

Online Multi-level aggregation with delays and stochastic arrivals

Mathieu Mari*, Michał Pałowski†, Runtian Ren‡, Piotr Sankowski§

Abstract

This paper presents a new research direction for online Multi-Level Aggregation (MLA) with delays. In this problem, we are given an edge-weighted rooted tree T , and we have to serve a sequence of requests arriving at its vertices in an online manner. Each request r is characterized by two parameters: its arrival time $t(r)$ and location $l(r)$ (a vertex). Once a request r arrives, we can either serve it immediately or postpone this action until any time $t > t(r)$. We can serve several pending requests at the same time, and the service cost of a service corresponds to the weight of the subtree that contains all the requests served and the root of T . Postponing the service of a request r to time $t > t(r)$ generates an additional delay cost of $t - t(r)$. The goal is to serve all requests in an online manner such that the total cost (i.e., the total sum of service and delay costs) is minimized. The current best algorithm for this problem achieves a competitive ratio of $O(d^2)$ (Azar and Touitou, FOCS'19), where d denotes the depth of the tree.

The MLA problem is a generalization of several well-studied problems, including TCP Acknowledgment (depth 1), Joint Replenishment (depth 2) and multi-level message aggregation (arbitrary depth). Although it appeared implicitly in many previous papers, it has been formalized by Bienkowski et al. (ESA'16).

Here, we consider a stochastic version of MLA where the requests follow a Poisson arrival process. We present a deterministic online algorithm which achieves a constant ratio of expectations, meaning that the ratio between the expected costs of the solution generated by our algorithm and the optimal offline solution is bounded by a constant. Our algorithm is obtained by carefully combining two strategies. In the first one, we plan periodic oblivious visits to the subset of frequent vertices, whereas in the second one, we greedily serve the pending requests in the remaining vertices. This problem is complex enough to demonstrate a very rare phenomenon that “single-minded” or “sample-average” strategies are not enough in stochastic optimization.

1 Introduction

Imagine the manager of a biscuit factory needs to deal with the issue of delivering products from the factory to the convenience stores. Once some products, say chocolate waffle, is in shortage at some store, then the store employee will inform the factory for replenishment. From the factory's perspective, each time a service is created to deliver the products, a truck has to travel from the factory to go to each store, and then come back to the factory. A cost proportional to the total traveling distance has to be paid for this service. For the purpose of saving delivery cost, it is

*LIRMM, University of Montpellier, Montpellier, France. mari.mathieu.06@gmail.com

†MIMUW, University of Warsaw and IDEAS NCBR, Warsaw, Poland. michal.pawlowski196@gmail.com

‡University of Wrocław and IDEAS NCBR, Wrocław, Poland. renruntian@gmail.com

§MIMUW, University of Warsaw, IDEAS NCBR, MIM Solutions, Warsaw, Poland. piotr.sankowski@gmail.com

beneficial to accumulate the replenishment requests from many stores and then deliver the ordered products altogether in one service. However, this accumulated delay of delivering products may cause the stores unsatisfied and the complaints will have negative influence on future contracts between the stores and the factory. Typically, for each request ordered from a store, the time gap between ordering the products and receiving the products, is known as delay cost (of this request). The goal of the factory manager, is to plan the delivery service schedule in an online manner, such that the total service cost and the total delay cost is minimized.

The above is an example of an online problem called Multi-level Aggregation (MLA) with linear delays. Formally, the input is an edge-weighted rooted tree T and a sequence of requests, with each request r specified by an arrival time $t(r)$ and a location at a particular vertex. Once a request r arrives, its service does not have to be processed immediately, but can be delayed to any time $t \geq t(r)$ at a delay cost of $t - t(r)$. The benefit of delaying requests is that several requests can be served together to save some service cost: to serve any set of requests R at time t , a subtree T' containing the tree root and all locations of requests R needs to be bought at a service cost equal to the total weight of edges in T' . The goal of MLA is to serve all requests in an online manner such that the total cost (i.e., the total service cost plus the total delay cost) is minimized.

The MLA problem is first formally introduced by Bienkowski et al. [16]. Due to many real-life applications ranging from logistic, supply chain management, data transmission in sensor network, this MLA problem has recently drawn considerable attentions [22, 16, 27, 13]. Besides, two classic problems, TCP-acknowledgment (also known as lot-sizing problem, from operation research community) and Joint Replenishment (JRP), as special cases of MLA with tree depths of 1 and 2 respectively, are studied by extensive previous works [32, 45, 62, 1, 47, 28, 21, 3, 59, 20]. Particularly for MLA, the state-of-the-art is as follows:

- the current best online algorithm, proposed by Azar and Touitou [13, Theorem IV.2], achieves a competitive ratio of $O(d^2)$, where d denotes the depth of the given tree (i.e., the maximum number of edges from any leaf vertex to the tree root);
- no online algorithm can achieve a competitive ratio less than 4 [16, Theorem 6.3] — this is the current best lower bound, even restricted to the case when the given tree is a path, and the root is an endpoint of the path.

Obviously, there is a huge gap between the upper bound and the lower bound on the competitiveness of MLA. Closing the gap remains an interesting open question.

In fact, it is often too pessimistic to assume no stochastic information on the input is available in practice — again, consider our delivery example. The factory knows all the historical orders and can estimate the request frequencies from the stores of all locations. It is reasonable to assume that the requests follow some stochastic distribution. Therefore, the following question is natural: *if stochastic information on the input is available, can we devise online algorithms for MLA with better performance guarantees?*

In this paper, we provide an affirmative answer to this question. We study a stochastic online version of MLA, assuming that the requests arrive following a Poisson arrival process. More precisely, the waiting time between any two consecutive requests arriving at the same vertex u follows an exponential distribution $\text{Exp}(\lambda(u))$ with parameter $\lambda(u)$. In this model, the goal is to minimize the expected cost produced by an algorithm ALG for a random input sequence generated in a long time interval $[0, \tau]$. In order to evaluate the performance of our algorithms on stochastic inputs,

we use the *ratio of expectations* (RoE), that corresponds to the ratio of the expected cost of the algorithm to the expected cost of the optimal offline solution (see Definition 2.8).

Our contribution. We prove that the performance guarantee obtained in the Poisson arrival model is significantly better compared with the current best competitiveness obtained in the adversarial model. More specifically, we propose a non-trivial deterministic online algorithm which achieves a constant ratio of expectations.

Theorem 1.1. *For MLA with linear delays in the Poisson arrival model, there exists a deterministic online algorithm which achieves a constant ratio of expectations.*

Our algorithm is obtained by synergistically merging two carefully crafted strategies. The first strategy incorporates periodic oblivious visits to a subset of frequently accessed vertices, while the second strategy employs a proactive, greedy approach to handle pending requests in the remaining vertices. The complexity of this problem unveils a rare phenomenon — the inadequacy of “single-minded” or “sample-average” strategies in stochastic optimization. In this paper we not only address this challenge but also point on to further complex problems that require similar approach in stochastic environments. We stress that it is open to obtain improved results for stochastic cases of facility location with delays [13] or online service with delays [9].

Previous works. The MLA problem has been only studied in the adversarial model. Bienkowski et al. [16] introduced a general version of MLA, assuming that the cost of delaying a request r by a duration t is $f_r(t)$. Here $f_r(\cdot)$ denotes the delay cost function of r and it only needs to be non-decreasing and satisfy $f_r(0) = 0$. They proposed an $O(d^4 2^d)$ -competitive online algorithm for this general delay cost version problem, where d denotes the tree depth [16, Theorem 4.2]. Later, the competitive ratio is further improved to $O(d^2)$ by Azar and Touitou [13, Theorem IV.2] (for this general delay cost version). However, no matching lower bound has been found for the delay cost version of MLA — the current best lower bound on MLA (with delays) is only 4 [16, Theorem 6.3], restricted to a path case with linear delays. Thus far, no previous work has studied MLA in the stochastic input model.

Organization. In Section 2, we give all the necessary notations and preliminaries. In Section 3, we study a special single-edge tree instance as a warm-up. We show that there are two different situations, one heavy case and one light case, and to achieve constant RoE, the ideas for the two cases are different. In Section 4, we give an overview of our deterministic online algorithm (Theorem 1.1). This algorithm is the combination of two different strategies for two different types of instances. In Section 5, we study the first type, called light instances, that are a generalization of light single-edge trees. In Section 6, we study the other type called heavy instances as a generalization of heavy single-edge trees. In Section 7, we prove Theorem 1.1. We finish the paper by detailing in Section 8 all other related works, and by discussing some future directions in Section 9.

2 Notations and Preliminaries

Weighted tree. Consider an edge-weighted tree T rooted at vertex $\gamma(T)$. We refer to its vertex set by $V(T)$ and its edge set by $E(T)$. When the context is clear, we denote the root vertex, vertex set, and edge set by γ , V , and E , respectively. We assume that each edge $e \in E$ has a positive

weight w_e . For any vertex $u \in V$, except for the root vertex γ , we denote its *parent vertex* as $\text{par}(u) \in V$, and $e_u = (u, \text{par}(u))$ as the edge connecting u and its parent. We also define T_u as the subtree of T rooted at vertex u . In addition to the edge weights, we use the term *vertex weight* to refer to $w_u := w(e_u)$, where $u \in V$ and $u \neq \gamma$. Given any two vertices $u, v \in V(T)$, we denote the path length from u to v in T by $d_T(u, v)$, i.e., it is the total weight of the edges along this path.¹ Finally, we use $T[U]$ to denote the forest induced by vertices of $U \subseteq V(T)$ in T .

Problem description. An MLA *instance* is characterized by a tuple (T, σ) , where T is a weighted tree rooted at γ and σ is a sequence of *requests*. Each request r is described by a tuple $(t(r), l(r))$ where $t(r) \in \mathbb{R}^+$ denotes r 's *arrival time* and $l(r) \in V(T)$ denotes r 's *location*. Thus, denoting by m the number of requests, we can rewrite $\sigma := (r_1, \dots, r_m)$ with the requests sorted in increasing order of their arrival times, i.e., $t(r_1) \leq t(r_2) \leq \dots \leq t(r_m)$. Given a sequence of requests σ , a *service* $s = (t(s), R(s))$ is characterized by the *service time* $t(s)$ and the set of requests $R(s) \subseteq \sigma$ it serves. A *schedule* S for σ is a sequence of services. We call schedule S *valid* for σ if each request $r \in \sigma$ is assigned a service $s \in S$ that does not precede r 's arrival. In other words, a valid S for σ satisfies (i) $\forall s \in S \forall r \in R(s) t(r) \geq t(s)$; (ii) $\{R(s) : s \in S\}$ forms a partition of σ . Given any MLA instance (T, σ) , an MLA algorithm ALG needs to produce a valid schedule S to serve all the requests σ . Particularly for an online MLA algorithm ALG, at any time t , the decision to create a service to serve a set of pending request(s) cannot depend on the requests arriving after time t .

For each request $r \in \sigma$, let $S(r)$ denote the service in S which serves r , i.e., for each $s \in S$, $S(r) = s$ if and only if $r \in R(s)$. Given a sequence of requests σ and a valid schedule S , the *delay cost* for a request $r \in \sigma$ is defined as $\text{delay}(r) := t(S(r)) - t(r)$. Using this notion, we define the *delay cost* for a service $s \in S$ and the *delay cost* for the schedule S as

$$\text{delay}(s) := \sum_{r \in R(s)} \text{delay}(r) \quad \text{and} \quad \text{delay}(S) := \sum_{s \in S} \text{delay}(s).$$

Besides, given any request $r \in \sigma$, if it is pending at time $t > t(r)$, let $\text{delay}(r, t) = t - t(r)$ denote the delay cost of r at this moment.

The *weight* (also called *service cost*) of a service $s \in S$, denoted by $\text{weight}(s, T)$, is defined as the weight of the minimal subtree of T that contains root γ and all locations of requests $R(s)$ served by s . The *weight* (or *service cost*) of a schedule S is defined as $\text{weight}(S, T) := \sum_{s \in S} \text{weight}(s, T)$. To compute the *cost* of a service s , we sum its delay cost and weight, i.e.,

$$\text{cost}(s, T) := \text{delay}(s) + \text{weight}(s, T).$$

Similarly, we define the *cost* (or *total cost*) of a schedule S for σ as

$$\text{cost}(S, T) := \text{delay}(S) + \text{weight}(S, T).$$

When the context is clear, we simply write $\text{cost}(S) = \text{cost}(S, T)$. Moreover, given an MLA instance (T, σ) , let $\text{ALG}(\sigma)$ denote the schedule of algorithm ALG for σ and let $\text{OPT}(\sigma)$ denote the optimal schedule for σ with minimum total cost. Note that without loss of generality, we can assume that no request in σ arrives at the tree root γ since such a request can be served immediately at its arrival with zero cost.

¹When the context is clear, we simply write $d(u, v)$ instead of $d_T(u, v)$. Furthermore, we stress that the order of vertices in this notation is not arbitrary — the second vertex (v) is always an ancestor of the first one (u).

Poisson arrival model. In this paper, instead of using an adversarial model, we assume that the requests arrive according to some stochastic process. A *stochastic instance* that we work with is characterized by a tuple (T, λ) , where T denotes an edge-weighted rooted tree, and $\lambda : V(T) \rightarrow \mathbb{R}_+$ is a function that assigns each vertex $u \in V(T)$ an *arrival rate* $\lambda(u) \geq 0$.² Formally, such a tuple defines the following process.

Definition 2.1 (Poisson arrival model). *Given any stochastic MLA instance (T, λ) and any value $\tau > 0$, we say that a (random) requests sequence σ follows a Poisson arrival model over time interval $[0, \tau]$, if (i) for each vertex $u \in V(T)$ with $\lambda(u) > 0$ the waiting time between any two consecutive requests arriving at u follows an exponential distribution with parameter $\lambda(u)$,³ (ii) variables representing waiting times are mutually independent; (iii) all the requests in σ arrive within time interval $[0, \tau]$. We denote this fact by writing $\sigma \sim (T, \lambda)^\tau$.*

Given any subtree T' of T , we use both $\lambda|_{T'}$ and $\lambda|_{V(T')}$ to denote the arrival rates restricted to the vertices of T' . Similarly, given a random sequence of requests $\sigma \sim (T, \lambda)^\tau$, we use $\sigma|_{T'} \subseteq \sigma$ and $\sigma|_I \subseteq \sigma$ for $I \subseteq [0, \tau]$ to denote the sequences of all requests in σ that arrive inside the subtree T' and within the time interval I , respectively. Note that the above Poisson arrival model satisfies the following properties (see Appendix A for the formal proof).

Proposition 2.2. *Given a subtree T' of T : (i) for any $\tau > 0$ and any sequence $\sigma \sim (T, \lambda)^\tau$, $\sigma|_{T'}$ follows the Poisson arrival model over the MLA instance restricted to T' , i.e. $\sigma|_{T'} \sim (T', \lambda|_{T'})^\tau$, (ii) the process determining arrivals inside T' is independent of the requests arriving in $T \setminus T'$.*

Proposition 2.3. *Given any stochastic MLA instance (T, λ) , $\tau = \sum_{i=1}^k \tau_i$ with $\tau_i > 0$ for each $i \in [k]$,⁴ and a family of random sequences $\{\sigma_i \sim (T, \lambda)^{\tau_i} : i \in [k]\}$, we merge them into one sequence defined over a τ -length time interval by postponing arrivals of requests in σ_i by $\sum_{j=1}^{i-1} \tau_j$ for all $i \in [k]$. Due to the memoryless property of exponential variables, this process results in a sequence σ that follows the Poisson arrival model over $[0, \tau]$, i.e., $\sigma \sim (T, \lambda)^\tau$.*

Intuitively, the first proposition gives us the freedom to select a sub-instance of the problem and focus on the requests arriving in a subtree $T' \subseteq T$. The second one allows us to split the time horizon into smaller intervals and work with shorter request sequences. The most important fact, though, is that both operations preserve the arrival model and are independent of the remaining part of the initial request sequence. Here, we also stress that in the following sections, we use the notation of $\lambda(T') := \sum_{v \in T'} \lambda(v)$ to denote the arrival rate for a given subtree $T' \subseteq T$.

Another equivalent characteristic of the Poisson arrival model gives us a more “centralized” perspective on how the request sequences are generated (see Appendix A for the formal proof).

Proposition 2.4. *Given any stochastic MLA instance (T, λ) and a random sequence of requests $\sigma \sim (T, \lambda)^\tau$, we have (i) the waiting time between any two consecutive requests in σ follows an exponential distribution with parameter $\lambda(T)$; (ii) for each vertex $u \in T$ and each request $r \in \sigma$ the probability of r being located at u equals $\lambda(u)/\lambda(T)$.*

²Without loss of generality we assume $\lambda(\gamma(T)) = 0$, i.e., no request arrives at the tree root.

³For the first request r arriving at u , we require that the waiting time from 0 to $t(r)$ follows this distribution $\text{Exp}(\lambda(u))$. Similarly, if we look at the last request r' arriving at u and let $W_{r'} \sim \text{Exp}(\lambda(u))$ denote the variable determining its waiting time, we require that $\tau - t(r') < W_{r'}$.

⁴For simplicity, we use $[k]$ to denote $\{1, 2, \dots, k\}$ everywhere in this paper.

In the following, we introduce three more properties of the Poisson arrival model. To simplify their statements, from now on we denote the random variable representing the number of requests in sequence $\sigma \sim (T, \boldsymbol{\lambda})^\tau$ by $N(\sigma)$. The first property describes the expected value of $N(\sigma)$ for a fixed time horizon τ . The second one describes our model's behavior under the assumption that we are given the value of $N(\sigma)$. Finally, the third one presents the value of the expected waiting time generated by all the requests arriving before a fixed time horizon. All the proofs can be found in [61].⁵ However, for completeness, we also include them in Appendix A.

