehat{\tau}\widehat{\lambda}^{t+1}$ . By invariants (P1c) and (P1d), no loss is incurred to any of the two instances of PBF. Hence, this ADD just shifts both sets  $\mathcal{I}$  and points  $\mathbf{x}$  and all invariants are preserved.
2.  $\lambda_{t+1} = \text{ADD}$  and this ADD is a hit in sequence  $\tau\lambda^{t+1}$ . We show that  $t + 1$  is of type T2. As there is no hit for the sequence  $\widehat{\tau}\widehat{\lambda}^{t+1}$ , invariants (P1c) and (P1d) guarantee no loss for both instances of PBF. This implies invariant (P2c). As  $d(\mathbf{x}(\tau\lambda^{t+1})) = d(\mathbf{x}(\widehat{\tau}\widehat{\lambda}^{t+1})) = d(\text{pb}(\mathcal{I}(\widehat{\tau}\widehat{\lambda}^{t+1}))) = d(\text{pb}(\mathcal{I}(\tau\lambda^{t+1}))) \pm 1$ , invariants (P2d)–(P2f) follow.
3.  $\lambda_{t+1} = \text{IDLE}$ . By the lemma assumption,  $\|\mathbf{x}(\tau\lambda^{t+1})\| \geq 0$ . In this case,  $\|\mathcal{I}\|$  and the number of packets in the buffers of PBF run on the original sequence decrease by 1. The number of packets in the buffers of PBF run on the modified sequence may decrease by a smaller amount if  $\|\mathbf{x}(\widehat{\tau}\widehat{\lambda}^{t+1})\| = 0$ . In either case, invariant (P1c) follows. Invariant (P1d) follows trivially.

*Preserving invariants when  $t$  is of type T2.*

Without loss of generality, we assume that at step  $t$ ,  $\mathcal{I}(\tau\lambda^t)$  is *right-shifted*, i.e.,  $\mathcal{I}(\tau\lambda^t) = \mathcal{I}(\widehat{\tau}\widehat{\lambda}^t) + (0, 1)$ . Recall that this means that both  $\mathcal{I}(\tau\lambda^t)$  and  $\mathcal{I}(\widehat{\tau}\widehat{\lambda}^t)$  touch upper 0-boundary and the horizontal distance between their pbs is exactly 1.

1.  $\lambda_{t+1} = \text{IDLE}$ . If  $\mathcal{I}(\widehat{\tau}\widehat{\lambda}^t)$  lies strictly above the main diagonal, then pbs remain intact. As both instances of PBF are trying to reduce their distance to the respective pbs, whose mutual horizontal distance is 1, invariants (P2d)–(P2f) are preserved. Invariant (P2c) follows trivially. If, however,  $\mathcal{I}(\widehat{\tau}\widehat{\lambda}^t)$  lies on the main diagonal, then step  $t + 1$  is of type T1, where invariant (P1c) follows trivially and invariant (P1d) follows by Lemma 3.
2.  $\lambda_{t+1} = \text{ADD}(0)$ . Such request is clearly a hit. If both  $\mathcal{I}$  are points on the boundary then this incurs a loss of a packet for OPT in both instances. The pbs remain in place, while  $\mathbf{x}$  corresponding to PBF instances may be shifted. This preserves invariants (P2d)–(P2f). Invariant (P2c) of step  $t$  states that there exists  $c \geq 0$  such that  $\|\mathbf{x}(\widehat{\tau}\widehat{\lambda}^t)\| + c = \|\mathbf{x}(\tau\lambda^t)\|$ . On the other hand, invariant (P2d) of step  $t$  states  $x_1(\widehat{\tau}\widehat{\lambda}^t) - x_0(\widehat{\tau}\widehat{\lambda}^t) \leq x_1(\tau\lambda^t) - x_0(\tau\lambda^t)$ . Therefore,  $x_0(\widehat{\tau}\widehat{\lambda}^t) \geq x_0(\tau\lambda^t) - c/2$ . By adding a packets to the buffer 0 in step  $t + 1$ , the loss incurred to the modified instance can be at most  $c/2$  smaller than the loss incurred to the original one. Therefore,  $\|\mathbf{x}(\widehat{\tau}\widehat{\lambda}^{t+1})\| + c/2 \leq \|\mathbf{x}(\tau\lambda^{t+1})\|$ , which yields invariant (P2c).

