

The Power of Choice in Priority Scheduling

Dan Alistarh
IST Austria & ETH Zurich

dan.alistarh@inf.ethz.ch

Jerry Li
Massachusetts Institute of Technology
jerryzli@mit.edu

Justin Kopinsky
Massachusetts Institute of Technology
jkopin@mit.edu

Giorgi Nadiradze
ETH Zurich
giorgi.nadiradze@inf.ethz.ch

August 22, 2018

Abstract

Consider the following random process: we are given n queues, into which elements of increasing labels are inserted uniformly at random. To remove an element, we pick two queues at random, and remove the element of lower label (higher priority) among the two. The *cost* of a removal is the *rank* of the label removed, among labels still present in any of the queues, that is, the distance from the optimal choice at each step. Variants of this strategy are prevalent in state-of-the-art concurrent priority queue implementations. Nonetheless, it is not known whether such implementations provide any rank guarantees, even in a sequential model.

We answer this question, showing that this strategy provides surprisingly strong guarantees: Although the single-choice process, where we always insert and remove from a single randomly chosen queue, has degrading cost, going to infinity as we increase the number of steps, in the two choice process, the expected rank of a removed element is $O(n)$ while the expected worst-case cost is $O(n \log n)$. These bounds are tight, and hold irrespective of the number of steps for which we run the process. The argument is based on a new technical connection between “heavily loaded” balls-into-bins processes and priority scheduling. Our analytic results inspire a new concurrent priority queue implementation, which improves upon the state of the art in terms of practical performance.

1 Introduction

The last two decades have seen tremendous progress in the area of concurrent data structures, to the point where scalable versions of many classic data structures are now known, including lists [4, 16], skip-lists [13, 18, 21], hash-tables [19, 22, 25, 34], or search trees [1, 9, 23]. In this context, significant research attention has been given to concurrent *producer-consumer* data structures, such as queues, stacks, or priority queues, which are critical in concurrent systems. On the one hand, several elegant designs, e.g. [11, 12, 27, 28], have pushed these data structures to the limits of scalability on current processor architectures. On the other, impossibility results [2, 10] show that such data structures providing strong ordering semantics also have inherent sequential bottlenecks which limit their scalability.

A popular approach for circumventing these impossibility results is to *relax* data structure semantics [35]. The first instance of this strategy is probably the relaxed PRAM priority queue structure of Karp and Zhang [20], designed for efficient parallel execution of branch-and-bound programs. A flurry of subsequent work, e.g. [3, 6, 15, 24, 29, 32, 33, 36, 38], has shown that relaxed semantics can yield extremely good practical results, and that it can be very useful in the context of higher-level constructs, such as concurrent graph algorithms [29].¹

Notably, the current state-of-the-art concurrent priority queue is obtained by the following *MultiQueue* strategy [14, 32]. We start from n queues, each protected by a lock. To *insert* an element, a processor picks one of the n queues uniformly at random, locks it, and inserts into it. To *remove* an element, the processor picks *two queues* uniformly at

¹In practice, it is often possible to offset the cost of *priority inversions* by performing *additional work*. For instance, in Dijkstra’s single-source shortest-paths algorithm, priority inversions can be offset by nodes being relaxed multiple times.

random, locks them, and removes the element of *higher priority* among their two top elements. Extensive testing [14, 32] has shown that this natural strategy can provide state-of-the-art throughput, and that the average number of priority inversions is relatively low. Further variants of this strategy are used to implement general priority schedulers [29] and relaxed concurrent queues [15].

Despite their significant practical success, randomized relaxed concurrent data structures still lack a theoretical foundation: it is not known whether distributed strategies such as the MultiQueue can provide any *guarantees* in terms of their inversion cost, or whether their performance can still be improved.

Contribution. In this paper, we take a first step towards addressing this challenge. We focus on the analysis of the following sequential process, inspired by the MultiQueue strategy: we are given n queues, into which a large number of consecutively labeled elements are inserted initially, uniformly at random. To remove an element, we pick two queues at random, and remove the element of lower label on top of either queue. We are interested in the *cost* of removals, defined as the *rank* of the removed label among labels still present in any of the queues, i.e., the “distance” from the optimal choice at each step, which would be the top element among all queues.

Our main technical contribution is showing that this process provides surprisingly strong rank guarantees: for any time $t \geq 0$ in the execution of the process, the expected rank of an element removed at time t is $O(n)$, while the expected worst-case rank removed at a step is $O(n \log n)$. These bounds are asymptotically tight, and hold for arbitrarily large t . Our analysis generalizes to a $(1 + \beta)$ extension of the process, where the algorithm deletes from the higher priority among two random queues with probability $0 < \beta < 1$, and from a single randomly chosen queue with probability $(1 - \beta)$. It also extends to show that the process is *robust to bias* in the insertion distribution towards some bins by a constant factor $\gamma \in (0, 1)$. By contrast, we show that the strategy which always removes from a *single* randomly chosen queue *diverges*, in the sense that its average rank guarantee evolves as $\Omega(\sqrt{tn \log n})$, for time $t \geq n \log n$. The $(1 + \beta)$ strategy can be extended to an efficient implementation, which improves upon the state-of-the-art in terms of practical performance.

Technical Overview. A tempting first approach at analysis is to reduce to classic “power of two choices” processes, e.g. [5, 30], in which balls are always inserted into the least loaded of two randomly chosen bins. A simple reduction between these two processes exists for the case where labels are inserted into the queues in *round-robin* fashion (please see Section A for details). However, this reduction breaks if elements are inserted uniformly at random, as is the case in real systems. Another important difference from the classic process is that elements are *labelled*, and so the state of the system at a given step is highly correlated with previous steps. For instance, given a queue Q whose top element has label ℓ , if we wish to characterize the probability that the next label on top of Q is a specific value $\ell' > \ell$, we need to know whether the history of elements examined by the algorithm (in other queues) contains element ℓ' . Such correlations make a standard step-by-step analysis extremely challenging.

We circumvent these issues by relating the original labelled process to a continuous *exponential process*, which reduces correlations by replacing integer labels with real-valued labels. In this process, on every “insertion,” we insert a new label whose value equals that of the latest inserted label in that bin (or 0 for the first insertion), plus an *exponential random variable* of mean n . Intuitively, we wish to simulate the integer label insertion process with the same mean n , while removing correlations but using continuous label values to prevent label clashes and reduce correlations. Once this insertion process has ended, we replace each (real-valued) label with its *rank* among all (real-valued) labels.

We relate these two processes by the following key claim: *the rank distributions for the original process and for the exponential process are identical*. More precisely, for any queue i and rank j , the event that, after insertion, the label with rank j is located in queue i has the same probability in the original process and in the exponential process (and all such events are independent in both processes).

Given this fact, we can couple the two processes as follows. We first couple the insertion steps such that each bin contains the same sequence of ranks in both processes. We then couple the removal steps so that they make exactly the same sequence of random choices. Under this coupling, the two processes will pay the same rank cost. Thus, it will suffice to characterize the rank cost of the exponential process under our removal scheme.

This is achieved via two steps. The first is a potential argument which carefully characterizes the average value of labels on top of the queues, and the maximum deviation from the average by a queue, at any time $t > 0$. This part of the proof is very technical, and generalizes an analytic approach developed for weighted balls-into-bins processes by Peres, Talwar and Wieder [30]. Specifically, if we let $x_i(t)$ be the difference between the label on top of queue i at time t and the mean top label across all queues, then the *total potential* of the system at time t is defined as

$$\sum_{i=1}^n [\exp(\alpha x_i/n) + \exp(-\alpha x_i/n)].$$

The key claim in the argument is showing that the expected value of this potential is bounded by $O(n)$ for any step t . This is done by a careful technical analysis of the expected change in potential at a step, which proves that the potential has supermartingale-like behavior: that is, it tends to decrease in expectation once it surpasses the $O(n)$ threshold. It is interesting to note that neither of the two exponential factors in the definition of the potential satisfies this bounded increase property in isolation, yet their sum can be bounded in this way.

The $O(n)$ bound on the total potential provides a strong handle on the maximum deviation from the mean in the exponential process. The last step of our argument builds on this characterization to show that, since label values cannot stray too far either above or below the mean, the *ranks* of these values among all values on top of queues cannot be too large. In particular, the average rank of elements on top of queues must be $O(n)$, and the maximum rank is $O(n \log n)$. The rank equivalence theorem implies our main claim.

The analytic framework we sketched above is quite general. It can be extended to showing that this implementation strategy is robust to insertion bias, i.e. that it guarantees similar rank bounds even if the insertion process is biased (within some constant $\gamma \ll 1$) towards some of the bins. Further, it also applies to the $(1 + \beta)$ variant of the process, where each removal considers two randomly chosen queues with probability $\beta < 1$, and a single queue otherwise. Formally, the expected cost at a step is $O(n/\beta^2)$, and the expected maximum cost at a step is $O(n \log n/\beta)$.

Applications. The process considered above is sequential, and has FIFO semantics, since we are inserting labels in strictly increasing order, and assumes that all labels are inserted initially. We relax each of these assumptions, and discuss sufficient conditions for extending the analysis to concurrent and general priority insertions. To validate the practicality of our results, we implemented a variant of the $(1 + \beta)$ priority queue in C++. Empirical results, given in Section 5, show that it improves upon the state-of-the-art implementation of Rihani et al. [32] by 20% in terms of throughput, while providing similar rank guarantees. Moreover, when used in Dijkstra’s algorithm, $(1 + \beta)$ priority queue reduces running time by 10% over the original MultiQueue, and by 40% over a deterministic implementation [38]. This suggests that the decrease in rank guarantees is well compensated by improved scalability.

2 Related Work

The first instance of a distributed data structure implementation using a similar approach to the one we consider is probably the parallel branch-and-bound framework by Karp and Zhang [20]. In their construction, each processor is assigned a queue, and elements are inserted randomly into one of the queues. Each processor removes from its own queue. This critically relies on the synchronicity of PRAM processors to achieve bounds on average rank and maximum rank difference. It is easy to see that processor delays can cause the rank difference to become unbounded. Quantitative relaxations of concurrent data structures have been previously formalized by Henzinger et al. [17], but only for *deterministic* relaxations.

The MultiQueue. The strategy we consider is based on the MultiQueue data structure of Rihani et al. [32]. They start from n sequential priority queues, each of which is protected by a lock. To *insert* an element, each processor repeatedly picks a queue at random, and tries to lock it. When successful, the processor inserts the element into the queue, and unlocks it. Otherwise, the process re-tries with a new randomly chosen lock.

To *delete* an element, the algorithm samples two queues uniformly, and reads the top element from both. It then locks the queue whose top element had higher priority, removes that element, and returns it. If the algorithm failed to obtain the lock on the chosen queue, then it restarts the whole process. (They also consider variations of this strategy, which are essentially equivalent from the point of view of the sequential version.)

In the paper, the authors show via a basic balls-into-bins argument that, initially, before any elements are removed, the expected rank cost is $O(n)$, and the max cost is $O(n \log n)$ with high probability. Our argument applies to a sequential version of the concurrent process described above. The argument is significantly more general, since it applies to any time in the execution of the process. Variants of this approach have also been considered in other contexts, for instance for relaxed queues [15] and priority schedulers [29]. Designing efficient concurrent priority queues is a very active research area, see e.g. [3, 6, 36, 38] for recent examples, and [3] for a survey.

Balls-into-bins Processes. In the classic two-choice balanced allocation process, at each step, one new ball is to be inserted into one of n bins; the ball is inserted into the less loaded among two random choices [5, 26]. In this setting,

it is known that the most loaded bin is at most $O(\log \log n)$ above the average. The literature studying extensions of this process is extremely vast; we direct the reader to [31] for a survey. Considerable effort has been dedicated to understanding guarantees in the “heavily-loaded” case, where the number of insertion steps is unbounded [7, 30], and in the “weighted” case, in which ball weights come from a probability distribution [8, 37]. In a tour-de-force, Peres, Talwar, and Wieder [30] gave a potential argument characterizing a general form of the heavily-loaded, weighted process on graphs. The second step in our proof, which characterizes the deviation of weights from the mean in the exponential process, builds on their approach. Specifically, we use a similar potential function, and the parts of the analysis are similar. However, we generalize their approach to the case of biased insert distributions, and our argument diverges in several technical points: in particular, the potential analysis under unbalanced conditions is different. The first and third steps of our argument are completely new, as we consider a more complex *labelled* allocation process.

3 Process Definition

We are given with n priority queues, labeled $1, \dots, n$. To insert, we choose queue i with probability $\pi_i \in (0, 1)$, and insert into that queue. We assume that $\sum_{i=1}^n \pi_i = 1$, and that the bias is bounded, i.e. there exists a constant $\gamma \in (0, 1)$ such that, for any queue i , $1 - \gamma \leq \frac{1}{n\pi_i} \leq 1 + \gamma$. To remove, with probability $0 < \beta \leq 1$, we pick two queues u.a.r., and remove the element of minimum label among the two choices, and with probability $(1 - \beta)$ remove from a random queue. We call such a process a $(1 + \beta)$ -sequential process. When omit β when unimportant or clear from context.

