



Stochastic coordination in heterogeneous load balancing systems

Guy Goren¹ · Shay Vargaftik² · Yoram Moses¹

Received: 11 November 2021 / Accepted: 25 August 2022 / Published online: 12 October 2022
© The Author(s), under exclusive licence to Springer-Verlag GmbH Germany, part of Springer Nature 2022

Abstract

Current-day data centers and high-volume cloud services employ a broad set of heterogeneous servers. In such settings, client requests typically arrive at multiple entry points, and dispatching them to servers is an urgent distributed systems problem. This paper presents an efficient solution to the load balancing problem in such systems that improves on and overcomes problems of previous solutions. The load balancing problem is formulated as a stochastic optimization problem, and an efficient algorithmic solution is obtained based on a subtle mathematical analysis of the problem. Finally, extensive evaluation of the solution on simulated data shows that it outperforms previous solutions. Moreover, the resulting dispatching policy can be computed very efficiently, making the solution practically viable.

Keywords Distributed load balancing · Heterogeneous load balancing · Stochastically coordinated dispatching · Parallel server model

1 Introduction

Load balancing in modern computer clusters is a challenging task. Unlike in the traditional parallel server model where all client requests arrive through a single centralized entry point, today's cluster designs are distributed [6,14,17,48]. In particular, they involve many dispatchers that serve as entry points to client requests and distribute these requests among a multitude of servers. The dispatchers' goal is to distribute the client requests in a balanced manner so that no server is overloaded or underutilized. This is particularly challenging due to two system design attributes: (1) the dispatchers must take decisions immediately upon arrival of requests, and independently from each other. This requirement is critical to adhere to the high rate of incoming client requests and to the extremely low required response times [11,46]. Indeed, even a small sub-second addition to response time in dynamic content websites can lead to a persistent loss of users and revenue [34,49]; (2) today's systems are heterogeneous with

different servers containing different generations of CPUs, various types of acceleration devices such as GPUs, FPGAs, and ASICs, with varying processing speeds [12,13,25,27,37]. This paper presents a new load balancing solution for distributed dispatchers in heterogeneous systems.

Most previous works on distributed cluster load balancing focused either on homogeneous systems or on systems in which dispatchers have limited information about server queue-lengths. In contrast, in today's heterogeneous systems the dispatchers (e.g., high-performance production L7 load balancers such as HAProxy [55] and NGINX [21]) typically have access to abundant queue-length information.

A series of recent works have shown that popular "join-the-shortest-queue (JSQ)"-based load-balancing policies behave poorly when the dispatchers' information is highly correlated [24,40,58,64]. In particular, such policies suffer from so-called "herd behavior," in which different dispatchers concurrently recognize the same set of less loaded servers and forward all their incoming requests to this set. The queue-lengths of the servers in this small set then grow rapidly, causing excessive processing delays, increased response times, and even dropped jobs. In severe cases, herding may even result in servers stalling and crashing.

Herding has been dubbed the "finger of death" in [39]. In order to avoid it, both in recent theory works [58,64] and in modern state-of-the-art production deployments [20,54], researchers introduced load balancing policies that break symmetry among the dispatchers via a random subsampling

✉ Guy Goren
sgoren@campus.technion.ac.il

Shay Vargaftik
shayv@vmware.com

Yoram Moses
moses@ee.technion.ac.il

¹ Technion, Haifa, Israel

² VMware Research, Palo Alto, USA

of queue-length information. Even large cloud service companies such as Netflix report making use of such limited queue-size information policies to avoid suffering from detrimental herd behavior effects [39,51]. They choose to do so despite the potential for degraded resource utilization and worse response times.

The better-information/worse-performance paradox manifested by herding was recently addressed for load balancing in homogeneous systems (in which servers are all equally powerful) by [24]. They demonstrated that the tradeoff between more accurate information and herding is not inherent. Rather, herding stems from the fact that currently employed dispatching techniques do not account for the fact that multiple dispatchers operate concurrently in the system. Based on this observation, they suggested a solution in which the dispatchers' policies address the presence and concurrent operation of multiple dispatchers by employing *stochastic coordination* (Sect. 1.1). While their solution provides superior performance in homogeneous systems, it performs poorly in heterogeneous systems, since it does not account for variation in server service rates.

Our work generalizes the approach of [24] and obtains an efficient load-balancing policy based on stochastic coordination for the more challenging heterogeneous case. We initially expected that mild adjustments to the scheme of [24] for homogeneous systems would yield similarly effective policies for the heterogeneous case. That turned out not to be the case. The generalization required overcoming non-trivial hurdles at the level of mathematical analysis, of the algorithmic treatment, and at the conceptual level. The solution that we obtained has several features that the policy presented in [24] does not have. In both cases, a randomized load-balancing policy is designed based on the solution of a stochastic optimization problem. In the homogeneous case, [24] have shown that this solution can be formulated in terms of a closed-form formula that can be efficiently computed in real-time by the dispatchers. In contrast, explicitly computing the solution obtained for the heterogeneous case requires exponential time. Even obtaining an approximate solution using standard optimization techniques requires cubic time, which is not feasible for dispatchers to perform. Using a novel analysis of the optimization problem in the heterogeneous case, we show that the search for an optimal solution can be made in an extremely efficient manner, which results in a practically feasible algorithm. We thus obtain a novel policy, we term *stochastically coordinated dispatching* (*SCD*) that greatly improves over existing load-balancing policies for heterogeneous systems, and is essentially as easy to compute as the simplest policies are.

To evaluate the performance of *SCD* we implemented it as well as 10 other algorithms (including both traditional techniques and recent state-of-the-art ones) in C++. Extensive evaluation results indicate that *SCD* consistently outperforms

all tested techniques over different systems, heterogeneity levels and metrics. For example, at high loads, *SCD* improves the 99th percentile delay of client requests by more than a factor of 2 in comparison to the second-best policy, and by more than an order of magnitude compared to the heterogeneity-oblivious solution in [24]. In terms of computational running time, *SCD* is competitive with currently employed techniques (e.g., *JSQ*). For example, even in a system with 100 servers, *SCD* requires only a few microseconds on a single CPU core to make dispatching decisions. Finally, our results are reproducible. Our *SCD* implementation is available on GitHub [23].

