

BIN PACKING WITH QUEUES

DEVAVRAT SHAH* ** AND

JOHN N. TSITSIKLIS,* *** *Laboratory for Information and Decision Systems, MIT*

Abstract

We study the best achievable performance (in terms of the average queue size and delay) in a stochastic and dynamic version of the bin-packing problem. Items arrive to a queue according to a Poisson process with rate 2ρ , where $\rho \in (0, 1)$. The item sizes are independent and identically distributed (i.i.d.) with a uniform distribution in $[0, 1]$. At each time unit, a single unit-size bin is available and can receive any of the queued items, as long as their total size does not exceed 1. Coffman and Stolyar (1999) and Gamarnik (2004) have established that there exist packing policies under which the average queue size is finite for every $\rho \in (0, 1)$. In this paper we study the precise scaling of the average queue size, as a function of ρ , with emphasis on the critical regime where ρ approaches 1. Standard results on the probabilistic (but static) bin-packing problem can be readily applied to produce policies under which the queue size scales as $O(h^2)$, where $h = 1/(1 - \rho)$, which raises the question of whether this is the best possible. We establish that the average queue size scales as $\Omega(h \log h)$, under any policy. Furthermore, we provide an easily implementable policy, which packs at most two items per bin. Under that policy, the average queue size scales as $O(h \log^{3/2} h)$, which is nearly optimal. On the other hand, if we impose the additional requirement that any two items packed together must have near-complementary sizes (in a sense to be made precise), we show that the average queue size must scale as $\Theta(h^2)$.

Keywords: Bin packing; queueing; heavy traffic

2000 Mathematics Subject Classification: Primary 90B36

Secondary 90B22

1. Introduction

Bin packing is an extremely well-studied combinatorial optimization problem, which arises naturally in many situations including, for example, bandwidth allocation (see, for example, [5]) and selection of multicast trees (see, for example, [20]). The classical version of the problem is as follows. We are given n items, of sizes X_1, \dots, X_n with $X_i \in [0, 1]$, and an unlimited supply of bins of unit size. Items i_1, \dots, i_p can be packed into the same bin if $X_{i_1} + \dots + X_{i_p} \leq 1$. The goal is to pack all n items in the smallest possible number of bins. In the offline setting, all item sizes are known up front, while in the online setting, items sizes are revealed one by one and each item must be placed in a bin as soon as it appears. It is well known that the decision version of the offline bin-packing problem is NP-complete. This has resulted in extensive work on polynomial-time (as a function of n) approximation algorithms for both the offline and online settings; see, for example, the extensive survey by Coffman *et al.* [6].

Received 5 December 2006; revision received 8 August 2008.

* Postal address: Laboratory for Information and Decision Systems, MIT, 77 Massachusetts Avenue, Cambridge, MA 02139, USA.

** Email address: devavrat@mit.edu

*** Email address: jnt@mit.edu

The probabilistic version of the bin-packing problem, in which the sizes X_1, \dots, X_n are assumed to be independent and identically distributed (i.i.d.) with a known distribution μ on $[0, 1]$, has also been extensively studied in both the offline and online settings. For example, Knodel [10] and Lueker [12] established that the expected waste in a stochastic offline setting is $\Theta(\sqrt{n})$; Shor [19] studied online algorithms (best fit and first fit); Bentley *et al.* [2] studied the first fit and first-fit decreasing algorithms; Karp *et al.* [9] studied multidimensional bin packing; Leighton and Shor [11] provided various results through a connection with grid-matching problems; and several results are given in a series of papers by Rhee and Talagrand [14], [15], [16], [17], [18]. A central theme in this literature is that, for both the offline and online settings, and for a large class of distributions, the expected number of required bins is $c(\mu)n + O(\sqrt{n})$; furthermore, with high probability, the number of required bins is $c(\mu)n + O(\sqrt{n \log n})$ [16], [17]. Here $c(\mu)$ is a constant determined by the distribution μ . For the case of the uniform distribution on $[0, 1]$, $c(\mu)$ is equal to $\frac{1}{2}$. The $O(\sqrt{n})$ term can be interpreted as the average waste caused by the randomness in the item sizes. However, in the online case, the waste is $\Omega(\sqrt{n \log n})$ with high probability [19], and there exist online algorithms [19] whose waste is $O(\sqrt{n \log^{3/4} n})$ with high probability.

In this paper we consider a different model, in which items arrive to a queue according to a Poisson process with rate 2ρ , where $\rho \in (0, 1)$, and the item sizes are i.i.d. with a uniform distribution in $[0, 1]$. At each time unit, a single unit-size bin is available and can receive any of the queued items, as long as their total size does not exceed 1.

We highlight some key differences between the online setting and our model. In the online setting, items need to be packed as soon as they arrive. In contrast, our model allows queueing, which provides additional flexibility, and can in principle result in smaller waste, albeit at the expense of delay. (For example, under our model, an item can be queued until the arrival of a ‘matching’ item, i.e. an item whose size is approximately 1 minus the size of the queued item.) On the other hand, the online setting assumes that all bins are always available. In contrast, our model assumes that a bin remains available for only one time unit, and this restriction may result in larger waste. Because of these two differences, the online model and our queueing model cannot be reduced to each other, existing results are not directly applicable, and a new analysis is required.

At a higher level, our work deals with the trade-off between throughput and delay. However, because of the stability constraint, high throughput is equivalent to low waste. Thus, we are in effect dealing with the trade-off between waste and delay.

Our model is motivated from contexts such as networking, where a bin corresponds to the available bandwidth during a single time slot, and an item corresponds to a packet or a file that can only be transmitted ‘unbroken’ (within one time slot). In an alternative interpretation, bins correspond to regularly scheduled fixed-capacity trucks, and items correspond to pieces of cargo that cannot be broken up into smaller pieces.

1.1. Our model

We provide here a precise formulation of the queueing version of the bin-packing problem, which is the subject of this paper. We assume that items arrive according to a Poisson process with rate 2ρ , where $\rho \in (0, 1)$. (Different arrival models are possible. For example, we could assume that the numbers of arrivals during each unit time interval are independent random variables with mean 2. Renewal arrival process models are also possible. Our results remain valid for more general arrival processes, under suitable assumptions; see the discussion in Section 3.6.) We index the items according to their arrival order, so that item i is the i th

arriving item after time 0. We let X_i be the size of the i th item, and assume that the random variables X_i are i.i.d., uniformly distributed in $[0, 1]$, and independent of the arrival process.

Arriving items join a queue and remain in queue until they are placed in a bin. At each positive integer time, a bin of size 1 becomes available (‘arrives’) and can receive any of the items that are in queue at that time, as long as their total size does not exceed 1. (In an alternative model, bins could be arriving according to a stochastic arrival process, for example, a Poisson process with rate 1. As will be discussed in Section 3.6, the results do not change under the assumption of Poisson bin arrivals.) The items that are packed are removed from the queue and are discarded forever, together with the corresponding bin. We let $Q(t)$ be the number of items in queue just before time t , i.e. without including the effects of an item or bin arrival at exactly time t . (Thus, $Q(t^-)$ would have been a more accurate, though more cumbersome, notation.)

A *policy* is a rule that, at each integer time t , selects the items (if any) to be placed in the available bin, as a function of the available items and the entire history of the process until just before time t . As before, let $F^\pi(\rho)$ be the value of $\limsup_{t \rightarrow \infty} E[Q(t)]$ under policy π , as a function of the arrival parameter ρ , starting from an empty queue. Let $F^*(\rho)$ be the infimum of $F^\pi(\rho)$ over all policies π . We are interested in the behavior of $F^*(\rho)$ as $\rho \uparrow 1$ and, more specifically, in terms of the parameter h defined by

$$h = \frac{1}{1 - \rho}.$$

For some preliminary insights into the problem, we make some observations. If we were to relax the problem and allow items to be broken into pieces that can be placed in different bins, then the system becomes almost (the reason for the qualification ‘almost’ is that services can only start at integer times) identical to a single-server M/U/1 queue with Poisson arrivals (arrival rate of 2ρ), uniformly distributed service times (mean service time of $\frac{1}{2}$), and utilization factor equal to $2\rho\frac{1}{2} = \rho$. The M/U/1 queue is stable if and only if $\rho < 1$, and the steady-state queue size is of order $\Theta(h)$ when $\rho < 1$. This readily implies that $F^*(\rho) = \infty$ when $\rho \geq 1$ and $F^*(\rho) = \Omega(h)$ when $\rho < 1$ (this follows from Kingman’s lower bound [21]). It is not hard to come up with stable policies and show that $F^*(\rho) < \infty$ for all $\rho < 1$. However, available results, for example, from Coffman and Stolyar [5] and Gamarnik [8], only lead to an upper bound of the form $F^*(\rho) = O(h^2)$. This motivates the central question to be addressed in this paper, namely, whether the behavior of $F^*(\rho)$ is of the form $\Theta(h)$, of the form $\Theta(h^2)$, or something in between.