At any point in the execution, we define the *rank* of any element to be the number of elements currently in the system which have lower label than it (including itself), so the minimal rank is 1. Our goal is to show that the ranks of the elements returned by the random process are small, throughout the entire execution. We will restrict our attention to executions which never inspect empty queues, and no priority inversions on inserts are visible to the removal process. Formally:

Definition 1. *An execution of the sequential process is prefixed if except with negligible probability (i.e. probability less than $n^{-\omega(1)}$), the sequential process never inspects an empty queue in a remove, and no remove sees an insert with lower priority than one that is inserted later in the same queue.*

There are many natural types of executions which are prefixed. For instance, a common strategy in practice is to insert a large enough “buffer” of elements initially, so as to minimize the likelihood of examining empty queues. More formally, since in each iteration we touch a queue with probability at most $\frac{2}{n} + O(\frac{1}{n^2})$ and only remove one element at a time, by Bernstein’s inequality, any execution of length T whose first $T/2 + \sqrt{n \cdot \omega(\log n)}$ operations are inserts, and which contains no inserts after this time, or which only inserts larger labels after this point, will be prefixed.

For the rest of the paper, we will tacitly assume that no remove inspects any empty queue in the sequential process, and that no priority inversions are visible to the removal process, and condition on the event that this does not happen. By Markov’s inequality, as long as we only consider prefixed executions, then for any sequence of operations of length $\text{poly}(n)$, this will change the max rank or average rank of the operations by at most a subconstant factor.

4 Analysis of the Sequential Process

The goal of this section is to prove the following theorem:

Theorem 1 (see Corollary 1 and Corollary 2). *Fix a bias bound γ . In any prefixed $(1 + \beta)$ sequential process, for $\beta = \Omega(\gamma)$, and for any time $t > 0$, the max rank of any element on top of a queue at time t is at most $O(\frac{1}{\beta}(n \log n + n \log \frac{1}{\beta}))$, and the average rank at time t is at most $O(n/\beta^2)$.*

The Exponential Process. For the purposes of analysis, it is useful to consider an alternative process, which we will call the *exponential process*, in which each bin independently generates real-valued labels by starting from 0, and adding an exponentially-distributed random variable with parameter $\lambda_i = 1/\pi_i$ to its previous label. More precisely, if $w_i(t)$ is the label present on top of bin i at time t , then the value of the next label is $w_i(t+1) = w_i(t) + \exp(1/\pi_i)$. Once we have enqueued all the items we wanted to enqueue, we proceed to remove items by the two-choice rule: we

always pick the element of minimum rank (label value) among two random choices. At any step, and for any $v \in \mathbb{R}$, we define $\text{rank}(v)$ to be the number of elements currently in the system with label at most v . At each step, we pay a cost equal to the rank of the element we just removed.

Proof Strategy. The argument proceeds in three steps. First, we show that, perhaps surprisingly, after all insertion steps have been performed, the rank distribution of the exponential process is *the same* as the rank distribution of the original process (Section 4.1). With this in place, we will perform the following coupling between the two processes:

Given an arbitrary number M of balls to be inserted into the n bins, we first generate M total weights across the n queues in the exponential process, so that each bin has exactly the same number of elements in both processes. We then replace each (real-valued) generated weight with its *rank* among all weights in the system. As we will see, for any rank j and queue i , the probability that the ball with rank j is in bin i is exactly the same as in the simple process, where queue i has probability of insertion π_i . More precisely, if we fix an increasing sequence of ranks r_1, r_2, \dots , the probability that bin i contains this exact sequence of ranks is the same across the two processes. Hence, we can couple the two processes to generate the same sequence of ranks in each bin.

For the removal step, we couple the two processes such that, in every step, they generate exactly the same β values and random choice indices. Since the ranks are the same, and the choices are the same, the two processes will remove from exactly the same bins at each step, and will pay the same rank cost in every step.

Hence, it will suffice to bound the expected rank cost paid at a step in the exponential process. Since this is difficult to do directly, we will first aim to characterize the *concentration* of the difference between the weight (value) on top of each bin and the average weight on top of bins, via a potential argument (Section 4.2). This will show that relatively few values can stray far from the mean. In turn, this will imply a bound on the average rank removed (Section 4.4).

In the following, we use the terms “queue” and “bin” interchangeably.

4.1 Equivalence between Rank Distributions

Let π_i be the probability that a ball is inserted into bin i . This section will be dedicated to showing that the distribution of ranks in the exponential process where bin i gets weights exponentially distributed with mean $1/\pi_i$ is identical to the distribution of labels in the original process, where the i th bin is chosen for insertion with probability π_i .

Theorem 2. *Let $I_{j \leftarrow i}$ be the event that the label with rank i is located in bin j , in either process. Let $\Pr_e[I_{j \leftarrow i}]$ be its probability in the exponential process, and $\Pr_o[I_{j \leftarrow i}]$ be its probability in the original process. Then for both processes $I_{j \leftarrow i}$ is independent from $I_{j' \leftarrow i'}$ for all $i \neq i'$ and*

$$\Pr_e[I_{j \leftarrow i}] = \Pr_o[I_{j \leftarrow i}] = \pi_j.$$

Proof. That $\Pr_o[I_{j \leftarrow i}] = \pi_j$ and that $I_{j \leftarrow i}, I_{j' \leftarrow i'}$ are independent for $i \neq i'$ in the original process both follow immediately by the construction of the sequential process.

We now require some notation. Let $\ell(i)$ be the label with rank i in the exponential process, and let $b(i)$ be the bin containing $\ell(i)$. We will employ the following *memoryless* property of exponential random variables:

Fact 1. *Let $X \sim \text{Exp}(r)$. Then, for all $s, t \geq 0$, we have $\Pr[X > s + t] = \Pr[X > s] \Pr[X > t]$.*

Fix an arbitrary $i \geq 1$, and let $L = \ell(i - 1)$ be the label of the $(i - 1)$ th element. For each bin j , we isolate the two consecutive weights between which L can be placed. More precisely, for each bin $1 \leq j \leq n$, denote the smallest label *larger* than L in bin j by $\ell_{j, > L}$ and let the largest label in B_j *smaller* than L be $\ell_{j, < L}$. By construction of the exponential process, $\ell_{j, > L} = \ell_{j, < L} + X_j$ where X_j is an exponentially distributed random variable with mean $1/\pi_j$. Furthermore, by assumption we have that $\ell_{j, > L} > L$, so $X_j > L - \ell_{j, < L}$. Then by memoryless-ness,

$$\begin{aligned} \Pr[\ell_{j, > L} > \ell_{j, < L} + s | \ell_{j, > L} > L] &= \Pr[X_j > s | X_j > (L - \ell_{j, < L})] \\ &= \Pr[X_j > (L - \ell_{j, < L}) + (\ell_{j, > L} - L) | X_j > (L - \ell_{j, < L})] \\ &= \Pr[X_j > \ell_{j, > L} - L] \end{aligned}$$

Let $\Delta_j = \ell_{j, > L} - L$, observing by the above that Δ_j is distributed exponentially with mean π_j . Furthermore, since the X_j are independent, the Δ_j are independent as well. Consider the event that $\ell_{j, > L}$ is the smallest label larger than

L in the system; that is $\ell(i) = \ell_{j, > L}$ and $b(i) = j$. This occurs if $\Delta_j < \Delta_{j'}$ for all $j' \neq j$, i.e. if Δ_j is the smallest such value among all bins. In turn, this occurs with probability

$$\begin{aligned} \int_0^\infty \Pr[\Delta_j = t] \prod_{k \neq j} \Pr[\Delta_k > t] dt &= \int_0^\infty \pi_j \exp(-t\pi_j) \prod_{k \neq j} \exp(-t\pi_k) dt \\ &= \pi_j \int_0^\infty \exp\left(-t \sum_{k=1}^n \pi_k\right) dt = \pi_j. \end{aligned}$$

□

4.2 Analysis of the Exponential Process

In the previous section, we have shown that the rank distributions are the same at the end of the insert process. The coupling described in Section 4, by which we inspect the same queues in each removal steps, implies that two processes will produce the same cost in terms of average rank removed. Hence, we focus on bounding the *rank* of the labels removed from the exponential process at each step t . Our strategy will be to first bound the deviation of the *weight* on top of each queue from the average, at every step t , and then re-interpret the deviation in terms of rank cost.

Notation. From this point on, we will assume for simplicity that, at the beginning of each step, queues are always ranked in *increasing order of their top label*. If p_i is the probability that we pick the i th ranked bin for a removal, and β is the two-choice probability, then it is easy to see that the $(1 + \beta)$ process guarantees that

$$p_i = (1 - \beta) \frac{1}{n} + \beta \left[\frac{2}{n} \left(1 - \frac{i-1}{n} \right) - \frac{1}{n^2} \right].$$

Further, notice that, for any $1 \leq m \leq n$, we have, ignoring the negligible $O(1/n^2)$ factor, that

$$\sum_{i=1}^m p_i \simeq \frac{m}{n} \left(1 + \beta - \frac{m}{n} \beta \right).$$

For any bin j and time t , let $w_j(t)$ be the label on top of bin j at time t , and let $x_j(t) = w_j(t)/n$ be the normalized label. Let $\mu(t) = \sum_{j=1}^n x_j(t)/n$ be the average normalized label at time t over the bins. Let $y_i(t) = x_i(t) - \mu(t)$, and let $\alpha < 1$ be a parameter we will fix later. Define

$$\Phi(t) = \sum_{j=1}^n \exp(\alpha y_i(t)), \text{ and } \Psi(t) = \sum_{j=1}^n \exp(-\alpha y_i(t)).$$

Finally, define the potential function

$$\Gamma(t) = \Phi(t) + \Psi(t).$$

Parameters and Constants. Define $\epsilon = \frac{\beta}{16}$. Recall that the parameter $\gamma > 0$ is such that, for every $1 \leq i \leq n$, $\frac{1}{n\pi_i} \in [1 - \gamma, 1 + \gamma]$. Next, let $c \geq 2$ be a constant, and δ be a parameter such that

$$1 + \delta = \frac{1 + \gamma + c\alpha(1 + \gamma)^2}{1 - \gamma - c\alpha(1 + \gamma)^2}. \quad (1)$$

In the following, we assume that the parameters α, β , and γ satisfy the inequality

$$\frac{\beta}{16} = \epsilon \geq \delta. \quad (2)$$

We assume that the insertion bias $\gamma \leq 1/2$ is small, and hence this is satisfied by setting $\beta = \Omega(\gamma)$ and $\alpha = \Theta(\beta)$. The rest of this section is dedicated to the proof of the following claim:

Theorem 3. Let $\vec{p} = (p_1, p_2, \dots, p_n)$ be the vector of probabilities, sorted in increasing order. Let $\epsilon = \beta/16$, and $c \geq 2$ be a small constant. Let $\alpha, \beta, \gamma, \delta$ be parameters as given above, such that $\epsilon \geq \delta$. Then there exists a constant $C(\epsilon) = \text{poly}(\frac{1}{\epsilon})$ such that, for any time $t \geq 0$, we have $\mathbb{E}[\Gamma(t)] \leq C(\epsilon)n$.

Potential Argument Overview. The argument proceeds as follows. We will begin by bounding the change of each potential function Ψ and Φ , at each queue in a step (Lemmas 1 and 6). We then use these bounds to show that, if not too many queues have weights above or below the mean ($y_{n/4} \leq 0$ and $y_{3n/4} \geq 0$), then Φ and Ψ respectively decrease in expectation (Lemmas 7 and 8). Unfortunately, this does not necessarily hold in general configurations. However, we are able to show the following claim: if the configuration is unbalanced (e.g. $y_{n/4} > 0$) and Φ does not decrease in expectation at a step, then either the symmetric potential Ψ is large, and will decrease on average, or the global potential function Γ must be in $O(n)$ (Lemma 9). We will prove a similar claim for Ψ . Putting everything together, we get that the global potential function Γ always decreases in expectation once it exceeds the $O(n)$ threshold, which implies Theorem 3.

Potential Change at Each Step. We begin by analyzing the expected change in potential for each queue from step to step. We first look at the change in the weight vector $\vec{y} = (y_1, y_2, \dots, y_n)$. (Recall that we always re-order the queues in increasing order of weight at the beginning of each step.) Below, let κ_j be the cost increase if the bin of rank j is chosen, which is an exponential random variable with mean $1/\pi_j$. We have that, for every rank i ,

$$y_i(t+1) = \begin{cases} y_i(t) + \kappa_i (1 - 1/n), & \text{with probability } p_i; \text{ i.e., the queue of rank } i \text{ is picked, and} \\ y_i(t) - \frac{\kappa_j}{n}, & \text{if some other bin } j \neq i \text{ is chosen.} \end{cases}$$

The Change in Φ at a Step. We prove the following bound on the expected increase in Φ at a step.