3.  $\lambda_{t+1} = \text{ADD}(1)$  and  $\mathcal{I}(\tau\lambda^t)$  does not touch the upper 1-boundary. As this request is not a hit, sets  $\mathcal{I}$  and points  $\mathbf{x}$  are just shifted and all invariants hold.
4.  $\lambda_{t+1} = \text{ADD}(1)$  and  $\mathcal{I}(\tau\lambda^t)$  touches the upper 1-boundary. Note that in such case  $d(\text{pb}(\mathcal{I}(\tau\lambda))) = 0$  and  $d(\text{pb}(\mathcal{I}(\widehat{\tau}\widehat{\lambda}))) = -1$ . We consider two subcases.
  - (a) If  $d(\mathbf{x}(\tau\lambda^t)) \leq -1$ , then we set  $\widehat{\lambda}^{t+1} = \text{ADD}(0)$ . For such choice, both  $\text{pbs}$  remain intact, and invariants (P2d)–(P2f) follow. Furthermore, this step incurs no loss on PBF on the original sequence, and therefore invariant (P2c) follows as well. Note that this step may incur a loss of OPT on the modified sequence in case when set  $\mathcal{I}(\widehat{\tau}\widehat{\lambda}^t)$  contains a single point (on the upper boundary). However, in such case  $\mathcal{I}(\tau\lambda)$  consists of a single point  $\{(B, B)\}$  and  $\lambda_{t+1}$  incurs a loss on OPT on the original sequence, too. Hence invariant (I1) follows.
  - (b) If  $d(\mathbf{x}(\tau\lambda^t)) \in (-1, 0]$ , then we perform swapping operation. Note that the operation does not change the shape of  $\mathcal{I}(\tau\lambda^t)$  as this set is buffer symmetric. It does however change the position of  $\mathbf{x}(\tau\lambda^t)$ , so that  $d(\mathbf{x}(\tau\lambda^t)) \in [0, 1)$  and  $\lambda_{t+1} = \text{ADD}(0)$ . In particular, after performing this request,  $d(\mathbf{x}(\tau\lambda_t)) < 0$ , and thus invariant (P2f) is preserved. In this case, we set  $\widehat{\lambda}^{t+1} = \text{ADD}(0)$  and apply the analysis from case 2.

After the input sequence ends, we analyze the performance ratio on the original and the modified sequence. Clearly,  $S(\widehat{\tau}\widehat{\lambda}) = S(\tau\lambda)$ . Within  $\lambda$ , PBF run on the original instance of PBF transmitted load 1 for any IDLE request and when run on the modified instance transmitted load at most 1. Hence, by a simple induction we obtain the following relation:  $\|\mathbf{x}(\tau\lambda)\| - \|\mathbf{x}(\tau)\| + \Delta_\tau \text{loss}_{\text{PBF}}(\lambda) \leq \|\mathbf{x}(\widehat{\tau}\widehat{\lambda})\| - \|\mathbf{x}(\widehat{\tau})\| + \Delta_{\widehat{\tau}} \text{loss}_{\text{PBF}}(\widehat{\lambda})$ . As  $\text{loss}_{\text{PBF}}(\tau) = \text{loss}_{\text{PBF}}(\widehat{\tau})$  and  $\|\mathbf{x}(\tau)\| = \|\mathbf{x}(\widehat{\tau})\|$ , it holds that  $\text{loss}_{\text{PBF}}(\widehat{\tau}\widehat{\lambda}) - \text{loss}_{\text{PBF}}(\tau\lambda) \geq \|\mathbf{x}(\tau\lambda)\| - \|\mathbf{x}(\widehat{\tau}\widehat{\lambda})\|$ .

If the sequence ends with a step of type T1, then  $\text{loss}_{\text{OPT}}(\widehat{\tau}\widehat{\lambda}) = \text{loss}_{\text{OPT}}(\tau\lambda)$  (by invariant (I1)) and  $\text{loss}_{\text{PBF}}(\widehat{\tau}\widehat{\lambda}) \geq \text{loss}_{\text{PBF}}(\tau\lambda)$  (by invariant (P1c)). Same relations hold if a sequence ends with a step of type T2 (by (I1) and (P2c)). Therefore,  $\mathcal{R}(\widehat{\tau}\widehat{\lambda}) \geq \mathcal{R}(\tau\lambda)$ .  $\square$

### A.3 Properties of Function $s()$

*Proof (of Lemma 10).* We fix any  $0 \leq a < b \leq 1$ . Let  $i$  be an integer, such that  $1 - 2^{-i} \leq a < 1 - 2^{-(i+1)}$ . Note that for any  $b > a$ , it holds that  $\frac{a+b}{2} < 1 - 2^{-(i+2)}$ . We consider two cases.