1.2. Comparison with other stochastic and dynamic combinatorial problems

A question similar to the one studied in this paper has been raised and studied for a stochastic and dynamic version of the traveling repairman problem (TRP) [3], [4]. In that problem, jobs arrive as a Poisson process with rate λ at random locations (for example, uniformly distributed) in the unit square. The different jobs have i.i.d. processing time requirements with mean $1/\mu$. A repairman moves from job to job at unit speed and spends time traveling, as well as for processing the jobs. Let $\rho = \lambda/\mu$. If the repairman could travel infinitely fast, we would be dealing with an M/G/1 queueing system, whose expected queue size is $\Theta(h)$, where again $h = 1/(1 - \rho)$. However, because of the finite speed, the repairman must also waste some time traveling between jobs, and the optimal average queue size (number of yet unprocessed jobs) turns out to be of the order of $1/h^2$.

From the analysis of the static probabilistic traveling salesman problem, we know that the travel time (‘waste’) to serve n jobs is of the order of \sqrt{n} . For the bin-packing problem,

the waste is also $\Theta(\sqrt{n})$ and, arguing by analogy, one is tempted to conjecture that queue sizes should also be $\Theta(h^2)$. However, as will be shown, there is no such simple relation between static probabilistic problems and the corresponding stochastic and dynamic problems, with the exception of a straightforward upper bound based on the following batching policy (cf. Section 2.1). This upper bound happens to be tight for the dynamic TRP, but far from tight for the dynamic bin-packing problem.

1.3. Our results

Our main result, Theorem 1.1, below, effectively shows that the penalty caused by the bin-packing constraints (as compared to the relaxed problem where the items can be broken up) is only a polylogarithmic function of $1/(1 - \rho)$.

Theorem 1.1. *We have*

$$F^*(\rho) = \Omega(h \log h),$$

where $h = 1/(1 - \rho)$. Furthermore, there exists a family $\pi(\cdot)$ of policies (one policy for each value of ρ) that packs at most two items in each bin and which satisfies

$$F^{\pi(\rho)}(\rho) = O(h \log^{3/2} h).$$

We will also consider a special class of packing policies, which we call *restricted-pair policies*. Under these policies (to be described in detail in Section 2.2), the unit interval is partitioned into pairs of subintervals, every bin receives at most two items, and if two items are packed together then they must belong to paired subintervals.

Theorem 1.2. *If $\pi(\cdot)$ is a family of restricted-pair policies (one policy for each value of ρ) then*

$$F^{\pi(\rho)}(\rho) = \Omega(h^2),$$

where $h = 1/(1 - \rho)$. Furthermore, there exists a family of restricted-pair policies such that

$$F^{\pi(\rho)}(\rho) = O(h^2).$$

1.4. Organization

The rest of the paper is organized as follows. In Section 2 we briefly discuss a batching policy and argue that the corresponding average queue size is $\Theta(h^2)$. We introduce restricted-pair policies and establish an $\Omega(h^2)$ lower bound for such policies. We also display a simple round-robin restricted-pair policy and a corresponding $O(h^2)$ upper bound. In Section 3 we introduce a more powerful policy and use results of [11] to derive an $O(h \log^{3/2} h)$ upper bound. Section 3.6 contains some brief remarks on possible extensions. Then, in Section 4 we build on results of [1] and [19] to establish an $\Omega(h \log h)$ lower bound for any policy. Finally, Section 5 contains our conclusions.

2. Two inadequate policies

In this section we examine the best performance that can be obtained with relatively simple policies. We consider two alternative policies and show that they both result in $\Omega(h^2)$ queue sizes; thus, justifying the development of more elaborate policies in the next section. The first policy is based on batching, which is a generic method, applicable to wide classes of stochastic combinatorial problems; despite its simplicity, it is often near-optimal, as in the dynamic TRP. Its $\Omega(h^2)$ performance in our context however motivates us to consider a more complex policy.

The simplest class of nontrivial, and possibly interesting policies, that comes to mind is one that packs at most two items per bin, but places restrictions on the pairs that can be packed together. We show that any policy in this class also results in $\Omega(h^2)$ queue sizes.

2.1. Batching

Batching is a routine approach for converting policies for static problems to policies for dynamic problems. In our case, a batching policy works as follows. All the items arriving during the interval $[(k - 1)t^*, kt^*)$ form the k th batch. Items in the k th batch can be packed into bins after time kt^* and after all items in the $(k - 1)$ th batch have already been packed. Let S_k be the number of bins used for the k th batch, which can be viewed as the ‘service time’ of that batch. The number of items in the k th batch is concentrated around $2\rho t^*$. From this, and the results of [10] and [12], it follows that $E[S_k] = \rho t^* + \Theta(\sqrt{t^*})$. For the expected ‘work’ in the system (number of bins required for the queued items) to remain bounded, the stability condition $\rho t^* + \Theta(\sqrt{t^*}) \leq t^*$ must hold, which yields $t^* = \Omega(1/(1 - \rho)^2) = \Omega(h^2)$. On the other hand, a typical arriving item must wait until the end of the interval $[(k - 1)t^*, kt^*)$ during which it arrives, resulting in an $\Omega(t^*)$ delay. This shows that the expected queue size associated with the batching policy we have described is $\Omega(h^2)$.

Conversely, it is not hard to show (using, for example, Kingman’s bound [21] and a concentration inequality on the waste, such as the one in [18]) that, by choosing $t^* = \Theta(h^2)$, the resulting expected queue size is $\Theta(h^2)$.

2.2. Restricted-pair policies

In this subsection we partition the range $[0, 1]$ of item sizes into subintervals of equal length. We consider policies that pack at most two items in each bin and require that any two items that are packed together belong to an allowed pair of subintervals.

The intuitive idea behind the approach in this subsection is that, with a given utilization rate ρ , there is a certain amount of slack, $1 - \rho$. If item sizes are discretized, with a discretization error bounded by $c(1 - \rho)$, where $c < 1$, the remaining slack is $(1 - c)(1 - \rho)$; so the optimal achievable performance, as a function of $(1 - \rho)$, should still be of the same order of magnitude. Discretizing the item sizes is the same as partitioning the interval $[0, 1]$ into subintervals of length $O(1 - \rho)$, and taking into account only the interval to which an item belongs, not its exact size. However, besides discretization, we will be placing additional restrictions on the allowed packing policies.

Given the value of ρ , we divide the unit interval into $2m + 1$ equal subintervals of length $\delta = 1/(2m + 1)$, where m is of the order of $1/(1 - \rho)$. For concreteness, we assume that

$$\frac{8}{1 - \rho} \leq 2m \leq \frac{8}{1 - \rho} + 1. \tag{2.1}$$

We note, for future reference, that, for every $\rho \geq \frac{1}{2}$, we have

$$m \geq \frac{2}{1 - \rho}, \quad \delta \geq \frac{1 - \rho}{9}, \quad \frac{2m}{m + 1} \geq 1 - \frac{1 - \rho}{8}. \tag{2.2}$$

We associate a separate queue with each subinterval. In particular, an arriving item whose size belongs to the interval $[(i - 1)\delta, i\delta)$, $1 \leq i \leq 2m$, joins the i th queue, while an item whose size belongs to the interval $[2m\delta, 1]$ joins the $(2m + 1)$ th queue. Since the item sizes are i.i.d. with a uniform distribution in $[0, 1]$, each arriving item is equally likely (independent of everything else) to join each of the $2m + 1$ queues. The well-known *splitting* property of

the Poisson process implies that the arrival processes to each one of the $2m + 1$ queues are independent Poisson processes with rate $\gamma = 2\rho / (2m + 1) = 2\rho\delta$.