Lemma 1. For any bin rank i ,

$$\mathbb{E}[\Delta\Phi \mid y(t)] \leq \sum_{i=1}^n \hat{\alpha} \left((1 + \delta)p_i - \frac{1}{n} \right) \exp(\alpha y_i(t)).$$

Proof. Let $\Phi_i(t) = \exp(\alpha y_i(t))$. We have two cases. If the bin is chosen for removal, then the change is:

$$\begin{aligned} \Delta\Phi_i &:= \Phi_i(t+1) - \Phi_i(t) = \exp\left(\frac{\alpha}{n}(y_i(t) + \kappa_i(1 - \frac{1}{n}))\right) - \exp\left(\frac{\alpha}{n}y_i(t)\right) \\ &= \exp\left(\frac{\alpha}{n}y_i(t)\right) \left(\exp\left(\frac{\alpha}{n}\kappa_i(1 - \frac{1}{n})\right) - 1\right). \end{aligned}$$

Taking expectations with respect to the random choices made on insertion, that is, the value of κ_j , we have

$$\mathbb{E}_i \left[\exp\left(\kappa_i \frac{\alpha}{n} \left(1 - \frac{1}{n}\right)\right) \right] \stackrel{(a)}{=} \frac{\pi_i}{\pi_i - (\frac{\alpha}{n}(1 - \frac{1}{n}))} = \frac{1}{1 - \frac{\alpha}{\pi_i n}(1 - \frac{1}{n})} \stackrel{(b)}{\leq} 1 + \frac{\alpha}{\pi_i n} \left(1 - \frac{1}{n}\right) + c \left(\frac{\alpha}{\pi_i n} \left(1 - \frac{1}{n}\right)\right)^2,$$

for some constant $c > 1$. Step (a) follows from the observation that the expectation we wish to compute is the moment-generating function of the exponential distribution at $\frac{\alpha}{n}(1 - \frac{1}{n})$, while (b) follows from the Taylor expansion. (We slightly abused notation in step (a) by denoting with π_i the insert probability for the i th ranked queue, according to the ranking in this step.)

The second step if if some other bin $j \neq i$ is chosen for removal, then the change is:

$$\Delta\Phi_i = \exp\left(\frac{\alpha}{n}(y_i(t) - \kappa_j \frac{1}{n})\right) - \exp\left(\frac{\alpha}{n}y_i(t)\right) = \exp\left(\frac{\alpha}{n}y_i(t)\right) \left(\exp\left(-\frac{\alpha\kappa_j}{n^2}\right) - 1\right).$$

Again taking expectations with respect to the random choices made on insertion, i.e. the value of κ_j , we have

$$\mathbb{E}_i \left[\exp\left(-\kappa_j \frac{\alpha}{n^2}\right) \right] = \frac{\pi_j}{\pi_j + \frac{\alpha}{n^2}} = \frac{1}{1 + \frac{\alpha}{\pi_j n^2}} \leq 1 - \frac{\alpha}{\pi_j n^2} + \left(\frac{\alpha}{\pi_j n^2}\right)^2 - \dots \leq 1 - \frac{\alpha}{\pi_j n^2} + \left(\frac{\alpha}{\pi_j n^2}\right)^2.$$

Therefore, we have that

$$\begin{aligned}
\mathbb{E}[\Delta\Phi_i]/\Phi_i(t) &\leq \left(1 + \frac{\alpha}{\pi_i n} \left(1 - \frac{1}{n}\right) + c \left(\frac{\alpha}{\pi_i n} \left(1 - \frac{1}{n}\right)\right)^2\right) p_i + \sum_{j \neq i} \left(1 - \frac{\alpha}{\pi_j n^2} + \left(\frac{\alpha}{\pi_j n^2}\right)^2\right) p_j - 1 \\
&\leq \left(\frac{\alpha}{\pi_i n} + c \left(\frac{\alpha}{\pi_i n}\right)^2\right) p_i - \sum_{j=1}^n \left(\frac{\alpha}{\pi_j n^2} - c \left(\frac{\alpha}{\pi_j n^2}\right)^2\right) p_j \\
&\leq \left(1 + \gamma + c\alpha(1 + \gamma)^2\right) \alpha p_i - \frac{\alpha}{n} \sum_{j=1}^n \left(1 - \gamma - c\alpha(1 + \gamma)^2\right) p_j.
\end{aligned}$$

If we denote for convenience

$$\hat{\alpha} = \alpha \left(1 - \gamma - c\alpha(1 + \gamma)^2\right), \text{ and recall that } \delta := \frac{1 + \gamma + c\alpha(1 + \gamma)^2}{1 - \gamma - c\alpha(1 + \gamma)^2} - 1,$$

then we can rewrite this as

$$\mathbb{E}[\Delta\Phi | y(t)] \leq \sum_{i=1}^n \hat{\alpha} \left(p_i(1 + \delta) - \frac{1}{n}\right) \Phi_i(t).$$

□

The Change in Ψ . Using a symmetric argument, we can prove the following about the expected change in Ψ .

Lemma 6.

$$\mathbb{E}[\Delta\Psi | y(t)] \leq \sum_{i=1}^n \hat{\alpha} \left((1 + \delta) \frac{1}{n} - p_i\right) \exp(-\alpha y_i(t)).$$

Bounds under Balanced Conditions. Let us briefly stop to examine the bounds in the above Lemmas. The terms $(p_i(1 + \delta) - \frac{1}{n})$ are *decreasing* in i , and in fact become *negative* as i increases. (The exact index where this occurs is controlled by β and δ .) The $\exp(\alpha y_i(t))$ terms are *increasing* in i . Bins whose weight is *below* the mean (i.e., $y_i(t) \leq 0$) have a negligible effect on Φ , since each of their contributions is at most 1. At the same time, notice that the contribution of bins of large index i will be negative. Hence, we can show that, if at least $n/4$ bins have weights below average, then the value of Φ tends to *decrease* on average.

Lemma 7. If $y_{n/4} \leq 0$, then we have that

$$\mathbb{E}[\Phi(t + 1) | y(t)] \leq \left(1 - \frac{\hat{\alpha}\epsilon}{3n}\right) \Phi(t) + 1.$$

A similar claim holds for Ψ , under the condition that there are at least $n/4$ bins with weight *larger* than average.

Lemma 8. If $y_{3n/4} \geq 0$, then we have that

$$\mathbb{E}[\Psi(t + 1) | y(t)] \leq \left(1 - \frac{\hat{\alpha}\epsilon}{3n}\right) \Psi(t) + 1.$$

Bounds under Unbalanced Conditions. We now analyze unbalanced configurations, where there are either many bins whose weights are above average (e.g., $y_{n/4} > 0$), or below average ($y_{3n/4} < 0$). The rationale we used to bound each potential function independently no longer applies. In particular, as shown in [30], we can have unbalanced settings where for example Φ does not decrease in expectation. However, we can show that, one of two things must hold: either the other potential Ψ is *larger* and *does decrease* in expectation, or the global potential Γ is in $O(n)$.

Lemma 9. Given ϵ as above, assume that $y_{n/4}(t) > 0$, and $\mathbb{E}[\Delta\Phi] \geq -\frac{\epsilon\hat{\alpha}}{3n}\Phi(t)$. Then either $\Phi < \frac{\epsilon}{4}\Psi$ or $\Gamma = O(n)$.

Proof. Fix $\lambda = 2/3 - 1/54$ for the rest of the proof. We can split the inequality in Lemma 1 as follows:

$$\mathbb{E} [\Delta \Phi | y(t)] \leq \sum_{i=1}^{\lambda n} \hat{\alpha} \left((1 + \delta) p_i - \frac{1}{n} \right) \exp(\alpha y_i(t)) + \sum_{i=\lambda n+1}^n \hat{\alpha} \left((1 + \delta) p_i - \frac{1}{n} \right) \exp(\alpha y_i(t)). \quad (3)$$

We bound each term separately. Since the probability terms are non-increasing and the exponential terms are non-decreasing, the first term is maximized when all p_i terms are equal. Since these probabilities are at most 1, we have

$$\sum_{i=1}^{\lambda n} \hat{\alpha} \left((1 + \delta) p_i - \frac{1}{n} \right) \exp(\alpha y_i(t)) \leq \frac{\hat{\alpha}}{n} \left((1 + \delta) \frac{1}{\lambda} - 1 \right) \Phi_{\leq \lambda n}. \quad (4)$$

The second term is maximized by noticing that the p_i factors are non-increasing, and are thus dominated by their value at λn . Noticing that we carefully picked λ such that

$$p_{\lambda n} \leq \frac{1}{n} - \frac{4\epsilon}{n},$$

we obtain, using the assumed inequality $\delta \leq \epsilon$, that

$$\sum_{i=\lambda n+1}^n \hat{\alpha} \left((1 + \delta) p_i - \frac{1}{n} \right) \exp(\alpha y_i(t)) \leq \hat{\alpha} (\delta - 4\epsilon) \frac{\Phi_{> \lambda n}}{n} \leq -\frac{3\epsilon \hat{\alpha}}{n} \Phi_{> \lambda n}. \quad (5)$$

By the case assumption, we know that $\mathbb{E} [\Delta \Phi | y(t)] \geq -\frac{\hat{\alpha}\epsilon}{3n} \Phi(t)$. Combining the bounds (18), (19), and (20), this yields:

$$\frac{\hat{\alpha}}{n} \left((1 + \delta) \frac{1}{\lambda} - 1 \right) \Phi_{\leq \lambda n} - \frac{3\epsilon \hat{\alpha}}{n} \Phi_{> \lambda n} \geq -\frac{\hat{\alpha}\epsilon}{3n} \Phi(t).$$

Substituting $\Phi_{> \lambda n} = \Phi - \Phi_{\leq \lambda n}$ yields:

$$(3\epsilon - \epsilon/3) \Phi \leq \left((1 + \delta) \frac{1}{\lambda} - 1 + 3\epsilon \right) \Phi_{\leq \lambda n}.$$

For simplicity, we fix $C(\epsilon) = \frac{(1+\delta)\frac{1}{\lambda}-1+3\epsilon}{3\epsilon-\epsilon/3} = O\left(\frac{1}{\epsilon}\right)$, to obtain

$$\Phi \leq C(\epsilon) \Phi_{\leq \lambda n}. \quad (6)$$

Let $B = \sum_{y_i > 0} y_i$. Since we are normalizing by the mean, it also holds that $B = \sum_{y_i < 0} (-y_i)$. Notice that

$$\Phi_{\leq \lambda n} \stackrel{y_i \text{ incr.}}{\leq} \lambda n \exp(\alpha y_{\lambda n}) \stackrel{y_i \text{ incr.}}{\leq} \lambda n \exp\left(\frac{\alpha B}{(1-\lambda)n}\right). \quad (7)$$

We put inequalities (21) and (22) together and get

$$\Phi(t) \leq \lambda n C(\epsilon) \exp\left(\frac{\alpha B}{(1-\lambda)n}\right), \quad (8)$$

Let us now lower bound the value of Ψ under these conditions. Since $y_{n/4} > 0$, all the costs below average must be in the first quarter of y . We can apply Jensen's inequality to the first $n/4$ terms of Ψ to get that

$$\Psi \geq \sum_{i=1}^{n/4} \exp(-\alpha y_i) \geq \frac{n}{4} \exp\left(-\alpha \frac{\sum_{i=1}^{n/4} y_i}{n/4}\right).$$

We now split the sum $\sum_{i=1}^{n/4} y_i$ into its positive part and its negative part. We know that the negative part is summing up to exactly $-B$, as it contains all the negative y_i 's and the total sum is 0. The positive part can be of size at most

$B/4$, since it is maximized when there are exactly $n - 1$ positive costs and they are all equal. Hence the sum of the first $n/4$ elements is at least $-3B/4$, which implies that the following bound holds:

$$\Psi \geq \frac{n}{4} \exp\left(-\alpha \frac{-3B/4}{n/4}\right) \geq \frac{n}{4} \exp\left(\alpha \frac{3B}{n}\right). \quad (9)$$

If $\Phi < \frac{\epsilon}{4}\Psi$, then there is nothing to prove. Otherwise, if $\Phi \geq \frac{\epsilon}{4}\Psi$, we get from (24) and (23) that

$$\frac{\epsilon}{4} \frac{n}{4} \exp\left(\alpha \frac{3B}{n}\right) \leq \frac{\epsilon}{4} \Psi \leq \Phi(t) \leq \lambda n C(\epsilon) \exp\left(\frac{\alpha B}{(1-\lambda)n}\right),$$

Therefore, we get that:

$$\exp\left(\alpha \frac{B}{n} \left(3 - \frac{1}{1-\lambda}\right)\right) \leq \frac{4\lambda}{\epsilon} C(\epsilon) = O\left(\frac{1}{\epsilon^2}\right).$$

Using the mundane fact that $3 - \frac{1}{1-\lambda} = 3/19$, we get that

$$\exp\left(\frac{\alpha B}{n}\right) \leq O\left(\frac{1}{\epsilon^{14}}\right). \quad (10)$$

To conclude, notice that (25) implies we can upper bound Γ in this case as:

$$\Gamma = \Phi + \Psi \leq \frac{4+\epsilon}{\epsilon} \Phi \leq \frac{4+\epsilon}{\epsilon} \lambda n C(\epsilon) \exp\left(\frac{\alpha B}{(1-\lambda)n}\right) \leq O\left(\frac{1}{\epsilon^{14/(1-\lambda)}}\right) \frac{4+\epsilon}{\epsilon} C(\epsilon) \lambda n = O(\text{poly}\left(\frac{1}{\epsilon}\right) n).$$

□

We can prove a symmetric claim for Ψ by a slightly different argument.