Proposition 2.5. *Given any stochastic MLA instance $(T, \boldsymbol{\lambda})$ and a random sequence of requests $\sigma \sim (T, \boldsymbol{\lambda})^\tau$, it holds that (i) $N(\sigma) \sim \text{Pois}(\boldsymbol{\lambda}(T) \cdot \tau)$; (ii) $\mathbb{E}[N(\sigma) \mid \sigma \sim (T, \boldsymbol{\lambda})^\tau] = \boldsymbol{\lambda}(T) \cdot \tau$; (iii) if $\boldsymbol{\lambda}(T) \cdot \tau \geq 1$, then $\mathbb{P}(N(\sigma) \geq \mathbb{E}[N(\sigma)]) \geq 1/2$.*

Proposition 2.6. *Given that n requests arrive during time interval $[0, \tau]$ according to Poisson arrival model, the n arrival times (in sequence) have the same distribution as the order statistics corresponding to n independent random variables uniformly distributed over $[0, \tau]$.*

Proposition 2.7. *Given any stochastic MLA instance $(T, \boldsymbol{\lambda})$ and $\sigma \sim (T, \boldsymbol{\lambda})^\tau$, the expected delay cost generated by all the requests arriving before τ is equal to*

$$\mathbb{E} \left[\sum_{i=1}^{N(\sigma)} (\tau - t(r_i)) \mid \sigma \sim (T, \boldsymbol{\lambda})^\tau \right] = \frac{\tau}{2} \cdot \mathbb{E}[N(\sigma) \mid \sigma \sim (T, \boldsymbol{\lambda})^\tau] = \frac{1}{2} \cdot \boldsymbol{\lambda}(T) \cdot \tau^2.$$

Benchmark description. For an online algorithm ALG that takes as input a random sequence $\sigma \sim (T, \boldsymbol{\lambda})^\tau$, let $\mathbb{E}_\sigma[\text{cost}(\text{ALG}(\sigma), T)]$ denotes the expected cost of the schedule it generates. To measure the performance of ALG in this stochastic version of MLA, we use the *ratio of expectations*.

Definition 2.8 (ratio of expectations). *An online MLA algorithm ALG achieves a ratio of expectations (RoE) $C \geq 1$, if for all stochastic MLA instances $(T, \boldsymbol{\lambda})$ we have*

$$\frac{\mathbb{E}[\text{cost}(\text{ALG}(\sigma), T) \mid \sigma \sim (T, \boldsymbol{\lambda})^\tau]}{\lim_{\tau \rightarrow \infty} \mathbb{E}[\text{cost}(\text{OPT}(\sigma), T) \mid \sigma \sim (T, \boldsymbol{\lambda})^\tau]} \leq C.$$

3 Warm-up: single edge instances

In this section, we study the case of a single-edge tree in the stochastic model. Thus, throughout this section, we fix a tree T that consists of a single edge $e = (u, \gamma)$ of weight $w > 0$, and denote the arrival rate of u by $\lambda > 0$. In such a setting, the problem of finding the optimal schedule to serve the requests arriving at vertex u is known as *TCP acknowledgment* (here, we consider the stochastic model). It is worth mentioning that in the adversarial setting, a 2-competitive deterministic and a $(1 - e^{-1})^{-1}$ -competitive randomized algorithms are known for this problem [32, 45].⁶

Let us stress that the goal of this section is not to improve the best-known competitive ratio for a single-edge case, but to illustrate the efficiency of two opposite strategies, and introduce the important concepts of this paper. The first strategy called the *instant strategy*, is to serve each

⁵The proof of the first proposition follows from Proposition 2.2.1, and Definition 2.1.1, while the proofs of the remaining two facts can be found in Theorem 2.3.1 and Example 2.3(A).

⁶To our best extent, no previous work studied this problem in the Poisson arrival model from a theoretical perspective, i.e., evaluating the performance of the algorithms using the ratio of expectations.

request as soon as it arrives. Intuitively, this approach is efficient when the requests are not so frequent, so that on average, the cost of delaying a request to the arrival time of the next request, is enough to compensate the service cost. The second strategy, called the *periodic approach* is meant to work in the opposite case where requests are frequent enough so that it is worth grouping several of them for the same service. In this way, the weight cost of a service can be shared between the requests served. Assuming that requests follow some stochastic assumptions, it makes sense to enforce that services are ordered at regular time intervals, where the time between any two consecutive services is a fixed number p , which depends only on the instance's parameters.

There are two challenges here. First, when should we use each strategy? Second, what should be the value of p that optimizes the performance of the periodic strategy? For the first question, we show that this depends on the value of $\pi := w\lambda$ that we call the *heaviness* of the instance. More precisely, we show that if $\pi > 1$, i.e., the instance is *heavy*, the periodic strategy is more efficient. On the other hand, if $\pi \leq 1$, then the instance is *light*, and the instant strategy is essentially better. For the second question, we show that the right value for the period, up to a constant in the ratio of expectations, is $p = \sqrt{2w/\lambda}$. Without loss of generality, in what follows we assume that the time horizon τ is always a multiple of the period chosen, which simplifies the calculation and does not affect the ratio of expectations.

Lemma 3.1. *Given a stochastic instance where the tree consists of a single edge of weight $w > 0$ and the leaf has an arrival rate $\lambda > 0$, let $\pi = w\lambda$ and let σ be a random sequence of requests of duration τ , for some $\tau > 0$. It holds that*

- (i) *the instant strategy on σ has the expected cost of $\tau \cdot \pi$;*
- (ii) *the periodic strategy on σ , with period $p = \sqrt{2w/\lambda}$, has the expected cost of $\tau \cdot \sqrt{2\pi}$.*

Proof. Notice that the instant strategy incurs an expected cost equal to the expected number of requests arriving within the time horizon τ times the cost of serving one. By Proposition 2.5, we have that on average $\lambda \cdot \tau$ requests arrive within the time interval $[0, \tau]$. Thus, since the cost of serving one equals w , the total expected cost is $\lambda \cdot \tau \cdot w = \tau \cdot \pi$.

Similarly, for the periodic strategy, we know that within each period $p = \sqrt{2w/\lambda}$, we generate the expected delay cost of $1/2 \cdot \lambda \cdot p^2 = w$ (Proposition 2.7). The service cost we pay at the end of each period equals w as well. Thus, the total expected cost within $[0, \tau]$ is equal to $\tau/p \cdot 2w = \tau \cdot \sqrt{2\lambda w} = \tau \cdot \sqrt{2\pi}$, which ends the proof. \square

We now compare these expected costs with the expected cost of the optimal offline schedule. The bounds obtained imply that the instant strategy has constant RoE when $\pi \leq 1$, and the periodic strategy (with $p = \sqrt{2w/\lambda}$) has a constant RoE when $\pi > 1$.

Lemma 3.2. *Given a stochastic instance where the tree consists of a single edge of weight $w > 0$ and the leaf has an arrival rate $\lambda > 0$, let $\pi = w\lambda$ and let σ be a random sequence of requests of duration τ , for some $\tau > 0$. The lower bounds for the optimal offline schedule for σ are as follows*

- (i) *if $\pi \leq 1$, then it has an expected cost of at least $1/2 \cdot (1 - e^{-1}) \cdot \tau \cdot \pi$;*
- (ii) *if $\pi > 1$, then it has an expected cost of at least $3/8\sqrt{2} \cdot \tau \cdot \sqrt{\pi}$.*

In the following subsection, we prove Lemma 3.2.

3.1 Lower bounding OPT

Let $\sigma \sim (T, \lambda)^\tau$ be a random sequence of requests defined for the given single-edge instance and some time horizon τ . In this instance, the edge has weight $w > 0$ and the vertex has arrival rate $\lambda > 0$. Now we lower bound the expected cost of the optimal offline algorithm OPT on σ .

The main idea is to partition the initial time horizon $[0, \tau]$ into a collection of shorter intervals $\{I_1, I_2, \dots, I_k\}$ of length p each, for some value p that is defined later. We denote $\sigma_i := \sigma|_{I_i}$ for $i \in [k]$. From Proposition 2.3, we know that all σ_i are independent and follow the same Poisson arrival model $(T, \lambda)^p$. Thus, we should be able to analyze them separately and combine the results to get the estimation of the total cost incurred by OPT over the initial sequence σ .

Let $D(\sigma_1)$ denote the total delay cost of σ_1 at time p when no services are issued during $[0, p]$. Note that OPT either serves some requests during $[0, p]$ and incurs the service cost of at least w , or issues no services during $[0, p]$ and pays the delay cost of $D(\sigma_1)$. Thus, the total cost of OPT within $[0, p]$ is at least $\min(w, D(\sigma_1))$. Since $\{\sigma_i : i \in [\tau/p]\}$ are i.i.d, we deduce the following bound:

$$\mathbb{E}[\text{cost}(\text{OPT}(\sigma))] \geq \frac{\tau}{p} \cdot \mathbb{E}[\min(w, D(\sigma_1)) \mid \sigma_1 \sim (T, \lambda)^p]. \quad (1)$$

Using Proposition 2.6, we can partition the right-hand side of the inequality further. Indeed, we know that when conditioned on the number of requests $N(\sigma_1) = n$, for some $n \in \mathbb{N}$, the arrival times in σ_1 follow the same distribution as the order statistics corresponding to n independent random variables uniformly distributed over $[0, p]$. Let us denote these variables by A_1, A_2, \dots, A_n . Consider any request $r_j \in \sigma_1$ that arrived at time $t(r_j) = A_j$ and is still pending at time p . It is easy to notice that the variable U_j representing the delay cost r_j incurred until p also follows a uniform distribution $[0, p]$ as it holds that $U_j = p - A_j$, i.e., $U_j \sim \mathcal{U}(p)$. Thus, when we condition on $N(\sigma_i) = n$, we can write $D(\sigma_1) = \sum_{j=1}^n U_j$ as the sum of n uniform variables representing the waiting times. This allows us to rewrite the right-hand side of (1) as

$$\mathbb{E}[\min(w, D(\sigma_1))] = \sum_{n=1}^{\infty} \mathbb{P}[N(\sigma_1) = n] \cdot \mathbb{E} \left[\min \left(w, \sum_{j=1}^n U_j \right) \mid N(\sigma_i) = n \right], \quad (2)$$

where the expectation on the right side is taken over all sequences $\{U_j : j \in \mathbb{N}\}$ of independent uniform variables in $[0, p]$. We now estimate these expectations.

Claim 3.3. *Given $p, w > 0$, an integer $n \geq 1$, and a sequence $\{U_j : j \in \mathbb{N}\}$ of independent uniform random variables defined over $[0, p]$, if it holds that $np \geq w$, then*

$$\mathbb{E} \left[\min \left(w, \sum_{j=1}^n U_j \right) \right] \geq w \left(1 - \frac{w}{2np} \right).$$

Proof. When $n = 1$, we have

$$\begin{aligned} \mathbb{E}[\min(w, U_1)] &= w(1 - \mathbb{P}[U_1 \leq w]) + \mathbb{P}[U_1 \leq w] \cdot \mathbb{E}[U_1 \mid U_1 \leq w] \\ &= w \left(1 - \frac{w}{p} \right) + \frac{w}{p} \cdot \frac{w}{2} = w \left(1 - \frac{w}{2p} \right). \end{aligned} \quad (3)$$

When $n \geq 2$, we notice that $\min(w, \sum_{j=1}^n U_j) \geq \sum_{j \in [n]} \min(w/n, U_j)$. Indeed, let us denote the variables on the right-hand side by B_j , i.e., $B_j := \min(w/n, U_j)$ for $j \in [n]$. Whenever the sum of

U_j realizations is smaller than w , the sum of B_j values cannot be larger as each of them is upper bounded by U_j . On the other hand, in case U_j sum up to something bigger than w , the sum of B_j s is not larger than w as there are n of them, each upper bounded by w/n . Thus, we can use the monotonicity of expectation. Moreover, since $w/n \leq p$, we can follow the steps of (3) to get

$$\mathbb{E} \left[\min \left(w, \sum_{j=1}^n U_j \right) \right] \geq \sum_{j \in [n]} \mathbb{E} \left[\min \left(\frac{w}{n}, U_j \right) \right] \geq \sum_{j \in [n]} \frac{w}{n} \left(1 - \frac{w}{2np} \right) = w \left(1 - \frac{w}{2np} \right).$$

As a result, we proved this claim. \square

Let $n_0 \in \mathbb{N}$. If $w \leq n_0 p$, then by applying the bound of the claim in equation (2), we obtain

$$\begin{aligned} \mathbb{E}[\min(w, D(\sigma_1))] &\geq \sum_{n=n_0}^{\infty} \mathbb{P}[N(\sigma_1) = n] \cdot w \left(1 - \frac{w}{2np} \right) \geq \sum_{n=n_0}^{\infty} \mathbb{P}[N(\sigma_1) = n] \cdot w \left(1 - \frac{w}{2n_0 p} \right) \\ &= \mathbb{P}[N(\sigma_1) \geq n_0] \cdot w \left(1 - \frac{w}{2n_0 p} \right). \end{aligned} \quad (4)$$

In order to obtain the desired bound on the expected cost of the optimal schedule, we now define the suitable values of p and n_0 depending on whether $w\lambda \leq 1$ or $w\lambda > 1$.

Case $w\lambda \leq 1$. We define $n_0 = 1$ and $p = 1/\lambda$. Then, $n_0 p = 1/\lambda \geq w$. Moreover, $\mathbb{P}(N(\sigma_1) \geq 1)$ equals $1 - e^{-p\lambda} = 1 - e^{-1}$. Thus, combining (1), (4) and the values of n_0 and p , we obtain

$$\mathbb{E}[\text{cost}(\text{OPT}(\sigma))] \geq \frac{\tau}{p} \left(1 - e^{p\lambda} \right) w \left(1 - \frac{w}{2p} \right) = \tau \lambda w \left(1 - e^{-1} \right) \left(1 - \frac{w\lambda}{2} \right) \geq \frac{1 - e^{-1}}{2} \tau \lambda w. \quad (5)$$

Case $w\lambda > 1$. We define $p = \sqrt{2w/\lambda}$ and $n_0 = \lceil \lambda p \rceil$. We have $n_0 p \geq \lambda p^2 = 2w > w$. Moreover, by Proposition 2.5, we get that $\mathbb{P}(N(\sigma_1) \geq n_0) \geq 1/2$. Thus, combining (1), (4) and the value of n_0 and p , we obtain

$$\mathbb{E}[\text{cost}(\text{OPT}(\sigma))] \geq \frac{\tau}{p} \mathbb{P}[N(\sigma_1) \geq n_0] w \left(1 - \frac{w}{2n_0 p} \right) \geq \frac{\tau}{2} \sqrt{\frac{\lambda}{2w}} w \left(1 - \frac{w}{2 \cdot 2w} \right) = \frac{3}{8\sqrt{2}} \tau \sqrt{w\lambda}.$$

This concludes the proof of Lemma 3.2.

4 Overview

We now give an overview of the following sections. Inspired by the two strategies for the single edge instance, we define two types of stochastic instances: the *light* instances for which the strategy of serving requests instantly achieves a constant RoE, and the *heavy* instances for which the strategy of serving requests periodically achieves a constant RoE. Heavy and light instances are defined precisely below (Definitions 4.1 and 4.3) and generalizes the notions of heavy and light single-edge trees studied in the previous section.

We first define the light instances by extending the notion of *heaviness* for an arbitrary tree.

Definition 4.1. An instance (T, λ) is called *light* if $\pi(T, \lambda) \leq 1$, where $\pi(T, \lambda)$ is

$$\pi(T, \lambda) := \sum_{u \in V(T)} \lambda(u) \cdot d(u, \gamma(T)),$$

is called the *heaviness* of the instance.

We show in Section 5 that for a light instance, serving the requests immediately at the their arrival time achieves a constant RoE. We refer to the schedule produced with this strategy (see Algorithm 1) on a sequence of requests σ by $\text{INSTANT}(\sigma)$.

Algorithm 1: INSTANT
<p>Input: A sequence σ of requests.</p> <p>Output: A valid schedule $\text{INSTANT}(\sigma)$ for σ.</p> <p>1 for each request $r \in \sigma$ at its arrival time $t(r)$ do</p> <p>2 Create a service $(t(r), \{r\})$.</p>

Notice that this algorithm does not require the knowledge of the arrival rates.

Theorem 4.2. *INSTANT* has a ratio of expectations of $16/(3 - 3e^{-1}) < 8.44$ for light instances.

We prove this theorem in Section 5. We now turn our attention to heavy instances. An instance (T, λ) is heavy if for every subtree $T' \subseteq T$, we have $\pi(T', \lambda) > 1$. By monotonicity of $\pi(\cdot, \lambda)$, we obtain the following equivalent definition. Recall that for a vertex $u \in V(U)$, w_u denotes the weight of the edge incident to u on the path from $\gamma(T)$ to u .

Definition 4.3. An instance (T, λ) is called *heavy* if $w_u \geq 1/\lambda(u)$ for all $u \in V(T)$ with $\lambda(u) > 0$.

To give some intuition, suppose that u is a vertex of an heavy instance, and r and r' are two consecutive (random) requests located on u . Then, the expected duration between their arrival times is $1/\lambda(u) < w_u$. This suggests that to minimize the cost, we should in average gather r and r' into the same service, in order to avoid to pay twice the weight cost w_u . Since we expect services to serve a group of two or more requests, our stochastic assumptions suggests to the services must follow some form of regularity.

In Section 6, we present an algorithm called PLAN, that given an heavy instance (T, λ) , computes for each vertex $u \in V(T)$ a period $p_u > 0$, and will serve u at every time that is a multiple of p_u . One intuitive property of these periods $\{p_u : u \in V(T)\}$ is that the longer the distance to the root, the longer the period. While losing only a constant fraction of the expected cost, we choose the periods to be powers of 2. This enables us to optimize the weights of the services on the long run. One interesting feature of our algorithm is that it acts “blindly”: the algorithm does not need to know the requests, but only the arrival rate of each point! Indeed, our algorithm may serve a vertex where there are no pending requests. For the detail of the PLAN algorithm, see Section 6.