Definition 2.1. A restricted-pair policy is a policy with the following two properties.

- (a) No more than two items can be packed into the same bin.
- (b) If an item from queue i is packed together with an item from queue j then $i + j = 2m + 1$.

Note that any two items in queues i and j , where $i + j = 2m + 1$, can always be packed together, because their total size is bounded by $i\delta + j\delta = (2m + 1)\delta = 1$. Furthermore, when two such items are packed together, the wasted space in the bin is bounded by $2\delta \approx (1 - \rho)/4$, which is consistent with our earlier discussion. Restricted-pair policies are simple enough to be tractable. On the other hand, they are restrictive, hence, potentially inefficient. Indeed, the best possible performance of such policies is $\Theta(h^2)$. Next, we present the details of the proof of Theorem 1.2.

2.3. Lower bound for restricted-pair policies

In this subsection we prove the $\Omega(h^2)$ lower bound in Theorem 1.2. Without loss of generality, we assume that $\rho \geq \frac{1}{2}$. For $i = 1, \dots, 2m + 1$, let $N_i(t)$ be the number of arrivals to queue i until time t . For $i = 1, \dots, m$, let $Q_i(t)$ be the total number of items in the two paired queues i and $2m + 1 - i$ at time t . Finally, let $S_i(t)$ be the number of bins packed with items from queues i and/or $2m + 1 - i$ until time t .

We have

$$Q_i(t) \geq \max\{N_i(t), N_{2m+1-i}(t)\} - S_i(t).$$

Let X and Y be independent Poisson random variables with mean $2\rho\delta t$, which is the same as the distribution of $N_i(t)$. Summing over i , using the property that $\sum_{i=1}^m S_i(t) \leq t$, and taking expectations, we obtain

$$\begin{aligned} E[Q(t)] &\geq \sum_{i=1}^m E[Q_i(t)] \\ &\geq \sum_{i=1}^m E[\max\{N_i(t), N_{2m+1-i}(t)\}] - t \\ &= m E[\max\{X, Y\}] - t \\ &= m E[X] + \frac{m}{2} E[|X - Y|] - t \\ &= \rho \frac{2m}{2m + 1} t + \frac{m}{2} E[|X - Y|] - t. \end{aligned} \tag{2.3}$$

Lemma 2.1. Consider a Poisson random variable X with mean λ . There exists a positive constant β , independent of λ , such that

$$P(X > \lambda + \sqrt{\lambda}) \geq \beta \quad \text{and} \quad P(X < \lambda - \sqrt{\lambda}) \geq \beta \quad \text{for all } \lambda \geq 1.$$

Proof. Note that $P(X > \lambda + \sqrt{\lambda})$ is a positive and continuous function of λ . Let β be its infimum over $\lambda \in [1, \infty)$. Suppose that $\beta = 0$. Then the infimum cannot be attained on a compact set (this would contradict continuity and positivity). Thus, the limit, as $\lambda \rightarrow \infty$, must be 0. Recalling that the variance of X is also λ , the central limit theorem implies that the

random variable $Z_\lambda = (X - \lambda)/\sqrt{\lambda}$ converges in distribution to the standard normal distribution. Therefore, $\lim_{\lambda \rightarrow \infty} P(Z_\lambda > 1) > 0$, which is a contradiction. This completes the proof of the lower bound on $P(X > \lambda + \sqrt{\lambda})$. The argument for $P(X < \lambda - \sqrt{\lambda})$ is similar.

Lemma 2.1 establishes that $E[|X - Y|] \geq \beta^2 \sqrt{2\rho\delta t}$, where β does not depend on δ and t , as long as $t = \Omega(1/\delta)$. Using (2.3), Lemma 2.1, and the bounds in (2.2), we obtain, after some straightforward algebra,

$$E[Q(t)] \geq \rho t - \frac{1 - \rho}{8} t + \frac{1}{1 - \rho} \frac{\beta^2}{3} \sqrt{(1 - \rho)t} - t.$$

Let us now set $t = c/(1 - \rho)^3$, where c is chosen such that $2c = \beta^2 \sqrt{c}/6$. Some more algebra yields

$$E[Q(t)] \geq \frac{2c}{(1 - \rho)^2},$$

which is the desired result.

2.4. Upper bound

In this subsection we consider a simple round-robin restricted-pair policy that results in $O(h^2)$ expected queue sizes. Thus, the lower bound of the preceding subsection is tight, completing the proof of Theorem 1.2.

Let m be as in (2.1), and consider a time interval of length $m + 1$. For $i \leq 2$, the i th bin in that interval serves (depending on availability) up to one item from queue i and one time from queue j , where $i + j = 2m + 1$. Finally, the $(m + 1)$ th bin serves up to one item from queue $2m + 1$. Let us focus on items in the i th queue. These items arrive as a Poisson process with rate $\gamma = 2\rho/(2m + 1)$, and one of them can be served every $m + 1$ time steps. This is equivalent to an M/D/1 queue with utilization rate

$$\rho' = 2\rho \frac{m + 1}{2m + 1} \leq 1 - c(1 - \rho)$$

for some absolute constant c , except for a restriction that service can only start at integer multiples of $m + 1$. (The latter restriction adds at most $m + 1$ to the expected delay.) Thus, using standard bounds on M/D/1 queues [21], the expected queue size for items in the i th interval is of order $O(m) = O(h)$. Since there are $O(h)$ queues/intervals, it follows that the expected queue size of the overall system is $O(h^2)$.

3. A near-optimal policy

In this section we present a packing policy, analyze its performance, and show that the average queue size grows not much faster than h . The policy we describe is randomized, although as discussed in Section 5, randomization is easily removed.

3.1. The policy

The policy is as follows. When the i th bin arrives, we independently generate a Bernoulli random variable B_i , with $P(B_i = 1) = \rho$, a Bernoulli random variable C_i , with $P(C_i = 1) = \frac{1}{2}$, and a random variable U_i which is uniformly distributed in the interval $[0, \frac{1}{2}]$. We assume that the random variables associated with each bin arrival are independent.

- (a) If $B_i = 1$, we select the largest-sized item whose size is in the range $[0, U_i]$ (if any) and the largest-sized item whose size is in the range $(\frac{1}{2}, 1 - U_i]$ (if any), and place them in the bin.

- (b) If $B_i = 0$ and $C_i = 0$, we place the largest-sized item whose size is in the range $[0, \frac{1}{2}]$ (if any) in the bin.
- (c) If $B_i = 0$ and $C_i = 1$, we place the largest-sized item whose size is in the range $(\frac{1}{2}, 1]$ (if any) in the bin.

We view the randomization above as splitting the bin into two pieces, dedicating one piece to ‘small’ items (size at most $\frac{1}{2}$) and the other to ‘large’ items (size larger than $\frac{1}{2}$). The dynamics for the small items are as follows. Small items arrive as a Poisson process with rate ρ . At each positive integer time, with probability ρ , a bin of size U_i (uniformly distributed in $[0, \frac{1}{2}]$) becomes available; with probability $(1 - \rho)/2$, a bin of size $\frac{1}{2}$ becomes available; with probability $(1 - \rho)/2$, no bin becomes available. At each time step, a single item (the largest available) is placed in the bin. The dynamics for the large items are essentially identical: the statistics of item and bin arrivals are exactly the same, except that item and bin sizes are larger by an additive factor of $\frac{1}{2}$. For this reason, it suffices to analyze a system involving small items and bins, and multiply the end result by 2 in order to obtain a result for the overall system. From now on, we restrict attention to a system involving only small item and bin sizes.