Lemma 10. *Given ϵ as above, assume that $y_{3n/4} < 0$, and that $\mathbb{E}[\Delta\Psi] \geq -\frac{\hat{\alpha}\epsilon}{3n}\Psi$. Then either $\Psi < \frac{\epsilon}{4}\Phi$ or $\Gamma = O(n)$.*

Endgame. We now finally have the required machinery to prove that Γ satisfies a supermartingale property:

Lemma 2. *There exists a constant ϵ such that*

$$\mathbb{E}[\Gamma(t+1)|y(t)] \leq \left(1 - \frac{\hat{\alpha}\epsilon}{4n}\right) \Gamma(t) + C, \text{ where } C \text{ is a constant in } O\left(\text{poly}\left(\frac{1}{\epsilon}\right)\right).$$

Proof. Case 1: If $y_{n/4} \leq 0$ and $y_{3n/4} \geq 0$, then the property follows by putting together Lemmas 7 and 8.

Case 2: If $y_{3n/4} \geq y_{n/4} \geq 0$. This means that the weight vector is unbalanced, in particular that there are few bins of low cost, and many bins of high cost. However, we can show that the expected decrease in Ψ can compensate this decrease, and the inequality still holds. Notice that Lemmas 1 and 6 imply that $\mathbb{E}[\Delta\Phi] \leq \frac{\hat{\alpha}\delta}{n}\Phi(t)$ and $\mathbb{E}[\Delta\Psi] \leq \frac{\hat{\alpha}\delta}{n}\Psi(t)$, respectively.

If $\mathbb{E}[\Delta\Phi] \leq -\frac{\epsilon\hat{\alpha}}{4n}\Phi(t)$, then the claim simply follows by putting this bound together with Lemma 8. Otherwise, if $\mathbb{E}[\Delta\Phi] \geq -\frac{\epsilon\hat{\alpha}}{4n}\Phi(t)$, we obtain from Lemma 9 that either $\Phi < \frac{\epsilon}{4}\Psi$ or $\Gamma = O(n)$. In the first sub-case, we have that

$$\begin{aligned} \mathbb{E}[\Delta\Gamma] &= \mathbb{E}[\Delta\Phi] + \mathbb{E}[\Delta\Psi] \leq \frac{\hat{\alpha}\delta}{n}\Phi(t) + 1 - \frac{\hat{\alpha}\epsilon}{3n}\Psi(t) \leq -\frac{\hat{\alpha}\epsilon}{3n}\Psi(t) + 1 + \frac{\epsilon\delta\hat{\alpha}}{4n}\Psi(t) \\ &\leq -\frac{\hat{\alpha}\epsilon}{3n}\left(1 - \frac{3\delta}{4}\right)\Psi(t) + 1 \leq -\frac{\hat{\alpha}}{3n}\left(1 + \frac{\epsilon}{4}\right)^{-1}\left(1 - \frac{3\delta}{4}\right)\Gamma(t) \leq -\frac{\hat{\alpha}}{4n}\Gamma(t), \text{ as claimed.} \end{aligned}$$

In the second sub-case, we know that $\Gamma \leq Cn$, for some constant C . Hence, we can get that

$$\mathbb{E}[\Delta\Gamma] = \mathbb{E}[\Delta\Phi] + \mathbb{E}[\Delta\Psi] \stackrel{(a)}{\leq} \frac{\hat{\alpha}\delta}{n}\Gamma \stackrel{\Gamma=O(n)}{\leq} 2C\hat{\alpha},$$

where in step (a) we used the upper bounds in Lemmas 1 and 6, respectively. On the other hand,

$$C - \frac{\hat{\alpha}\epsilon}{4n}\Gamma(t) \geq C - \frac{\hat{\alpha}\epsilon}{4n}Cn = C\left(1 - \frac{\hat{\alpha}\epsilon}{4}\right) \geq 3C\hat{\alpha} \geq \mathbb{E}[\Delta\Gamma].$$

Case 3: If $y_{n/4} \leq y_{3n/4} < 0$. This case is symmetric to the one above.

□

The intuition behind the above bound is that Γ will always tend to decrease once it surpasses the $\Theta(n)$ threshold. This implies a bound on the expected value of Γ , which completes the proof of Theorem 3.

Lemma 3. For any $t \geq 0$, $\mathbb{E}[\Gamma(t)] \leq \frac{4C}{\hat{\alpha}\epsilon}n$.

Proof. By induction. This holds for $t = 0$ by a direct computation: see Lemma 2 for the argument. Then, we have

$$\mathbb{E}[\Gamma(t+1)] = \mathbb{E}[\mathbb{E}[\Gamma(t+1)|\Gamma(t)]] \leq \mathbb{E}\left[\left(1 - \frac{\hat{\alpha}\epsilon}{4n}\right)\Gamma(t) + C\right] \leq \frac{4C}{\hat{\alpha}\epsilon}n\left(1 - \frac{\hat{\alpha}\epsilon}{4n}\right) + C \leq \frac{4C}{\hat{\alpha}\epsilon}n.$$

□

4.3 Guarantees on Max Rank

We can use the characterization of the exponential process to prove the following:

Lemma 4. If $w_{\max}(t)$ is the maximum bin weight at time t and $w_{\min}(t)$ is the minimum, then

$$\mathbb{E}[w_{\max}(t) - w_{\min}(t)] = O\left(\frac{1}{\alpha}n(\log n + \log C)\right). \quad (11)$$

Proof. Let $x_{\max}(t) = w_{\max}(t)/n$ and $x_{\min}(t) = w_{\min}(t)/n$. By definition, we have $\exp(\alpha(x_{\max}(t) - \mu(t))) \leq \Gamma(t)$ and $\exp(\alpha(\mu(t) - x_{\min}(t))) \leq \Gamma(t)$. Therefore $\alpha(x_{\max}(t) - x_{\min}(t)) \leq 2 \log \Gamma(t)$. Thus, we have

$$\mathbb{E}[\alpha(x_{\max}(t) - x_{\min}(t))] \leq 2 \mathbb{E}[\log \Gamma(t)] \stackrel{(a)}{\leq} 2 \log(\mathbb{E}[\Gamma(t)]) \stackrel{(b)}{=} O(\log n + \log C),$$

where (a) follows from Jensen's inequality and (b) follows from Theorem 3. Simplifying yields the desired claim. □

In Appendix E we show that this implies the following theorem:

Theorem 4. For all $t \geq 0$, we have $\mathbb{E}[\max \text{rank}(t)] = O\left(\frac{1}{\alpha}n(\log n + \log C)\right)$.

By plugging in constants as in (1) and (2), we get:

Corollary 1. For all $t \geq 0$, and any $\beta = \Omega(\gamma)$, we have $\mathbb{E}[\max \text{rank}(t)] = O\left(\frac{1}{\beta}n(\log n + \log 1/\beta)\right)$.

4.4 Guarantees on Average Rank

We now focus on at the rank cost paid in a step by the algorithm. Let $A = \log C/\alpha$. For real values $s \geq 0$, we “stripe” the bins according to their top value, denoting by $b_{>s}(t)$ the number of bins with $w_j(t) \geq (s + A)n + \mu$ at time t , and let $b_{<-s}(t)$ be the number of bins with $w_j(t) \leq \mu - (s + A)n$ at time t . For any bin j and interval I , we also let $\ell_{j,I}(t)$ be the number of elements in j at time t with label in I . Finally, let $p_{j,\mu}$ denote the PDF of $w_j(t)$ given μ_t .

We discuss the proof strategy at a high level here, deferring the full proofs to Appendix F. First, we use the bounds on Γ to obtain the following bounds on the quantities defined previously:

Lemma 5. For any time t , we have that $\mathbb{E}[b_{>s}] \leq n \exp(-\alpha s)$ and that $\mathbb{E}[b_{<-s}] \leq n \exp(-\alpha s)$.

Proof. Recall that $\Phi(t) = \sum_{i=1}^n \exp(\alpha(x_i(t) - \mu))$ and that $\mathbb{E}[\Phi(t)] \leq Cn$. By linearity of expectation, we have

$$\mathbb{E}[\Phi(t)] \geq \mathbb{E}[b_{>s} \exp(\alpha(s + A))] = C \exp(\alpha s) \cdot \mathbb{E}[b_{>s}],$$

which implies the claim. The converse claim follows from the bound on $\Psi(t)$. □

This lemma gives us strong bounds on the tail behavior of the w_j . We show that this implies a bound on the average rank. We will need the following technical result.

Lemma 11. For any interval $I = [a, b]$ which may depend on μ , we have $\mathbb{E}[\text{rank}(b) - \text{rank}(a)|\mu] \leq n(b - a + 1)$.

We then show a bound on the rank of μ :

Lemma 12. *For all t , we have $E[\text{rank}(\mu(t))] \leq O((A + 1/\alpha^2)n)$.*

The proof of this lemma, deferred to Appendix F, works as follows. We divide up the interval $(-\infty, \mu]$ into infinitely many intervals of constant length. We will count the number of $w_j(t)$ within each interval separately. Within each such interval, we can give a crude upper bound the number of weights that could have ever been in that interval at any point in the execution. We then use Lemma 5 to show that at the current time in the execution, the number of bins with elements in each interval decays exponentially as we take intervals which are further and further away from μ . Summing these values up gives the desired bound. Finally, this implies:

Theorem 5. *For all t , we have $\mathbb{E}[\frac{1}{n} \sum_{i=1}^n \text{rank}_j(t)] = O(A + \frac{1}{\alpha^2})n$.*

Proof. By Lemma 12 and Lemma 11, we have that

$$\begin{aligned} \mathbb{E}[\text{rank}_j(t) | w_j(t) \leq \mu(t) + An] &\leq O((A + 1/\alpha^2)n) + (A + 1)n \\ &= O((A + 1/\alpha^2)n). \end{aligned}$$

Thus, we have

$$\begin{aligned} \mathbb{E}\left[\sum_{i=1}^n \text{rank}_j(t)\right] &= \mathbb{E}\left[\sum_{w_j(t) \leq \mu + An} \mathbb{E}[\text{rank}_j(t) | \mu] + \sum_{w_j(t) > \mu + An} \mathbb{E}[\text{rank}_j(t) | \mu]\right] \\ &= \mathbb{E}\left[\sum_{w_j(t) \leq \mu + An} \mathbb{E}[\text{rank}_j(t) | \mu] + \sum_{j=1}^n \sum_{w_j(t) > \mu + An} \mathbb{E}[\text{rank}_j(\mu + An) + \ell_{j,(\mu, w_j(t))} | \mu, w_j(t)]\right] \\ &\stackrel{(a)}{\leq} O\left(\left(A + \frac{1}{\alpha^2}\right)n^2\right) + \sum_{j=1}^n \mathbb{E}\left[\sum_{w_j(t) > \mu + An} \mathbb{E}[\ell_{j,(\mu, w_j(t))} | \mu, w_j(t)]\right] \\ &\stackrel{(b)}{\leq} O\left(\left(A + \frac{1}{\alpha^2}\right)n^2\right) + \sum_{j=1}^n \mathbb{E}\left[\sum_{w_j(t) > \mu + An} (w_j(t) - \mu + 1)\right] \\ &= O\left(\left(A + \frac{1}{\alpha^2}\right)n^2\right) + \sum_{j=1}^n \sum_{k=0}^{\infty} \mathbb{E}\left[\sum_{w_j(t) \in [\mu + (A+k)n, \mu + (A+k+1)n]} (w_j(t) - \mu + 1)\right] \\ &\leq O\left(\left(A + \frac{1}{\alpha^2}\right)n^2\right) + n \sum_{k=0}^{\infty} (A + k + 2)n \mathbb{E}[b_{>k}] \\ &\stackrel{(c)}{\leq} O\left(\left(A + \frac{1}{\alpha^2}\right)n^2\right) + O\left(\frac{1}{\alpha^2}n^2\right). \end{aligned}$$

where (a) follows by Lemma 12, (b) follows by Lemma 11, and (c) follows from Lemma 5. Thus

$$\mathbb{E}\left[\frac{1}{n} \sum_{i=1}^n \text{rank}_j(t)\right] = O\left(A + \frac{1}{\alpha^2}\right)n$$

as claimed. \square

We now show how these imply bounds for our removal processes. The actual rank choice at time t is always better than a uniform choice in expectation, since it uses power of two choices. If we consider an $(1 + \beta)$ process, where we only do two choices with probability $\beta = \Omega(\gamma)$, we obtain the following, by setting parameters as in (1) and (2):

Corollary 2. *For all t , if we let $\beta = \Omega(\gamma)$, and we let $r(t)$ denote the rank of the removed element at time t , then*

$$\mathbb{E}[r(t)] = O\left(\left(\frac{\log C}{\alpha} + \frac{1}{\alpha^2}\right)n\right) = O\left(\frac{n}{\beta^2}\right).$$




Figure 1: Throughput comparison for the $(1 + \beta)$ priority queue with $\beta = 0.5$ and 0.75 , versus the original MultiQueues, the Linden-Jonsson implementation, and kLSM. Higher is better.