Theorem 4.4. *PLAN* has a ratio of expectations $64/3 < 21.34$ for heavy instances.

We remark that light instances and heavy instances are not complementary: there are instances that are neither light nor heavy.⁷ In Section 7, we focus on the general case of arbitrary instances.

⁷Furthermore, there exists a stochastic MLA instance where the ratio of expectations are both unbounded if INSTANT or PLAN is directed applied to deal with. See Appendix B for details.

The strategy there is to partition the tree (and the sequence of requests) into two groups of vertices (two groups of requests), so that the first group corresponds to a light instance where we can apply the instant strategy while the second group corresponds to a heavy instance where we can apply a periodic strategy. However, this correspondence for the heavy group is not straightforward. For this, we need to define an *augmented tree* that is a copy of the original tree, with the addition of some carefully chosen vertices. Each new vertex is associated with a subset of vertices of the original tree called *part*. We then define an arrival rate for each of these new vertices that is equal to the sum of the arrival rates of the vertices in the corresponding part. We show that this defines an heavy instance on which we can apply the algorithm PLAN. For each service made by PLAN on each of this new vertices, we serve all the pending requests in the corresponding part. The full description of this algorithm, called GEN, is given in Section 7. We show that this algorithm achieves a constant ratio of expectations.

Theorem 4.5. *GEN has a ratio of expectations of 210 for an arbitrary stochastic instance.*

5 Light instances

In this section, we prove Theorem 4.2 that we recall below.

Theorem 4.2. *INSTANT has a ratio of expectations of $16/(3 - 3e^{-1}) < 8.44$ for light instances.*

Recall that an instance (T, λ) is light if $\pi(T, \lambda) := \sum_{u \in V(T)} (\lambda(u) \cdot d(u, \gamma(T))) \leq 1$. When the context is clear we simply write $\pi(T) = \pi(T, \lambda)$. The proof of the theorem easily follows from the two following lemmas, that respectively estimates the expected cost of the algorithm, and give a lower bound on the expected cost of the optimal offline schedule.

Lemma 5.1. *Let (T, λ) be a light instance, and $\tau > 0$. Then,*

$$\mathbb{E}[\text{cost}(\text{INSTANT}(\sigma)) \mid \sigma \sim (T, \lambda)^\tau] = \tau \cdot \pi(T, \lambda).$$

Lemma 5.2. *Let (T, λ) be a light instance, and $\tau > 0$. Then,*

$$\mathbb{E}[\text{cost}(\text{OPT}(\sigma)) \mid \sigma \sim (T, \lambda)^\tau] \geq \frac{3}{16}(1 - e^{-1}) \cdot \tau \cdot \pi(T, \lambda).$$

Proof of Lemma 5.1. Let σ be a sequence of requests for T of duration τ and let $u \in V(T)$. For each request located on u , the algorithm issues a service of cost $d(u, \gamma(T))$ (notice that the delay cost is equal to zero). Let $N(\sigma|_u)$ denotes the number of requests in σ that are located at u . By Proposition 2.5, we know that $\mathbb{E}[N(\sigma|_u)] = \tau \lambda(u)$. Thus, we have

$$\begin{aligned} \mathbb{E}[\text{cost}(\text{INSTANT}(\sigma))] &= \mathbb{E}\left[\sum_{u \in V(T)} N(\sigma|_u) \cdot d(u, \gamma(T))\right] = \sum_{u \in V(T)} \tau \cdot \lambda(u) \cdot d(u, \gamma(T)) \\ &= \tau \cdot \pi(T, \lambda), \end{aligned} \tag{6}$$

which concludes the proof. \square

Proof of Lemma 5.2. Without loss of generality, we assume that τ is a multiple of $1/\lambda$. The plan of the proof is to associate to (T, λ) a specific family of single-edge light instances. We then apply the bounds proved in Section 3 to establish the bound of the lemma.

For each integer $j \in \mathbb{Z}$, we denote

$$V_j := \{u \in V(T) : 2^j \leq d_T(u, \gamma(T)) < 2^{j+1}\}.$$

Let σ be a sequence of requests for T . We denote

$$\sigma_j := \sigma|_{V_j} = \{r \in \sigma : \ell(r) \in V_j\}.$$

For each $j \in \mathbb{Z}$, we create a single-edge stochastic instance as follows. Let e_j denote a single-edge tree of weight $w_j = 2^{j-1}$. Let $\lambda_j := \boldsymbol{\lambda}(V_j)$ denote the arrival rate of the vertex in e_j (that is not the root). We construct σ'_j as a sequence of requests for e_j with the same arrival times as requests in σ_j . Let $S(\sigma)$ denote a schedule for σ . We create a schedule $S(\sigma'_j)$ for σ'_j as follows: for each service $s \in S(\sigma)$ that serves at least one request from V_j , add a service s'_j in $S(\sigma'_j)$ with the same service time to serve the corresponding requests in σ'_j . It is clear from the construction that if $S(\sigma)$ is a valid schedule for σ , then for each $j \in \mathbb{Z}$, $S(\sigma'_j)$ is a valid schedule for σ'_j . Further, we have the following inequality on the cost of these schedules.

$$\text{cost}(S(\sigma), T) \geq \sum_{j \in \mathbb{Z}} \text{cost}(S(\sigma'_j), e_j). \quad (7)$$

Indeed, first notice that the delay costs are the same, i.e., $\text{delay}(S(\sigma)) = \sum_j \text{delay}(S(\sigma'_j))$. We now focus on the weight of these schedules. Let s be a service in $S(\sigma)$ which serves a subset of requests $R(s) \subseteq \sigma$ at time t . Let $k \in \mathbb{Z}$ be the largest index such that $R(s) \cap \sigma_k \neq \emptyset$. This means that s serves a request that is at distance at least 2^k from the root, and then

$$\text{weight}(s) \geq 2^k \geq \sum_{j \leq k} 2^{j-1} \geq \sum_{j \leq k} \text{weight}(s'_j) = \sum_{j \in \mathbb{Z}} \text{weight}(s'_j),$$

where s'_j denotes the service created at time t that serves the requests in σ'_j that correspond to $R(s) \cap \sigma_j$. As a result,

$$\text{cost}(S(\sigma), T) \geq \sum_{j \in \mathbb{Z}} \text{cost}(S(\sigma'_j), e_j) \geq \sum_{j \in \mathbb{Z}} \text{cost}(\text{OPT}(\sigma'_j), e_j),$$

where $\text{OPT}(\sigma'_j)$ denotes the optimal schedule for σ'_j . Since this holds for any valid schedule $S(\sigma)$, we obtain

$$\text{cost}(\text{OPT}(\sigma), T) \geq \sum_j \text{cost}(\text{OPT}(\sigma'_j), e_j).$$

Due to the equivalence between the distributed and centralized Poisson arrival model, we know that $\sigma \sim (T, \boldsymbol{\lambda})^\tau$ implies $\sigma'_j \sim (e_j, \lambda_j)^\tau$. By taking expectation over all the random sequences $\sigma \sim (T, \boldsymbol{\lambda})^\tau$, we have

$$\mathbb{E}[\text{cost}(\text{OPT}(\sigma), T)] \geq \sum_{j \in \mathbb{Z}} \mathbb{E}[\text{cost}(\text{OPT}(\sigma'_j), e_j) \mid \sigma'_j \sim (e_j, \lambda_j)^\tau].$$

We now show that for each $i \in \mathbb{Z}$, (e_j, λ_j) is a light instance. Indeed, we have

$$\pi(e_j, \lambda_j) = w_j \cdot \lambda_j = 2^{j-1} \cdot \lambda_j = 2^{j-1} \cdot \sum_{u \in V_j} \boldsymbol{\lambda}(u) = \frac{1}{2} \sum_{u \in V_j} \boldsymbol{\lambda}(u) \cdot 2^j \leq \frac{1}{2} \sum_{u \in V_j} \boldsymbol{\lambda}(u) \cdot d(u, \gamma(T))$$

$$\leq \frac{1}{2} \sum_{u \in V(T)} \lambda(u) \cdot d(u, \gamma(T)) = \frac{1}{2} \cdot \pi(T, \lambda) \leq \frac{1}{2},$$

since (T, λ) is light. Thus, for each $i \in \mathbb{Z}$, by plugging $p = 1/\lambda_j$ in equation (5) in the proof of Lemma 3.2, we obtain

$$\begin{aligned} \mathbb{E} [\text{cost}(\text{OPT}(\sigma'_j), e_j) \mid \sigma'_j \sim (e_j, \lambda_j)^\tau] &\geq \frac{\tau}{p} (1 - e^{p\lambda}) w_j (1 - \frac{w_j}{2p}) \geq (1 - \frac{w_j \lambda_j}{2}) (1 - e^{-1}) \tau w_j \lambda_j \\ &\geq \frac{3}{4} (1 - e^{-1}) \tau w_j \lambda_j. \end{aligned}$$

On the other hand, we have

$$w_j \cdot \lambda_j = 2^{j-1} \sum_{u \in V_j} \lambda(u) = \frac{1}{4} \sum_{u \in V_j} 2^{j+1} \lambda(u) \geq \frac{1}{4} \sum_{u \in V_j} d(u, \gamma(T)) \cdot \lambda(u).$$

Putting everything together we obtain the desired bound:

$$\begin{aligned} \mathbb{E}[\text{cost}(\text{OPT}(\sigma))] &\geq \sum_{j \in \mathbb{Z}} \mathbb{E} [\text{cost}(\text{OPT}(\sigma'_j), e_j) \mid \sigma'_j \sim (e_j, \lambda_j)^\tau] \geq \sum_{j \in \mathbb{Z}} \frac{3}{4} (1 - e^{-1}) \tau w_j \lambda_j \\ &\geq \frac{3}{4} (1 - e^{-1}) \cdot \frac{1}{4} \cdot \tau \cdot \sum_j \sum_{u \in V_j} d(u, \gamma(T)) \cdot \lambda(u) \\ &= \frac{3}{16} (1 - e^{-1}) \cdot \tau \cdot \sum_{u \in V(T)} d(u, \gamma(T)) \cdot \lambda(u) = \frac{3}{16} (1 - e^{-1}) \cdot \tau \cdot \pi(T, \lambda). \end{aligned}$$

This concludes the proof. □

6 Heavy instances

In this section, we focus on heavy MLA instances. Let us first recall their definition.

Definition 4.3. *An instance (T, λ) is called heavy if $w_u \geq 1/\lambda(u)$ for all $u \in V(T)$ with $\lambda(u) > 0$.*

To solve the problem, when restricted to this class of MLA instances, we define a new algorithm PLAN. In the main theorem of this section, we prove that within this class, PLAN achieves a constant ratio of expectations.

Our approach can be seen as a generalization of the algorithm for a single-edge case. Once again, we serve the requests periodically, although this time, we may assign different periods for different vertices. Intuitively, vertices closer to the root and having a greater arrival rate should be served more frequently. For this reason, the PLAN algorithm generates a partition P of a given tree T into a family of subtrees (clusters) and assigns a specific period to each of them.

The partition procedure allows us to analyze each cluster separately. Thus, we can assume that from now on, we are restricted to a given subtree $T' \in P$. To lower bound the cost generated by OPT on T' , we split the weight of T' among its vertices using a saturation procedure. This action allows us to say that for each vertex v , the optimal algorithm either covers the delay cost of all the requests arriving at v within a given time horizon or pays some share of the service cost. The last step is to round the periods assigned to the subtrees in P to minimize the cost of PLAN. In what follows, we present the details of our approach.

6.1 Periodical algorithm PLAN

As mentioned before, the main idea is to split tree T rooted at vertex γ into a family of subtrees and serve each of them periodically. In other words, we aim to find a partition $P = \{T_1, T_2, \dots, T_k\}$ of T where each subtree T_i besides the one containing γ is rooted at the leaf vertex of another subtree. At the same time, we assign each subtree T_i some period p_i . To decide how to choose the values of p_i s, let us recall how we picked the period for a single-edge case. In that setting, for the period p we had an equality between the expected delay cost $\lambda/2 \cdot p^2$ at the leaf u and the weight w of the edge. Thus, the intuition behind the PLAN algorithm is as follows.

We start by assigning each vertex $v \in T$ a process that saturates the edge connecting it to the parent at the pace of $\lambda(v)/2 \cdot t^2$, i.e., within the time interval $[t, t + \epsilon]$ it saturates the weight of $\lambda(v)/2 \cdot ((t + \epsilon)^2 - t^2)$. Whenever an edge gets saturated, the processes that contributed to this outcome join forces with the processes that are still saturating the closest ascendant edge. As the saturation procedure within the whole tree T reaches the root γ , we cluster all the vertices corresponding to the processes that made it possible into the first subtree T_1 . Moreover, we set the period of T_1 to the time it got saturated. After this action, we are left with a partially saturated forest having the leaves of T_1 as the root nodes. The procedure, however, follows the same rules, splitting the forest further into subtrees T_2, \dots, T_k .

To simplify the formal description of our algorithm, we first introduce some new notations. Let $p(v)$ denote the saturation process defined for a given vertex v . As mentioned before, we define it to saturate the parent edge at the pace of $\lambda(v)/2 \cdot t^2$. Moreover, we extend this notation to the subsets of vertices, i.e., we say that $p(S)$ is the saturation process where all the vertices in S cooperate to cover the cost of an edge. The pace this time is equal to $\lambda(S)/2 \cdot t^2$. To trace which vertices cooperate at a given moment and which edge they saturate, we denote the subset of vertices that v works with by $S(v)$ and the edge they saturate by $e(v)$. We also define a method `join`(u, v) that takes as the arguments two vertices and joins the subsets they belong to. It can be called only when the saturation process of $S(u)$ reaches v . Formally, at this moment the `join` method merges subset $S(u)$ with $S(v)$ and sets $e(v)$ as the outcome of the function e on all the vertices in the new set. It also updates the saturation pace of the new set. Now, we present the pseudo-code for PLAN as Algorithm 2. For a visual support to illustrate the saturating process, see Figure 1.

We start by listing some properties of the partition generated by this algorithm.

Proposition 6.1. *Let (T, γ, λ, w) be a heavy instance and let $P = \{T_1, T_2, \dots, T_k\}$ be the partition generated on it by Algorithm 2. We denote the period corresponding to T_i by p_i . Assuming that T_i s are listed in the order they were added to P , it holds that:*

1. each T_i is a rooted subtree of T ;
2. the periods are increasing, i.e., $1 \leq p_1 \leq p_2 \leq \dots \leq p_k$;
3. each vertex $v \in T_i$ saturated exactly $\lambda(v)/2 \cdot p_i^2$ along the path to the root of T_i .

Proof. To show that the first property is satisfied, we proceed by induction. Initially, we have that each subset $S(v)$ for $v \in T$ is a single vertex and thus it forms a subtree. Then, in line 9, we can merge two subtrees only if an edge connects them. Thus, the `join` call also creates a new subset that induces a subtree. Finally, we notice that we cluster a subset only as it reaches a vertex from the set R . It becomes the root of this subtree, which implies the desired property.

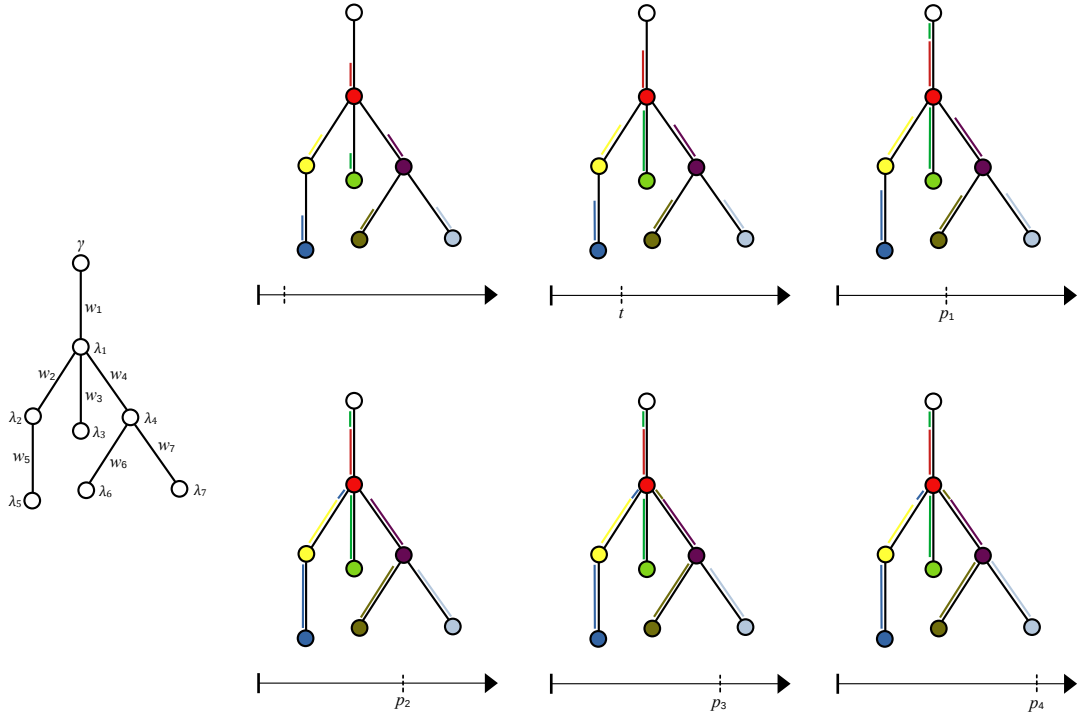


Figure 1: Here is an example to show how Algorithm 2 works on an heavy instance. Given the tree consisting of 7 vertices (with $w_i \geq 1/\lambda_i$ for each vertex $i \in [7]$ marked in different color), we use the length of the colored line to denote the saturated amount (i.e., $\lambda_i/2 \cdot t^2$) of a vertex i at any time t . At time p_1 , the subtree T_1 including vertices 1 and 3 is determined; similarly, T_2 includes vertices 2 and 5 at time t_2 ; T_3 includes vertices 4 and 6 at time p_3 ; and T_4 includes vertex 7 at time p_4 .