1.1 Related work

In a traditional computer cluster with a centralized design and a single dispatcher that takes all the decisions, a centralized algorithm such as *JSQ*, that assigns each arriving request to the currently shortest queue, offers favorable performance and strong theoretical guarantees [15,61,63]. However, in a distributed design, where each dispatcher independently follows *JSQ*, the aforementioned herding phenomenon occurs [39,51,58,64]. As a result, both researchers and system designers often resort to traditional techniques that were originally designed either for centralized systems with a single dispatcher or for systems based on limited queue state information. For example, in the *power-of-d-choices* policy (denoted *JSQ(d)*) [36,41,59], when requests arrive, a dispatcher samples d servers uniformly at random and employs *JSQ* considering only the d sampled servers. This policy alleviates the herding phenomenon for sufficiently low d values since it is only with a low probability that different dispatchers will sample the same good server(s) at a given point in time. However, low d values often come at the price of longer response times and low resource utilization [16], whereas herding does occur for higher d values. Moreover, in heterogeneous systems, *JSQ(d)* may even result in *instability*¹ for any d value strictly below the number of servers.

Recently, to account for server heterogeneity, *shortest-expected-delay* (*SED*) policies were proposed. These policies operate similarly to *JSQ* and *JSQ(d)* but, instead of ranking servers according to their queue-lengths, servers are compared according to their *normalized* queue-lengths, i.e., their queue length divided by their processing capacity. This way, a job is sent not to the server with the shortest queue but to the server with the shortest expected wait time. Indeed these policies were shown to outperform their

¹ A queue is unstable when its size can grow with no bound. A load balancing system is unstable if at least one of its queues is unstable. Instability in heterogeneous load balancing systems usually occurs when the faster queues are constantly idling because they do not receive enough requests, whereas slower servers receive too many requests, and their queues continue to grow.

heterogeneity-unaware JSQ -based counterparts in heterogeneous systems (for homogeneous systems these policies coincide) [18,19,28,50]. However, in a setting with multiple, distributed dispatchers, the SED policies suffer from the same herding phenomena.

The first dispatching policy that was designed specifically for the multi-dispatcher case, called *join-the-idle-queue* (JIQ) [34,42,52,53,56], was originally introduced by Microsoft [35]. In JIQ , a dispatcher sends requests only to idle servers. If there are no idle servers, the requests are forwarded to randomly chosen ones. JIQ significantly outperforms $JSQ(d)$ when the system operates at low loads. However, its performance quickly deteriorates when the load increases and the dispatching approaches a random one. In fact, as is the case with $JSQ(d)$, the JIQ policy may exhibit instability in the presence of high loads [4,65]. A recent work [18] considered an improvement to JIQ that accounts for server heterogeneity by adapting the server sampling probabilities to account for their processing rate. While this policy restores stability when the load is high, it is significantly outperformed by policies such as SED and even JSQ when queue-length information is available [64].

Recent state-of-the-art techniques include the local-shortest-queue (LSQ) [58] and the *local-estimation-driven* (LED) [64] policies. They address the limitations of both $JSQ(d)$ and JIQ by maintaining a local array of server queue-lengths at each dispatcher. Each dispatcher updates its local array by randomly choosing servers and querying them for their queue-length. Consequently, the dispatchers have different views of the server queue-lengths. However, the performance of LSQ and LED depends on the dispatchers' local arrays being weakly correlated. When the dispatchers' queue length information is even partially correlated, both LSQ and LED incur herding [24].

As discussed in the Introduction, the seemingly paradoxical herding behavior was recently addressed in the case of homogeneous systems in [24]. Specifically, that work introduced the *tidal-water-filling* policy (denoted TWF). In TWF , the dispatching policy of a dispatcher is defined by the probabilities at which it sends each arriving request to each server. The main idea that enables utilizing accurate server queue-lengths information without incurring herding relies on stochastic coordination of the dispatchers — i.e., setting these probabilities such that all the dispatchers' decisions combined result in a balanced state. Nevertheless, TWF does not account for server service rates. Consequently, its performance significantly degrades in a heterogeneous system, resulting in reduced resource utilization and excessively long response times. This hinders the applicability of TWF to modern computer clusters.

A related line of work dealing with load balancing challenges in distributed systems is based on the balls-into-bins model [1], including extensions to dynamic or heteroge-

neous settings (e.g., [5,7]). In the balls-into-bins model, it is often possible to obtain more precise theoretical guarantees. Indeed, common approaches include regret minimization (e.g., [31]) and adaptive techniques (e.g., [33]). Nevertheless, this model is not aligned with our model (e.g., we consider multiple dispatchers, stochastic arrivals at each dispatcher and stochastic departures at each server). As a result, their analysis does not apply in our model, and vice versa.

A seminal work deals with parallel accesses of CPUs to memory regions by trying to minimize collisions [29]. Their solutions rely on hash functions to prevent memory accesses from becoming the bottleneck for system performance. In our setting, using hash functions resembles random allocation which is known to be sub-optimal.

Another related line of work concerns static load balancing, where the goal is to assign jobs to servers in a distributed manner to optimize some balance metric and with a minimal number of communication rounds among the participants (usually in the CONGEST or LOCAL model) [3,10,26]. This model is not aligned with ours since we consider dynamic systems with the demand for immediate and independent decision making among the dispatchers to sustain the high incoming rate of client requests and the demand for low latency. In particular, in our model, the dispatchers do not interact.²

2 Model

We consider a system with a set \mathcal{S} of n servers and a set \mathcal{D} of m dispatchers. The system operates over discrete and synchronous rounds $t \in \mathbb{N}$. Each server $s \in \mathcal{S}$ has its own FIFO queue³ of pending client requests. We denote by $q_s(t)$ the number of client requests at server's s queue at the beginning of round t . We assume that the values of $q_s(t)$ for all servers s are available in round t to all dispatchers, for all $t \in \mathbb{N}$. Each round consists of three phases:

1. **Arrivals.** Each dispatcher $d \in \mathcal{D}$ has its own stochastic, independent and unknown client request arrival process. We use $a^{(d)}(t) \in \mathbb{N}$ to denote the number of new client requests that exogenously arrive at dispatcher d in the first phase of round t .

² In practice, any communication among the distributed dispatchers introduces additional processing and, more importantly, possibly unpredictable network delay. Therefore, the dispatchers' high rate of incoming jobs makes interaction among them highly undesirable and even not feasible, especially when the dispatchers are not co-located (i.e., not on the same machine). As a result, assuming that dispatchers do not interact is standard practice in our L7 cluster load balancing model [20,24,34,52–54,58,60,64].

³ FIFO stands for first-in-first-out. Namely, client requests at each server are processed in the order at which they arrive at the queue. The order among client requests that arrive at the same time is arbitrary.