3.2. Queue dynamics

We will say that t is an *event time* if $t = 0$, or if a bin arrives at time t , or if an item arrives at time t . Note that, for $t > 0$, the probability of two such events occurring simultaneously equals 0 because the arrival process is Poisson and bins arrive at integer times. Thus, we can associate each nonzero event time with a single event (item or bin arrival). Furthermore, because the Poisson process involves a finite number of arrivals during a finite interval, the event times can be ordered in a sequence. For $t > 0$ and $x \in [0, \frac{1}{2}]$, we will say that event $(t, x, +)$ has occurred if an item of size x arrives at time t and that event $(t, x, -)$ has occurred if a bin of size x arrives at time t . Finally, we introduce an event $(0, x, +)$ for every item of size x in queue at time 0.

We now develop equations that describe the evolution of the queue. For any $x \in [0, \frac{1}{2}]$ and $t > 0$, we use $Q_x(t)$ to denote the number of items of size greater than x that are found in the queue just before time t , i.e. before the effects of an item or bin arrival at time t are accounted for. Note that the collection of variables $\{Q_x(t) \mid x \in [0, \frac{1}{2}]\}$ is a complete description of the ‘state’ of the system at time t . For $x = 0$, we use the simpler notation $Q(t)$ in place of $Q_0(t)$. Even though we are interested in the evolution of the queue starting with an empty system, for the purposes of our development here, it will be convenient to allow for the presence of some queued items at time 0.

Suppose that $t > 0$ is the first event time after time 0. Since nothing happens between times 0 and t , we have

$$Q_x(t) = Q_x(0) \quad \text{for all } x \in [0, \frac{1}{2}] \text{ if no event occurred during } (0, t). \tag{3.1}$$

Now suppose that $t > 0$ is an event time, but not the first one. Suppose that the previous event occurred at time $\tau < t$ and that it involved an arrival of an item or of a bin of size y , $y \leq x$. Such an arrival does not affect the number of items in queue of size greater than x ; thus,

$$Q_x(t) = Q_x(\tau) \quad \text{if the previous event was } (\tau, y, +) \text{ or } (\tau, y, -), \text{ where } y \leq x. \tag{3.2}$$

Next suppose that the previous event was the arrival (at time $\tau < t$) of an item of size y , $y > x$. This item is added to the queue and

$$Q_x(t) = Q_x(\tau) + 1 \quad \text{if the previous event was } (\tau, y, +), \text{ where } y > x. \tag{3.3}$$

Finally, suppose that the previous event was the arrival (at time τ) of a bin of size y , $y > x$. This bin cannot be used with an item of size greater than y ; thus, $Q_y(\tau)$ such items are still in queue at time t . However, the number of items of size in the range $(x, y]$, which is $Q_x(\tau) - Q_y(\tau)$ if it is positive, decreases by 1. Therefore,

$$Q_x(t) = Q_y(\tau) + \max\{0, Q_x(\tau) - Q_y(\tau) - 1\}$$

or

$$Q_x(t) = \max\{Q_y(\tau), Q_x(\tau) - 1\} \quad \text{if the previous event was } (\tau, y, -), \text{ where } y > x. \quad (3.4)$$

Equations (3.1)–(3.4) for all $x \in [0, \frac{1}{2}]$ are a complete description of the queue dynamics.

3.3. Discrepancies and their relation to queue sizes

In this subsection we relate queue sizes to a certain ‘discrepancy’ measure. Let us fix a time $t > 0$ and consider the rectangle $[0, t) \times [0, \frac{1}{2}]$. We place a ‘+’ at $(0, z)$ if there is an item of size z in queue at time 0. We place a ‘+’ at (τ, z) if an item of size z arrives at time $\tau > 0$. We place a ‘-’ at (τ, z) if a bin of size z arrives at time $\tau > 0$. Now consider a continuous curve that starts at $(0, \frac{1}{2})$, ends at (t, x) , and consists entirely of horizontal segments (moving to the right) and vertical segments (moving downward); see Figure 1. We use this curve to define a region (subset of the rectangle) as follows. A point (τ, z) in the rectangle is in the region if and only if $0 \leq \tau < t$ and there exists a point (τ, w) on the curve with $w < z$. We note some consequences of this definition. If the initial queue contains an item of size z and if the first segment of the curve is vertical and ends at a point $(0, y)$ with $y < z$, then the point $(0, z)$ associated with this item is included in the region. Points of the form $(\tau, \frac{1}{2})$ (which correspond to bins of size $\frac{1}{2}$) are included only if $\tau < t$ and τ is greater than or equal to the time of the first vertical segment; see Figure 1.

Let \mathcal{A} be a region constructed as above, to be called an *admissible* region. We add the number of plusses in the region and subtract the number of minusses to arrive at a quantity to be denoted by $D(\mathcal{A})$. Let $D_x(t)$ be the maximum possible value of $D(\mathcal{A})$, where the maximum is taken over all admissible regions defined by curves that end at (t, x) . It turns out that $Q_x(t)$ is always equal to $D_x(t)$. For our purposes, we only need to state and prove one direction of this fact. The intuition is the same as for Loynes’ description of the G/G/1 queue in terms of the maximum of a random walk: queues build up only when there is a discrepancy between the arrival and departure processes; hence, small discrepancies imply small queue sizes.




FIGURE 1: Two types of curve and resulting admissible regions. In both cases, the points marked with a ‘Y’ are included in the region, but the points marked with an ‘N’ are not.

Proposition 3.1. *For every $t > 0$ and every $x \in [0, \frac{1}{2}]$, we have $Q_x(t) \leq D_x(t)$.*

Proof. Suppose that $t > 0$ is the first event time after time 0. Consider the curve that starts at $(0, \frac{1}{2})$, moves down to $(0, x)$, and then moves horizontally to (t, x) . For the resulting region \mathcal{A} , the quantity $D(\mathcal{A})$ equals the number of points $(0, z)$ crossed by the vertical segment, which is $Q_x(0)$. Since $D(\mathcal{A}) \leq D_x(t)$, we obtain

$$Q_x(0) \leq D_x(t) \quad \text{for all } x \in [0, \frac{1}{2}] \text{ if no event occurred during } (0, t). \quad (3.5)$$

Now suppose that $t > 0$ is an event time, but not the first one. Suppose that the previous event occurred at time $\tau < t$ and that it involved an arrival of an item or of a bin of size y , $y \leq x$. Consider a curve that ends at (τ, x) such that the resulting region \mathcal{A}' satisfies $D(\mathcal{A}') = D_x(\tau)$. Extend the curve horizontally until (t, x) , forming a new region \mathcal{A} . The event at time τ is outside the region \mathcal{A} , and, therefore, $D_x(t) \geq D(\mathcal{A}) = D(\mathcal{A}') = D_x(\tau)$, leading to

$$D_x(t) \geq D_x(\tau) \quad \text{if the previous event was } (\tau, y, +) \text{ or } (\tau, y, -), \text{ where } y \leq x. \quad (3.6)$$

Next suppose that the previous event was the arrival (at time τ) of an item of size y , $y > x$. Consider a curve that ends at (τ, x) such that the resulting region \mathcal{A}' satisfies $D(\mathcal{A}') = D_x(\tau)$. Extend the curve horizontally until (t, x) . The newly formed region \mathcal{A} includes the event $(\tau, y, +)$, so that $D_x(t) \geq D(\mathcal{A}) = D(\mathcal{A}') + 1 = D_x(\tau) + 1$, leading to

$$D_x(t) \geq D_x(\tau) + 1 \quad \text{if the previous event was } (\tau, y, +), \text{ where } y > x. \quad (3.7)$$

Finally, suppose that the previous event was the arrival (at time τ) of a bin of size y , $y > x$. Consider the following two options for coming up with a curve that ends at (t, x) .

- (i) Form a curve that ends at (τ, y) and which attains $D_y(\tau)$; move horizontally to (t, y) and then down to (t, x) . The resulting region \mathcal{A} satisfies $D(\mathcal{A}) = D_y(\tau)$.
- (ii) Form a curve that ends at (τ, x) and which attains $D_x(\tau)$; move horizontally to (t, x) . The resulting region \mathcal{A}' includes the point $(\tau, y, -)$ and, therefore, satisfies $D(\mathcal{A}') = D_x(\tau) - 1$.