Algorithm 2: PLAN (part I)

Input: an heavy instance (T, λ) with tree T rooted at γ
Output: a partition $P = \{T_1, T_2, \dots, T_k\}$ of T with each subtree T_i assigned a period p_i

- 1 let R be the set of roots, initially $R = \{\gamma\}$ \triangleright at the beginning we work with the whole tree
- 2 **for** each vertex $v \in V(T)$ **do**
- 3 define the saturation process $p(v)$ as described before
- 4 set $S(v) := \{v\}$ and $e(v) := \text{par}(v)$
- 5 start the clock at time 0
- 6 **while** there exist some unclustered vertices in T **do**
- 7 wait until the first time t_e when an edge $e = (u, v)$ gets saturated
- 8 **if** $v \notin R$ **then**
- 9 join(u, v)
- 10 **else**
- 11 add cluster $C := S(u) \cup \{v\}$ to partition P \triangleright we reached the root vertex v
- 12 set the period p for C to be equal to t_e
- 13 set the saturation pace for C to 0
- 14 extend R by the leaves in C \triangleright leaves of C are the roots of the remaining forest

The second property follows straight from the assumption that we started the clock at time 0 in line 5 and we process the edges in order they get saturated, i.e., there is no going back in time. Similarly, the last property is implied by the definition of the saturation process. \square

6.2 Lower bounding OPT

In this subsection, we lower bound the total cost incurred by the optimal offline schedule OPT on a heavy instance. Let us first consider each subtree T_i generated by Algorithm 2 separately.

Lemma 6.2. *Let (T, λ) be a heavy instance. We denote the partition generated for it by Algorithm 2 by $P = \{T_1, T_2, \dots, T_k\}$ and the period corresponding to T_i by p_i for all $i \in [k]$. Let T_i be any subtree in P , and let us define σ_i as a random sequence of requests arriving within the MLA instance restricted to T_i over a time horizon τ . We assume that τ is a multiple of p_i . It holds that*

$$\mathbb{E}[\text{cost}(\text{OPT}(\sigma_i), T_i) \mid \sigma_i \sim (T_i, \lambda|_{T_i})^\tau] \geq \frac{3}{16} \cdot w(T_i) \cdot \frac{\tau}{p_i}.$$

Proof. We use the same approach as in Section 3 and first focus on lower bounding the cost incurred by OPT within a shorter time interval — for now, we set the horizon to p_i . By Proposition 6.1, we have that by time p_i each vertex $v \in T_i$ saturates the weight of $\hat{w}_v := \lambda(v)/2 \cdot p_i^2$ along the path to the root. Thus, whenever OPT issues a service that contains v , we can distribute the service cost among the served vertices and say that v needs to cover \hat{w}_v share.

By the definition of PLAN, we know that the sum of \hat{w}_v over all the vertices $v \in T_i$ is equal to the weight $w(T_i)$, as T_i is fully saturated at moment p_i . Moreover, by the definition of a heavy instance, we have that $w_v \geq 1/\lambda(v)$ for each $v \in T_i$. Combining it all together gives us

$$\sum_{v \in T_i} \frac{\lambda(v) \cdot p_i^2}{2} = \sum_{v \in T_i} \hat{w}_v = \sum_{v \in T_i} w_v \geq \sum_{v \in T_i} \frac{1}{\lambda(v)}. \quad (8)$$

Now, we apply the single-edge case analysis for some of the vertices in T_i . To be more precise, we focus on the case where the product of the arrival rate and the weight of a given vertex was at least 1. The crucial assumption there was to guarantee that the parameter for the Poisson arrival variable, i.e., the product of the arrival rate and the period, was at least 1. Thus, in the current scenario, we need to check for which vertices $v \in T_i$ it holds that $\lambda(v)p_i \geq 1$.

Here, we use a different approach and upper bound the contribution to the total saturated cost $\sum_{v \in T_i} \hat{w}_v$ incurred by the vertices that do not satisfy this property. Let us denote the set of such vertices by L_i . We have that $p_i \leq 1/\lambda(v)$ for each $v \in L_i$. Thus, combining this with inequality (8), implies that

$$\sum_{v \in L_i} \hat{w}_v = \sum_{v \in L_i} \frac{\lambda(v) \cdot p_i^2}{2} \leq \frac{1}{2} \sum_{v \in L_i} \frac{\lambda(v)}{\lambda^2(v)} = \frac{1}{2} \sum_{v \in L_i} \frac{1}{\lambda(v)} \leq \frac{1}{2} \sum_{v \in T_i} \frac{1}{\lambda(v)} \leq \frac{1}{2} \sum_{v \in T_i} w_v, \quad (9)$$

which proves that at least half of the saturation cost comes from the heavy vertices.

As we apply the single-edge case analysis for all the vertices in $T_i \setminus L_i$, saying that within each period p_i , OPT has to pay either the service cost of \hat{w}_v or the total delay cost generated at vertex v for each $v \in T_i$, we obtain

$$\mathbb{E}[\text{cost}(\text{OPT}(\sigma_i))] \geq \sum_{v \in T_i \setminus L_i} \frac{3}{8} \cdot \hat{w}_v \cdot \frac{\tau}{p_i} = \frac{3}{8} \sum_{v \in T_i \setminus L_i} \hat{w}_v \cdot \frac{\tau}{p_i} \geq \frac{3}{8} \cdot \frac{1}{2} \sum_{v \in T_i} w_v \cdot \frac{\tau}{p_i} = \frac{3}{16} \cdot w(T_i) \cdot \frac{\tau}{p_i},$$

which ends the proof. \square

6.3 Cost analysis for PLAN

Let us start by assuming that we serve all the subtrees T_i generated by Algorithm 2 periodically according to the periods p_i . In this setting, to serve any cluster besides the one containing the root vertex γ , not only we need to cover the service cost of the cluster vertices but also the cost of the path connecting them to γ . Since we only know how to lower bound the cost incurred by OPT on the clusters, we improve the PLAN algorithm to get rid of this problem. The idea is to round the periods p_i s to be of form $2^{e_i}p_1$ for some positive integers e_i . Thus, whenever we need to serve some cluster S_i , we know that we get to serve all the clusters generated before it as well. Formally, our approach is presented in Algorithm 3.

Algorithm 3: PLAN (part II)

Input: increasing sequence of periods (p_1, p_2, \dots, p_k) obtained from Algorithm 2

Output: new sequence of periods $(\hat{p}_1, \hat{p}_2, \dots, \hat{p}_k)$

- 1 **for** $i \in \{2, 3, \dots, k\}$ **do**
- 2 find $e_i \in \mathbb{Z}_+$ such that $2^{e_i}p_1 \leq p_i < 2^{e_i+1}p_1$
- 3 set $\hat{p}_i := 2^{e_i}p_1$

Finally, we define the algorithm PLAN to serve the requests periodically, according to the new periods (see Algorithm 4).

To conclude this section, we prove Theorem 4.4 (restated below).

Theorem 4.4. *PLAN has a ratio of expectations $64/3 < 21.34$ for heavy instances.*

Algorithm 4: PLAN (part III)

Input: heavy instance (T, λ) and a sequence of requests σ in $[0, \tau]$
Output: valid schedule $\text{PLAN}(\sigma)$ for σ

- 1 let $P = \{T_1, T_2, \dots, T_k\}$ be the partition generated by Algorithm 2;
- 2 let $(\hat{p}_1, \hat{p}_2, \dots, \hat{p}_k)$ be a sequence of periods obtained from Algorithm 3;
- 3 **for** $i \in [k]$ **do**
- 4 **for** $j \in \llbracket \tau / \hat{p}_i \rrbracket$ **do**
- 5 | schedule a service that serves T_i at time $j\hat{p}_i$
- 6 **if** τ is not a multiple of \hat{p}_i **then**
- 7 | schedule a service that serves T_i at time τ

Proof. Let (T, λ) be a heavy instance with tree T rooted at γ and let $P = \{T_1, T_2, \dots, T_k\}$ be the partition generated for it by Algorithm 2. Moreover, let (p_1, p_2, \dots, p_k) and $(\hat{p}_1, \hat{p}_2, \dots, \hat{p}_k)$ denote the periods obtained from Algorithm 2 and Algorithm 3, respectively. Here, we analyze the cost generated by PLAN (Algorithm 4) on a random sequence of requests $\sigma \sim (T, \lambda)^\tau$, where the time horizon τ is a multiple of $2\hat{p}_k$.

Notice that since we align the periods to be of form $2^l \hat{p}_1$ for some positive integer l , whenever PLAN serves some tree T_i , it serves all the trees containing the path from T_i to γ at the same time. Thus, the service cost can be estimated on the subtree level. Moreover, since for each $i \in [k]$ it holds that period $\hat{p}_i \leq p_i$, the expected delay cost incurred within $[0, \hat{p}_i]$ does not exceed $w(T_i)$. Thus, denoting $\sigma_i \sim (T_i, \lambda|_{T_i})^{\hat{p}_i}$ for $i \in [k]$, by Propositions 2.2 and 2.3, we have that

$$\mathbb{E}[\text{cost}(\text{PLAN}(\sigma), T)] = \sum_{i=1}^k \frac{\tau}{\hat{p}_i} \mathbb{E}[\text{cost}(\text{PLAN}(\sigma_i), T_i)] = \sum_{i=1}^k \frac{\tau}{\hat{p}_i} \cdot 2w(T_i) = 2 \sum_{i=1}^k \frac{\tau}{\hat{p}_i} w(T_i).$$

Now, let us lower bound the expected cost for the optimal offline schedule for σ . By Proposition 2.2, we have that

$$\mathbb{E}[\text{cost}(\text{OPT}(\sigma), T)] = \sum_{i=1}^k \mathbb{E}[\text{cost}(\text{OPT}(\sigma|_{T_i}), T_i)] \geq \sum_{i=1}^k \frac{\tau}{p_i} \frac{3}{16} w(T_i).$$

Since by definition of Algorithm 3, it holds that $p_i < 2\hat{p}_i$ for $i \in [k]$, we can rewrite the above as

$$\mathbb{E}[\text{cost}(\text{OPT}(\sigma), T)] > \sum_{i=1}^k \frac{\tau}{2\hat{p}_i} \frac{3}{16} w(T_i) = \frac{3}{32} \sum_{i=1}^k \frac{\tau}{\hat{p}_i} w(T_i).$$

Thus, the ratio between the expected costs of OPT and PLAN algorithms is upper bounded by $64/3$. However, to simplify the calculations, until now we assumed that the time horizon τ is a multiple of $2\hat{p}_k$. Nonetheless, this implies that PLAN achieves a ratio of expectations equal to $64/3$, since with the value of τ going to infinity, the marginal contribution of extra cost generated by the last service (line 7 of Algorithm 4) goes to 0. \square

7 General instances

In this section, we describe our algorithm GEN for an arbitrary stochastic instance $(T, \gamma, \boldsymbol{\lambda}, \mathbf{w})$ and prove that it achieves a constant RoE. The main idea is to distinguish two types of requests, and apply a different strategy for each type. The first type are the requests that are located close to the root. These requests will be served immediately at their arrival times, i.e., we apply INSTANT to the corresponding sub-sequence. The second type includes all remaining requests and they are served in a periodic manner. To determine the period of these vertices, we will use the algorithm PLAN on a specific heavy instance $(T', \boldsymbol{\lambda}^h)$ constructed in Section 7.2. The construction of this heavy instance relies on a partition of the vertices of T into *balanced parts*, whose definition and construction is given in Section 7.1. Intuitively, a part is balanced when it is light (or close to being light), but if we merge all vertices of the part into a single vertex whose weight corresponds to the average distance to the root of the part, then we obtain an heavy edge. This “merging” process is captured by the construction of the *augmented tree* T' , which is part of the heavy instance. The augmented tree is essentially a copy of T with the addition of one (or two) new vertices for each balanced part. See Section 7.2 for the formal description.

Once we have determined the corresponding heavy instance, we can compute the periods of each vertex of the heavy instance with the PLAN algorithm. The period of a vertex in the original instance is equal to the period of the corresponding vertex in the heavy instance. The full description of the GEN algorithm is given in Section 7.3.

The main challenge of this section is to analyze the ratio of expectations of this algorithm, and in particular to establish good lower bounds on the expected cost of the optimal offline schedule. This is done in Section 7.4.2, where we prove two lower bounds (Lemmas 7.8 and 7.9) that depend on the heaviness of each part of the balanced partition.

In the entire section, we assume without loss of generality that γ has only one child. To see that this is possible, consider an MLA instance (T, σ) with tree T rooted at vertex γ that has at least $k \geq 2$ children. For each $i \in [k]$ let T_i be a tree obtained from T by pruning all the children of γ except the i th one, and denote the sequence of requests in σ arriving at T_i by σ_i . Finding a schedule for instance (T, σ) is then equivalent to finding a family of schedules for (T_i, σ_i) since each service s for (T, σ) can be partitioned into a set of services for each instance (T_i, σ_i) . Further the sum of the weights of the new services is the same as the weight of s .

Moreover, recall that we assume $\boldsymbol{\lambda}(\gamma) = 0$, as each request arriving at γ can be served immediately with zero weight cost.

Notations. Recall that given the edge-weighted tree T rooted at γ and a set of vertices $U \subseteq V(T)$, $T[U]$ denotes the forest induced on U in T . We say that a subset $U \subseteq V(T)$ is *connected* if $T[U]$ is connected (i.e., $T[U]$ is a subtree of T but not a forest). If $U \subseteq V(T)$ is connected, we write $\gamma(U) = \gamma(T[U])$ to denote the root vertex of $T[U]$, i.e., the vertex in U which has the shortest path length to γ in the original tree T . Given any vertex $u \in V(T)$, let $V_u \subseteq V(T)$ denote all the descendant vertices of u in T (including u). For simplification, set $w_{\gamma(T)} = \infty$. Given $T = (V, E)$ and $\boldsymbol{\lambda} : V(T) \rightarrow \mathbb{R}_+$, and $U \subseteq V$, we denote $\boldsymbol{\lambda}|_U : U \rightarrow \mathbb{R}_+$ such that $\boldsymbol{\lambda}|_U(u) = \boldsymbol{\lambda}(u)$ for each $u \in U$. For a sequence of requests $\sigma \sim (T, \boldsymbol{\lambda})$, we use $\sigma|_U = \{r \in \sigma : \ell(r) \in U\}$ to denote the corresponding sequence for $T[U]$.




Figure 2: An example of a balanced partition (Definition 7.2). The weight of each edge is shown in black, and the arrival rate of each vertex is shown in red. Green subsets corresponds to parts of type-I while purple ones correspond to parts of type-II. Some value of π are shown for the top-left type-I part and for the top-right type-II part.

7.1 Balanced partition of $V(T)$

Given $\lambda : V(T) \rightarrow \mathbb{R}_+$, recall that $\pi(T, \lambda) = \sum_{u \in V(T)} \lambda(u) \cdot d(u, \gamma(T))$. When the context is clear we simply write $\pi(T) = \pi(T, \lambda)$, and for a connected subset $U \subseteq V(T)$ we simply write $\pi(U) := \pi(T[U], \lambda|_U)$.

Definition 7.1. Given a stochastic instance (T, λ) , we say that $U \subseteq V(T)$ is balanced if U is connected and if one of the following conditions holds:

- (1) $\pi(U) \leq 1$, and either $\gamma(U) = \gamma(T)$ or $\pi(U \cup \{\text{par}(\gamma(U))\}) > 1$. In this case, we say that U is of type-I.
- (2) $\pi(U) > 1$, and for each child vertex y of $\gamma(U)$ in $T[U]$, we have $\pi(\{\gamma(U)\} \cup (U \cap V_y)) < 1$. In this case, we say that U is of type-II.

Remark. Note that the root $\gamma(U)$ of a balanced part U of type-II is necessarily to have at least two children in $T[U]$.

Definition 7.2. Given a stochastic instance (T, λ) and a partition \mathcal{P} of the vertices $V(T)$, we say that \mathcal{P} is a balanced partition of T if every part $U \in \mathcal{P}$ is balanced.

See Figure 2 for an example of a balanced partition. If \mathcal{P} is a *balanced partition* of T , then the part $U \in \mathcal{P}$ containing $\gamma(T)$ is called the *root part* in \mathcal{P} . Since we assumed that $\gamma(T)$ has only one child vertex, we deduce from the previous remark that the root part is necessarily of type-I. Given a balanced partition \mathcal{P} , we denote $\mathcal{P}^* := \mathcal{P} \setminus \{\gamma(\mathcal{P})\}$, we denote $\mathcal{P}_1 \subseteq \mathcal{P}$ the set of parts of type-I, $\mathcal{P}_1^* := \mathcal{P}_1 \cap \mathcal{P}^*$, and $\mathcal{P}_2 \subseteq \mathcal{P}$ the set of parts of type-II.

Lemma 7.3. Given any stochastic instance (T, λ) , there exists a balanced partition of T . Moreover, such a partition can be computed in $O(|V(T)|^2)$ time.⁸

⁸we assume that basic operations on numbers can be done in constant time.