- (i) If  $\frac{a+b}{2} < 1 - 2^{-(i+1)}$ , then both  $a$  and  $\frac{a+b}{2}$  belong to the same linear segment of the function  $s$ . Then,

$$\begin{aligned}
s\left(\frac{a+b}{2}\right) - s(a) &= \left(i + \frac{\frac{a+b}{2} - (1 - 2^{-i})}{2^{-(i+1)}}\right) - \left(i + \frac{a - (1 - 2^{-i})}{2^{-(i+1)}}\right) \\
&= 2^{i+1} \cdot \left(\frac{a+b}{2} - a\right) \\
&= 2^i \cdot (b - a) \\
&\leq 2^i \cdot [b - (1 - 2^{-i})] \\
&= (2^i - 1) \cdot (b - 1) + b \\
&\leq b .
\end{aligned}$$



- (ii) If  $1 - 2^{-(i+1)} \leq \frac{a+b}{2} < 1 - 2^{-(i+2)}$ , then  $a$  and  $\frac{a+b}{2}$  belong to the two consecutive linear fragments of the function  $s$ . In this case,

$$\begin{aligned}
s\left(\frac{a+b}{2}\right) - s(a) &= \left( (i+1) + \frac{\frac{a+b}{2} - (1 - 2^{-(i+1)})}{2^{-(i+2)}} \right) - \left( i + \frac{a - (1 - 2^{-i})}{2^{-(i+1)}} \right) \\
&= 1 + \frac{(a+b) - (2 - 2^{-i})}{2^{-(i+1)}} - \frac{a - (1 - 2^{-i})}{2^{-(i+1)}} \\
&= 1 + 2^{i+1} \cdot (b-1) \\
&= (2^{i+1} - 1) \cdot (b-1) + b \\
&\leq b.
\end{aligned}$$

Thus, in both cases it holds that  $s(\frac{a+b}{2}) - s(a) \leq b$ .  $\square$

*Proof (of Lemma 11).* Let  $g(a) = \frac{16}{3} \cdot a - 2$ . It suffices to show that  $s(a) \geq g(a)$  for any  $a \in [0, 1]$ . The plot of both functions is depicted in Fig. 4, but the relation between  $s(a)$  and  $g(a)$  can be also proved analytically.

For  $a \in [0, \frac{1}{2}]$ ,  $s(a) = 2a$ , and for  $a \in [\frac{1}{2}, \frac{3}{4}]$ ,  $s(a) = 4a - 1$ . It can be checked that for these cases  $s(a) \geq g(a)$ . On the other hand,  $s(\frac{3}{4}) = g(\frac{3}{4})$  and for  $a \geq \frac{3}{4}$ ,  $s(a)$  is growing faster than  $g(a)$ . To show this, it is sufficient to observe that the function  $s$  is continuous, differentiable on the whole  $(0, 1)$  except for the points  $1 - 2^{-i}$ , and its first derivative on  $(\frac{3}{4}, 1)$  is at least 4, which is more than  $\frac{16}{3}$ , the first derivative of  $g(a)$ .  $\square$


#### A.4 Randomization

*Proof (of Lemma 14).* We prove the lemma by induction on the number of steps. At the beginning, the invariant is trivially fulfilled, as both algorithms start with empty buffers. Assume that  $\mathbf{x}^{\text{PB}}(\sigma^{t-1}) = \mu(\mathbf{x}^{\text{PBF}}(\sigma^{t-1}))$ . Fix  $k_0, k_1 \in \mathbb{N}$ ,  $a, b \in [0, 1]$ , such that  $\mathbf{x}^{\text{PBF}}(\sigma^{t-1}) = (k_0 + a, k_1 + b)$ . To keep the description of a random variable concise, we introduce the following notation. We write  $\mathbf{x}^{\text{PB}} = \{(t_1, s_1) : p_1, (t_2, s_2) : p_2, \dots\}$  meaning that  $\mathbf{x}^{\text{PB}}$  is equal to  $(t_i, s_i)$  with probability  $p_i$ . Furthermore, we call a triangle corresponding to the case  $a + b \geq 1$  *upper*; the other ones are called *lower*.