Since $D_x(t) \geq \max\{D(\mathcal{A}), D(\mathcal{A}')\}$, we obtain

$$D_x(t) \geq \max\{D_y(\tau), D_x(\tau) - 1\} \quad \text{if the previous event was } (\tau, y, -), \text{ where } y > x. \quad (3.8)$$

Note that inequalities (3.5)–(3.8) describing the evolution of $D_x(t)$ are satisfied by the variables $Q_x(t)$ with equality (cf. (3.1)–(3.4)). By moving along the sequence of event times, an easy induction shows that $Q_x(t) \leq D_x(t)$ for every (t, x) associated with an event. For any (t, x) that does not correspond to an event, we use one step of the recursive equality or inequality, starting from the last event, to obtain the desired result.

3.4. Proof of the upper bound

For any admissible region \mathcal{A} , let $\Delta(\mathcal{A})$ be the difference between the plusses and minuses in the region (as in the definition of $D(\mathcal{A})$), but without counting points on the top horizontal boundary (corresponding to bins of size exactly $\frac{1}{2}$) and without counting points on the left vertical boundary (corresponding to items in the initial queue). Let $\Delta(t) = \max_{\mathcal{A}} |\Delta(\mathcal{A})|$, where the maximum is taken over all admissible regions \mathcal{A} in the rectangle $[0, t] \times [0, \frac{1}{2}]$.

Let us consider a curve such that the resulting region \mathcal{A} satisfies $D(\mathcal{A}) = D_x(t)$. First suppose that the curve starts with a vertical segment. Then the region \mathcal{A} includes up to $Q_0(0)$ points from the initial queue (they are all plusses), and all items on the top boundary (they are all minusses—let $Y(t)$ be their number). Thus, for this case, $D_x(t) = D(\mathcal{A}) \leq \Delta(\mathcal{A}) + Q_0(0) - Y(t)$. Next suppose that the curve starts with a horizontal segment. Then none of the items in the initial queue contribute to $D(\mathcal{A})$. In this case, $D_x(t)$ equals $\Delta(\mathcal{A})$ minus the contribution of (some) events of the form $(\tau, \frac{1}{2})$, so that $D_x(t) \leq \Delta(\mathcal{A})$. Combining the two cases, and using Proposition 3.1, we obtain

$$Q(t) \leq \max\{\Delta(t), \Delta(t) + Q(0) - Y(t)\}.$$

Using the inequality $\max^2\{a, b\} \leq a^2 + b^2$, we obtain

$$E[Q(t)^2] \leq E[\Delta^2(t)] + E[Q^2(0)] + E[\Delta^2(t)] + E[Y^2(t)] + 2E[Q(0)]E[\Delta(t) - Y(t)], \tag{3.9}$$

where we used the independence of $Q(0)$ from the future of the process and omitted the negative term $-2E[\Delta(t)Y(t)]$. In order to use this inequality, we need some information on the moments of $\Delta(t)$. This is provided by the following lemma, based on results from [11]. So as not to disrupt continuity, the proof of Lemma 3.1 is deferred to Section 3.5.

Lemma 3.1. *Suppose that $\rho \geq \frac{1}{2}$. There exist positive constants c_1 and c_2 such that, for all $\rho \geq \frac{1}{2}$ and all $t \geq 2/(1 - \rho)$, we have*

$$E[\Delta(t)] \leq c_1\sqrt{t} \log^{3/4} t, \quad E[\Delta^2(t)] \leq c_2 t \log^{3/2} t.$$

In order to bound the moments of $Y(t)$, recall that bins of size exactly $\frac{1}{2}$ are produced at each time step with probability $(1 - \rho)/2$. Thus, $Y(t)$ is the number of successes in t independent Bernoulli trials with the above success probability (we are assuming here that t is an integer). Therefore, $E[Y(t)] = (1 - \rho)t/2$, which shows that $E[Y(t)]$ increases (with t) faster than $E[\Delta(t)]$. Furthermore, $\text{var}(Y(t)) \leq (1 - \rho)t/2$ and

$$E[Y^2(t)] \leq \frac{(1 - \rho)^2 t^2}{4} + \frac{(1 - \rho)t}{2} \leq (1 - \rho)^2 t^2.$$

Let $t^* = \alpha h^2 \log^{3/2} h$, where α is a positive constant and $h = 1/(1 - \rho)$. By choosing α large enough, and using some straightforward algebra, it follows, from (3.9), Lemma 3.1, and the bounds for $E[Y(t)]$ and $E[Y^2(t)]$, that there exist positive constants γ_1 and γ_2 such that

$$E[\Delta(t^*) - Y(t^*)] \leq -\gamma_1 h \log^{3/2} h, \quad E[\Delta^2(t^*)] + E[Y^2(t^*)] \leq \gamma_2 h^2 \log^3 h. \tag{3.10}$$

From (3.9), with time 0 replaced by kt^* , and (3.10), we obtain, for $k \geq 0$,

$$E[Q^2((k + 1)t^*)] \leq E[Q^2(kt^*)] + 2\gamma_2 h^2 \log^3 h - 2\gamma_1 E[Q(kt^*)]h \log^{3/2} h. \tag{3.11}$$

For the policy under consideration, it is not hard to see that $E[Q(t) \mid Q(0) = 0]$ is non-decreasing with t . Thus, the limit $q = \lim_{t \rightarrow \infty} E[Q(t) \mid Q(0) = 0]$ exists. It then follows from (3.11) that $2\gamma_1 q h \log^{3/2} h \leq 2\gamma_2 h^2 \log^3 h$, which implies that $q \leq (\gamma_2/\gamma_1)h \log^{3/2} h$. This completes the proof of the upper bound.

3.5. Proof of Lemma 3.1

Let us fix some time t and some admissible region $\mathcal{A} \subset (0, t) \times [0, \frac{1}{2})$. Let $N_+(\mathcal{A})$ and $N_-(\mathcal{A})$ be the number of plusses and minusses, respectively, inside \mathcal{A} . Note that $E[N_+(\mathcal{A})] = 2\rho\lambda(\mathcal{A})$, where $\lambda(\mathcal{A})$ is the area of \mathcal{A} . Also, note that $\Delta(\mathcal{A}) = N_+(\mathcal{A}) - N_-(\mathcal{A})$ with probability 1 (this is because the events $U_i = \frac{1}{2}$ or $X_i = \frac{1}{2}$ for some i have zero probability). We decompose $\Delta(\mathcal{A})$ into two terms, involving the discrepancies in the item and bin arrival processes.

$$\begin{aligned} \Delta(t) &= \max_{\mathcal{A}} |\Delta(\mathcal{A})| \\ &\leq \max_{\mathcal{A}} |N_+(\mathcal{A}) - E[N_+(\mathcal{A})]| + \max_{\mathcal{A}} |N_-(\mathcal{A}) - E[N_+(\mathcal{A})]|. \end{aligned} \tag{3.12}$$

In order to bound $E[\Delta^2(t)]$, it suffices to obtain a bound on the mean-square value of each one of the two terms on the right-hand side of (3.12).

3.5.1. *Item arrival discrepancies.* We consider here the first term in (3.12). Let N_+ be the total number of points in the rectangle $(0, t) \times [0, \frac{1}{2})$, which is a Poisson random variable with mean ρt . We will be using the following decomposition:

$$\begin{aligned} \max_{\mathcal{A}} |N_+(\mathcal{A}) - E[N_+(\mathcal{A})]| \\ \leq \max_{\mathcal{A}} |N_+(\mathcal{A}) - E[N_+(\mathcal{A}) \mid N_+]| + \max_{\mathcal{A}} |E[N_+(\mathcal{A}) \mid N_+] - E[N_+(\mathcal{A})]|. \end{aligned}$$

We will bound separately the mean square of the two terms on the right-hand side above.