Figure 3: Construction of a balanced partition in the proof of Lemma 7.3. The weights of the edges and the arrival rates of the vertices are the same as in Figure 2. The numbers represent an ordering of the vertices. The gray sets corresponds to U_i , for $i \in [40]$. We illustrate the step $i = 30$ of the algorithm. We have $C_{30} = \{24, 26\}$ and $U_{30} = \{u_{30}\} \cup U_{24} \cup U_{30} = \{u_{30}, u_{24}, u_{19}, u_{20}, u_{26}\}$. Since $\pi(U_{30} \cup \{u_{25}\}) = 2.65 > 1$, we add U_{30} into $\mathcal{P}^{(29)}$ to create $\mathcal{P}^{(30)}$ (red stroke). We remark that U_{30} is a balanced subset of type-II and U_{27} is a balanced subset of type-II, while U_{25} is not a balanced subset.

Proof. We first describe an algorithm that constructs a partition \mathcal{P} , and then argue that this partition is balanced. See Figure 3 for a visual support. We sort the vertices u_1, \dots, u_n of T (where $n = |V(T)|$) by decreasing distances to the tree root, i.e., for $1 \leq i < j \leq n$, $d(u_i, \gamma(T)) \geq d(u_j, \gamma(T))$. Let $\mathcal{P}^{(0)} = \emptyset$. For each $i \in [n]$, do the following. Let $C_i \subseteq \{1, \dots, n-1\}$ be the subset of indexes j , such that (i) u_j is a child of u_i ; (ii) $u_j \notin \bigcup_{U \in \mathcal{P}^{(i-1)}} U$. Define $U_i := \{u_i\} \cup \bigcup_{j \in C_i} U_j$ recursively. If $i = n$ (i.e., if u_i is the root of T) or $\pi(U_i \cup \{\text{par}(u_i)\}) > 1$, then define $\mathcal{P}^{(i)} := \mathcal{P}^{(i-1)} \cup \{U_i\}$. Otherwise, define $\mathcal{P}^{(i)} := \mathcal{P}^{(i-1)}$.

We first show that for each $i \in [n]$, the set U_i is connected. This shows the correctness of the algorithm⁹ and will also be useful later. We proceed by induction. Given $i \in [n]$, suppose that all U_j with $j < i$ are connected. Since the vertices are ordered by decreasing distance to the root of T , it holds that for every $j \in C_i$, we have $j < i$, and hence U_j is connected. The set $U_i = \{u_i\} \cup \bigcup_{j \in C_i} U_j$ is connected since for each $j \in C_i$, u_j and u_i are adjacent in T , $u_j \in U_j$, and U_j is connected.

Let $\mathcal{P} = \mathcal{P}^{(n)}$ be the set of subsets of $V(T)$ at the end of the algorithm. Now we show that \mathcal{P} is a partition of $V(T)$. For each $i \in [n]$, let $b(i)$ be the smallest integer j such that (i) u_j is on the path from $\gamma(T)$ to u_i ; (ii) $\mathcal{P}^{(j)} \neq \mathcal{P}^{(j-1)}$. Note that such index $b(i)$ exists and is unique. Besides, for each $i \in [n]$, $u_i \in U_{b(i)}$ and $U_{b(i)} \in \mathcal{P}$. Hence, \mathcal{P} is indeed a partition of $V(T)$.

It remains to show that each part $W \in \mathcal{P}$ is either of type-I or of type-II. Let i be the index such that $U_i = W$. Note that $\mathcal{P}^{(i)} \neq \mathcal{P}^{(i-1)}$ by the definition of our algorithm. This implies that either $u_i = \gamma(T)$ or $\pi(U_i \cup \{\text{par}(u_i)\}) > 1$. If $\pi(W) \leq 1$, then W is of type-I. Otherwise, $\pi(W) > 1$. We show that in this case W is of type-II, i.e., for each child vertex y of $\gamma(W)$ in $T[W]$, we have $\pi(\{\gamma(W)\} \cup (W \cap V_y)) < 1$. Notice that $\gamma(W) = u_i = \text{par}(y)$, and that for each child

⁹ $\pi(\cdot)$ is only defined for connected sets.

vertex y of u_i in $T[W]$, there is an index $j \in C_i$ such that $y = u_j$. Furthermore, by construction we have $W \cap V_{u_j} = U_j$. Thus, $\{\gamma(W)\} \cup (W \cap V_y) = U_j \cup \{\text{par}(u_j)\}$. Now, by definition of C_i , we know that $u_j \notin \bigcup_{U \in \mathcal{P}^{(i-1)}} U$, and in particular, $u_j \notin \bigcup_{U \in \mathcal{P}^{(j)}} U$, since $j \leq i - 1$. This implies that $\mathcal{P}^{(j)} = \mathcal{P}^{(j-1)}$, and hence that $\pi(U_j \cup \{\text{par}(u_j)\}) \leq 1$, which is what we wanted to prove. To summarize, \mathcal{P} is indeed a balanced partition.

For the running time to produce \mathcal{P} , recall that $n = |V(T)|$ denotes the number of vertices. On one hand, sorting all the tree vertices in the order of their distances to the tree root takes a running time at most $O(n^2)$. On the other hand, there are n iterations for producing the balanced partition \mathcal{P} . In each iteration $i \in [n]$, determining C_i needs $O(n)$ time and calculating $\pi(U_i \cup \{\text{par}(u_i)\})$ needs $O(n)$ time. As such, the total running time is $O(n^2)$.

This concludes the proof of Lemma 7.3. \square

7.2 The heavy instance

Given a stochastic instance (T, λ) , and a balanced partition \mathcal{P} of T , we construct a tree T' that we call the *augmented tree* of T . This tree is essentially a copy of T with additional one or two vertices for each part of \mathcal{P}^* .¹⁰ Then, we define arrival rates λ^h on T' in a way that the stochastic instances (T', λ^h) is heavy. Finally, we construct from a request sequence σ , the corresponding *heavy sequence* σ^h for the augmented tree.

Construction of the augmented tree. We define $T' = (V', E')$ where

$$V' = V(T) \cup \{z_U, z'_U : U \in \mathcal{P}^*\},$$

and the edge set E' is constructed based on $E(T)$ as follows (see Figure 4).

- First, for each $U \in \mathcal{P}_1^*$, replace the edge $(\gamma(U), \text{par}(\gamma(U)))$ (of length $w_{\gamma(U)}$) by two edges $(\gamma(U), z'_U)$ and $(z'_U, \text{par}(\gamma(U)))$ of respective lengths $(1 - \pi(U))/\lambda(U)$ and $w_{\gamma(U)} - (1 - \pi(U))/\lambda(U)$; where $\text{par}(\gamma(U))$ denotes the parent of $\gamma(U)$ in T . Then, add an edge (z_U, z'_U) of weight $1/\lambda(U)$.
- For each $U \in \mathcal{P}_2$, set $z'_U = \gamma(U)$, and add an edge (z_U, z'_U) of weight $\pi(U)/\lambda(U)$.

This completes the construction of the augmented tree. Notice that if a part $U = \{u\}$ in \mathcal{P} that contains only one vertex, then we have $\pi(U) = 0$ and thus U is necessarily of type-I. To simplify, in the following we identify vertices in T with their copy in T' , and consider that $V(T)$ is a subset of $V(T')$.

Arrival rates for the heavy instance. Recall that $\mathcal{P}^* = \mathcal{P} \setminus \{\gamma(\mathcal{P})\}$, where $\gamma(\mathcal{P})$ denotes the part in \mathcal{P} containing the root $\gamma(T)$. We define $\lambda^h : V(T') \rightarrow \mathbb{R}_+$ as follows: for each $U \in \mathcal{P}^*$, set $\lambda^h(z_U) = \lambda(U)$; and $\lambda^h(u) = 0$ otherwise.

Proposition 7.4. (T', λ^h) is heavy.

Proof. For each $U \in \mathcal{P}_1^*$ we have $\lambda^h(z_U) \cdot w(z_U) = \lambda(U) \cdot (1/\lambda(U)) = 1$. For each $U \in \mathcal{P}_2^*$ we have $\lambda^h(z_U) \cdot w(z_U) = \lambda(U) \cdot (\pi(U)/\lambda(U)) = \pi(U) > 1$. Hence (T', λ^h) is heavy; see Definition 4.3. \square

¹⁰Recall that $\mathcal{P}^* = \mathcal{P} \setminus \{\gamma(\mathcal{P})\}$. Here $\gamma(\mathcal{P})$ denotes the particular part in \mathcal{P} which includes the tree root $\gamma(T)$.




Figure 4: The construction of the augmented tree associated with the instance and the balanced partition of Figure 2. The new edges and vertices are shown in red. The illustrate the calculation of the length of these edges for a part U_1 of type-I and for a part U_2 of type-II. For each part U of the partition, we indicate the values of $\lambda(U)$ and $\pi(U)$. For simplicity, we have rounded the values to their second decimal.

The heavy sequence.

Definition 7.5. Given a stochastic instance (T, λ) , a balanced partition \mathcal{P} of T , the corresponding augmented tree T' , and a sequence of request $\sigma \sim (T, \lambda)^\tau$, we construct the heavy sequence associated with σ for T' and denoted by σ^h as follows: for each request $r = (t, u) \in \sigma$ that is located on some part $U \in \mathcal{P}^*$ (i.e., $u \in U$), there is a request (z_U, t) in σ^h .

Remark. It is important to notice that σ^h can be constructed in an online fashion: for any time t , the restriction of the σ^h to the requests that arrives before t only depends on the requests that arrives before t in σ .

Proposition 7.6. Given a stochastic instance (T, λ) and the corresponding heavy instance (T', λ^h) , let $\sigma' \sim (T', \lambda^h)^\tau$ denote a random sequence for T' and let $X(\sigma')$ be a random variable depending on σ' . Then, for any $\tau > 0$, it holds that

$$\mathbb{E} \left[X(\sigma') \mid \sigma' \sim (T', \lambda^h)^\tau \right] = \mathbb{E} \left[X(\sigma^h) \mid \sigma \sim (T, \lambda)^\tau \right],$$

where σ^h is the heavy sequence associated with $\sigma \sim (T, \lambda)^\tau$ (see Definition 7.5).

Proof. Suppose we are given a random sequence of requests $\sigma \sim (T, \lambda)^\tau$. We show that $\sigma^h \sim (T', \lambda^h)^\tau$. According to Proposition 2.4, (i) the waiting time between two consecutive requests in σ follows exponential distribution $\text{Exp}(\lambda(T))$; (ii) once a request $r \in \sigma$ arrives, the probability that it is located at a vertex $u \in V(T)$ is $\lambda(u)/\lambda(T)$. This implies that (i) the waiting time between any two consecutive requests arriving at vertices $\bigcup_{U \in \mathcal{P}^*} U$ follows exponential distribution $\text{Exp}(\lambda(\mathcal{P}^*))$; ¹¹ (ii) the probability that a request is located at some vertex of $U \in \mathcal{P}^*$ is $\lambda(U)/\lambda(\mathcal{P}^*)$. Note that

¹¹For simplicity, denote $\lambda(\mathcal{P}^*) := \sum_{U \in \mathcal{P}^*} \sum_{u \in U} \lambda(u)$ in this section.

when a request $r \in \sigma$ arrives at $\bigcup_{U \in \mathcal{P}^*} U$, then a corresponding request in σ^h arrives at vertex $z_U \in V(T')$. As a result, we know that σ^h follows the centralized Poisson arrival model (see Proposition 2.4). Due to the equivalence between the centralized and distributed Poisson arrival model, we thus have $\sigma^h \sim (T', \lambda^h)^\tau$. Since $\sigma' \sim (T', \lambda^h)^\tau$ also follows Poisson arrivals model generated over T' with arrival rates λ^h , we thus have $\mathbb{E}[X(\sigma') \mid \sigma \sim (T', \lambda^h)^\tau] = \mathbb{E}[X(\sigma^h) \mid \sigma \sim (T, \lambda)^\tau]$. \square

7.3 The algorithm

The algorithm GEN works as follows. It is given a stochastic instance (T, λ) , known in advance, and a sequence of requests σ for T , revealed over time. In the pre-processing step, the algorithm computes a balanced partition \mathcal{P} of T (Lemma 7.3), and computes a light instance $(T[\gamma(\mathcal{P})], \sigma|_{\gamma(\mathcal{P})})$ and the heavy instance (T', σ^h) (Section 7.2). Upon arrival of each request, the algorithm updates the sequences of requests $\sigma|_{\gamma(\mathcal{P})}$ and σ^h as described in the previous paragraph.

The algorithm runs PLAN (Algorithm 2) on input (T', σ^h) . Suppose that PLAN serves at time t a set of vertices $\{z_U, U \in \mathcal{P}'\} \subseteq V(T')$ for some subset $\mathcal{P}' \subseteq \mathcal{P}^*$. Then, GEN serves at time t all pending requests on vertices $(\bigcup_{U \in \mathcal{P}'} U) \subseteq V(T)$.

In parallel, the algorithm runs INSTANT (Algorithm 1) on input $(T[\gamma(\mathcal{P})], \sigma|_{\gamma(\mathcal{P})})$, and performs the same services. This finishes the description of the algorithm GEN.

Algorithm 5: GEN

<p>Input: stochastic instance (T, λ) and a random sequence of requests $\sigma \sim (T, \lambda)^\tau$ Output: a valid schedule of σ</p> <ol style="list-style-type: none"> 1 — pre-processing the given instance — 2 produce a balanced partition \mathcal{P} for T (see Lemma 7.3); 3 construct the heavy instance (T', λ^h) (according to Section 7.2); 4 use PLAN (Algorithm 2) to determine the period of the vertices of T'; 5 — Serve the requests — 6 for each request $r \in \sigma$ do 7 if r arrives in $\gamma(\mathcal{P})$ then 8 serve r immediately. 9 if r arrives in a vertex of $U \in \mathcal{P}^*$ then 10 serve r at time $t(r')$ where $r' \in \sigma^h$ is the corresponding request located on z_U and $t(r')$ is the time at which r' is served by PLAN(σ^h).

Observation 7.7. GEN(σ) is a valid schedule for any sequence of requests σ .

Proof. Each time a request $r \in \sigma$ arrives, if r is classified into light sequence $\sigma|_{\gamma(\mathcal{P})}$, then it is served immediately; if r is classified into heavy sequence σ^h , then it is served periodically. \square

7.4 Analysis on GEN's ratio of expectations

In this section, we analyze the ratio of expectations of GEN for arbitrary stochastic instances and show that it can be bounded by 210. We first present in Section 7.4.1 two lower bounds on the expectation of optimal schedule (Lemma 7.8 and Lemma 7.9), then in Section 7.4.2, we upper bound

the expected cost of the schedule produced by our algorithm GEN (Lemma 7.10), and finally we combine the three results in Section 7.4.3 to prove GEN achieves a constant RoE (Theorem 4.5).

In this section, we use the following notation. Given a balance partition \mathcal{P} of a stochastic instance $(T, \boldsymbol{\lambda})$, for each part $U \in \mathcal{P}$ we set $\pi'(U) = 1$, if $U \in \mathcal{P}_1^*$, and $\pi'(U) = \pi(U)$, otherwise.

7.4.1 Lower bounds on the cost of the optimal schedule

Lemma 7.8. *Given a stochastic instance $(T, \boldsymbol{\lambda})$, a balanced partition \mathcal{P} for this instance, and $\tau > 0$, it holds that*

$$\mathbb{E}[\text{cost}(\text{OPT}(\sigma), T)] \geq \frac{3}{16}(1 - e^{-1}) \cdot \tau \cdot \left(\sum_{U \in \mathcal{P}} \pi'(U) \right),$$

where $\text{OPT}(\sigma)$ denotes an optimal schedule for σ and the expectation is taken over all random sequences $\sigma \sim (T, \boldsymbol{\lambda})^\tau$.

Proof. To prove this bound, we define a family $\tilde{\mathcal{T}}$ of edge-disjoint light subtrees of T' . This family $\tilde{\mathcal{T}}$ is defined as follows:

0. we add $T'[\gamma(\mathcal{P})]$;
1. for each $U \in \mathcal{P}_1^*$, we add the subtree $T'[U \cup \{z'_U\}]$;
2. for each $U \in \mathcal{P}_2$, and each child y of $\gamma(U)$ in $T'[U]$, we add $T'[(U \cap V_y) \cup \{\gamma(U)\}]$;

Notice that these subtrees are pairwise edge-disjoint, but two subtrees that corresponds to the same part of type-II share the same root vertex.

For each subtree $\tilde{T} \in \tilde{\mathcal{T}}$ we construct an arrival rate function $\boldsymbol{\lambda}_{\tilde{T}} : V(\tilde{T}) \rightarrow \mathbb{R}_+$. We fix an arbitrary strict total order on the vertices of T . In cases 0 and 1, we simply define $\boldsymbol{\lambda}_{\tilde{T}} := \boldsymbol{\lambda}|_{\tilde{T}}$. In case 2, if \tilde{T} corresponds to a part $U \in \mathcal{P}_2$ and a vertex y that is *not* the smallest¹² child of $\gamma(U)$, then we set $\boldsymbol{\lambda}_{\tilde{T}}(\gamma(U)) = 0$ and $\boldsymbol{\lambda}_{\tilde{T}}(u) = \lambda(u)$ for $u \in V(\tilde{T}) \setminus \{\gamma(U)\}$. If y is the smallest child of $\gamma(U)$, then we simply set $\boldsymbol{\lambda}_{\tilde{T}} := \boldsymbol{\lambda}|_{\tilde{T}}$.