2. **Dispatching.** In the second phase of a round, each dispatcher immediately and independently chooses a destination server for each received request and forwards the request to the chosen server queue for processing. We denote by $\bar{a}_s^{(d)}(t)$ the number of requests dispatcher d forwards to server s in the second phase of round t , and by $\bar{a}_s(t) = \sum_{d \in \mathcal{D}} \bar{a}_s^{(d)}(t)$ the total number of requests server s receives from all dispatchers.
3. **Departures.** During the third phase, each server performs work and possibly completes requests. Completed requests immediately depart from the system. We denote by $c_s(t)$ the number of requests that server s can complete during the third phase of round t , provided it has that many requests to process. We assume that $c_s(t)$ is determined by an unknown independent stochastic process. We only assume that each server has some inherent expected time invariant processing rate (i.e., speed), and denote $\mathbb{E}[c_s(t)] = \mu_s$.

Note that any processing system must adhere to the requirement that, on average, the sum of server processing rates must be sufficient to accommodate the sum of arriving client requests. We mathematically express this additional demand in context when appropriate. Finally, we use the terms *job* and *client request* interchangeably.

3 Solving the dispatching problem

Our goal is to devise an algorithm that leads to short response times and high resource utilization. In this section, we define a notion of an “ideal” assignment for an online dispatching algorithm. Intuitively, the quality of a dispatching assignment can be assessed by comparing it with the ideal one. A key element in our distributed solution will be solving an optimization problem whose goal is to approximate the ideal assignment as closely as possible.

3.1 Ideally balanced assignment

For each round t we are given the current sizes of the queues at the servers $(q_1(t), \dots, q_n(t))$ and the arrivals at the dispatchers $(a^{(1)}(t), \dots, a^{(m)}(t))$, and we need to compute an assignment of the incoming jobs to the servers $(\bar{a}_1(t), \dots, \bar{a}_n(t))$. (From here on, we omit the round notation t when clear from context.) We would like to distribute the incoming jobs in a manner that balances the *load*, which we think of as the amount of work that each server has after the assignment. Since servers have different processing rates μ_s , the load on a server s does not correspond to the number of jobs $q_s + \bar{a}_s$ in its queue (at the end of the round). Rather, the load is taken to be $\frac{q_s + \bar{a}_s}{\mu_s}$, i.e., the expected amount of time it would take s to process the jobs that are

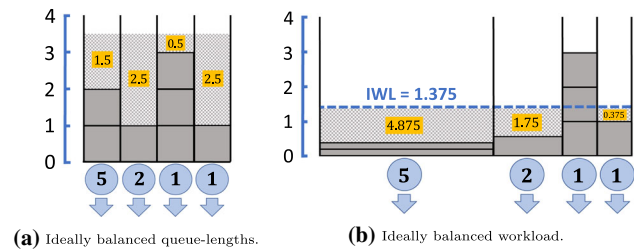


Fig. 1 Illustrating the difference between balancing the number of jobs and balancing the workload at the servers. An example with 4 servers with rates [5, 2, 1, 1] (from left to right), 7 queued jobs at the servers [2, 1, 3, 1] and 7 new arrivals. An ideally balanced assignment is [4.875, 1.75, 0, 0.375], which differs from [1.5, 2.5, 0.5, 2.5] that balances the number of jobs per server

in its queue. (See Figure 1 for illustration.) In general, no assignment that would completely balance the load among all servers will necessarily exist, since server queue-lengths may vary considerably at the start of a round. We can, however, aim at minimizing the difference between the load of the most loaded and the least loaded servers. If the units of incoming work were continuous, this would be achieved by an assignment $\{\bar{a}_s\}_{s \in \mathcal{S}}$ that solves:

$$\begin{aligned} & \max \min_{s \in \mathcal{S}} \frac{q_s + \bar{a}_s}{\mu_s} \\ & \text{s.t. } \sum_{s \in \mathcal{S}} \bar{a}_s = \sum_{d \in \mathcal{D}} a^{(d)} \quad \text{and} \quad \forall s \in \mathcal{S}: \bar{a}_s \geq 0. \end{aligned} \tag{1}$$

We call an assignment that satisfies Eq. (1) an *ideally balanced assignment* (IBA for short), and the value of the target function the *ideal workload* (IWL). An illustrative example for an IBA is presented in Figure 1b. The IBA is an idealized goal in the sense that it is not always possible to achieve. This is because jobs are discrete and cannot be split among different servers. Therefore, a realistic load balancing algorithm should strive to assign jobs in a manner that is as close as possible to the IBA by some distance measure (e.g., the Euclidean norm).

In a *centralized dispatching system* (with a single dispatcher) minimizing the Euclidean norm distance from the IWL can be achieved in a straightforward manner. The dispatcher sends each job to the server that is expected to process it the earliest. Namely, to the server with the currently minimal $\frac{q_s + 1}{\mu_s}$. It then updates the server’s queue length and moves on to the next job. In a *distributed dispatching system* (i.e., with multiple dispatchers), however, the solution is considerably more complex. In particular, finding the best approximation to the IBA in systems with multiple dispatchers requires exact coordination among the dispatchers. However, dispatchers must make immediate and independent decisions. They have no time to communicate and so such exact coordination is impossible.

Following [24], we deal with the need to coordinate dispatchers' decisions while keeping them independent by randomizing their decisions. In essence, a dispatcher computes a probability distribution $P = [p_1, \dots, p_n]$ in each round. Each job's destination is then drawn according to P , thus making the decisions independent. The probabilities in P take into consideration both server loads and what other dispatchers might draw. In [24], the authors name this approach stochastic coordination. The main challenge in stochastic coordination is identifying the right probabilities, and computing them.

We now turn our focus to computing the ideal workload (IWL). Note that the IWL determines the IBA ($= \{\bar{a}_s\}_{s \in \mathcal{S}}$) since for all $s \in \mathcal{S}$:

$$\bar{a}_s = \mu_s \cdot \max \left\{ \frac{q_s}{\mu_s}, \text{IWL} \right\} - q_s. \tag{2}$$

We develop an algorithm that computes the IWL in a heterogeneous system. The pseudocode appears in Algorithm 1. Roughly speaking, it works as follows. It starts with the current state of the system and iteratively assigns work that increases the minimal load, i.e., $\min \frac{q_s + \bar{a}_s}{\mu_s}$ until it reaches the IWL. At each iteration, unassigned work is assigned to the least loaded servers until they reach the closest higher load. The algorithm returns when no unassigned work remains. Moreover, COMPUTEIDEALWORKLOAD is efficient. That is, it runs in $O(n)$ time, if the values of $\frac{q_s}{\mu_s}$ at the servers are pre-sorted (otherwise, the algorithm's complexity is dominated by the task of sorting these n values). In the algorithm and henceforth, we often use 'a' as a shorthand for the total sum of arrivals: $a \triangleq \sum_{d \in \mathcal{D}} a^{(d)}$.