For the second term, we note that

$$E[N_+(\mathcal{A}) \mid N_+] = \frac{\lambda(\mathcal{A})}{t/2} N_+.$$

Therefore,

$$E[N_+(\mathcal{A}) \mid N_+] - E[N_+(\mathcal{A})] = \frac{2\lambda(\mathcal{A})}{t} (N_+ - \rho t).$$

To maximize over all admissible regions \mathcal{A} , we let \mathcal{A} be the entire rectangle whose area is $t/2$. It follows that

$$E\left[\max_{\mathcal{A}} (E[N_+(\mathcal{A}) \mid N_+] - E[N_+(\mathcal{A})])^2\right] = \text{var}(N_+) = \rho t.$$

We now deal with the first term. Let B be the event $|N_+ - \rho t| \geq \sqrt{t} \log^{3/4} t$, and let B^c be its complement. Let $Z = \max_{\mathcal{A}} |N_+(\mathcal{A}) - E[N_+(\mathcal{A}) \mid N_+]|$. We have

$$E[Z^2] = E[Z^2 \mid B]P(B) + E[Z^2 \mid B^c]P(B^c).$$

We have $P(B^c) \leq 1$. We recall the following well-known tail bound for Poisson and binomial distributions (see [7, pp. 27–35] and [13, p. 72]).

Lemma 3.2. *Let X be a Poisson random variable with mean λ . Then,*

$$P(X > \lambda + \delta) \leq \exp\left\{-\frac{\delta^2}{4\lambda}\right\} \text{ for all } \delta \in (0, \lambda/2).$$

Similarly, let Y be a binomial random variable with parameters m and p , so that its mean is mp . Then,

$$P(|Y - mp| \geq \delta) \leq 2 \exp\left\{-\frac{\delta^2}{4mp}\right\} \text{ for all } \delta \in (0, mp).$$

We apply Lemma 3.2 to the Poisson random variable N_+ and obtain

$$P(B) \leq c_4 \exp\{-c_5 \log^{3/2} t\} \leq c_6 t^{-2}$$

for some constants c_4, c_5 , and c_6 independent of t . Let $B_1 = \{N_+ \leq \rho t - \sqrt{t} \log^{3/4} t\}$ and $B_2 = \{N_+ \geq \rho t + \sqrt{t} \log^{3/4} t\}$, so that $B = B_1 \cup B_2$. In order to bound the term $E[Z^2 | B]$, we will use the fact that $Z \leq N_+$ and the easily derived inequalities $E[N_+^2 | N_+ \geq a] \leq E[(N_+ + a)^2] \leq 2E[N_+^2] + 2a^2$. We have

$$\begin{aligned} E[Z^2 | B] &\leq E[N_+^2 | B] \\ &\leq \max\{E[N_+^2 | B_1], E[N_+^2 | B_2]\} \\ &= E[N_+^2 | B_2] \\ &= E[N_+^2 | N_+ \geq \rho t + \sqrt{t} \log^{3/4} t] \\ &\leq E[(N_+ + \rho t + \sqrt{t} \log^{3/4} t)^2] \\ &\leq c_7 t^2 \end{aligned}$$

for some constant c_7 independent of t . Therefore, $E[Z^2 | B]P(B)$ is bounded by a constant.

We finally deal with the term $E[Z^2 | B^c]$. Given any particular value n of N_+ (in the range allowed by the event B^c), the n points are independent and uniformly distributed. We will apply the following result of [11].

Lemma 3.3. ([11, Theorem 3].) *Consider a set of n points uniformly and independently distributed in the $\sqrt{n} \times \sqrt{n}$ square. Let \mathcal{R} be the set of all simply connected subsets of the square. There exist constants c and n_0 (independent of n) such that*

$$\begin{aligned} P(\text{there exists } R \in \mathcal{R} : |N(R) - E[N(R)]| > cp(R) \log^{3/4} n + c \log^{3/2} n) \\ \leq \frac{1}{n\sqrt{\log n}} \text{ for all } n \geq n_0, \end{aligned}$$

where $p(r)$ is the perimeter of R and $N(R)$ is the number of points in R .

Note that our model involves the rectangle $(0, t) \times [0, \frac{1}{2})$, whereas Lemma 3.3 involves the rectangle $[0, \sqrt{n}] \times [0, \sqrt{n}]$. This is not an issue because we can just rescale the dimensions of our rectangle. With this rescaling, an admissible region $\mathcal{A} \subset (0, t) \times [0, \frac{1}{2})$ is mapped to a region $\mathcal{A}' \subset [0, \sqrt{n}] \times [0, \sqrt{n}]$ with $N(\mathcal{A}') = N_+(\mathcal{A})$. Because we are interested in only admissible regions \mathcal{A} , it is easily seen that $p(\mathcal{A}') \leq 4\sqrt{n}$. Thus, ‘with high probability’, Z is bounded above by $c_8 \sqrt{n} \log^{3/4} n$. With the remaining (small) probability, it is bounded above by n , which gives negligible contribution to the second moment of Z . The values of n allowed by the event B^c are of order $\Theta(t)$. Putting everything together, we conclude that $E[Z^2 | B] \leq c_9 t \log^{3/2} t$ for some constant c_9 . In particular, the mean square of the first term on the left-hand side of (3.12) is $O(t \log^{3/2} t)$.

3.5.2. *Bin arrival discrepancies.* We now turn to the second term on the right-hand side of (3.12). If bins of size less than $\frac{1}{2}$ were arriving as a Poisson process, the argument would be identical to the one for the case of item arrivals. However, in our case, the process of small bin arrivals is different, and some additional work is needed. For simplicity, we will be assuming for the rest of this subsection that t is an integer.

The *actual* bin arrival process can be described as follows. At each integer time k , $0 < k < t$, a potential bin arrives whose size U_k is uniform in $[0, \frac{1}{2})$. An independent coin with success probability ρ is flipped, and the bin is retained if the coin flip is a success. Thus, the total number of bin arrivals, N_- , is a binomial random variable with parameters $t - 1$ and ρ .

Let us now consider an *alternative* bin arrival process, coupled with the actual one. We generate $t - 1$ arrival times, independently and uniformly in $(0, t)$, sort them in increasing order, and let \hat{T}_k be the time of the k th arrival. We let the k th arriving bin in the sorted sequence have the same size U_k as in the above description of the actual process. Finally, we use the same coins as in the actual bin arrival process to thin out the arrivals. We will use a ‘hat’ to indicate quantities associated with the alternative bin arrival process. Let $C = \max_k |k - \hat{T}_k|$. The points in $(0, t) \times [0, \frac{1}{2})$ corresponding to the alternative process are the same as the points corresponding to the actual process, except that each point is shifted horizontally by a random amount bounded by C . Using Lemma 3.3 (restricted to sets R of the form $(0, s) \times [0, \frac{1}{2})$), it is easily shown that $E[C^2] \leq c_{10} t \log^{3/2} t$ for some absolute constant c_{10} .

For any admissible region $\hat{\mathcal{A}}$, let $\hat{N}_-(\hat{\mathcal{A}})$ be the number of points in $\hat{\mathcal{A}}$ obtained from the alternative process. Given an admissible region \mathcal{A} , we can shift horizontally the vertical segments of the curve that defines it, to obtain a new admissible region $\hat{\mathcal{A}}$ such that the k th point in the actual process is in \mathcal{A} if and only if the k th point in the alternative process is in $\hat{\mathcal{A}}$; in particular, $\hat{N}_-(\hat{\mathcal{A}}) = N_-(\mathcal{A})$. Furthermore, the shifts in the vertical segments are bounded by C . It follows that $|\lambda(\hat{\mathcal{A}}) - \lambda(\mathcal{A})| \leq C$.