We claim that for each $\tilde{T} \in \tilde{\mathcal{T}}$, the stochastic instance $(\tilde{T}, \boldsymbol{\lambda}_{\tilde{T}})$ is light (see Definition 4.1). For subtrees that are associated with parts of type-II, or with the root part, this simply follows from the definition of balanced parts (Definition 7.1). Now consider a part $U \in \mathcal{P}_1^*$ (of type-I), and its associated subtree $\tilde{T} = T'[U \cup \{z'_U\}]$. Notice that z'_U is the root in \tilde{T} and its only child in \tilde{T} is $\gamma(U)$. Thus, for each $u \in U$, we have $d(u, z'_U) = d(u, \gamma(U)) + w(\gamma(U), z'_U) = d(u, \gamma(U)) + (1 - \pi(U)) / \boldsymbol{\lambda}(U)$. We now calculate $\pi(\tilde{T}, \boldsymbol{\lambda}_{\tilde{T}})$ and show that it is equal to 1, which means that $(\tilde{T}, \boldsymbol{\lambda}_{\tilde{T}})$ is light:

$$\begin{aligned} \pi(\tilde{T}, \boldsymbol{\lambda}_{\tilde{T}}) &= \boldsymbol{\lambda}(z'_U) \cdot d(z'_U, z'_U) + \sum_{u \in U} \boldsymbol{\lambda}(u) \cdot d(u, z'_U) = 0 + \sum_{u \in U} \boldsymbol{\lambda}(u) \left(d(u, \gamma(U)) + \frac{1 - \pi(U)}{\boldsymbol{\lambda}(U)} \right) \\ &= \pi(U) + \boldsymbol{\lambda}(U) \cdot \frac{1 - \pi(U)}{\boldsymbol{\lambda}(U)} = 1. \end{aligned}$$

Thus, by Theorem 5.2, for each $\tilde{T} \in \tilde{\mathcal{T}}$,

$$\mathbb{E}[\text{cost}(\text{OPT}(\sigma'), \tilde{T})] \geq \frac{3}{16}(1 - e^{-1}) \cdot \tau \cdot \pi(\tilde{T}, \boldsymbol{\lambda}_{\tilde{T}}), \quad (10)$$

¹²according to the fixed order.

where $\text{OPT}(\sigma')$ denote the optimal schedule for the sequence of request σ' in T' , and the expectation is taken over all the random sequences $\sigma' \sim (\tilde{T}, \lambda|_{\tilde{T}})^\tau$.

Now, let σ be a sequence of requests for T of duration τ . For each $\tilde{T} \in \tilde{\mathcal{T}}$, we define a request sequence $\sigma_{\tilde{T}}$ for \tilde{T} : for each request $(t, u) \in \sigma$, there is a request (t, u) in $\sigma_{\tilde{T}}$ if and only if $u \in V(\tilde{T})$ and $\lambda_{\tilde{T}}(u) > 0$.

Let $S = \text{OPT}(\sigma)$ denote an optimal schedule for σ . For each $\tilde{T} \in \tilde{\mathcal{T}}$, we define a schedule $S_{\tilde{T}}$ for $\sigma_{\tilde{T}}$ as follows. For each service $s = (R(s), t) \in S$, there is a service $s_{\tilde{T}} := (R(s) \cap V(\tilde{T}), t)$ in $S_{\tilde{T}}$. It is not difficult to see that

- for each $\tilde{T} \in \tilde{\mathcal{T}}$, $S_{\tilde{T}}$ is a valid schedule for $\sigma_{\tilde{T}}$, and particularly $\text{cost}(S_{\tilde{T}}, \sigma_{\tilde{T}}) \geq \text{cost}(\text{OPT}(\sigma_{\tilde{T}}), \tilde{T})$.
- $\text{delay}(S) = \sum_{\tilde{T} \in \tilde{\mathcal{T}}} \text{delay}(S_{\tilde{T}})$.

We now argue that $\text{weight}(S, T) \geq \sum_{\tilde{T} \in \tilde{\mathcal{T}}} \text{weight}(S_{\tilde{T}}, \tilde{T})$. Indeed, since subtrees in $\tilde{\mathcal{T}}$ are pairwise edge-disjoint, it holds that for each service $s \in S$, we have: $\text{weight}(s, T) \geq \sum_{\tilde{T} \in \tilde{\mathcal{T}}} \text{weight}(s_{\tilde{T}}, \tilde{T})$, which implies what we want.

Finally, we show that $\sum_{\tilde{T} \in \tilde{\mathcal{T}}} \pi(\tilde{T}, \lambda_{\tilde{T}}) = \sum_{U \in \mathcal{P}} \pi'(U)$. Let us first consider the root part, and its associated subtree $\tilde{T} = T'[\gamma(\mathcal{P})]$. We have

$$\pi(\tilde{T}, \lambda_{\tilde{T}}) = \pi(T[\gamma(\mathcal{P})], \lambda) = \pi'(\gamma(\mathcal{P})).$$

Now, consider a part $U \in \mathcal{P}_2$ (of type-II). Let $Y \subseteq V(T)$ denote the children of $\gamma(U)$. The part U is associated with subtrees $T[(U \cap V_y) \cup \{\gamma(U)\}]$, for $y \in Y$. The family $\{U \cap V_y, y \in Y\}$ forms a partition of U , and thus we have

$$\begin{aligned} \pi'(U) &= \pi(U) = \sum_{u \in U} \lambda(u) \cdot d(u, \gamma(U)) = \sum_{y \in Y} \sum_{u \in U \cap V_y} \lambda_{\tilde{T}}(u) \cdot d(u, \gamma(U)) \\ &= \sum_{u \in U \cap V_y} \pi(T[(U \cap V_y) \cup \{\gamma(U)\}], \lambda_{\tilde{T}}). \end{aligned}$$

Given $U \in \mathcal{P}_1^*$, let \tilde{T} be the subtree associated. We have proved before that in this case $\pi(\tilde{T}, \lambda_{\tilde{T}}) = 1 = \pi'(U)$ by definition of π' . Finally, we have

$$\begin{aligned} \text{cost}(\text{OPT}(\sigma), T) &= \text{cost}(S, T) = \text{delay}(S) + \text{weight}(S, T) \geq \sum_{\tilde{T} \in \tilde{\mathcal{T}}} \text{delay}(S_{\tilde{T}}) + \sum_{\tilde{T} \in \tilde{\mathcal{T}}} \text{weight}(S_{\tilde{T}}, \tilde{T}) \\ &= \sum_{\tilde{T} \in \tilde{\mathcal{T}}} \text{cost}(S_{\tilde{T}}, \tilde{T}) \geq \sum_{\tilde{T} \in \tilde{\mathcal{T}}} \text{cost}(\text{OPT}(\sigma_{\tilde{T}}), \tilde{T}). \end{aligned}$$

We now take expectation over all the random sequences $\sigma \sim (T, \lambda)^\tau$. It is not difficult to see (the proof is similar as the proof of Proposition 7.6) that for $\tilde{T} \in \tilde{\mathcal{T}}$, $\sigma \sim (T, \lambda)^\tau \implies \sigma_{\tilde{T}} \sim (\tilde{T}, \lambda_{\tilde{T}})^\tau$. Thus,

$$\begin{aligned} \mathbb{E} \left[\text{cost}(\text{OPT}(\sigma), T) \mid \sigma \sim (T, \lambda)^\tau \right] &= \sum_{\tilde{T} \in \tilde{\mathcal{T}}} \mathbb{E} \left[\text{cost}(\text{OPT}(\sigma_{\tilde{T}}), \tilde{T}) \mid \sigma_{\tilde{T}} \sim (\tilde{T}, \lambda_{\tilde{T}})^\tau \right] \\ &\geq \sum_{\tilde{T} \in \tilde{\mathcal{T}}} \frac{3}{16} (1 - e^{-1}) \tau \cdot \pi(\tilde{T}, \lambda_{\tilde{T}}) = \frac{3}{16} (1 - e^{-1}) \tau \sum_{U \in \mathcal{P}} \pi'(U). \end{aligned}$$

This concludes the proof of Lemma 7.8. □

Lemma 7.9. *Let (T, λ) be a stochastic instance, \mathcal{P} a balanced partition for this instance, (T', λ^h) the corresponding heavy instance, and $\tau > 0$. It holds that*

$$\mathbb{E}[\text{cost}(\text{OPT}(\sigma), T)] + \tau \cdot \sum_{U \in \mathcal{P}^*} \pi'(U) \geq \mathbb{E}[\text{cost}(\text{OPT}(\sigma'), T') \mid \sigma' \sim (T', \lambda^h)],$$

where the expectation on the left side is taken over all random sequences $\sigma \sim (T, \lambda)^\tau$.

Proof. Let σ be a sequence of requests for T of duration τ and let S be a schedule for σ . Let σ^h be the corresponding sequence for the heavy instance. We construct a schedule S_h for σ^h in T' as follows. For each service $s \in S$, that serves requests located on $R(s)$, create a service in S_h that serves the points $R(s) \cup \{z_U \mid U \in \mathcal{P}^*, U \cap R(s) \neq \emptyset\}$, with the same service time as s .

It is clear that S_h is a valid schedule for σ^h . It is also clear that $\text{delay}(S, \sigma) = \text{delay}(S_h, \sigma^h)$. We now claim that

$$\text{weight}(S, \sigma) + \sum_{U \in \mathcal{P}^*} N(\sigma|_U) \cdot \frac{\pi'(U)}{\lambda(U)} \geq \text{weight}(S_h, \sigma^h), \quad (11)$$

where $N(\sigma|_U)$ is the number of requests in $\sigma|_U$, i.e., the number of request in σ that are located at some vertex in U . Indeed, for a service $s \in S$, the weight of the rooted tree induced by $R(s) \cup \{z_U \mid U \in \mathcal{P}^*, U \cap R(s) \neq \emptyset\}$ is equal to the weight of the rooted tree induced by $R(s)$ plus the sum of the weights of the edges (z_U, z'_U) for each $U \in \mathcal{P}^*$ such that $U \cap R(s) \neq \emptyset$. To see this, notice that for a service $s \in S$, and a part $U \in \mathcal{P}^*$ such that $U \cap R(s) \neq \emptyset$, the vertex z'_U is contained in the rooted subtree served by s . Now, the weight of (z_U, z'_U) is equal to $\pi'(U)/\lambda(U)$, and for $U \in \mathcal{P}^*$, we have $U \cap R(s) \neq \emptyset$ if and only if $N(\sigma|_U) \geq 1$. This shows inequality (11).

Note that $\text{cost}(S, \sigma) = \text{delay}(S, \sigma) + \text{weight}(S, \sigma)$ and $\text{delay}(S, \sigma) = \text{delay}(S_h, \sigma^h)$. Combining these two inequalities together with (11), we obtain

$$\text{cost}(S, \sigma) + \sum_{U \in \mathcal{P}^*} N(\sigma|_U) \cdot \frac{\pi'(U)}{\lambda(U)} \geq \text{cost}(S_h, \sigma^h) \geq \text{cost}(\text{OPT}(\sigma^h)).$$

In expectation, when $\sigma \sim (T, \lambda)$, this becomes, using linearity of expectation:

$$\mathbb{E}[\text{cost}(\text{OPT}(\sigma), T)] + \sum_{U \in \mathcal{P}^*} \mathbb{E}[N(\sigma|_U)] \cdot \frac{\pi'(U)}{\lambda(U)} \geq \mathbb{E}[\text{cost}(\text{OPT}(\sigma^h))].$$

Besides, the expected number of requests located in a part $U \in \mathcal{P}^*$ is $\mathbb{E}[N(\sigma|_U)] = \tau \cdot \lambda(U)$. Finally, thanks to Proposition 7.6, we know that for any random variable $X(\sigma')$ that depends on a sequence σ' for T' of duration τ , we have $\mathbb{E}[X(\sigma^h) \mid \sigma \sim (T, \lambda)^\tau] = \mathbb{E}[X(\sigma') \mid \sigma' \sim (T', \lambda^h)^\tau]$. This finishes the proof of Lemma 7.9. \square

7.4.2 Upper bound on the cost of GEN

Lemma 7.10. *Let (T, λ) be a stochastic instance, \mathcal{P} the balanced partition computed by the algorithm GEN, (T', λ^h) the corresponding heavy instance, and $\tau > 0$. It holds that*

$$\mathbb{E}[\text{cost}(\text{GEN}(\sigma), T)] \leq \mathbb{E}[\text{cost}(\text{PLAN}(\sigma'), T') \mid \sigma' \sim (T', \lambda^h)^\tau] + \tau \cdot \sum_{U \in \mathcal{P}} \pi'(U),$$

where the expectation on the left side is taken over all sequences $\sigma \sim (T, \lambda)^\tau$.

Proof. Given $\sigma \sim (T, \boldsymbol{\lambda})^\tau$, let $S_I \subseteq \text{GEN}(\sigma)$ be the set of services to corresponds of applying INSTANT during the execution of GEN on σ . Let $S' = \text{GEN}(\sigma) \setminus S_I$ be the set of services ordered by PLAN. It is clear that $\text{cost}(\text{GEN}(\sigma)) = \text{cost}(S_I) + \text{cost}(S')$.

To prove the lemma, we show the two following bounds:

$$\mathbb{E}[\text{cost}(S_I)] = \tau \cdot \pi'(\gamma(\mathcal{P})), \quad (12)$$

$$\mathbb{E}[\text{cost}(S')] \leq \mathbb{E}[\text{cost}(\text{PLAN}(\sigma^h), T') \mid \sigma' \sim (T', \boldsymbol{\lambda}^h)^\tau] + \tau \cdot \sum_{U \in \mathcal{P}^*} \pi'(U), \quad (13)$$

where the expectations are taken over all sequences $\sigma \sim (T, \boldsymbol{\lambda})^\tau$. It is easy to see that these two results implies the statement of the Lemma.

Proof of (12). By definition of $\pi'(\cdot)$, we have $\pi'(\gamma(\mathcal{P})) = \pi(\gamma(\mathcal{P}))$. Recall that $\pi(\gamma(\mathcal{P})) = \sum_{u \in \gamma(\mathcal{P})} \boldsymbol{\lambda}(u) \cdot d(u, \gamma(T))$ is the expected distance from the root to the location of the requests in $\sigma \sim (T[\gamma(\mathcal{P})], \boldsymbol{\lambda}|_{\mathcal{P}})$, which in turn is the expected weight cost of the schedule produced by INSTANT. Since this algorithm serves each request in $\gamma(\mathcal{P})$ immediately when it arrives, the delay cost is 0. Hence, we obtain (12).

Proof of (13). Let σ^h be the heavy sequence for T' associated with σ . We show that

$$\text{cost}(S', T) \leq \text{cost}(\text{PLAN}(\sigma^h), T') + \sum_{U \in \mathcal{P}^*} \sum_{r \in \sigma|_U} d(\ell(r), z'_U), \quad (14)$$

It is easy to check that $\text{delay}(S', T) = \text{delay}(\text{PLAN}(\sigma^h), T')$. Now we claim that

$$\text{weight}(S', T) \leq \text{weight}(\text{PLAN}(\sigma^h), T') + \sum_{U \in \mathcal{P}^*} \sum_{r \in \sigma|_U} d(\ell(r), z'_U), \quad (15)$$

which together with the previous equation on the delay cost implies (14). The only thing left now is to prove (15).

Let s_h be a service in $\text{PLAN}(\sigma^h)$ and let s_g denote the corresponding service in S' . Let $U \in \mathcal{P}^*$ such that $z_U \in R(s_h)$. Let $t' < t(s_h)$ be the latest time at which $\text{PLAN}(\sigma^h)$ served z_U . It is clear by the definition of GEN that all requests $(t, u) \in \sigma$ with $u \in U$ and $t \leq t'$ have been served by GEN at time t' , and that all requests $(t, u) \in \sigma$ with $u \in U$ and $t' < t \leq t(s_h)$ are served by s_g .

Let Q_h denote the subtree in T served by s_h and let Q_g be the subtree of T' served by s_g . For a request $r \in \sigma$ located in some $U \in \mathcal{P}^*$, let P_r denote the path in T' from $\ell(r)$ to z'_U . If $z_U \in Q_h$, then it is necessarily $z'_U \in Q_h$, so $Q_h \cup P_r$ is connected. Thus, the subtree

$$Q'_g := Q_h \cup \left(\bigcup_{U \in \mathcal{P}^*, z_U \in R(s_h)} \bigcup_{r \in R(s_g), \ell(r) \in U} P_r \right)$$

contains all the requests served by s_g , which implies that

$$w(Q_g) \leq w(Q'_g) \leq w(Q_h) + \sum_{U \in \mathcal{P}^*, z_U \in R(s_h)} \sum_{r \in R(s_g), \ell(r) \in U} d(\ell(r), z'_U).$$

We obtain (15) by summing this inequality over all services $\sigma^h \in \text{PLAN}(\sigma^h)$. Hence, we have proved (14).