Figure 2: Mean rank returned (log scale) for the $(1 + \beta)$ priority queue, for various values of β on 8 queues and 8 threads. Lower is better.

5 Experimental Results

Setup and Methodology. We implemented a $(1 + \beta)$ priority queue based on the MultiQueue implementation from the priority queue benchmark framework of [14], and benchmarked it against the original MultiQueue ($\beta = 1$), the Linden-Jonsson [24] skiplist-based implementation, and the kLSM deterministic-relaxed data structure [38], with a relaxation factor of 256, which has been found to perform best. The MultiQueue uses efficient sequential priority queues from the boost library. Tests were performed on a recent Intel(R) Xeon(R) CPU E7-8890 (Haswell architecture), with 18 hardware threads, each running at 2.5GHz. The tests for mean rank returned use coherent timestamps to record the times when elements are returned at each thread. We use these in a post-processing step to count rank inversions. This methodology might not be 100% accurate, since the use of timestamps might change the schedule; however, we believe results should be reasonably close to the true values.

For the throughput experiments, we consider executions consisting of alternating insert and deleteMin operations, for 10 seconds. Experiment outputs are averaged over 10 trials. Removals on empty queues do not count towards throughput. Since we are interested in the regime where queues are never empty, we insert 10 million elements initially. Threads are pinned to cores, and memory allocation is affinitized. The single-source shortest paths benchmark is a version of Dijkstra’s algorithm, running on a weighted, directed California road network graph.

Results. Figure 1 illustrates the throughput differential between the various implementations. As previously stated, the MultiQueue variants are superior to other implementations (except at very low thread counts). Of note, the variants with $\beta < 1$ improve on the standard implementation by up to 20%. Since throughput figures are not conclusive in isolation, we also benchmarked the average rank cost in Figure 2. (The y axis is logarithmic.) Note that the increase in average cost due to the further β relaxation is relatively limited. Results are conformant with our analysis for $\beta \geq 0.5$. The apparent inflection point at around $\beta = 0.5$ could be explained by the $\epsilon \geq \delta$ bias assumptions breaking down after this point, or by non-trivial correlations in the actual execution which mean that our analysis no longer applies.

Finally, Figure 3 gives a running times for a single-source shortest path benchmark, using a parallel version of Dijkstra’s algorithm. We note that the relaxed versions with $\beta < 1$ can be superior in terms of running time to the version with $\beta = 1$, by up to 10%. The version with $\beta = 0$ (not shown) is the fastest at low thread counts, but then loses performance at thread counts ≥ 8 , probably because of excessive relaxation.




Figure 3: Running times for single-source shortest path benchmark, using various versions of the priority queue, and kLSM. Lower is better.

6 Discussion and Future Work

We have provided tight rank guarantees for a practically-inspired randomized priority scheduling process. Moreover, we showed that this strategy is robust in terms of its bias and randomness requirements. Intuitively, our results show that, given biased random insertions into the queues, the preference towards lower ranks provided by the two-choice process is enough to give strong linear bounds on the average rank removed. We extended our analysis to a practical algorithm which improves on the state-of-the-art MultiQueues in Section 5.

Tightness. The bounds we provide for the two-choice process are asymptotically tight. To see that the $O(n)$ bound is tight, it suffices to notice that, even in a best-case scenario, the rank of the k th least expensive queue is at least k . Hence, the process has to have expected rank cost $\Omega(n)$. The tightness argument for $\Theta(n \log n)$ expected worst-case cost is more complex. In particular, it is known [30, Example 2] that the gap between the most loaded and average bin load in a weighted balls-into-bins process with weights coming from an exponential distribution of mean 1 is $\Theta(\log n)$ in expectation. That is, there exists a queue j from which we have removed $\Theta(\log n)$ fewer times than the average. We can extend the argument in Section 4.4 to prove that there exist $\Theta(n)$ queues which have $\Theta(\log n)$ elements of higher label than the top element of queue j . This implies that the *rank cost* of queue j is $\Theta(n \log n)$, as claimed. We conjecture that the dependency in β for the expected rank bound on the $(1 + \beta)$ process can be improved to *linear*.

Relation to Concurrent Processes. An important question is whether we can show that the practical *concurrent* implementations provide similar bounds. This claim holds if we assume that, in the concurrent implementation, the steps comparing the two top elements of the randomly chosen queues and removing the better top element are all performed *atomically*, e.g. through hardware transactional memory. In this case, we can show that the probability distributions of ranks returned by the concurrent and sequential processes are identical, and hence our guarantees apply. We formalize the required property, called *distributional linearizability*, in Appendix C.

This property does not appear to hold in general for fine-grained lock-based or lock-free implementations. Appendix C provides a detailed discussion of the difficulties and limitations of extending this approach to a concurrent setting. In particular, due to complex correlations implied by concurrent execution, our bounds might not hold for some locking strategies for subtle reasons, and it appears highly non-trivial to rigorously extend them to any particular fine-grained strategy.

In future work, we plan to examine whether exist *relaxations* of distributional linearizability satisfied by real implementations, which would allow our bounds to be extended to concurrent implementations without additional assumptions. The results in Section 5 suggest that existing implementations do satisfy strong rank guarantees.

Future Work. An extremely intriguing question regards the impact of using such relaxed priority queue structures in the context of existing parallel algorithms, e.g. for shortest-paths computation. In particular, it would be interesting to

bound the amount of extra work caused by relaxation versus the benefits of parallelism.

Another interesting question in terms of future work is whether our results can also be extended to processes on *graphs*. More precisely, imagine we are given a strongly connected, undirected graph with n vertices. Insertions of increasing integer labels are performed at uniformly random nodes. In every removal step, we pick a random edge, remove the higher priority label among its endpoints, and pay the rank of this label among all labels. It is not hard to see that, for graph families with good expansion, our analytic framework can be extended to this process as well. We plan to explore tight bounds for such graphical processes in future work.

References

- [1] Yehuda Afek, Haim Kaplan, Boris Korenfeld, Adam Morrison, and Robert E. Tarjan. Cbtree: A practical concurrent self-adjusting search tree. In *Proceedings of the 26th International Conference on Distributed Computing, DISC'12*, pages 1–15, Berlin, Heidelberg, 2012. Springer-Verlag.
- [2] Dan Alistarh, James Aspnes, Keren Censor-Hillel, Seth Gilbert, and Rachid Guerraoui. Tight bounds for asynchronous renaming. *J. ACM*, 61(3):18:1–18:51, June 2014.
- [3] Dan Alistarh, Justin Kopinsky, Jerry Li, and Nir Shavit. The spraylist: A scalable relaxed priority queue. In *20th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, PPOPP 2015*, San Francisco, CA, USA, 2015. ACM.
- [4] Hillel Avni, Nir Shavit, and Adi Suissa. Leaplist: Lessons learned in designing tm-supported range queries. In *Proceedings of the 2013 ACM Symposium on Principles of Distributed Computing, PODC '13*, pages 299–308, New York, NY, USA, 2013. ACM.
- [5] Yossi Azar, Andrei Z Broder, Anna R Karlin, and Eli Upfal. Balanced allocations. *SIAM journal on computing*, 29(1):180–200, 1999.
- [6] Dmitry Basin, Rui Fan, Idit Keidar, Ofer Kiselov, and Dmitri Perelman. CafÉ: Scalable task pools with adjustable fairness and contention. In *Proceedings of the 25th International Conference on Distributed Computing, DISC'11*, pages 475–488, Berlin, Heidelberg, 2011. Springer-Verlag.
- [7] Petra Berenbrink, Artur Czumaj, Angelika Steger, and Berthold Vöcking. Balanced allocations: The heavily loaded case. In *Proceedings of the Thirty-second Annual ACM Symposium on Theory of Computing, STOC '00*, pages 745–754, New York, NY, USA, 2000. ACM.
- [8] Petra Berenbrink, Tom Friedetzky, Zengjian Hu, and Russell Martin. On weighted balls-into-bins games. *Theor. Comput. Sci.*, 409(3):511–520, December 2008.
- [9] Trevor Brown, Faith Ellen, and Eric Ruppert. A general technique for non-blocking trees. *SIGPLAN Not.*, 49(8):329–342, February 2014.
- [10] Faith Ellen, Danny Hendler, and Nir Shavit. On the inherent sequentiality of concurrent objects. *SIAM J. Comput.*, 41(3):519–536, 2012.
- [11] Panagiota Fatourou and Nikolaos D. Kallimanis. Revisiting the combining synchronization technique. *SIGPLAN Not.*, 47(8):257–266, February 2012.
- [12] Panagiota Fatourou and Nikolaos D. Kallimanis. Highly-efficient wait-free synchronization. *Theory Comput. Syst.*, 55(3):475–520, 2014.
- [13] Mikhail Fomitchev and Eric Ruppert. Lock-free linked lists and skip lists. In *Proceedings of the 23rd annual ACM symposium on Principles of Distributed Computing (PODC'04)*, pages 50–59, New York, NY, USA, 2004. ACM Press.

- [14] Jakob Gruber, Jesper Larsson Träff, and Martin Wimmer. Brief announcement: Benchmarking concurrent priority queues. In *Proceedings of the 28th ACM Symposium on Parallelism in Algorithms and Architectures*, SPAA '16, pages 361–362, New York, NY, USA, 2016. ACM.
- [15] Andreas Haas, Michael Lippautz, Thomas A. Henzinger, Hannes Payer, Ana Sokolova, Christoph M. Kirsch, and Ali Sezgin. Distributed queues in shared memory: multicore performance and scalability through quantitative relaxation. In *Computing Frontiers Conference, CF'13, Ischia, Italy, May 14 - 16, 2013*, pages 17:1–17:9, 2013.
- [16] Tim L. Harris. A pragmatic implementation of non-blocking linked-lists. In *Proceedings of the International Conference on Distributed Computing (DISC)*, pages 300–314, 2001.
- [17] Thomas A. Henzinger, Christoph M. Kirsch, Hannes Payer, Ali Sezgin, and Ana Sokolova. Quantitative relaxation of concurrent data structures. *SIGPLAN Not.*, 48(1):317–328, January 2013.
- [18] Maurice Herlihy, Yossi Lev, Victor Luchangco, and Nir Shavit. A simple optimistic skiplist algorithm. In *Proceedings of the 14th international conference on Structural information and communication complexity*, SIROCCO'07, pages 124–138, Berlin, Heidelberg, 2007. Springer-Verlag. <http://dl.acm.org/citation.cfm?id=1760631.1760646>.
- [19] Maurice Herlihy, Nir Shavit, and Moran Tzafrir. Hopscotch hashing. In *Proceedings of the 22nd international symposium on Distributed Computing*, DISC '08, pages 350–364, Berlin, Heidelberg, 2008. Springer-Verlag. <http://dl.acm.org/citation.cfm?id=1432316>.
- [20] R. M. Karp and Y. Zhang. Parallel algorithms for backtrack search and branch-and-bound. *Journal of the ACM*, 40(3):765–789, 1993.
- [21] Doug Lea, 2007. <http://java.sun.com/javase/6/docs/api/java/util/concurrent/ConcurrentSkipListMap.html>.
- [22] Doug Lea, 2007. <http://g.oswego.edu/dl/jsr166/dist/docs/java/util/concurrent/ConcurrentHashMap.html>.
- [23] Justin J Levandoski, David B Lomet, and Sudipta Sengupta. The bw-tree: A b-tree for new hardware platforms. In *Data Engineering (ICDE), 2013 IEEE 29th International Conference on*, pages 302–313. IEEE, 2013.
- [24] Jonatan Lindén and Bengt Jonsson. A skiplist-based concurrent priority queue with minimal memory contention. In *International Conference On Principles Of Distributed Systems*, pages 206–220. Springer, 2013.
- [25] Maged M Michael. High performance dynamic lock-free hash tables and list-based sets. In *Proceedings of the fourteenth annual ACM symposium on Parallel algorithms and architectures*, pages 73–82. ACM, 2002.
- [26] Michael Mitzenmacher. The power of two choices in randomized load balancing. *IEEE Transactions on Parallel and Distributed Systems*, 12(10):1094–1104, 2001.
- [27] Adam Morrison. Scaling synchronization in multicore programs. *Commun. ACM*, 59(11):44–51, October 2016.
- [28] Adam Morrison and Yehuda Afek. Fast concurrent queues for x86 processors. *SIGPLAN Not.*, 48(8):103–112, February 2013.
- [29] Donald Nguyen, Andrew Lenharth, and Keshav Pingali. A lightweight infrastructure for graph analytics. In *Proceedings of the Twenty-Fourth ACM Symposium on Operating Systems Principles*, SOSP '13, pages 456–471, New York, NY, USA, 2013. ACM.
- [30] Yuval Peres, Kunal Talwar, and Udi Wieder. Graphical balanced allocations and the $1 + \beta$ -choice process. *Random Struct. Algorithms*, 47(4):760–775, December 2015.
- [31] Andrea W Richa, M Mitzenmacher, and R Sitaraman. The power of two random choices: A survey of techniques and results. *Combinatorial Optimization*, 9:255–304, 2001.