Using also the facts that $E[N_+(\mathcal{A})] = 2\rho\lambda(\mathcal{A})$ and $E[\hat{N}_-(\hat{\mathcal{A}})] = 2\rho\lambda(\hat{\mathcal{A}})$, we find that, for every admissible \mathcal{A} , there exists an admissible $\hat{\mathcal{A}}$ such that

$$\begin{aligned} |N_-(\mathcal{A}) - E[N_+(\mathcal{A})]| &= |\hat{N}_-(\hat{\mathcal{A}}) - E[N_+(\mathcal{A})]| \\ &\leq |\hat{N}_-(\hat{\mathcal{A}}) - 2\rho\lambda(\hat{\mathcal{A}})| + 2\rho|\lambda(\mathcal{A}) - \lambda(\hat{\mathcal{A}})| \\ &\leq |\hat{N}_-(\hat{\mathcal{A}}) - E[\hat{N}_-(\hat{\mathcal{A}})]| + C. \end{aligned}$$

It follows that

$$\max_{\mathcal{A}} |N_-(\mathcal{A}) - E[N_+(\mathcal{A})]| \leq \max_{\hat{\mathcal{A}}} |\hat{N}_-(\hat{\mathcal{A}}) - E[\hat{N}_-(\hat{\mathcal{A}})]| + C.$$

To obtain a bound for the mean square of $\max_{\mathcal{A}} |\hat{N}_-(\hat{\mathcal{A}}) - E[\hat{N}_-(\hat{\mathcal{A}})]|$, we proceed exactly as in the case of item arrival discrepancies, where we obtained a bound on the mean square of $\max_{\mathcal{A}} |\hat{N}_+(\mathcal{A}) - E[\hat{N}_+(\mathcal{A})]|$. The only difference is that with item arrivals we were dealing with a Poisson number of points uniformly distributed in the rectangle, whereas here we are dealing with a binomial number of points. However, using a binomial tail bound analogous to the one in Lemma 3.2, the proof goes through without change.

3.6. Discussion and extensions

The upper bound for the policy we have considered remains valid for several variations in our assumptions on the item and bin arrival processes. Our results were developed for the case of item sizes uniformly distributed in $[0, 1]$, Poisson item arrivals, and deterministic bin arrivals.

The Poisson item arrival assumption can be relaxed to an assumption of i.i.d. interarrival times with finite second moment. For example, for an upper bound, we can delay the item arrivals by an $O(h)$ amount and introduce $O(1 - \rho)$ additional artificial arrivals, so that the new arrival process is Poisson with rate $(1 + \rho)/2$, and then apply the policy we have developed. Similarly, the assumption that bin arrivals are deterministic can be relaxed. However, the assumption on the item-size distribution is more restrictive. We believe that our analysis extends to any item-size distribution that is symmetric around $\frac{1}{2}$. Moving beyond this simple extension appears to be an interesting future research problem.

We make a brief remark about our policy. The policy we have analyzed uses randomization. As an alternative, consider the following deterministic policy. Assume, for simplicity, that $1/(1 - \rho)$ is an integer, and let $\delta = (1 - \rho)/4\rho$. Instead of letting the sizes of the bins allocated to the small items be random variables, we let them cycle through the values $0, \delta, 2\delta, \dots, \frac{1}{2}, \frac{1}{2}$ (the last value, $\frac{1}{2}$, is repeated twice). The bin arrival discrepancies are now easier to bound (everything is deterministic). While the proof for this deterministic policy is somewhat simpler, we find the randomized policy to be more elegant.

4. Lower bound

As discussed in the introduction, if we relax the problem and allow items to be broken into pieces that can be placed in different bins, we obtain a tractable queueing system, resulting in a $\Theta(h)$ queue size. This readily yields an $\Omega(h)$ lower bound for the dynamic bin-packing problem. This lower bound does not match the upper bound $O(h \log^{3/2} h)$, which leaves the question of whether a better policy might result in a $\Theta(h)$ queue size. The $\Omega(h \log h)$ lower bound to be proved in this section shows that this is not the case, although it still leaves a small gap from the upper bound. The proof below is inspired from the proof of a lower bound for online bin packing [19].

Given ρ , let us define

$$t^* = \frac{\alpha}{(1 - \rho)^2} \log \frac{1}{1 - \rho},$$

where α is an absolute constant. Consider the items that arrive during the interval $[0, t^*]$. Let B be the number of such items whose size belongs to $(\frac{1}{2}, 1)$. Let N_- be the number of such items whose size belongs to $(\frac{1}{3}, \frac{1}{2})$. Let N_+ be the number of such items whose size belongs to $(\frac{1}{2}, \frac{2}{3})$. Note that $E[N_-] = E[N_+] = 2\rho t^*/6 = \rho t^*/3$. Let

$$n = \frac{\rho t^*}{3} - 2\sqrt{t^*},$$

assumed for simplicity to be an integer. Let G be the event $\{N_+ \geq n, N_- \geq n\}$. Note that, for ρ sufficiently close to 1, t^* is sufficiently large and the Chebychev inequality yields $P(G) \geq \frac{1}{2}$.

Conditioned on the event G having occurred, let us select at random n of the N_- arriving items of size in $(\frac{1}{3}, \frac{1}{2})$ and n of the N_+ arriving items of size in $(\frac{1}{2}, \frac{2}{3})$, and call them small and large special items, respectively. A probabilistically equivalent way of generating the small special items is as follows (the process is similar for large special items). We generate the value of the Poisson random variable N_- and then generate n independent random points, uniformly distributed in the rectangle $[0, t^*] \times (\frac{1}{3}, \frac{1}{2})$.

Let us now fix a policy, assume that the event G has occurred, and let $Q(t)$ be the number of items in queue at some time t under that policy. For any small special item i_- , there are two possibilities: (a) item i_- is eventually placed in a bin (either before or after t^*) together

with a large special item i_+ (plus possibly other items), in which case we say that i_- and i_+ are ‘matched’; (b) the bin to which item i_- is eventually placed does not contain a large special item. In case (b) we say that item i_- is ‘unmatched’. Let K be the number of unmatched small special items.

For $i = 1, \dots, n$, let t_i^- and t_i^+ be the arrival times of the pair of items i_- and i_+ that are matched together. Clearly, one of these two items will have to be in queue for at least $|t_i^- - t_i^+|$ time units. Therefore,

$$\int_0^{t^*} Q(t) dt \geq \sum_{i=1}^n |t_i^- - t_i^+|. \tag{4.1}$$

Another useful bound is provided by the next lemma.

Lemma 4.1. *We have*

$$E[Q(t^*) \mid G] \geq n + \frac{2}{3}\rho t^* + \frac{E[K \mid G]}{2} - t^*.$$

Proof. Let us consider the number of bins that the policy uses for the items that arrive in $[0, t^*]$. There are n large special items, and they have to be in separate bins. Furthermore, the expected number of items whose size is in $[\frac{2}{3}, 1]$ is $2\rho t^*/3$, and these must also be in separate bins. Finally, any unmatched small special item has size at least $\frac{1}{3}$ and cannot be together with an item of size in $[\frac{2}{3}, 1]$; also, it cannot be together with a large special item, by the definition of unmatched items. Since at most two small special items can fit in the same bin, at least $K/2$ bins will be required. Thus, an expected number of $n + 2\rho t^*/3 + E[K \mid G]/2$ bins will be required, but only t^* bins are available until time t^* . The difference provides a lower bound on the expected queue size at time t^* .

Combining (4.1) and Lemma 4.1, and using the definition of n , we obtain

$$\frac{1}{t^*} \int_0^{t^*} E[Q(t) \mid G] dt + E[Q(t^*) \mid G] \geq c E \left[K + \frac{1}{t^*} \sum_{i=1}^n |t_i^- - t_i^+| \mid G \right] - (1 - \rho)t^* - 2\sqrt{t^*}, \tag{4.2}$$

where c is a positive absolute constant. (The above argument involves an interchange of integration and expectation, which is justified by the Fubini theorem.)