To obtain (13) from (14), we take expectation over all random sequences $\sigma \sim (T, \boldsymbol{\lambda})^\tau$. First, by Proposition 7.6, we know that

$$\mathbb{E}[\text{cost}(\text{PLAN}(\sigma^h), T') \mid \sigma^h \sim (T', \boldsymbol{\lambda}^h)^\tau] = \mathbb{E}[\text{cost}(\text{PLAN}(\sigma), T') \mid \sigma \sim (T, \boldsymbol{\lambda})^\tau].$$

Finally, in expectation we have

$$\begin{aligned} \mathbb{E} \left[\sum_{U \in \mathcal{P}^*} \sum_{r \in \sigma|_U} d(\ell(r), z'_U) \right] &= \sum_{U \in \mathcal{P}^*} \sum_{u \in U} \mathbb{E}[\text{number of requests of } \sigma \text{ located on } u] \cdot d(u, z'_U) \\ &= \sum_{U \in \mathcal{P}^*} \sum_{u \in U} (\boldsymbol{\lambda}(u) \cdot \tau) \cdot d(u, z'_U) = \tau \cdot \sum_{U \in \mathcal{P}^*} \pi'(U). \end{aligned}$$

This concludes the proof of (13), and thus the proof of the lemma. \square

7.4.3 Proof of Theorem 4.5

Let $(T, \boldsymbol{\lambda})$ be a stochastic instance and let \mathcal{P} be the balanced partition of T . Let $(T', \boldsymbol{\lambda}^h)$ denote the corresponding heavy instance, and let $\tau > 0$. Taking expectation over all the sequences $\sigma \sim (T, \boldsymbol{\lambda})^\tau$, we obtain:

$$\begin{aligned} &\mathbb{E}_\sigma^\tau[\text{cost}(\text{GEN}(\sigma), T)] \\ &\leq \mathbb{E} \left[\text{cost}(\text{PLAN}(\sigma'), T') \mid \sigma' \sim (T', \boldsymbol{\lambda}^h)^\tau \right] + \tau \sum_{U \in \mathcal{P}} \pi'(U) && \text{(Lemma 7.10)} \\ &\leq \frac{64}{3} \cdot \mathbb{E} \left[\text{cost}(\text{OPT}(\sigma'), T') \mid \sigma' \sim (T', \boldsymbol{\lambda}^h)^\tau \right] + \tau \sum_{U \in \mathcal{P}} \pi'(U) && \text{(Thm. 4.4 and Prop. 7.4)} \\ &\leq \frac{64}{3} \left(\mathbb{E}[\text{cost}(\text{OPT}(\sigma), T)] + \tau \sum_{U \in \mathcal{P}^*} \pi'(U) \right) + \tau \sum_{U \in \mathcal{P}} \pi'(U) && \text{(Lemma 7.9)} \\ &\leq \frac{64}{3} \cdot \mathbb{E}[\text{cost}(\text{OPT}(\sigma), T)] + \left(\frac{64}{3} + 1 \right) \left(\tau \cdot \sum_{U \in \mathcal{P}} \pi'(U) \right) \\ &\leq \frac{64}{3} \cdot \mathbb{E}[\text{cost}(\text{OPT}(\sigma), T)] + \frac{67}{3} \cdot \frac{16}{3(1-e^{-1})} \cdot \mathbb{E}[\text{cost}(\text{OPT}(\sigma), T)] && \text{(Lemma 7.8)} \\ &< 210 \cdot \mathbb{E}[\text{cost}(\text{OPT}(\sigma), T)]. \end{aligned}$$

This proves the theorem.

8 Other related works

The MLA problem was first introduced by Bienkowski et al. [16] and they study a more general version in their paper, where the cost of delaying a request r by a duration t is $f_r(t)$. Here $f_r(\cdot)$ denotes the delay cost function of r and it only needs to be non-decreasing and satisfy $f_r(0) = 0$. An $O(d^4 2^d)$ -competitive online algorithm is proposed for this general delay cost version problem, where d denotes the depth of the given tree. Besides, a deadline version of MLA is also considered

in [16], where each request r has a time window (between its arrival and its deadline) and it has to be served no later than its deadline. The target is to minimize the total service cost for serving all the requests. For this deadline version problem, they proposed an online algorithm with better competitive ratio of $d^2 2^d$. Later, the competitiveness of MLA are further improved to $O(d^2)$ by Azar and Touitou [13] for the general delay cost version and to $O(d)$ by Buchbinder et al. [27] for the deadline version.¹³ However, for the delay cost version of MLA, no matching lower bound has been found thus far — the current best lower bound on MLA (with delays) is only 4 [16, 17, 18], restricted to a path case with linear delays. In the offline setting, MLA is NP-hard in both delay and deadline versions [3, 15] and a 2-approximation algorithm was proposed by Becchetti et al. [15] for the deadline version. For a special path case of MLA with the linear delay, Bienkowski et al. [22] proved that the competitiveness is between 3.618 and 5, improving on an earlier 8-competitive algorithm given by Brito et al. [26]. Thus far, no previous work has studied MLA in the stochastic input model, no matter the delay or deadline versions.

Two special cases of MLA with linear delays, one TCP-acknowledgment (equivalent to MLA with tree being an edge, i.e. $d = 1$) and one Joint Replenishment (abbrev. JRP, equivalent to MLA with tree being a star, i.e. $d = 2$) are of particular interests. This is because, TCP-acknowledgment (a.k.a. single item lot-sizing problem in operation research community, see e.g. [25, 41, 60, 29, 44]) models the data transmission issue from sensor networks (see e.g. [67, 51]), while JRP models the inventory control issue from supply chain management (see e.g. [5, 37, 42, 63, 48]). For TCP-acknowledgment, in the online setting there exists an optimal 2-competitive deterministic algorithm [32] and an optimal $\frac{1}{1-e^{-1}}$ -competitive randomized algorithm [45, 62]; in the offline setting, the problem can be solved in $O(n \log n)$ time, where n denotes the number of requests [1]. For JRP, there exists a 3-competitive online algorithm based on primal-dual method proposed by Buchbinder et al. [28], and no online algorithm achieves competitive ratio less than 2.754 [21]. In the offline setting, JRP is NP-hard [3] and also APX-hard [59, 20]. The current best approximation ratio for JRP is 1.791 (Bienkowski et al. [21]), improving on earlier results given by Levi et al. [53, 54, 52]. For a deadline version of JRP, Bienkowski et al. [21] proposed an optimal 2-competitive online algorithm. For the stochastic version, unfortunately, to our best extent, we did not find previous works on these two problems with requests following Poisson arrival model from a theoretical perspective, i.e., proposing online algorithm with their performances evaluated using RoE.

Another problem, called online service with delays (OSD), first introduced by Azar et al. [9], is closely related to MLA (with linear delays). In this OSD problem, a n -points metric space is given as input. The requests arrive at metric points over time and a server is available to serve the requests. The target is to serve all the requests in an online manner such that their total delay cost plus the total distance travelled by the server is minimized. Note that MLA can be seen as a special case of OSD when the given metric is a tree and the server has to always come back to a particular tree vertex immediately after serving some requests elsewhere. For OSD, Azar et al. [9] proposed a $O(\log^4 n)$ -competitive online algorithm in their paper. Later, the competitive ratio for OSD is improved from $O(\log^2 n)$ (by Azar and Touitou [13]) to $O(\log n)$ (by Touitou [66]).

We remark here that besides MLA and OSD, many other online problems with delays/deadline have also drawn a lot of attentions recently, such as online matching with delays [33, 6, 4, 24, 23, 34, 10, 31, 55, 12, 58, 56, 49], facility location with delays/deadline [19, 13, 14], Steiner tree with delays/deadline [14], bin packing with delays [8, 35, 36, 2], set cover with delays [7, 64, 50], paging with delays/deadline [38, 39], list update with delays/deadline [11], and many others [58, 30, 65,

¹³Later, McMahan [57] further improve the competitive ratio to d for the deadline version of MLA.

40, 43, 46].

9 Concluding remarks

In this paper, we studied MLA with additional stochastic assumptions on the sequence of the input requests. In the case where the requests follow a Poisson arrival process, we presented a deterministic online algorithm with constant RoE. In the following text, we briefly discuss some potential future directions.

Does the greedy algorithm achieve a constant RoE? An intuitive heuristic algorithm for MLA is *Greedy*, which works as follows: *each time when a set of requests R arriving at vertices $U \subseteq V(T)$ have the total delay cost equal to the weight of the minimal subtree of T including γ and U , serve all the requests R .* Does this greedy algorithm achieves a constant RoE?

Generalize MLA with edge capacity and k tree roots. One practical scenario on MLA is that each edge has a capacity on the maximum number of requests served in one service if this edge is used, such as [60, 44, 63]. We conjecture that some $O(1)$ -RoE online algorithm can be proposed for this generalized MLA with edge capacity. Another generalized version of MLA is to assume k tree roots available for serving requests concurrently. That is, a set of pending requests can be served together by connecting to any of k servers. The question is, how to design an online algorithm for this k -MLA problem? Does there exist $O(1)$ -RoE algorithm still?

Other online network design problems with delays in the Poisson arrival model.

Recall that the online problems of service with delays (and its generalization called k -services with delays), facility location with delays, Steiner tree/forest with delays are all closely related to MLA with delays. Does there exist online algorithm with $O(1)$ -RoE for each problem?

Acknowledgement

This work is supported by ERC CoG grant TUGBOAT 772346, NCN grant 2020/37/B/ST6/04179 and NCN grant 2022/45/B/ST6/00559.

References

- [1] A. AGGARWAL AND J. K. PARK, *Improved algorithms for economic lot size problems*, Operations research, 41 (1993), pp. 549–571.
- [2] L. AHLROTH, A. SCHUMACHER, AND P. ORPONEN, *Online bin packing with delay and holding costs*, Operations Research Letters, 41 (2013), pp. 1–6.
- [3] E. ARKIN, D. JONEJA, AND R. ROUNDY, *Computational complexity of uncapacitated multi-echelon production planning problems*, Operations research letters, 8 (1989), pp. 61–66.
- [4] I. ASHLAGI, Y. AZAR, M. CHARIKAR, A. CHIPLUNKAR, O. GERI, H. KAPLAN, R. MAKHIJANI, Y. WANG, AND R. WATTENHOFER, *Min-cost bipartite perfect matching with delays*, in Proc. APPROX / RANDOM, 2017, pp. 1:1–1:20.

- [5] Y. ASKOY AND S. S. ERENGUK, *Multi-item inventory models with coordinated replenishment: a survey*, International Journal of Operations and Production Management, 8 (1988), pp. 63–73.
- [6] Y. AZAR, A. CHIPLUNKAR, AND H. KAPLAN, *Polylogarithmic bounds on the competitiveness of min-cost perfect matching with delays*, in Proc. SODA, 2017, pp. 1051–1061.
- [7] Y. AZAR, A. CHIPLUNKAR, S. KUTTEN, AND N. TOUITOU, *Set cover with delay-clairvoyance is not required*, in Proc. ESA, 2020, pp. 8:1–8:21.
- [8] Y. AZAR, Y. EMEK, R. VAN STEE, AND D. VAINSTEIN, *The price of clustering in bin-packing with applications to bin-packing with delays*, in Proc. SPAA, 2019, pp. 1–10.
- [9] Y. AZAR, A. GANESH, R. GE, AND D. PANIGRAHI, *Online service with delay*, in Proc. STOC, 2017, pp. 551–563.
- [10] Y. AZAR AND A. JACOB-FANANI, *Deterministic min-cost matching with delays*, Theory of Computing Systems, 64 (2020), pp. 572–592.
- [11] Y. AZAR, S. LEWKOWICZ, AND D. VAINSTEIN, *List update with delays or time windows*, in Proc. ICALP, 2024, pp. 15:1–15:20.
- [12] Y. AZAR, R. REN, AND D. VAINSTEIN, *The min-cost matching with concave delays problem*, in Proc. SODA, 2021, pp. 301–320.
- [13] Y. AZAR AND N. TOUITOU, *General framework for metric optimization problems with delay or with deadlines*, in Proc. FOCS, 2019, pp. 60–71.
- [14] ———, *Beyond tree embeddings—a deterministic framework for network design with deadlines or delay*, in Proc. FOCS, 2020, pp. 1368–1379.
- [15] L. BECCHETTI, A. MARCHETTI-SPACCAMELA, A. VITALETTI, P. KORTEWEG, M. SKUTELLA, AND L. STOUGIE, *Latency-constrained aggregation in sensor networks*, ACM Transactions on Algorithms, 6 (2009), pp. 1–20.
- [16] M. BIENKOWSKI, M. BÖHM, J. BYRKA, M. CHROBAK, C. DÜRR, L. FOLWARCZNY, Ł. JEŻ, J. SGALL, N. K. THANG, AND P. VESELÝ, *Online algorithms for multi-level aggregation*, in Proc. ESA, 2016, pp. 12:1–12:17.
- [17] ———, *Online algorithms for multilevel aggregation*, Operations Research, 68 (2020), pp. 214–232.
- [18] ———, *New results on multi-level aggregation*, Theoretical Computer Science, 861 (2021), pp. 133–143.
- [19] M. BIENKOWSKI, M. BÖHM, J. BYRKA, AND J. MARCINKOWSKI, *Online facility location with linear delay*, in Proc. APPROX/RANDOM, 2022, pp. 45:1–45:17.
- [20] M. BIENKOWSKI, J. BYRKA, M. CHROBAK, N. DOBBS, T. NOWICKI, M. SVIRIDENKO, G. ŚWIRSZCZ, AND N. E. YOUNG, *Approximation algorithms for the joint replenishment problem with deadlines*, Journal of Scheduling, 18 (2015), pp. 545–560.

- [21] M. BIENKOWSKI, J. BYRKA, M. CHROBAK, Ł. JEŻ, D. NOGNENG, AND J. SGALL, *Better approximation bounds for the joint replenishment problem*, in Proc. SODA, 2014, pp. 42–54.
- [22] M. BIENKOWSKI, J. BYRKA, M. CHROBAK, Ł. JEŻ, J. SGALL, AND G. STACHOWIAK, *Online control message aggregation in chain networks*, in Proc. WADS, 2013, pp. 133–145.
- [23] M. BIENKOWSKI, A. KRASKA, H.-H. LIU, AND P. SCHMIDT, *A primal-dual online deterministic algorithm for matching with delays*, in Proc. WAOA, 2018, pp. 51–68.
- [24] M. BIENKOWSKI, A. KRASKA, AND P. SCHMIDT, *A match in time saves nine: Deterministic online matching with delays*, in Proc. WAOA, 2017, pp. 132–146.
- [25] N. BRAHIMI, S. DAUZERE-PERES, N. M. NAJID, AND A. NORDLI, *Single item lot sizing problems*, European Journal of Operational Research, 168 (2006), pp. 1–16.
- [26] C. F. BRITO, E. KOUTSOPIAS, AND S. VAYA, *Competitive analysis of organization networks or multicast acknowledgment: How much to wait?*, Algorithmica, 64 (2012), pp. 584–605.
- [27] N. BUCHBINDER, M. FELDMAN, J. NAOR, AND O. TALMON, *O (depth)-competitive algorithm for online multi-level aggregation*, in Proc. SODA, 2017, pp. 1235–1244.
- [28] N. BUCHBINDER, T. KIMBRELT, R. LEVI, K. MAKARYCHEV, AND M. SVIRIDENKO, *Online make-to-order joint replenishment model: primal dual competitive algorithms*, in Proc. SODA, 2008, pp. 952–961.
- [29] M. A. BUSHUEV, A. GUIFFRIDA, M. JABER, AND M. KHAN, *A review of inventory lot sizing review papers*, Management Research Review, 38 (2015), pp. 283–298.
- [30] R. CHEN, J. KHATKAR, AND S. W. UMBOH, *Online weighted cardinality joint replenishment problem with delay*, in Proc. ICALP, 2022.
- [31] L. DERYCKERE AND S. W. UMBOH, *Online matching with set and concave delays*, in Proc. APPROX/RANDOM, 2023, pp. 17:1–17:17.
- [32] D. R. DOOLY, S. A. GOLDMAN, AND S. D. SCOTT, *On-line analysis of the tcp acknowledgment delay problem*, Journal of the ACM, 48 (2001), pp. 243–273.
- [33] Y. EMEK, S. KUTTEN, AND R. WATTENHOFER, *Online matching: haste makes waste!*, in Proc. STOC, 2016, pp. 333–344.
- [34] Y. EMEK, Y. SHAPIRO, AND Y. WANG, *Minimum cost perfect matching with delays for two sources*, Theoretical Computer Science, 754 (2019), pp. 122–129.
- [35] L. EPSTEIN, *On bin packing with clustering and bin packing with delays*, Discrete Optimization, 41 (2021), p. 100647.
- [36] ———, *Open-end bin packing: new and old analysis approaches*, Discrete Applied Mathematics, 321 (2022), pp. 220–239.
- [37] S. K. GOYAL AND A. T. SATIR, *Joint replenishment inventory control: deterministic and stochastic models*, European journal of operational research, 38 (1989), pp. 2–13.

- [38] A. GUPTA, A. KUMAR, AND D. PANIGRAHI, *Caching with time windows*, in Proc. STOC, 2020, pp. 1125–1138.
- [39] ———, *A hitting set relaxation for k -server and an extension to time-windows*, in Proc. FOCS, 2022, pp. 504–515.
- [40] S. IM, B. MOSELEY, C. XU, AND R. ZHANG, *Online dynamic acknowledgement with learned predictions*, in Proc. INFOCOM, 2023, pp. 1–10.
- [41] R. JANS AND Z. DEGRAEVE, *Modeling industrial lot sizing problems: a review*, International Journal of Production Research, 46 (2008), pp. 1619–1643.
- [42] D. JONEJA, *The joint replenishment problem: new heuristics and worst case performance bounds*, Operations Research, 38 (1990), pp. 711–723.
- [43] N. KAKIMURA AND T. NAKAYOSHI, *Deterministic primal-dual algorithms for online k -way matching with delays*, in Proc. ICCV, 2023, pp. 238–249.
- [44] B. KARIMI, S. F. GHOMI, AND J. WILSON, *The capacitated lot sizing problem: a review of models and algorithms*, Omega, 31 (2003), pp. 365–378.
- [45] A. R. KARLIN, C. KENYON, AND D. RANDALL, *Dynamic tcp acknowledgement and other stories about $e/(e-1)$* , in Proc. STOC, 2001, pp. 502–509.
- [46] Y. KAWASE AND T. NAKAYOSHI, *Online matching with delays and size-based costs*, arXiv preprint arXiv:2408.08658, (2024).
- [47] S. KHANNA, J. S. NAOR, AND D. RAZ, *Control message aggregation in group communication protocols*, in Proc. ICALP, 2002, pp. 135–146.
- [48] M. KHOUJA AND S. GOYAL, *A review of the joint replenishment problem literature: 1989–2005*, European journal of operational Research, 186 (2008), pp. 1–16.
- [49] T.-W. KUO, *Online deterministic minimum cost bipartite matching with delays on a line*, arXiv preprint arXiv:2408.02526, (2024).
- [50] N. M. LE, S. WILLIAM UMBOH, AND N. XIE, *The power of clairvoyance for multi-level aggregation and set cover with delay*, in Proc. SODA, 2023, pp. 1594–1610.
- [51] K.-C. LEUNG, V. O. LI, AND D. YANG, *An overview of packet reordering in transmission control protocol (tcp): problems, solutions, and challenges*, IEEE Transactions on Parallel and Distributed Systems, 18 (2007), pp. 522–535.
- [52] R. LEVI, R. ROUNDY, D. SHMOYS, AND M. SVIRIDENKO, *A constant approximation algorithm for the one-warehouse multiretailer problem*, Management Science, 54 (2008), pp. 763–776.
- [53] R. LEVI, R. ROUNDY, AND D. B. SHMOYS, *Primal-dual algorithms for deterministic inventory problems*, in Proc. STOC, 2004, pp. 353–362.
- [54] R. LEVI AND M. SVIRIDENKO, *Improved approximation algorithm for the one-warehouse multi-retailer problem*, in Proc. APPROX-RANDOM, 2006, pp. 188–199.