- [32] Hamza Rihani, Peter Sanders, and Roman Dementiev. Brief announcement: Multiqueues: Simple relaxed concurrent priority queues. In *Proceedings of the 27th ACM Symposium on Parallelism in Algorithms and Architectures*, SPAA '15, pages 80–82, New York, NY, USA, 2015. ACM.
- [33] P. Sanders. Randomized priority queues for fast parallel access. *Journal Parallel and Distributed Computing, Special Issue on Parallel and Distributed Data Structures*, 49:86–97, 1998.
- [34] Ori Shalev and Nir Shavit. Split-ordered lists: Lock-free extensible hash tables. *J. ACM*, 53:379–405, May 2006. <http://doi.acm.org/10.1145/1147954.1147958>.
- [35] Nir Shavit. Data structures in the multicore age. *Commun. ACM*, 54(3):76–84, 2011.
- [36] Nir Shavit and Itay Lotan. Skiplist-based concurrent priority queues. In *Parallel and Distributed Processing Symposium, 2000. IPDPS 2000. Proceedings. 14th International*, pages 263–268. IEEE, 2000.
- [37] Kunal Talwar and Udi Wieder. Balanced allocations: The weighted case. In *Proceedings of the Thirty-ninth Annual ACM Symposium on Theory of Computing*, STOC '07, pages 256–265, New York, NY, USA, 2007. ACM.
- [38] Martin Wimmer, Jakob Gruber, Jesper Larsson Träff, and Philippas Tsigas. The lock-free k-lsm relaxed priority queue. In *ACM SIGPLAN Notices*, volume 50, pages 277–278. ACM, 2015.

Appendix

The Appendix proceeds as follows. We give a reduction between classic power-of-two-choices processes and our sequential process, under round-robin insertions, in Section A. We prove that the single random choice process diverges in terms of rank cost in Section B. We cover the definition of distributional linearizability and relation to concurrent processes in Section C. We give complete proofs in Sections D, E, and F.

A Reduction to Two-Choices Process for Round-Robin Insertions

Process Description. For clarity, we re-define the process for round-robin insertions: we are given n queues, into which consecutively labeled elements are inserted in round-robin order: at the t th insertion step, we insert the element with label t into bin $t \pmod n$. To remove an element, we pick two queues at random, and remove the element of lower label on top of either queue. We again measure the cost of a removal as the rank of the removed label among labels still present in any of the queues.

Reduction. We will associate a “virtual” bin V_i with each queue i . Whenever an element is removed from the queue i , it gets immediately placed into the corresponding virtual bin. Notice that in every step we are removing the element of minimum label, which, by round-robin insertion, corresponds to the queue having been removed from less times than the other choice. Alternatively, we are inserting the removed element into the *less loaded* virtual bin, out of two random choices (implicitly breaking ties by bin ID). This is the classic definition of the two-choices process [5], which has been analyzed for the long-lived case in e.g. [7]. One can prove bounds on the rank of elements removed using the existing analyses of this process, e.g. [7, 30].

The Reduction Breaks for Random Insertions. Notice that the critical step in the above argument is the fact that, out of two random choices i and j , the less loaded virtual bin corresponds to the queue which has lower label on top. This does not hold in the random insertion case; for instance, it is possible for a bin to get two elements with consecutive labels, which breaks the above property. Upon closer examination, we can observe that in the random insertion case, the correlation between the event that queue i has lower label than queue j and the event that queue i has been removed from less times than queue j exists, but is too weak to allow a direct reduction.

B The Single Choice Process Diverges

In this section, we prove the following claim.

Theorem 6. *The expected maximum rank guarantee of the process which inserts and removes from uniform random queues in each time step t evolves as $\Omega(\sqrt{tn \log n})$, for $t = \Omega(n \log n)$.*

Proof. First, notice that we can apply the reduction in Section A to obtain that the maximum number of elements removed from a queue is the same as the maximum load of a bin in the classic long-lived random-removal process. (This holds irrespective of the insertion process, since we are performing uniform random removals.) The maximum load of a bin in the classic process is known to be $\Theta\left(t/n + \sqrt{\frac{t}{n} \log n}\right)$, with high probability, for $t > n \log n$ [30].

Let the queue with the most removals up to time t be Q . For any other queue Q_i with less removals, let $\ell(Q_i)$ be the difference between the number of elements removed from Q and the number of elements removed from Q_i . Notice that the expected number of elements removed from a queue up to t is t/n . Hence, it is easy to prove that, with probability at least $1/2$, there are at least $n/4$ queues which have had at most t/n elements removed. Hence, for any such queue Q_i , $\ell(Q_i) \geq \sqrt{\frac{t}{n} \log n}$. Hence, in expectation, $\sum_i \ell(Q_i) = \Omega(\sqrt{tn \log n})$.

We say that an element is *light* if its label is *smaller* than the label of the top element of Q . By symmetry, we know that each element counted in $\sum_i \ell(Q_i)$ is light with probability $\geq 1/2$. Putting everything together, we get that, in expectation, we have $\Omega(\sqrt{tn \log n})$ elements which are *light*, i.e. of smaller label than the top label of Q . This implies that the expected rank of Q is $\Omega(\sqrt{tn \log n})$, as claimed. \square

C Relationship to Concurrent Processes

On first glance, it might seem that a simple lock-based strategy should be linearizable to the sequential process we define in Section 3. For example, we could lock both queues that are to be examined (locking in order of their index to avoid deadlock and restarting the operation on failure) and declare the linearization point to be the point at which the second lock is grabbed. While this does linearize to *some* relaxed sequential process, it turns out that our upper bounds fail to hold for subtle reasons when concurrency is introduced.

Consider the extreme execution in which process 0 grabs the locks on some two queues, say Q_i, Q_j , and then hangs for a long time. Meanwhile, all the other processes complete many operations while process 0 holds these locks, and all operations will have to retry if they try to grab locks on Q_i or Q_j . In this case, many delete operations will be performed, none of which can delete from Q_i or Q_j . Such an execution could produce arbitrarily bad rank errors. We can formalize the property we would like a correct concurrent strategy to have as *distributional linearizability*:

Definition 2 (Distributional Linearizability). *A randomized data structure Q is distributionally linearizable to a sequential data structure S if for any parallel asynchronous execution, there exists a linearization of the operations of Q to operations of S so that the outputs of each operation of Q are distributionally equivalent to those S .*

The simple locking strategy fails to be distributionally linearizable due to, e.g., the counter-example execution above. We would like to ask if there is a distributionally linearizable strategy, yet there appears to be an inherent limitation. Consider three processes i, j, k performing DELETETMIN operations in lock-step. No matter what strategy is used, if some two processes, say i, j , try to delete from the same queue Q , at least one, say j , will necessarily be delayed. As a result i and k will finish their operations while j takes longer, causing i and k to be linearized before j (this can be forced by e.g. an insertion to Q while j is delayed). As a result, an additional constraint has been introduced which requires that the first two DELETETMIN operations to complete in this execution cannot have deleted from the same queue, a constraint which is not present for the sequential process.

In this way, the distribution of outputs of any concurrent process is affected by subtle timing issues in ways which seem hard to circumvent. We conjecture that there is no concurrent algorithm which is distributionally linearizable to the sequential process. That said, it may be possible that one can produce a simple strategy satisfying a weaker condition than distributional linearizability such as an analogous variant of sequential consistency or quiescent consistency.

D Omitted Proofs from Section 4.2

Here, we give a complete version of the potential argument.

Lemma 6.

$$\mathbb{E} [\Delta \Psi \mid y(t)] \leq \sum_{i=1}^n \hat{\alpha} \left((1 + \delta) \frac{1}{n} - p_i \right) \exp(-\alpha y_i(t)).$$

Proof. **Case 1:** If the bin is chosen, then the change is:

$$\begin{aligned} \Delta \Psi_i &= \Psi_i(t+1) - \Psi_i(t) = \exp \left(-\frac{\alpha}{n} (y_i(t) + \Delta_i \left(1 - \frac{1}{n} \right)) \right) - \exp \left(-\frac{\alpha}{n} y_i(t) \right) \\ &= \exp \left(-\frac{\alpha}{n} y_i(t) \right) \left(\exp \left(-\frac{\alpha}{n} \Delta_i \left(1 - \frac{1}{n} \right) \right) - 1 \right) \end{aligned}$$

Taking expectations with respect to the random choices made on insertion, we have

$$\begin{aligned} \mathbb{E}_i \left[\exp \left(-\Delta_i \frac{\alpha}{n} \left(1 - \frac{1}{n} \right) \right) \right] &= \frac{1}{1 + \frac{\alpha}{\pi_i n} \left(1 - \frac{1}{n} \right)} \\ &= 1 - \frac{\alpha}{\pi_i n} \left(1 - \frac{1}{n} \right) + \left(\frac{\alpha}{\pi_i n} \left(1 - \frac{1}{n} \right) \right)^2 - \dots \\ &\leq 1 - \frac{\alpha}{\pi_i n} \left(1 - \frac{1}{n} \right) + \left(\frac{\alpha}{\pi_i n} \left(1 - \frac{1}{n} \right) \right)^2. \end{aligned}$$

Case 2: If some other bin $j \neq i$ is chosen, then the change is:

$$\begin{aligned} \Delta \Psi_i &= \exp \left(-\frac{\alpha}{n} (y_i(t) - \Delta_j \frac{1}{n}) \right) - \exp \left(-\frac{\alpha}{n} y_i(t) \right) = \\ &= \exp \left(-\frac{\alpha}{n} y_i(t) \right) \left(\exp \left(\frac{\alpha \Delta_j}{n^2} \right) - 1 \right). \end{aligned}$$

Again taking expectations with respect to the random choices made on insertion, we have

$$\mathbb{E}_i \left[\exp \left(\Delta_i \frac{\alpha}{n^2} \right) \right] = \frac{\pi_j}{\pi_j - \frac{\alpha}{n^2}} = \frac{1}{1 - \frac{\alpha}{\pi_j n^2}} = 1 + \frac{\alpha}{\pi_j n^2} + \left(\frac{\alpha}{\pi_j n^2} \right)^2 + \dots \leq 1 + \frac{\alpha}{\pi_j n^2} + c \left(\frac{\alpha}{\pi_j n^2} \right)^2.$$

Therefore, we have that

$$\begin{aligned} \mathbb{E} [\Delta \Psi_i] &\leq \left(1 - \frac{\alpha}{n \pi_i} \left(1 - \frac{1}{n} \right) + \left(\frac{\alpha}{n \pi_i} \left(1 - \frac{1}{n} \right) \right)^2 \right) p_i - 1 + \sum_{j \neq i} \left(1 + \frac{\alpha}{n^2 \pi_j} + c \left(\frac{\alpha}{n^2 \pi_j} \right)^2 \right) p_j \\ &\leq \left(-\frac{\alpha}{n \pi_i} + c \left(\frac{\alpha}{n \pi_i} \right)^2 \right) p_i + \sum_{i=1}^n \left(\frac{\alpha}{n^2 \pi_j} + c \left(\frac{\alpha}{n^2 \pi_j} \right)^2 \right) p_j \\ &\leq -\alpha p_i (1 - \gamma - c\alpha(1 + \gamma)^2) p_i + \frac{\alpha}{n} (1 + \gamma + \alpha c(1 + \gamma)^2) \\ &\leq \hat{\alpha} \left((1 + \delta) \frac{1}{n} - p_i \right). \end{aligned}$$