The quantity inside the expectation on the right-hand side of (4.2) can be interpreted in terms of a certain matching problem, which we now describe. We have n i.i.d. points uniformly distributed in $[0, t] \times (\frac{1}{2}, \frac{2}{3})$, marked by a ‘+’, and n i.i.d. points uniformly distributed in $[0, t] \times (\frac{1}{3}, \frac{1}{2})$, marked by a ‘-’, corresponding to the large and small special items, respectively. Now, let us ‘fold’ the rectangle $[0, t] \times (\frac{1}{3}, \frac{2}{3})$ along the $x = \frac{1}{2}$ line, so that a ‘+’ point at (τ, x) gets moved to (τ, x') , where $x' = 1 - x$. As a result, we now have n ‘+’ points and n ‘-’ points, all uniformly and independently distributed in $[0, t] \times (\frac{1}{3}, \frac{1}{2})$. A ‘+’ point at (τ, x') can be matched together with a ‘-’ point at (s, x) (i.e. be placed in the same bin) only if $x' \geq x$ (i.e. if the ‘+’ point lies above the ‘-’ point). For such a pair of matched points, there is a penalty equal to their horizontal distance, $|s - \tau|$. For any unmatched ‘-’ point, there is a penalty of 1. Lemma 1 of [19] (building on the results of [1]) shows that the expected value of the optimal cost in such a matching problem on the unit rectangle is $\Omega(\sqrt{n \log n})$. (Actually, for an unmatched ‘-’ point at (s, x) , [19] used a penalty equal to $1 - x$; our unit penalty is at least as large, so the lower bound still applies.) Our rectangle has height $\frac{1}{3}$ instead of 1, but this does not matter. Also, our rectangle has length t^* (instead of 1), but this is taken care of

by our dividing by t^* in (4.2). Using also the facts that $P(G) \geq \frac{1}{2}$ and $n = \Theta(t^*)$, and taking expectations of both sides, we obtain

$$\frac{1}{t^*} \int_0^{t^*} E[Q(t)] dt + E[Q(t^*)] \geq c_1 \sqrt{t^* \log t^*} - (1 - \rho)t^* - 2\sqrt{t^*}. \quad (4.3)$$

Recalling the definition of t^* , some elementary algebra shows that if α is chosen suitably small, and if ρ is sufficiently close to 1, the right-hand side of (4.3) is lower bounded by $ch \log h$ (c is some absolute constant).

Now consider the same length time interval $[s, s + t^*]$ instead of the time interval $[0, t^*]$. What is different now is that there may be some items already in queue at the beginning of the interval $[s, s + t^*]$. However, the presence of such items cannot reduce the minimum possible value of the left-hand side of (4.3). It follows that

$$\frac{1}{2t^*} \int_s^{s+t^*} E[Q(t)] dt + \frac{1}{2} E[Q(t^*)] \geq ch \log h \quad \text{for all } t \geq 0.$$

This implies that $\limsup_{t \rightarrow \infty} E[Q(t)] \geq ch \log h$, concluding the proof of the lower bound in Theorem 1.1.

5. Conclusion

In this paper we have studied a dynamic version of the bin-packing problem that involves queueing (as opposed to the previously studied online problem, in which decisions need to be made as soon as an item arrives). Whereas some past work on our problem has addressed the stability question, we focused on the scaling of the expected queue size as the load factor ρ approaches its stability limit. (By Little's law, this also addresses the scaling of the expected delay.) We showed that, as ρ approaches 1, there exists a policy under which the expected queue size scaling is very close to (within a logarithmic factor of) the $\Theta(1/(1 - \rho))$ scaling associated with an M/U/1 queueing system in which items can be broken into pieces that can be placed into different bins. While the logarithmic factor may be aesthetically unappealing, we also showed that this is unavoidable. Our upper and lower bounds are not tight as far as the logarithmic factor is concerned ($\log^{3/2} h$ versus $\log h$), but a method for closing this gap is not apparent.

Our results were developed for the case of item sizes uniformly distributed in $[0, 1]$, Poisson item arrivals, and deterministic bin arrivals. The Poisson item arrival assumption can be relaxed to an assumption of i.i.d. interarrival times with finite second moment. For example, for an upper bound, we can delay the item arrivals by an $O(h)$ amount and introduce $O(1 - \rho)$ additional artificial arrivals, so that the new arrival process is Poisson with rate $(1 + \rho)/2$, and then apply the policy we have developed. Similarly, the assumption that bin arrivals are deterministic can be relaxed. However, the assumption on the item-size distribution is more restrictive. We believe that our analysis extends to any item-size distribution that is symmetric around $\frac{1}{2}$. Moving beyond this simple extension appears to be an interesting future research problem.

Acknowledgements

The second author would like to acknowledge discussions with Dimitris Bertsimas on this problem, about fifteen years ago. This research was partially supported by NSF grants ECS-0426453 and CNS-0546590.

References

- [1] AJTAI, M., KOMÓS, J. AND TUSNÁDY, G. (1983). On optimal matchings. *Combinatorica* **4**, 259–264.
- [2] BENTLEY, J. L. *et al.* (1984). Some unexpected expected behavior results for bin packing. In *Proc. 16th Annual ACM Symp. Theory Comput.*, ACM, New York, pp. 279–288.
- [3] BERTSIMAS, D. J. AND VAN RYZIN, G. (1993). Stochastic and dynamic vehicle routing in the Euclidean plane with multiple-capacitated vehicles. *Operat. Res.* **41**, 60–76.
- [4] BERTSIMAS, D. J. AND VAN RYZIN, G. (1993). Stochastic and dynamic vehicle routing with general demand and interarrival time distributions. *Adv. Appl. Prob.* **20**, 947–978.
- [5] COFFMAN, E. G. AND STOLYAR, A. L. (2001). Bandwidth packing. *Algorithmica* **29**, 70–88.
- [6] COFFMAN, E. G., JR., GALAMBOS, G., MARTELLO, S. AND VIGO, D. (1998). Bin packing approximation algorithms: combinatorial analysis. In *Handbook of Combinatorial Optimization*, eds D.-Z. Du and P. M. Pardalos, Kluwer, Boston.
- [7] DEMBO, A. AND ZEITOUNI, O. (1998) *Large Deviations Techniques and Applications*, 2nd edn. Springer, New York.
- [8] GAMARNIK, D. (2004). Stochastic bandwidth packing process: stability conditions via Lyapunov function technique. *Queueing Systems* **48**, 339–363.
- [9] KARP, R. M., LUBY, M. AND MARCHETTI-SPACCAMELA, A. (1984). Probabilistic analysis of multidimensional bin packing problems. In *Proc. 16th Annual ACM Symp. Theory Comput.*, ACM, New York, pp. 289–298.
- [10] KNODEL, W. (1981). A bin-packing algorithm with complexity $O(n \log n)$ and performance 1 in the stochastic limit. In *Proc. Math. Foundations Comput. Sci.* Springer, London, pp. 369–378.
- [11] LEIGHTON, T. AND SHOR, P. (1986). Tight bounds for minimax grid matching, with applications to the average case analysis of algorithms. In *Proc. 18th Annual ACM Symp. Theory Comput.*, ACM, New York, pp. 91–103.
- [12] LUEKER, G. S. (1983). Bin packing with items uniformly distributed over intervals $[a, b]$. In *Proc. 24th Annual Symp. Foundations Comput. Sci.*, IEEE Computer Society, Washington, DC, pp. 289–297.
- [13] MOTWANI, R. AND RAGHAVAN, P. (1995). *Randomized Algorithms*. Cambridge University Press.
- [14] RHEE, W. T. (1988). Optimal bin-packing with items of random sizes. *Math. Operat. Res.* **13**, 140–151.
- [15] RHEE, W. T. (1990). A note on optimal bin packing and optimal bin covering with items of random sizes. *SIAM J. Comput.* **19**, 705–710.
- [16] RHEE, W. T. AND TALAGRAND, M. (1989). Optimal bin packing with items of random sizes. II. *SIAM J. Comput.* **18**, 139–151.
- [17] RHEE, W. T. AND TALAGRAND, M. (1989). Optimal bin packing with items of random sizes. III. *SIAM J. Comput.* **18**, 473–486.
- [18] RHEE, W. T. (1993). Optimal bin packing of items of size uniformly distributed over $[0, 1]$. *Math. Operat. Res.* **18**, 694–704.
- [19] SHOR, P. W. (1986). The average case analysis of some on-line algorithms for bin packing. *Combinatorica* **6**, 179–200.
- [20] VEMPALA, S. AND VOCKING, B. (1999). Approximating multicast congestion. In *Proc. 10th Int. Symp. Algorithms and Computation*, Springer, London, pp. 367–372.
- [21] WOLFF, R. W. (1989). *Stochastic Modeling and the Theory of Queues*. Prentice-Hall, Englewood Cliffs, NJ.