- [55] X. LIU, Z. PAN, Y. WANG, AND R. WATTENHOFER, *Impatient online matching*, in Proc. ISAAC, vol. 123, 2018, pp. 62:1–62:12.
- [56] M. MARI, M. PAWŁOWSKI, R. REN, AND P. SANKOWSKI, *Online matching with delays and stochastic arrival times*, in Proc. AAMAS, 2023, p. 976–984.
- [57] J. MCMAHAN, *A d -competitive algorithm for the multilevel aggregation problem with deadlines*, arXiv preprint arXiv:2108.04422, (2021).
- [58] D. MELNYK, Y. WANG, AND R. WATTENHOFER, *Online k -way matching with delays and the h -metric*, arXiv preprint arXiv:2109.06640, (2021).
- [59] T. NONNER AND A. SOUZA, *Approximating the joint replenishment problem with deadlines*, Discrete Mathematics, Algorithms and Applications, 1 (2009), pp. 153–173.
- [60] D. QUADT AND H. KUHN, *Capacitated lot-sizing with extensions: a review*, Operation Research, 6 (2008), pp. 61–83.
- [61] S. M. ROSS, *Stochastic processes*, vol. 2, Wiley New York, 1996.
- [62] S. S. SEIDEN, *A guessing game and randomized online algorithms*, in Proc. STOC, 2000, pp. 592–601.
- [63] S. SINDHUCHAO, H. E. ROMELJN, E. AKÇALI, AND R. BOONDISKULCHOK, *An integrated inventory-routing system for multi-item joint replenishment with limited vehicle capacity*, Journal of Global Optimization, 32 (2005), pp. 93–118.
- [64] N. TOUITOU, *Nearly-tight lower bounds for set cover and network design with deadlines/delay*, in Proc. ISAAC, 2021, pp. 53:1–53:16.
- [65] ———, *Frameworks for nonclairvoyant network design with deadlines or delay*, in Proc. ICALP, 2023, pp. 105:1–105:20.
- [66] ———, *Improved and deterministic online service with deadlines or delay*, in Proc. STOC, 2023, pp. 761–774.
- [67] W. YUAN, S. V. KRISHNAMURTHY, AND S. K. TRIPATHI, *Synchronization of multiple levels of data fusion in wireless sensor networks*, in Proc. GLOBECOM, vol. 1, 2003, pp. 221–225.

A Missing Proofs in Section 2

We first introduce the two well-known properties of the exponential distribution, which will be used to prove Proposition 2.3, Proposition 2.3 and Proposition 2.4.

Proposition A.1 (memoryless property). *If X is an exponential variable with parameter λ , then for all $s, t \geq 0$, we have*

$$\mathbb{P}(X > s + t \mid X > s) = \mathbb{P}(X > t) = e^{-\lambda t}.$$

Proposition A.2. *Given n independent exponential variables $X_i \sim \text{Exp}(\lambda_i)$ for $i \in [n]$, let $Z := \min\{X_1, X_2, \dots, X_n\}$ and let $\lambda := \sum_{i=1}^n \lambda_i$. It holds that*

1. $Z \sim \text{Exp}(\lambda)$,
2. $\mathbb{P}(Z = X_i) = \lambda_i/\lambda$,
3. $Z \perp \{Z = X_i\}$,

where \perp denotes independence.

To illustrate the equivalence between the distributed Poisson arrival model (Definition 2.1) and the centralized Poisson arrival model (Proposition 2.4), see Figures 5 and 6 for visual supports.




Figure 5: Example showing the correspondence between distributed Poisson arrival model and exponential timers without resets. Y_x^i and $T_x^{x,i}$ both represent the waiting time between the $(i-1)$ -th and the i -th request arriving at any vertex $x \in V(T)$. The graph on the right highlights in blue the waiting times between the consecutive arrivals from the perspective of the tree vertices.




Figure 6: Example showing the correspondence between the centralized Poisson arrival model and timers with resets. On the left, the waiting time between appearances of the i -th and the $(i+1)$ -th requests is determined by the variable $Y_{i+1} \sim \text{Exp}(\lambda(T))$. On the right, a double-headed arrow represents a timer, while a single-headed arrow means that we had to reset a timer.

A.1 Proof of Proposition 2.4

Consider the distributed model where for each vertex $u \in V(T)$, we define an exponential variable Y_u^1 representing the time of arrival of the first request located at u . If we look at the whole tree, the time of arrival of the first request r is determined by the minimum of all these variables, $\min_{u \in V(T)} Y_u^1$. We denote this variable by Y_1 . By Proposition A.2, we know that Y_1 follows an exponential distribution with parameter $\lambda(T)$ being the sum of components' parameters. Moreover, by the second property presented in this proposition, we know that the probability of r arriving at vertex u equals $\lambda(u)/\lambda(T)$ for each $u \in V(T)$.

At time Y_1 when the first request arrives, we associate each vertex $u \neq l(r)$ with a new independent exponential random variable $Z_u^1 \sim \text{Exp}(\boldsymbol{\lambda}(u))$. By the memoryless property from Proposition A.1, for each vertex u the arrival time determined by $t(r) + Z_v^1$ follows the same distribution as Y_u^1 conditioned on being greater than $t(r)$. This shows that we can look at the first request arrival as it was defined by the centralized model and the consequent requests still follow the distributed model. We can continue this process to transform the distributed model into a centralized one.

A.2 Proof of Proposition 2.3

We only need to prove that, given $\sigma_1 \sim (T, \boldsymbol{\lambda})^{\tau_1}$ and $\sigma_2 \sim (T, \boldsymbol{\lambda})^{\tau_2}$, denoting by σ a merged sequence of σ_1 and σ_2 (each request in σ_2 has its arrival time delayed by a duration τ_1), we have $\sigma \sim (T, \boldsymbol{\lambda})^{\tau_1 + \tau_2}$. Indeed, given any sequence $\sigma_2 \sim (T, \boldsymbol{\lambda})^{\tau_2}$ and increasing the arrival time of each request by τ_1 , we obtain a sequence σ'_2 generated according to Poisson arrivals during time interval $[\tau_1, \tau_1 + \tau_2]$. Let the arrival time of the last request in $\sigma_1 \sim (T, \boldsymbol{\lambda})^{\tau_1}$ be t and we thus know that no request is generated during time interval $[t, \tau_1]$. By the memoryless property from Proposition A.1, we thus know that σ'_2 also follows Poisson arrivals during time interval $[t, \tau_1 + \tau_2]$. As a result, by Definition 2.1, we can conclude that $\sigma = \sigma_1 \sigma'_2$ follows Poisson arrival model during time interval $[0, \tau_1 + \tau_2]$ and hence this proposition.

A.3 Proof of Proposition 2.6

To prove this proposition, we show that the joint density function of the n requests' arrival times is identical to the joint density function of the order statistics corresponding to n independent random variables uniformly distributed over $[0, \tau]$. Let X_1, X_2, \dots, X_n denote the n independent identical variables, each with the same density function $f(x)$. Let $X_{(1)}, X_{(2)}, \dots, X_{(n)}$ denote these n variables in an increasing order, i.e., for each $i \in [n]$, $X_{(i)}$ is the i -th smallest variable among X_1, X_2, \dots, X_n . Denoting by $x_1 < x_2 < \dots < x_n$, then the joint density of $X_{(1)}, X_{(2)}, \dots, X_{(n)}$ is

$$f(x_1, x_2, \dots, x_n) = n! \prod_{i=1}^n f(x_i).$$

As a result, given n independent variables U_1, U_2, \dots, U_n , each of which is uniformly drawn from $[0, \tau]$ and hence has a density function of $g(t) = \frac{1}{\tau}$, the density function of $U_{(1)}, U_{(2)}, \dots, U_{(n)}$ is

$$g(t_1, \dots, t_n) = \frac{n!}{\tau^n}.$$

Now, let W_i denote the waiting time between the $(i-1)$ -th request's arrival and the i -th request's arrival, which follows an exponential distribution $\text{Exp}(\lambda)$. Let $S_i = \sum_{j=1}^i W_j$ denote the arrival time of the i -th request, according to Poisson arrival model. Now we derive the joint density function of $S_{(1)}, \dots, S_{(n)}$ (denoted by $f(t_1, \dots, t_n)$), given that n requests are generated within $[0, \tau]$ (i.e., $N(\sigma) = n$). Given $0 < t_1 < t_2 < \dots < t_n < t_{n+1} = \tau$, for each $i \in [n]$, let $h_i > 0$ be a value small enough so that $t_i + h_i < t_{i+1}$. By definition of Poisson arrival process, given that the Poisson arrival rate λ , the probability to generate $n \geq 0$ requests during any interval of length ℓ is equal to

$$e^{-\lambda \ell} \cdot \frac{(\lambda \ell)^n}{n!}.$$

Therefore, we have

$$\begin{aligned}
& \mathbb{P}(t_i \leq S_i \leq t_i + h_i \text{ for each } i \in [n] \mid N(\sigma) = n) \\
&= \frac{\prod_{i=1}^n \mathbb{P}(\text{exactly 1 request generated in } [t_i, t_i + h_i])}{\mathbb{P}(N(\sigma) = n)} \cdot \mathbb{P}(\text{no request arrives elsewhere in } [0, \tau]) \\
&= \frac{\prod_{i=1}^n \lambda h_i \cdot e^{-\lambda h_i}}{e^{-\lambda \tau} \cdot \frac{(\lambda \tau)^n}{n!}} \cdot e^{\lambda(\tau - h_1 - \dots - h_n)} = \frac{n!}{\tau^n} \cdot h_1 \cdot h_2 \cdot \dots \cdot h_n.
\end{aligned}$$

As a result,

$$\frac{\mathbb{P}(t_i \leq S_i \leq t_i + h_i \text{ for each } i \in [n] \mid N(\tau) = n)}{h_1 \cdot h_2 \cdot \dots \cdot h_n} = \frac{n!}{\tau^n}.$$

By letting $h_i \rightarrow 0$ for each i , the conditional joint density of S_1, \dots, S_n given $N(\tau) = n$ becomes

$$f(t_1, \dots, t_n) = \frac{n!}{\tau^n}.$$

Since $f(t_1, \dots, t_n) = g(t_1, \dots, t_n)$, we thus have Proposition 2.6.

A.4 Proof of Proposition 2.7

Again, let W_i denote the waiting time between the $(i-1)$ -th request's arrival time and the i -th request's arrival time, which follows an exponential distribution $\text{Exp}(\lambda)$. Let $S_i = \sum_{j=1}^i W_j$ denote the i -th request's arrival time. We first calculate the expected total delay cost under the condition that $N(\sigma) = n$ requests are generated within $[0, \tau]$.

$$\mathbb{E}_\sigma^\tau \left[\sum_{i=1}^{N(\sigma)} (\tau - S_i) \mid N(\sigma) = n \right] = n\tau - \mathbb{E}_\sigma^\tau \left[\sum_{i=1}^{N(\sigma)} S_i \mid N(\sigma) = n \right].$$

Let U_1, U_2, \dots, U_n denote the n uniform variables drawn from $[0, \tau]$ and let $U_{(1)}, U_{(2)}, \dots, U_{(n)}$ denote these variables in an increasing order. We thus have $\sum_{i=1}^n U_i = \sum_{i=1}^n U_{(i)}$ and

$$\mathbb{E} \left[\sum_{i=1}^n U_i \right] = \mathbb{E} \left[\sum_{i=1}^n U_{(i)} \right].$$

By Proposition 2.6, we further have

$$\mathbb{E}_\sigma^\tau \left[\sum_{i=1}^{N(\sigma)} S_i \mid N(\sigma) = n \right] = \mathbb{E} \left[\sum_{i=1}^n U_{(i)} \right].$$

As a result, we have

$$\mathbb{E}_\sigma^\tau \left[\sum_{i=1}^{N(\sigma)} (\tau - S_i) \mid N(\sigma) = n \right] = n\tau - \mathbb{E} \left[\sum_{i=1}^n U_i \right] = n\tau - \sum_{i=1}^n \mathbb{E}[U_i] = n\tau - \frac{n\tau}{2} = \frac{n\tau}{2}$$

and hence,

$$\mathbb{E}_\sigma^\tau \left[\sum_{i=1}^{N(\sigma)} (\tau - S_i) \right] = \frac{\tau}{2} \cdot \mathbb{E}_\sigma^\tau[N(\sigma)] = \frac{\tau}{2} \cdot \boldsymbol{\lambda}(T) \cdot \tau = \frac{1}{2} \cdot \boldsymbol{\lambda}(T) \cdot \tau^2.$$

B Missing contexts in Section 4

Lemma B.1. *There exists a stochastic instance (T, λ) such that both $\frac{\mathbb{E}[\text{cost}(\text{INSTANT}(\sigma)) | \sigma \sim (T, \lambda)^\tau]}{\mathbb{E}[\text{cost}(\text{OPT}(\sigma)) | \sigma \sim (T, \lambda)^\tau]}$ and $\frac{\mathbb{E}[\text{cost}(\text{PLAN}(\sigma)) | \sigma \sim (T, \lambda)^\tau]}{\mathbb{E}[\text{cost}(\text{OPT}(\sigma)) | \sigma \sim (T, \lambda)^\tau]}$ are unbounded.*

Consider such an MLA instance (T, λ) where T has a depth of 2. The root vertex γ has only one child u with $\lambda(u) = 0$ and u has n child vertices v_1, \dots, v_n , with $\lambda(v_i) = \frac{1}{n}$ for each $i \in [n]$. The edge (γ, u) has a weight of \sqrt{n} and each edge (u, v_i) has a weight of 1. By definition, such an instance is neither light nor heavy. Note that if INSTANT is applied to deal with this instance, then to serve each request, only a weight cost equal to $\sqrt{n} + 1$ is incurred and hence the expected cost of INSTANT is

$$(\sqrt{n} + 1) \cdot \mathbb{E}_\sigma^\tau[N(\sigma)] = (\sqrt{n} + 1) \cdot \lambda(T) \cdot \tau = \Theta(\sqrt{n}) \cdot \tau.$$

If PLAN algorithm is applied, then the period for each vertex v_i is determined as

$$\sqrt{\frac{2 \cdot (\sqrt{n} + n \cdot 1)}{\lambda(T)}} = \sqrt{2(n + \sqrt{n})}.$$

For each period of length $\sqrt{2(n + \sqrt{n})}$, a weight cost of $\sqrt{n} + n$ is incurred (since the whole tree T is bought) and the expected delay cost produced is $\frac{1}{2} \cdot \lambda(T) \cdot (\sqrt{2(n + \sqrt{n})})^2 = n + \sqrt{n}$ also. The expected cost PLAN is thus equal to

$$\frac{\tau}{\sqrt{2(n + \sqrt{n})}} \cdot 2(n + \sqrt{n}) = \sqrt{2(n + \sqrt{n})} \cdot \tau = \Theta(\sqrt{n}) \cdot \tau.$$

However, consider the following algorithm ALG:

- the edge (γ, u) is bought periodically with period being $n^{\frac{1}{4}}$;
- if at one request located at v_i is pending at time j , serve this request at this moment.

In this way, for each period of length $n^{\frac{1}{4}}$,

- the expected number of requests generated within this period is $\lambda(T) \cdot n^{\frac{1}{4}} = n^{\frac{1}{4}}$ — the expected weight cost is thus equal to $\sqrt{n} + 1 \cdot n^{\frac{1}{4}} = \sqrt{n} + n^{\frac{1}{4}}$;
- the expected delay cost is $\frac{1}{2} \cdot \lambda(T) \cdot (n^{\frac{1}{4}})^2 = \frac{\sqrt{n}}{2}$.

As a result, the expected cost produced in each period is equal to $\sqrt{n} + n^{\frac{1}{4}} + \frac{\sqrt{n}}{2} = 1.5\sqrt{n} + n^{\frac{1}{4}}$ and the expected cost of this algorithm is equal to

$$\frac{\tau}{n^{\frac{1}{4}}} \cdot (1.5\sqrt{n} + n^{\frac{1}{4}}) = (1.5 \cdot n^{\frac{1}{4}} + 1) \cdot \tau = \Theta(n^{\frac{1}{4}}) \cdot \tau.$$

Notice that

$$\lim_{\tau \rightarrow \infty} \frac{\mathbb{E}_\sigma^\tau[\text{cost}(\text{INSTANT}(\sigma))]}{\mathbb{E}_\sigma^\tau[\text{cost}(\text{ALG}(\sigma))]} = \Theta(n^{\frac{1}{4}}) \text{ and } \lim_{\tau \rightarrow \infty} \frac{\mathbb{E}_\sigma^\tau[\text{cost}(\text{PLAN}(\sigma))]}{\mathbb{E}_\sigma^\tau[\text{cost}(\text{ALG}(\sigma))]} = \Theta(n^{\frac{1}{4}}).$$

By letting $n \rightarrow \infty$, we can conclude that both INSTANT and PLAN achieve unbounded RoEs.