3.2 Distributed load balancing as a Stochastic optimization problem

Our distributed load balancing algorithm will be based on the solution to a stochastic optimization problem. The first step in the statement of the optimization problem is to define an appropriate error function that we seek to minimize. A job assignment is measured against the ideally balanced assignment (IBA). Assume that at each time slot t we are given the arrivals $(a^{(1)}, \dots, a^{(m)})$ in the current round,⁴ and we know the current sizes of the queues at the servers (q_1, \dots, q_n) . We can deduce the ideal workload (IWL) of a centralized IBA algorithm at that point (see Algorithm 1). Intuitively, a large deviation from the IWL in server s corresponds either to large job delays and slow response times (for a higher load than ideal), or lesser resource utilization (for a lower load than

ideal). Since we seek to avoid long delay tails and wasted server capacities, the error for a large deviation from the IWL should be higher than the error for several small deviations. On the other hand, minimizing the worst-case assignment, that is, focusing too much on the worst possible deviation, can damage the mean performance. We measure the distance a solution offers from the ideal solution in terms of the L_2 norm (squared distances). This balances the algorithm's mean and worst-case performance and, moreover, is amenable to formal analysis. Thus, the individual error of server s is defined by

Algorithm 1: Computing the ideal workload.

```

1 Function ComputeIdealWorkLoad( $\mathcal{S}, \{q_s\}_{s \in \mathcal{S}}, \{\mu_s\}_{s \in \mathcal{S}}$ ,
  a):
2    $\mathcal{S}_{\geq \text{IWL}} \leftarrow \mathcal{S}; \mu_{\text{tot}} \leftarrow 0; l \leftarrow a$ 
3    $r \leftarrow \arg \min_{s \in \mathcal{S}_{\geq \text{IWL}}} \frac{q_s}{\mu_s}$ 
4    $\text{IWL} \leftarrow q_r / \mu_r$ 
5   while  $l > 0$  do
6      $\mu_{\text{tot}} \leftarrow \mu_{\text{tot}} + \mu_r$ 
7      $\mathcal{S}_{\geq \text{IWL}} \leftarrow \mathcal{S}_{\geq \text{IWL}} \setminus \{r\}$ 
8     if  $\mathcal{S}_{\geq \text{IWL}} = \emptyset$  then
9       return  $\text{IWL} + \frac{l}{\mu_{\text{tot}}}$ 
10     $r \leftarrow \arg \min_{s \in \mathcal{S}_{\geq \text{IWL}}} \frac{q_s}{\mu_s}$ 
11     $\Delta \leftarrow \frac{q_r}{\mu_r} - \text{IWL}$ 
12    if  $\Delta \cdot \mu_{\text{tot}} \geq l$  then
13      return  $\text{IWL} + \frac{l}{\mu_{\text{tot}}}$ 
14     $l \leftarrow l - \Delta \cdot \mu_{\text{tot}}$ 
15     $\text{IWL} \leftarrow \text{IWL} + \Delta$ 

```

$$\text{error}_s \triangleq \left(\frac{q_s + \bar{a}_s}{\mu_s} - \text{IWL} \right)^2. \tag{3}$$

Finally, to account for the variability in processing power, a server's error is weighted by multiplying it by the server's processing speed. For example, an error of +1 in the workload of a server with a processing rate of $\mu_s = 10$ results in 10 jobs that will now wait for an expected extra round. The same +1 error for a server with $\mu_s = 1$ affects only a single job. The resulting total error of an assignment $\{\bar{a}_1, \dots, \bar{a}_n\}$ is

$$\begin{aligned} \text{error} &= \sum_{s \in \mathcal{S}} \mu_s \cdot \text{error}_s \\ &= \sum_{s \in \mathcal{S}} \mu_s \left(\frac{q_s + \bar{a}_s}{\mu_s} - \text{IWL} \right)^2 \\ &= \sum_{s \in \mathcal{S}} \frac{(q_s + \bar{a}_s - \mu_s \cdot \text{IWL})^2}{\mu_s}. \end{aligned} \tag{4}$$

⁴ We relax the assumption that the entire vector of arrivals is available in Sect. 5.

Recall that we employ randomness to determine job destinations. Therefore, an assignment $\{\bar{a}_s\}_{s \in \mathcal{S}}$ and the corresponding error are random variables. In particular, we seek the probabilities $P \triangleq [p_1, \dots, p_n]$ that minimize the expected error function given the current queue lengths $[q_1, \dots, q_n]$ and the new arrivals $[a^{(1)}, \dots, a^{(m)}]$. Formally,

$$\begin{aligned} & \arg \min_P \mathbb{E}[\text{error}] \\ &= \arg \min_P \mathbb{E} \left[\sum_{s \in \mathcal{S}} \frac{(\bar{a}_s + (q_s - \mu_s \text{IWL}))^2}{\mu_s} \right] \\ &= \arg \min_P \left(\sum_{s \in \mathcal{S}} \mathbb{E} \left[\frac{\bar{a}_s^2}{\mu_s} \right] \right. \\ & \quad \left. + 2 \sum_{s \in \mathcal{S}} \mathbb{E} \left[\frac{\bar{a}_s \cdot (q_s - \mu_s \text{IWL})}{\mu_s} \right] \right. \\ & \quad \left. + \sum_{s \in \mathcal{S}} \mathbb{E} \left[\frac{(q_s - \mu_s \text{IWL})^2}{\mu_s} \right] \right). \end{aligned} \tag{5}$$

In the error function only the values of $\{\bar{a}_s\}$ are affected by P . Therefore, we can drop all additive constants and take the expectation on the $\{\bar{a}_s\}$ terms only. This yields the simplified form

$$\begin{aligned} \arg \min_P \mathbb{E}[\text{error}] &= \arg \min_P \left(\sum_{s \in \mathcal{S}} \frac{1}{\mu_s} \mathbb{E}[\bar{a}_s^2] \right. \\ & \quad \left. + 2 \sum_{s \in \mathcal{S}} \frac{q_s - \mu_s \text{IWL}}{\mu_s} \mathbb{E}[\bar{a}_s] \right). \end{aligned} \tag{6}$$