□

Lemma 7. If $y_{n/4} \leq 0$, then we have that

$$\mathbb{E} [\Phi(t+1) \mid y(t)] \leq \left(1 - \frac{\hat{\alpha} \epsilon}{3n} \right) \Phi(t) + 1.$$

Proof. We start from the inequality

$$\mathbb{E} [\Delta \Phi | y(t)] \leq \sum_{i=1}^n \hat{\alpha} \left((1 + \delta) p_i - \frac{1}{n} \right) \exp(\alpha y_i(t)) \leq \hat{\alpha} (1 + \delta) \sum_{i=1}^n p_i \exp(\alpha y_i(t)) - \frac{\hat{\alpha}}{n} \Phi(t). \quad (12)$$

We will now focus on bounding the first term (without the constants). We can rewrite it as:

$$\sum_{i=1}^n p_i \exp(\alpha y_i(t)) = \sum_{i=1}^{n/4-1} p_i \exp(\alpha y_i(t)) + \sum_{i=n/4}^n p_i \exp(\alpha y_i(t)). \quad (13)$$

Since $y_{n/4} \leq 0$, the first term is upper bounded by 1. For the second term, notice that

$$\sum_{i=n/4}^n p_i \exp(\alpha y_i(t)) = \sum_{j=1}^{3n/4} p_{n-j+1} \exp(\alpha y_{n-j+1}(t)). \quad (14)$$

The p terms are non-decreasing in j , while the y terms are non-increasing in j . Further, note that

$$\sum_{j=1}^{3n/4} \exp(\alpha y_{n-j+1}(t)) \leq \Phi. \quad (15)$$

The whole sum is maximized when these non-increasing terms are equal. We therefore are looking to bound

$$\sum_{i=n/4}^n p_i \exp(\alpha y_i(t)) \leq \frac{4\Phi}{3n} \sum_{i=n/4}^n p_i. \quad (16)$$

Notice that

$$\sum_{i=n/4}^n p_i = 1 - \sum_{i=1}^{n/4-1} p_i \leq 1 - \left(\frac{1}{4} + \epsilon \right) = \frac{3}{4} - \epsilon. \quad (17)$$

Hence we have that

$$\begin{aligned} \mathbb{E} [\Delta \Phi | y(t)] &\leq \hat{\alpha} \left[(1 + \delta) \sum_{i=1}^n p_i \exp(\alpha y_i(t)) - \frac{1}{n} \Phi(t) \right] \\ &\leq \hat{\alpha} \left[(1 + \delta) \frac{4\Phi}{3n} \left(\frac{3}{4} - \epsilon \right) - \frac{1}{n} \Phi(t) + (1 + \delta) \right] \\ &\leq \hat{\alpha} \left[\frac{1 + \delta}{n} \left(1 - \frac{4\epsilon}{3} \right) \Phi(t) - \frac{1}{n} \Phi(t) + (1 + \delta) \right] \\ &= \hat{\alpha} \left[\frac{\Phi(t)}{n} \left[\left(1 - \frac{4\epsilon}{3} \right) (1 + \delta) - 1 \right] + (1 + \delta) \right] \\ &\leq -\hat{\alpha} \left[\frac{4\epsilon}{3} (1 + \delta) - \delta \right] \frac{\Phi(t)}{n} + 1 \leq -\hat{\alpha} (1 + \delta) \frac{\epsilon}{3} \frac{\Phi(t)}{n} + 1, \end{aligned}$$

where in the last step we have used the fact that $\delta \leq \epsilon$. Hence, we get that

$$\mathbb{E} [\Phi(t + 1) | y(t)] \leq \left(1 - \frac{\hat{\alpha}\epsilon}{3n} \right) \Phi(t) + 1.$$

□

Lemma 8. *If $y_{3n/4} \geq 0$, then we have that*

$$\mathbb{E}[\Psi(t+1) | y(t)] \leq \left(1 - \frac{\hat{\alpha}\epsilon}{3n}\right) \Psi + 1,$$

Proof. We start from the inequality

$$\begin{aligned} \mathbb{E}[\Delta\Psi | y(t)] &\leq \sum_{i=1}^n \hat{\alpha} \left((1+\delta)\frac{1}{n} - p_i \right) \exp(-\alpha y_i(t)) \\ &\leq \sum_{i=1}^{3n/4} \hat{\alpha} \left((1+\delta)\frac{1}{n} - p_i \right) \exp(-\alpha y_i(t)) + \sum_{i=3n/4+1}^n \hat{\alpha} \left((1+\delta)\frac{1}{n} - p_i \right) \exp(-\alpha y_i(t)). \end{aligned}$$

The second term is upper bounded as

$$\sum_{i=3n/4+1}^n \hat{\alpha} \left((1+\delta)\frac{1}{n} - p_i \right) \leq \hat{\alpha}(1+\delta)/4 - \hat{\alpha} \sum_{i=3n/4+1}^n p_i \leq 1.$$

We therefore now want to bound

$$\sum_{i=1}^{3n/4} \hat{\alpha} \left((1+\delta)\frac{1}{n} - p_i \right) \exp(-\alpha y_i(t)).$$

We again notice that the first factor is non-decreasing, whereas the second one is non-increasing. Hence, the sum is maximized when the non-decreasing terms are equal. At the same time, we have that

$$\sum_{i=1}^{3n/4} \exp(-\alpha y_i(t)) \leq \Psi.$$

Hence, it holds that

$$\begin{aligned} \sum_{i=1}^{3n/4} \hat{\alpha} \left((1+\delta)\frac{1}{n} - p_i \right) \exp(-\alpha y_i(t)) &\leq \frac{4\Psi}{3n} \sum_{i=1}^{3n/4} \hat{\alpha} \left((1+\delta)\frac{1}{n} - p_i \right) = \hat{\alpha} \frac{1+\delta}{n} \Psi - \hat{\alpha} \frac{4\Psi}{3n} \sum_{i=1}^{3n/4} p_i \\ &= \hat{\alpha} \frac{1+\delta}{n} \Psi - \hat{\alpha} \frac{1}{n} \Psi \left(1 + \frac{4\epsilon}{3} \right) = \frac{\hat{\alpha}}{n} \Psi \left((1+\delta) - \left(1 + \frac{4\epsilon}{3} \right) \right) \leq -\frac{\hat{\alpha}}{n} \frac{\epsilon}{3} \Psi, \end{aligned}$$

where in the last step we have used the fact that $\delta \leq \epsilon$. Therefore, we get that

$$\mathbb{E}[\Psi(t+1) | y(t)] \leq \left(1 - \frac{\hat{\alpha}\epsilon}{3n}\right) \Psi + 1,$$

as claimed. □

Lemma 9. *Given $\epsilon \in (0, 1)$ as defined, assume that $y_{n/4} > 0$ at t , and that $\mathbb{E}[\Delta\Phi] \geq -\frac{\epsilon\hat{\alpha}}{3n}\Phi(t)$. Then either $\Phi < \frac{\epsilon}{4}\Psi$, or $\Gamma = O(n)$.*

Proof. Fix $\lambda = 2/3 - 1/54$ for the rest of the proof. We can split the inequality in Lemma 1 as follows:

$$\mathbb{E}[\Delta\Phi | y(t)] \leq \sum_{i=1}^{\lambda n} \hat{\alpha} \left((1+\delta)p_i - \frac{1}{n} \right) \exp(\alpha y_i(t)) + \sum_{i=\lambda n+1}^n \hat{\alpha} \left((1+\delta)p_i - \frac{1}{n} \right) \exp(\alpha y_i(t)). \quad (18)$$

We bound each term separately. Since the probability terms are non-increasing and the exponential terms are non-decreasing, the first term is maximized when all p_i terms are equal. Since these probabilities are at most 1, we have

$$\sum_{i=1}^{\lambda n} \hat{\alpha} \left((1+\delta)p_i - \frac{1}{n} \right) \exp(\alpha y_i(t)) \leq \frac{\hat{\alpha}}{n} \left((1+\delta)\frac{1}{\lambda} - 1 \right) \Phi_{\leq \lambda n}. \quad (19)$$

The second term is maximized by noticing that the p_i factors are non-increasing, and are thus dominated by their value at λn . Noticing that we carefully picked λ such that

$$p_{\lambda n} \leq \frac{1}{n} - \frac{4\epsilon}{n},$$

we obtain, using the assumed inequality $\delta \leq \epsilon$, that

$$\sum_{i=\lambda n+1}^n \hat{\alpha} \left((1+\delta)p_i - \frac{1}{n} \right) \exp(\alpha y_i(t)) \leq \hat{\alpha} (\delta - 4\epsilon) \frac{\Phi_{>\lambda n}}{n} \leq -\frac{3\epsilon\hat{\alpha}}{n} \Phi_{>\lambda n}. \quad (20)$$

By the case assumption, we know that $\mathbb{E}[\Delta\Phi | y(t)] \geq -\frac{\hat{\alpha}\epsilon}{3n} \Phi(t)$. Combining the bounds (18), (19), and (20), this yields:

$$\frac{\hat{\alpha}}{n} \left((1+\delta)\frac{1}{\lambda} - 1 \right) \Phi_{\leq\lambda n} - \frac{3\epsilon\hat{\alpha}}{n} \Phi_{>\lambda n} \geq -\frac{\hat{\alpha}\epsilon}{3n} \Phi(t).$$

Substituting $\Phi_{>\lambda n} = \Phi - \Phi_{\leq\lambda n}$ yields:

$$(3\epsilon - \epsilon/3) \Phi \leq \left((1+\delta)\frac{1}{\lambda} - 1 + 3\epsilon \right) \Phi_{\leq\lambda n}.$$

For simplicity, we fix $C(\epsilon) = \frac{(1+\delta)\frac{1}{\lambda} - 1 + 3\epsilon}{3\epsilon - \epsilon/3} = O\left(\frac{1}{\epsilon}\right)$, to obtain

$$\Phi \leq C(\epsilon) \Phi_{\leq\lambda n}. \quad (21)$$

Let $B = \sum_{y_i > 0} y_i$. Since we are normalizing by the mean, it also holds that $B = \sum_{y_i < 0} (-y_i)$. Notice that

$$\Phi_{\leq\lambda n} \stackrel{y_i \text{ incr.}}{\leq} \lambda n \exp(\alpha y_{\lambda n}) \stackrel{y_i \text{ incr.}}{\leq} \lambda n \exp\left(\frac{\alpha B}{(1-\lambda)n}\right). \quad (22)$$

We put inequalities (21) and (22) together and get

$$\Phi(t) \leq \lambda n C(\epsilon) \exp\left(\frac{\alpha B}{(1-\lambda)n}\right), \quad (23)$$

Let us now lower bound the value of Ψ under these conditions. Since $y_{n/4} > 0$, all the costs below average must be in the first quarter of y . We can apply Jensen's inequality to the first $n/4$ terms of Ψ to get that

$$\Psi \geq \sum_{i=1}^{n/4} \exp(-\alpha y_i) \geq \frac{n}{4} \exp\left(-\alpha \frac{\sum_{i=1}^{n/4} y_i}{n/4}\right).$$

We now split the sum $\sum_{i=1}^{n/4} y_i$ into its positive part and its negative part. We know that the negative part is summing up to exactly $-B$, as it contains all the negative y_i 's and the total sum is 0. The positive part can be of size at most $B/4$, since it is maximized when there are exactly $n-1$ positive costs and they are all equal. Hence the sum of the first $n/4$ elements is at least $-3B/4$, which implies that the following bound holds:

$$\Psi \geq \frac{n}{4} \exp\left(-\alpha \frac{-3B/4}{n/4}\right) \geq \frac{n}{4} \exp\left(\alpha \frac{3B}{n}\right). \quad (24)$$

If $\Phi < \frac{\epsilon}{4} \Psi$, then there is nothing to prove. Otherwise, if $\Phi \geq \frac{\epsilon}{4} \Psi$, we get from (24) and (23) that

$$\frac{\epsilon}{4} \frac{n}{4} \exp\left(\alpha \frac{3B}{n}\right) \leq \frac{\epsilon}{4} \Psi \leq \Phi(t) \leq \lambda n C(\epsilon) \exp\left(\frac{\alpha B}{(1-\lambda)n}\right),$$

Therefore, we get that:

$$\exp\left(\alpha \frac{B}{n} \left(3 - \frac{1}{1-\lambda}\right)\right) \leq \frac{4\lambda}{\epsilon} C(\epsilon) = O\left(\frac{1}{\epsilon^2}\right).$$

Using the mundane fact that $3 - \frac{1}{1-\lambda} = 3/19$, we get that

$$\exp\left(\frac{\alpha B}{n}\right) \leq O\left(\frac{1}{\epsilon^{14}}\right). \quad (25)$$