Assume that  $\sigma_t = \text{ADD}(0)$  (for  $\sigma_t = \text{ADD}(1)$ , the reasoning is analogous). If  $\mathbf{x}_1^{\text{PBF}}(\sigma^{t-1}) \leq B - 1$ , then adding a packet to the buffer of PBF and PB just changes the triangle, without violating the relations between  $\mathbf{x}^{\text{PBF}}$  and  $\mathbf{x}^{\text{PB}}$ . The only problem may occur if  $\mathbf{x}_1^{\text{PBF}}(\sigma^{t-1}) > B - 1$ . So let  $k_0 = B - 1$  and assume first that  $\mathbf{x}^{\text{PBF}}(\sigma^{t-1}) = (B - 1 + a, k_1 + b)$  lies in an upper triangle. Hence,  $\mathbf{x}^{\text{PB}}(\sigma^{t-1}) = \{(B, k_1) : 1 - b, (B - 1, k_1 + 1) : 1 - a, (B, k_1 + 1) : a + b - 1\}$ . After adding a packet to the first buffer,  $\mathbf{x}^{\text{PBF}}(\sigma^t) = (B, k_1 + b)$ . On the other hand, PB adds the packet only if it happens to be in the state  $(B - 1, k_1 + 1)$ , in other cases it just loses this packet. Therefore,  $\mathbf{x}^{\text{PB}}(\sigma^t) = \{(B, k_1) : 1 - b, (B, k_1 + 1) : b\} = \mu(\mathbf{x}^{\text{PBF}}(\sigma^t))$  and  $\text{loss}_{\text{PB}}(\sigma^t) = \text{loss}_{\text{PBF}}(\sigma^t)$ . Similar reasoning can be applied if PBF lies in a lower triangle.

The reasoning for  $\sigma_t = \text{IDLE}$  is more complicated. Essentially, we have to show that it is possible to choose a strategy of (randomly) transmitting packets, such that the probability distribution of PB changes from  $\mu(\mathbf{x}^{\text{PBF}}(\sigma^{t-1}))$  to  $\mu(\mathbf{x}^{\text{PBF}}(\sigma^t))$ . Obviously, each possible state of PB can transmit once from a single buffer only.

First, we note that PBF is work-conserving, i.e., the total mass of packets sent in step  $t$  is equal to 1. The only exception of this rule occurs if PBF ends step  $t$  with the buffers completely empty. This may occur only if  $\mathbf{x}^{\text{PBF}}(\sigma^{t-1})$  belongs to the lower triangle with vertices  $(0, 0)$ ,  $(0, 1)$ , and  $(1, 0)$ . By simply transmitting from a non-empty buffer if PB is at state  $(0, 1)$  or  $(1, 0)$ , we may assure that  $\mathbf{x}^{\text{PB}}(\sigma^t) = 0 = \mathbf{x}^{\text{PBF}}(\sigma^t)$ .



**Fig. 7.** State transitions of PB

For the remaining part of the proof, we consider two cases, depicted on Fig. 7a and Fig 7b.  $\mathbf{x}^{\text{PBF}}(\sigma^{t-1})$  belongs either to an upper triangle  $S$  or to a lower triangle  $S'$ . Therefore,  $\mathbf{x}^{\text{PBF}}(\sigma^t)$  may belong to one of the three triangles  $A$ ,  $B$ , or  $C$  (or respectively  $A'$ ,  $B'$ , or  $C'$  if it starts from a lower triangle). A set of possible  $\mathbf{x}^{\text{PBF}}(\sigma^t)$  positions is an interval depicted with a thick line in the figure. We fix  $k'_0, k'_1 \in \mathbb{N}$ ,  $a', b' \in [0, 1]$ , such that  $\mathbf{x}^{\text{PBF}}(\sigma^t) = (k'_0 + a', k'_1 + b')$ . The behavior of PB is described by the table below. First two columns denote the starting and ending triangle for  $\mathbf{x}^{\text{PBF}}$ . Instead of writing from which buffer PB has to transmit a packet we write in which state it has to end (we guarantee that to achieve that it has to transmit a packet from exactly one buffer). A rule  $u_i \rightarrow u_j$  means that if PB is in state  $u_i$  then it has to end in state  $u_j$  (see Fig. 7). A rule  $u_i \rightarrow (u_j : p_j, u_k : p_k)$  means that if PB is in state  $u_i$ , it should end in state  $u_j$  with probability  $p_j$  and in state  $u_k$  with probability  $p_k$ .