Since job destinations are drawn independently according to P we have that \bar{a}_s is a binomial random variable with

$$\begin{aligned} \mathbb{E}[\bar{a}_s] &= p_s \sum_{d \in \mathcal{D}} a^{(d)} = ap_s, \quad \text{and} \\ \mathbb{E}[\bar{a}_s^2] &= a \cdot p_s(1 - p_s) + a^2 p_s^2. \end{aligned} \tag{7}$$

Plugging the above in Eq. (6) yields

$$\begin{aligned} & \arg \min_P \mathbb{E}[\text{error}] \\ &= \arg \min_P \left(\sum_{s \in \mathcal{S}} \frac{1}{\mu_s} (ap_s - ap_s^2 + a^2 p_s^2) \right. \\ & \quad \left. + 2 \sum_{s \in \mathcal{S}} \frac{q_s - \mu_s \text{IWL}}{\mu_s} \cdot ap_s \right) \\ &= \arg \min_P \left(a(a-1) \sum_{s \in \mathcal{S}} \frac{1}{\mu_s} p_s^2 \right. \\ & \quad \left. + a \sum_{s \in \mathcal{S}} \frac{2(q_s - \mu_s \text{IWL}) + 1}{\mu_s} \cdot p_s \right). \end{aligned} \tag{8}$$

When $a = 1$ the first term is eliminated, and we obtain:

$$\begin{aligned} & \arg \min_P \mathbb{E}[\text{error}] \\ &= \arg \min_P \sum_{s \in \mathcal{S}} \frac{2(q_s - \mu_s \text{IWL}) + 1}{\mu_s} \cdot p_s \\ &= \arg \min_P \sum_{s \in \mathcal{S}} \left(\frac{2q_s + 1}{\mu_s} - 2\text{IWL} \right) \cdot p_s \end{aligned} \tag{9}$$

The solution is to divide the probabilities among the servers that have the minimal value of $\frac{2q_s+1}{\mu_s}$. (The division can be arbitrary; any division of probabilities will do.)

We turn to the general case in which $a > 1$. Dividing the target function at Eq. (8) by a , we obtain a constrained optimization problem with the following standard form.

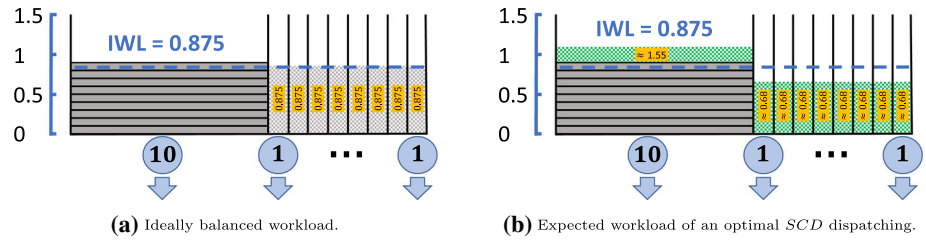
$$\begin{aligned} & \text{minimize}_P \quad f(P) = (a-1) \sum_{s \in \mathcal{S}} \frac{1}{\mu_s} \cdot p_s^2 \\ & \quad + \sum_{s \in \mathcal{S}} \frac{2(q_s - \mu_s \text{IWL}) + 1}{\mu_s} \cdot p_s \\ & \text{subject to} \quad \sum_{s \in \mathcal{S}} p_s - 1 = 0, \\ & \quad p_s \geq 0 \quad \forall s \in \mathcal{S}. \end{aligned} \tag{10}$$

Using the solution $P = [p_1, \dots, p_n]$ to the above problem yields an optimal greedy online algorithm, in which the best stochastic per-round decisions are made. Interestingly, an optimal P can contain positive probabilities even to servers that are above the IWL. This is in stark contrast to [24], where, underlying the formal analysis is the basic assumption that servers that have a current load of more than the IWL⁵ should have 0 probability of receiving any jobs. Figure 2b provides an illustrative example of what an optimal P is expected to yield. This example shows an optimal P in which a server that is above the IWL receives a positive probability of ≈ 0.221 (i.e., $\approx \frac{1.55}{7}$).

The optimization problem in Eq. (10) is convex quadratic with affine constraints and has a non-empty set of feasible solutions. Known general algorithms computing exact solutions to this problem incur an exponential in n (worst case) time complexity. Using well-established algorithms for quadratic programming such as the interior-point method and the ellipsoid method [9], it is possible to compute approximate solutions in $\Omega(n^3)$. However, this is still not good enough to be used for dispatching in the high-volume settings that we are targeting.

⁵ The work of [24] uses a term called *water-level*. In the special case of homogeneous systems, the water level coincides with the IWL.

Fig. 2 An example of a system with one fast ($\mu = 10$) server, and 8 slower servers ($\mu = 1$). The system state is 9 jobs queued at the fast server, empty queues at the remaining servers, and 7 incoming jobs



4 Deriving a computationally efficient solution

Real-time load balancing systems are reluctant to deploy algorithms that might incur exponential time complexity. Even reaching an approximate solution, by naïvely applying generic approximation-algorithms, might become impractical for systems with a few hundred servers. Thus, we seek a specialized solution that considers our specific problem. As we will show, it is possible to identify particular properties of our optimization problem Eq. (10), and utilize them to design a highly efficient algorithm.

4.1 The probable set and its ordering

Recall that $a = 1$ corresponds to a single job entering the system in the current round. In this case, no coordination between dispatchers is necessary, and indeed, as we have shown, the dispatching problem can be solved in a straightforward manner. The distributed problem arises only when $a > 1$, in which case we need to solve the optimization problem of Eq. (10). The Lagrangian function corresponding to Eq. (10) is

$$L(P, \Lambda) = (a - 1) \sum_{s \in \mathcal{S}} \frac{1}{\mu_s} \cdot p_s^2 + \sum_{s \in \mathcal{S}} \frac{2(q_s - \mu_s IWL) + 1}{\mu_s} \cdot p_s - \sum_{s \in \mathcal{S}} \Lambda_s p_s + \Lambda_0 (\sum_{s \in \mathcal{S}} p_s - 1), \tag{11}$$

where Λ_0 is the Lagrange multiplier that corresponds to the equality constraint and $\{\Lambda_s\}_{s \in \mathcal{S}}$ correspond to the inequality constraints. Since the problem is convex with affine constraints, the Karush-Kuhn-Tucker (KKT) [30,32] theorem states that the following conditions are necessary and sufficient for $P^* = [p_1^*, \dots, p_n^*]$ to be an optimal solution. For all $s \in \mathcal{S}$:

$$(Stationarity) \quad \frac{\partial L}{\partial p_s}(p_s^*) = 0 = 2(a-1) \frac{1}{\mu_s} p_s^* + \frac{2(q_s - \mu_s IWL) + 1}{\mu_s} - \Lambda_s + \Lambda_0$$

$$(Primal feasibility) \quad \sum_{s \in \mathcal{S}} p_s^* - 1 = 0, \quad p_s^* \geq 0$$

$$(Dual feasibility) \quad \Lambda_s \geq 0$$

$$(Complementary slackness) \quad p_s^* \Lambda_s = 0$$

(12)