To conclude, notice that (25) implies we can upper bound Γ in this case as:

$$\Gamma = \Phi + \Psi \leq \frac{4+\epsilon}{\epsilon} \Phi \leq \frac{4+\epsilon}{\epsilon} \lambda n C(\epsilon) \exp\left(\frac{\alpha B}{(1-\lambda)n}\right) \leq O\left(\frac{1}{\epsilon^{14/(1-\lambda)}}\right) \frac{4+\epsilon}{\epsilon} C(\epsilon) \lambda n = O(\text{poly}\left(\frac{1}{\epsilon}\right) n).$$

□

Lemma 10. *Given ϵ as above, assume that $y_{3n/4} < 0$, and that $\mathbb{E}[\Delta\Psi] \geq -\frac{\alpha\epsilon}{3n}\Psi$. Then either $\Psi > \frac{\epsilon}{4}\Phi$ or $\Gamma = O(n)$.*

Proof. Fix $\lambda = \lambda_2 = 1/3 + 1/54$. We start from the general bound on $\Delta\Psi$ from Lemma 6. We had that

$$\begin{aligned} \mathbb{E}[\Delta\Psi \mid y(t)] &\leq \sum_{i=1}^n \hat{\alpha} \left((1+\delta) \frac{1}{n} - p_i \right) \exp(-\alpha y_i(t)) \\ &= \sum_{i=1}^{\lambda n-1} \hat{\alpha} \left((1+\delta) \frac{1}{n} - p_i \right) \exp(-\alpha y_i(t)) + \sum_{i=\lambda n}^n \hat{\alpha} \left((1+\delta) \frac{1}{n} - p_i \right) \exp(-\alpha y_i(t)) \\ &\leq \frac{\hat{\alpha}}{n} ((1+\delta) - (1+4\epsilon)) \Psi_{<\lambda n} + \frac{\hat{\alpha}}{n} (1+\delta) \Psi_{\geq \lambda n} \\ &\leq -\frac{3\hat{\alpha}\epsilon}{n} \Psi_{<\lambda n} + \frac{\hat{\alpha}}{n} (1+\delta) \Psi_{\geq \lambda n}. \end{aligned}$$

Using the assumption, we have that

$$-\frac{\hat{\alpha}\epsilon}{3n} \Psi \leq \mathbb{E}[\Delta\Psi \mid y(t)] \leq -\frac{3\hat{\alpha}\epsilon}{n} (\Psi - \Psi_{\geq \lambda n}) + \frac{\hat{\alpha}}{n} (1+\delta) \Psi_{\geq \lambda n}.$$

We can re-write this as

$$\Psi \leq \Psi_{\geq \lambda n} \frac{1+\delta+3\epsilon}{3\epsilon-\epsilon/3} = C(\epsilon) \Psi_{\geq \lambda n}.$$

However, we also have that

$$\Psi_{\geq \lambda n} \leq (1-\lambda)n \exp(-\alpha y_{\lambda n}) \leq (1-\lambda)n \exp\left(\frac{\alpha B}{n\lambda}\right).$$

At the same time, since $y_{3n/4} < 0$, we have that

$$\Phi \geq \frac{n}{4} \exp\left(\frac{3\alpha B}{n}\right).$$

If $\Psi < \frac{\epsilon}{4}\Phi$, we can conclude. Let us examine the case where $\Psi \geq \frac{\epsilon}{4}\Phi$. Putting everything together, we get

$$\frac{\epsilon}{4} \frac{n}{4} \exp\left(\frac{3\alpha B}{n}\right) \leq \frac{\epsilon}{4} \Phi \leq \Psi \leq C(\epsilon) \Psi_{\geq \lambda n} \leq (1-\lambda)n C(\epsilon) \exp\left(\frac{\alpha B}{n\lambda}\right).$$

Alternatively,

$$\exp\left(\frac{\alpha B}{n} \left(3 - \frac{1}{\lambda}\right)\right) \leq \frac{16}{\epsilon} C(\epsilon)(1 - \lambda) = \frac{16}{\epsilon} (1 - \lambda) \frac{1 + \delta + 3\epsilon}{3\epsilon - \epsilon/3} = O\left(\frac{1}{\epsilon^2}\right).$$

Therefore,

$$\exp\left(\frac{\alpha B}{n}\right) \leq O\left(\frac{1}{\epsilon^{13}}\right).$$

To complete the argument, we bound:

$$\Gamma = \Psi + \Phi \leq \left(1 + \frac{4}{\epsilon}\right) \Psi \leq (1 - \lambda) C(\epsilon) \exp\left(\frac{\alpha B}{n\lambda}\right) n \leq O\left(\frac{n}{\epsilon^{22}}\right).$$

□

E Proof of Theorem 4

Before we prove Theorem 4 we need the following fact about exponential distributions:

Fact 2. Let X_1, X_2, \dots be independent and $X_i \sim \text{Exp}(\lambda)$ for all i . Let $Y_i = \sum_{k \leq i} X_k$. Fix any interval $I \subseteq [0, \infty)$ of length m . Then, $\#\{i : Y_i \in I\} \sim \text{Poi}(m\lambda)$.

Proof of Theorem 4. Let $I(t) = [w_{\min}(t), w_{\max}(t)]$, and let $L_j(t)$ be the number of elements in bin j in $I(t)$ at time t . Then, by memorylessness and Fact 2, for all bins j except the bin j_{\max} containing $w_{\max}(t)$, the number of labels in $I(t)$ in bin j before any deletions occur is distributed as $\text{Poi}(|I|/n)$. In particular, the expected number of elements after t rounds is bounded by $\mathbb{E}_{X \sim \text{Poi}(|I|/n)}[X] = |I|/n$. Moreover, for the bin containing $w_{\max}(t)$, by definition, at time t it contains exactly one element in I , namely, $w_{\max}(t)$. Hence, we have

$$\begin{aligned} \mathbb{E}[\max \text{rank}(t)] &\leq \mathbb{E}_{I(t)} \left[\mathbb{E} \left[\sum_{j=1}^n L_j(t) \middle| I(t) \right] \right] \\ &\leq \mathbb{E}_{I(t)} [1 + (n-1)|I|/n] \\ &= O\left(\frac{1}{\alpha} n(\log n + \log C)\right), \end{aligned}$$

by Lemma 4. □

F Omitted Proofs from Section 4.4

We first require the following technical lemma:

Lemma 11. For any interval $I = [a, b]$ which may depend on μ , we have $\mathbb{E}[\text{rank}(b) - \text{rank}(a) | \mu] \leq n(b - a + 1)$.

Proof. We first show that if I does not depend on μ , then $\mathbb{E}[\text{rank}(b) - \text{rank}(a) | \mu] \leq n(b - a)$. By Fact 2 we have

$$\begin{aligned} \mathbb{E}[\text{rank}(b) - \text{rank}(a) | \mu] &\leq \mathbb{E} \left[\sum_{j=1}^n \ell_{j, [a, b]} \middle| \mu \right] \\ &= \mathbb{E} \left[\sum_{j=1}^n \mathbb{E}_{X \sim \text{Poi}(b-a)}[X] \middle| \mu \right] \\ &= n(b - a). \end{aligned}$$

To conclude the proof, we now observe that μ depends on at most n elements, namely, those on top of the queues, and that if we remove those elements, then the remaining elements behave just as above. Thus, by conditioning on μ , we increase the rank by at most an additional factor of n . \square

Lemma 12. *For all t , we have $E[\text{rank}(\mu(t))] \leq O((A + 1/\alpha^2)n)$.*

Proof. For any bin j , we have

$$\begin{aligned} \mathbb{E}[\ell_{j,(-\infty, \mu]}(t)|\mu] &= \int_{-\infty}^{\mu} \mathbb{E}[\ell_{j,[x, \mu]}|\mu, w_j(t) = x] \cdot p_{j, \mu}(x) dx \\ &\stackrel{(a)}{\leq} \int_{-\infty}^{\mu} \mathbb{E}_{X \sim \text{Poi}(\mu-x)/n} [1 + X|\mu, w_j(t) = x] \cdot p_{j, \mu}(x) dx \\ &= \int_{-\infty}^{\mu} \frac{1}{n} (1 + \mu - x) \cdot p_{j, \mu}(x) dx = \frac{1}{n} \mathbb{E}[(1 + \mu - w_j(t))1_{w_j(t) \leq \mu}|\mu] \\ &\leq \frac{1}{n} (1 + \mathbb{E}[(\mu - w_j(t))1_{w_j(t) \leq \mu}|\mu]) . \end{aligned}$$

where (a) follows because after we've conditioned on the value of $w_j(t)$, the values of the remaining labels in bin j is independent of μ and we can apply Fact 2. Therefore, we have

$$\begin{aligned} \mathbb{E}[\text{rank}(\mu)] &\stackrel{(a)}{\leq} \mathbb{E}[\text{rank}(\mu - An)] + (A + 1)n = \mathbb{E} \left[\sum_{j=1}^n \mathbb{E}[\ell_{j,(-\infty, \mu - An]}(t)|\mu] \right] + (A + 1)n \\ &= \mathbb{E} \left[\sum_{j=1}^n \mathbb{E}[\ell_{j, (w_j(t), \mu - An]}(t)|\mu, w_j(t)] \right] + (A + 1)n \\ &\stackrel{(b)}{\leq} 1 + \frac{1}{n} \mathbb{E}_{\mu} \left[\sum_{j=1}^n (\mu - w_j(t))1_{w_j(t) \leq \mu - An} \right] + (A + 1)n \\ &\stackrel{(c)}{\leq} 1 + \frac{1}{n} \sum_{k=1}^{\infty} \mathbb{E}_{\mu} \left[\sum_{w_j(t) \in [\mu - (k+1+A)n, \mu - (k+A)n]} (\mu - w_j(t) + 1) \right] + (A + 1)n \\ &\leq 1 + \frac{1}{n} \sum_{k=1}^{\infty} (k + 2)n \mathbb{E}_{\mu} [b_{<-i}] + (A + 1)n \stackrel{(d)}{\leq} 1 + \sum_{k=1}^{\infty} (k + 2)n \exp(-\alpha k) + (A + 1)n \\ &\leq 1 + \frac{e^{\alpha}}{(e^{\alpha} - 1)^2} n + (A + 1)n , \end{aligned}$$

where (a) follows from Lemma 11, (b) follows from Fact 2, (c) follows from Lemma 11, and (d) follows from Lemma 5. Notice that for α small we have $\frac{e^{\alpha}}{(e^{\alpha} - 1)^2} = O(1/\alpha^2)$. Simplifying the above then yields the claimed bound. \square

F.1 Other Proofs

Lemma 13. *We have that $\Gamma(0) = O(n)$.*

Proof. In the initial state, the values $x_i(0)$ are independent exponentially distributed random variables with mean

$1/(n\pi_i)$. Now,

$$\begin{aligned}
\mathbb{E}[\Phi(0)] &= \mathbb{E} \left[\sum \exp \left(\alpha \left(\frac{w_i(0)}{n} - \mu(0) \right) \right) \right] \\
&= \sum \mathbb{E} \left[\exp \left(\alpha \left(x_i(0) - \frac{1}{n} \sum x_j \right) \right) \right] \\
&= \sum \mathbb{E} \left[\exp \left(\alpha x_i(0) \left(1 - \frac{1}{n} \right) \right) \prod_{j \neq i} \exp \left(-\frac{\alpha}{n} x_j(0) \right) \right] \\
&= \sum \mathbb{E} \left[\exp \left(\alpha x_i(0) \left(1 - \frac{1}{n} \right) \right) \right] \mathbb{E} \left[\prod_{j \neq i} \exp \left(-\frac{\alpha}{n} x_j(0) \right) \right],
\end{aligned}$$

where the last line follows from the (pairwise) independence of the x_i . All terms in the right hand side are now exactly moment generating functions of the w_i evaluated at some point. Since the moment generating function of an exponential with parameter λ evaluated at t is well known to be $\lambda/(\lambda - t)$, we can compute:

$$\begin{aligned}
\mathbb{E}[\Phi(0)] &= \sum \frac{\frac{1}{n\pi_i}}{\frac{1}{n\pi_i} - \alpha(1 - \frac{1}{n})} \prod_{j \neq i} \frac{\frac{1}{n\pi_j}}{\frac{1}{n\pi_j} + \frac{\alpha}{n}} \\
&\leq \sum \frac{1 + \gamma}{1 - \gamma - \alpha} \prod_{j \neq i} \frac{1}{1 + \alpha\pi_j} \\
&\approx \sum \frac{1 + \gamma}{1 - \gamma - \alpha} \left(1 - \sum_{j \neq i} \alpha\pi_j \right) \\
&= O(n).
\end{aligned}$$

The argument for $\Psi(0)$ is symmetric, replacing α with $-\alpha$.

□