$\mathbf{x}^{\text{PBF}}(\sigma^{t-1})$	$\mathbf{x}^{\text{PBF}}(\sigma^t)$	Rule	Notes
S	A	$u_5 \rightarrow u_2; u_6 \rightarrow u_4; u_4 \rightarrow (u_1 : \frac{1-a'}{1-a}, u_2 : \frac{a'-a}{1-a})$	$a' \geq a$
	B	$u_6 \rightarrow u_5; u_4 \rightarrow u_2; u_5 \rightarrow (u_3 : \frac{1-b'}{1-b}, u_2 : \frac{b'-b}{1-b})$	$b' \geq b$
	C	$u_4 \rightarrow u_2; u_5 \rightarrow u_2; u_6 \rightarrow (u_5 : \frac{a'}{a+b-1}, u_2 : \frac{b'}{a+b-1})$	$a' + b' = a + b - 1$
S'	A'	$u_4 \rightarrow u_1; u_7 \rightarrow u_4; u_6 \rightarrow (u_3 : \frac{b'}{b}, u_4 : \frac{b-b'}{b})$	$b' \leq b$
	B'	$u_4 \rightarrow u_2; u_6 \rightarrow u_4; u_7 \rightarrow (u_5 : \frac{a'}{a}, u_4 : \frac{a-a'}{a})$	$a' \leq a$
	C'	$u_6 \rightarrow u_4; u_7 \rightarrow u_4; u_4 \rightarrow (u_1 : \frac{1-a'}{1-a-b}, u_2 : \frac{1-b'}{1-a-b})$	$a' + b' = a + b + 1$

To illustrate the concept, we consider the case  $\mathbf{x}^{\text{PBF}}(\sigma^{t-1}) \in S$  and  $\mathbf{x}^{\text{PBF}}(\sigma^t) \in A$  more thoroughly. First, we note that  $a' \geq a$ , and therefore the probabilities occurring in the PB rule are legal. Additionally,  $a + b = a' + b'$ . Second, we compute PB distribution at the end of step  $t$ .

$$\begin{aligned} \Pr[\mathbf{x}^{\text{PB}}(\sigma^t) = u_1] &= \frac{1-a'}{1-a} \cdot \Pr[\mathbf{x}^{\text{PB}}(\sigma^{t-1}) = u_4] = 1 - a' , \\ \Pr[\mathbf{x}^{\text{PB}}(\sigma^t) = u_4] &= \Pr[\mathbf{x}^{\text{PB}}(\sigma^{t-1}) = u_6] = a + b - 1 = a' + b' - 1 , \end{aligned}$$

which implies that

$$\Pr[\mathbf{x}^{\text{PB}}(\sigma^t) = u_2] = 1 - (1 - a') - (a' + b' - 1) = 1 - b' .$$

The proof for the remaining five cases is analogous.  $\square$

### A.5 Why naive randomization does not work?

Let us take a look at a naive approach. When PBF transmits a fraction  $\delta$  from the first buffer and  $1 - \delta$  from the other, then the naive randomized algorithm PBN chooses a transmitting

buffer randomly with probabilities  $\delta$  and  $1 - \delta$ . Such approach does not guarantee that the expected numbers of packets in PBN buffers are equal to the fractional amounts of packets in the corresponding buffers of PB.

For example, consider a sequence  $\text{ADD}(0), \text{ADD}(1), \text{IDLE}, \text{IDLE}$  for  $B = 1$ . PBF ends with empty buffers, whereas after the prefix  $\text{ADD}(0), \text{ADD}(1), \text{IDLE}$ , state  $\mathbf{x}^{\text{PBN}}$  is equal to  $(0, 1)$  or  $(1, 0)$ , both with probability  $1/2$ . For the last  $\text{IDLE}$  request, PBN tries to transmit a packet from an empty buffer with probability  $1/2$ , which results in expected non-zero number of packets in the buffers.

Similar situation occurs for  $\text{ADD}$  requests: consider a sequence  $\text{ADD}(0)^2, \text{ADD}(1)^2, \text{IDLE}^2, \text{ADD}(0)$  for  $B = 2$ . Obviously, PBF serves this sequence without losing any packets. On the other hand, after processing everything but the last  $\text{ADD}(0)$ , PBN has non-zero probability of having two packets in the first buffer, which leads to non-zero expected loss due to the last  $\text{ADD}(0)$  request.