From Stationarity we can deduce that

$$p_s^* = \frac{-2(q_s - \mu_s IWL) - 1 + \mu_s \Lambda_s - \mu_s \Lambda_0}{2(a - 1)}. \tag{13}$$

We call the set of servers with positive probabilities in the optimal solution the *probable set* and denote it by \mathcal{S}^+ . Formally the probable set is $\mathcal{S}^+ \triangleq \{s \in \mathcal{S} \mid p_s^* > 0\}$. As we shall see, \mathcal{S}^+ plays an important role in our derivations. In particular, the Complementary slackness condition from Eq. (12) implies that $\Lambda_s = 0$ for every $s \in \mathcal{S}^+$. Thus,

$$p_s^* = \frac{-2(q_s - \mu_s IWL) - 1 - \mu_s \Lambda_0}{2(a-1)}, \quad \forall s \in \mathcal{S}^+. \tag{14}$$

We use the Primal feasibility condition from Eq. (12) together with Eq. (14) to obtain

$$1 = \sum_{s \in \mathcal{S}} p_s^* = \sum_{s \in \mathcal{S}^+} p_s^* = \sum_{s \in \mathcal{S}^+} \frac{-2(q_s - \mu_s IWL) - 1 - \mu_s \Lambda_0}{2(a - 1)}, \tag{15}$$

from which we can isolate Λ_0

$$\Lambda_0 = \frac{2 \sum_{s \in \mathcal{S}^+} (\mu_s IWL - q_s) - \sum_{s \in \mathcal{S}^+} 1 - 2(a-1)}{\sum_{s \in \mathcal{S}^+} \mu_s}. \tag{16}$$

The KKT conditions enabled us to derive Eq. (14) and Eq. (16), which show that identifying \mathcal{S}^+ provides an analytical solution to our optimization problem Eq. (10). (By first calculating Λ_0 and then each of the positive probabilities.) Hence, finding an optimal solution reduces to finding the probable set. This result is a complementary instance of the generic active set method for quadratic programming [47]. However, the active set method provides, in general, a worst case running-time complexity of 2^n . And while for the specific instance of the problem with $\mu_s = \mu$ for all servers (i.e., an homogeneous system) we can derive

Algorithm 2: Find probabilities in $O(n^2)$ time.

```

1 Function ComputeProbabilities( $\mathcal{S}, \{q_s\}_{s \in \mathcal{S}}, \{\mu_s\}_{s \in \mathcal{S}}, a,$ 
  IWL):
2    $\mathcal{O} \leftarrow \emptyset; val^* \leftarrow \infty$ 
3    $p_s \leftarrow 0 \forall s \in \mathcal{S}$ 
4   while  $\mathcal{S} \setminus \mathcal{O} \neq \emptyset$  do
5      $r \leftarrow \arg \min_{s \in \mathcal{S} \setminus \mathcal{O}} \frac{2q_s + 1}{\mu_s}$ 
6      $\mathcal{O} \leftarrow \mathcal{O} \cup r$  /*  $\mathcal{O}$  is the candidate for  $\mathcal{S}^+$ 
      */
7      $\Lambda_0 \leftarrow \frac{2 \sum_{s \in \mathcal{O}} (\mu_s IWL - q_s) - \sum_{s \in \mathcal{O}} 1 - 2(a-1)}{\sum_{s \in \mathcal{O}} \mu_s}$ 
      /* according to Eq. (16) */
8     for  $s \in \mathcal{O}$  do
9        $p_s \leftarrow \frac{-2(q_s - \mu_s IWL) - 1 - \mu_s \Lambda_0}{2(a-1)}$  /* according to
      Eq. (14) */
10      if  $p_s < 0$  then
11        go to line 4 /* the solution is
        infeasible; continue to next  $r$  */
12       $val \leftarrow (a-1) \sum_{s \in \mathcal{O}} \frac{1}{\mu_s} \cdot p_s^2 + \sum_{s \in \mathcal{O}} \frac{2(q_s - \mu_s IWL) + 1}{\mu_s} \cdot p_s$ 
      /* according to Eq. (10) */
13      if  $val < val^*$  then
14         $val^* \leftarrow val$ 
15         $P^* \leftarrow \{p_s\}_{s \in \mathcal{S}}$ 
16  return  $P^*$ 

```

$\mathcal{S}^+ = \{s \in \mathcal{S} \mid \frac{q_s}{\mu_s} < IWL\}$ analytically, the example in Figure 2b shows that this is no longer true in the general, heterogeneous, case.

Instead of deriving \mathcal{S}^+ analytically, we turn to find an algorithmic solution. A trivial algorithm is to examine each of the 2^n possible subsets of servers. For each candidate subset: first calculate Λ_0 according to Eq. (16) — this guarantees that the sum of probabilities is 1, then test whether all the probabilities are indeed positive, and finally calculate the objective function. Clearly, its exponential computation complexity renders this method infeasible. To overcome the exponential nature of searching in a domain of size 2^n , we must reduce the size of the domain we search in. Lemma 1 formulates a property of the objective function in Eq. (10) that holds the key to reducing the size of the search.

Lemma 1 *Let r be a server in the probable set \mathcal{S}^+ . For every server u in the set \mathcal{S} , if $\frac{2q_r + 1}{\mu_r} \geq \frac{2q_u + 1}{\mu_u}$ then u is in \mathcal{S}^+ as well.*

Proof Let r and u satisfy the assumption that $\frac{2q_r + 1}{\mu_r} \geq \frac{2q_u + 1}{\mu_u}$, and let $P^* = \{p_1^*, \dots, p_n^*\}$ be the optimal solution for Eq. (10). Assume by contradiction that $p_r^* > 0$ but $p_u^* = 0$. We show that there exists a feasible solution P that obtains a lower value of the objective function, thus contradicting the optimality of P^* . Specifically, we show that for the positive constant

$$z = \min \left\{ \frac{\mu_u(2q_r + 1) - \mu_r(2q_u + 1) + 2\mu_r\mu_u p_r^*}{(a-1)(\mu_r + \mu_u)}, p_r^* \right\},$$

any $0 < \epsilon < z$ and a different solution $P = \{p_1, \dots, p_n\}$ with $p_u = \epsilon$, $p_r = p_r^* - \epsilon$, and $p_s = p_s^*$ for all other servers, it holds that P is a better feasible solution than P^* .

First, we show that P is feasible. Since P^* is feasible, it holds that $\sum_{s \in \mathcal{S}} p_s^* = 1$, and $0 \leq p_s^* \leq 1$ for all $s \in \mathcal{S}$. Accordingly, for P it similarly holds that

$$\sum_{s \in \mathcal{S}} p_s = \sum_{s \in \mathcal{S} \setminus \{r, u\}} p_s^* + p_r + p_u = \sum_{s \in \mathcal{S}} p_s^* - \epsilon + \epsilon = 1.$$

The condition $0 \leq p_s \leq 1$ trivially follows from P^* 's feasibility and the fact that $0 < \epsilon < p_r^* \leq 1$.

Now we show that the new solution P obtains in a lower objective function's value (i.e., by Eq. (10)) than the optimal solution P^* , leading to a contradiction.

Denote $\text{diff} \triangleq f(P^*) - f(P)$. We now turn to show that $\text{diff} > 0$. By Eq. (10) we have

$$\begin{aligned} \text{diff} &= \left((a-1) \sum_{s \in \mathcal{S}} \frac{1}{\mu_s} \cdot (p_s^*)^2 \right. \\ &\quad \left. + \sum_{s \in \mathcal{S}} \frac{2(q_s - \mu_s IWL) + 1}{\mu_s} \cdot p_s^* \right) \\ &\quad - \left((a-1) \sum_{s \in \mathcal{S}} \frac{1}{\mu_s} \cdot p_s^2 \right. \\ &\quad \left. + \sum_{s \in \mathcal{S}} \frac{2(q_s - \mu_s IWL) + 1}{\mu_s} \cdot p_s \right) \\ &= \sum_{s \in \mathcal{S}} \left(\frac{a-1}{\mu_s} (p_s^*)^2 + \frac{2(q_s - \mu_s IWL) + 1}{\mu_s} p_s^* \right) \\ &\quad - \sum_{s \in \mathcal{S}} \left(\frac{a-1}{\mu_s} p_s^2 + \frac{2(q_s - \mu_s IWL) + 1}{\mu_s} p_s \right). \end{aligned}$$

Next, we split each summation term to a sum over $\mathcal{S} \setminus \{r, u\}$ and a sum over $\{r, u\}$

$$\begin{aligned} \text{diff} &= \sum_{s \in \mathcal{S} \setminus \{r, u\}} \left(\frac{a-1}{\mu_s} (p_s^*)^2 + \frac{2(q_s - \mu_s IWL) + 1}{\mu_s} p_s^* \right) \\ &\quad + \sum_{s \in \{r, u\}} \left(\frac{a-1}{\mu_s} (p_s^*)^2 + \frac{2(q_s - \mu_s IWL) + 1}{\mu_s} p_s^* \right) \\ &\quad - \sum_{s \in \mathcal{S} \setminus \{r, u\}} \left(\frac{a-1}{\mu_s} p_s^2 + \frac{2(q_s - \mu_s IWL) + 1}{\mu_s} p_s \right) \\ &\quad - \sum_{s \in \{r, u\}} \left(\frac{a-1}{\mu_s} p_s^2 + \frac{2(q_s - \mu_s IWL) + 1}{\mu_s} p_s \right), \end{aligned}$$

and since $p_s = p_s^*$ for every server in $\mathcal{S} \setminus \{r, u\}$, the sums over those sets cancel out. Thus, we obtain

$$\begin{aligned} \text{diff} &= \sum_{s \in \{r, u\}} \left(\frac{a-1}{\mu_s} (p_s^*)^2 + \frac{2(q_s - \mu_s \text{IWL}) + 1}{\mu_s} p_s^* \right) \\ &\quad - \sum_{s \in \{r, u\}} \left(\frac{a-1}{\mu_s} p_s^2 + \frac{2(q_s - \mu_s \text{IWL}) + 1}{\mu_s} p_s \right) \\ &= \sum_{s \in \{r, u\}} \left(\frac{a-1}{\mu_s} \left((p_s^*)^2 - p_s^2 \right) \right. \\ &\quad \left. + \frac{2(q_s - \mu_s \text{IWL}) + 1}{\mu_s} (p_s^* - p_s) \right) \\ &= \left(\frac{a-1}{\mu_r} \left((p_r^*)^2 - p_r^2 \right) \right. \\ &\quad \left. + \frac{2(q_r - \mu_r \text{IWL}) + 1}{\mu_r} (p_r^* - p_r) \right) \\ &\quad + \left(\frac{a-1}{\mu_u} \left((p_u^*)^2 - p_u^2 \right) \right. \\ &\quad \left. + \frac{2(q_u - \mu_u \text{IWL}) + 1}{\mu_u} (p_u^* - p_u) \right). \end{aligned}$$

We replace $(p_r^* - p_r)$ and $(p_u - p_u^*)$ by ϵ . Similarly, we replace $((p_r^*)^2 - p_r^2)$ by $(2p_r^*\epsilon - \epsilon^2)$, and $((p_u^*)^2 - p_u^2)$ by $(-\epsilon^2)$ and obtain

$$\begin{aligned} \text{diff} &= \left(\frac{a-1}{\mu_r} (2p_r^*\epsilon - \epsilon^2) + \frac{2(q_r - \mu_r \text{IWL}) + 1}{\mu_r} \cdot \epsilon \right) \\ &\quad + \left(\frac{a-1}{\mu_u} (-\epsilon^2) + \frac{2(q_u - \mu_u \text{IWL}) + 1}{\mu_u} (-\epsilon) \right) \\ &= \epsilon \left(\frac{2(q_r - \mu_r \text{IWL}) + 1}{\mu_r} - \frac{2(q_u - \mu_u \text{IWL}) + 1}{\mu_u} + 2p_r^* \right) \\ &\quad - \epsilon^2 \left(\frac{a-1}{\mu_r} + \frac{a-1}{\mu_u} \right) \\ &= \epsilon \underbrace{\left(\frac{2q_r + 1}{\mu_r} - \frac{2q_u + 1}{\mu_u} + 2p_r^* \right)}_x - \epsilon^2 \underbrace{\left(\frac{a-1}{\mu_r} + \frac{a-1}{\mu_u} \right)}_y. \end{aligned}$$

Since $\frac{2q_r + 1}{\mu_r} \geq \frac{2q_u + 1}{\mu_u}$ and $p_r^* > 0$, it must hold that $x > 0$. It must also be true that $y > 0$ since all μ -s are positive and we consider only $a > 1$. We therefore obtain that $0 < \frac{x}{y} = \frac{\mu_u(2q_r + 1) - \mu_r(2q_u + 1) + 2\mu_r\mu_u p_r^*}{(a-1)(\mu_r + \mu_u)}$. Hence, there exists $0 < \epsilon < \min\{\frac{x}{y}, p_r^*\}$, for which it holds that

$$\text{diff} = \epsilon x - \epsilon^2 y = \epsilon(x - \epsilon y) > \epsilon(x - \frac{x}{y} y) = 0.$$

Therefore, solutions exist that are both feasible and result in a lower objective function value than that of the optimal solution — a contradiction. This concludes the proof. \square

Lemma 1 imposes a very strict constraint on the structure of the probable set \mathcal{S}^+ :

Corollary 1 *Let $s_{i_1}, s_{i_2}, \dots, s_{i_n}$ be a listing of the servers in \mathcal{S} in non-decreasing sorted order of $\frac{2q_s + 1}{\mu_s}$. Moreover, denote $\mathcal{S}_j = \{s_{i_1}, \dots, s_{i_j}\}$ for $j = 1, \dots, n$. Then $\mathcal{S}^+ = \mathcal{S}_j$ for some $j \leq n$.*

Corollary 1 allows us to find the optimal probabilities in polynomial time. To be precise, both sorting \mathcal{S} and computing the IWL take $O(n \log n)$ time. So we are left with computing for each of the n subsets \mathcal{S}_j : (1) whether it satisfies that all probabilities are non-negative, and (2) What the value of the subset's objective function is. Steps (1) and (2) can be implemented in $O(n)$ time complexity each. Finally, we extract the subset with the minimal objective function value from those that respect (1). This can be done in $O(n^2)$ complexity, as presented in Algorithm 2.

4.2 $n \log n$ complexity

Given an already sorted data structure, the quadratic complexity of Algorithm 2 is caused by the $\Omega(n)$ cost per iteration of the calculation in line 7, the **for** loop in lines 8-11, and the summation in line 12. Since the outcome of an iteration depends on the results of past iterations, we can employ

Algorithm 3: Find probabilities in $O(n \log n)$ time; $O(n)$ given the servers' ordering.

```

1 Function ComputeProbabilities( $\mathcal{S}, \{q_s\}_{s \in \mathcal{S}}, \{\mu_s\}_{s \in \mathcal{S}}, a, \text{IWL}$ ):
2    $\mathcal{O} \leftarrow \emptyset; \text{val}^* \leftarrow \infty$ 
3    $\Lambda_{0,n} \leftarrow -2(a-1); \Lambda_{0,d} \leftarrow 0$ 
4    $v_1 \leftarrow 0; v_2 \leftarrow 0$ 
5   while  $\mathcal{S} \setminus \mathcal{O} \neq \emptyset$  do
6      $r \leftarrow \arg \min_{s \in \mathcal{S} \setminus \mathcal{O}} \frac{2q_s + 1}{\mu_s}$ 
7      $\mathcal{O} \leftarrow \mathcal{O} \cup r$ 
8      $\Lambda_{0,n} \leftarrow \Lambda_{0,n} + 2(\mu_r \text{IWL} - q_r) - 1$ 
9      $\Lambda_{0,d} \leftarrow \Lambda_{0,d} + \mu_r$ 
10     $\Lambda_0 \leftarrow \frac{\Lambda_{0,n}}{\Lambda_{0,d}}$  /* according to Eq. (15) */
11    if  $2\text{IWL} - \frac{2q_r + 1}{\mu_r} < \Lambda_0$  then
12      Continue /* the solution is
13      infeasible; go to line 6 */
14     $v_1 \leftarrow v_1 + \frac{\mu_r}{4(a-1)}$ 
15     $v_2 \leftarrow v_2 + \frac{(2(q_r - \mu_r \text{IWL}) + 1)^2}{4\mu_r(a-1)}$ 
16     $\text{val} \leftarrow v_1 \Lambda_0^2 - v_2$  /* according to Eq. (10)
17    */
18    if  $\text{val} < \text{val}^*$  then
19       $\text{val}^* \leftarrow \text{val}$ 
20       $\Lambda_0^* \leftarrow \Lambda_0$ 
21   $P^* \leftarrow \left\{ \max\{0, \frac{-2(q_s - \mu_s \text{IWL}) - 1 - \mu_s \Lambda_0^*}{2(a-1)}\} \right\}_{s \in \mathcal{S}}$ 
22  /* according to Eq. (14) */
23  return  $P^*$ 

```

a dynamic programming approach and design an algorithm with an $O(n)$ complexity given the data is presorted, and $O(n \log n)$ otherwise. The pseudocode in Algorithm 3 implements this. We next explain the necessary steps in deriving this algorithm.

Replacing the calculation of Λ_0 in line 7 of Algorithm 2 by an $O(1)$ computation per iteration is quite straightforward. We simply calculate the numerator and denominator sums separately, and then divide them. Each sum is computed by adding the current element to the sum from the previous iteration. The **for** loop in lines 8-11 of Algorithm 2 tests whether the computed probabilities are indeed non-negative. Since the denominator is a positive constant, it is sufficient to test the numerator for each server in \mathcal{O} . That is, whether it holds that $-2(q_s - \mu_s \text{IWL}) - 1 - \mu_s \Lambda_0 \geq 0$. The rates μ_s of the servers are positive. Therefore, dividing by μ_s and rearranging yields

$$\text{IWL} - \Lambda_0 \geq \frac{2q_s + 1}{\mu_s}. \tag{17}$$

In turn, observe that Eq. (17) holds for all servers in \mathcal{O} iff it holds for the server with the highest $\frac{2q_s+1}{\mu_s}$ value in \mathcal{O} , which is server r in each iteration. We thus replace the $\Omega(n)$ complexity **for** loop of lines 8-11 in Algorithm 2 by this single $O(1)$ complexity test.

Next, we address the summation in line 12 of Algorithm 2 which computes the value of the objective function $f(P)$ from Eq. (10). Since P contains $\Theta(n)$ elements, computing P costs $\Omega(n)$ time. However, we can make the computation of P more efficient by using the following Lemma.

Lemma 2 *The objective function in Eq. (10) satisfies*

$$f(P) = \underbrace{\Lambda_0^2 \sum_{s \in \mathcal{S}^+} \frac{\mu_s}{4(a-1)}}_{v_1} - \underbrace{\sum_{s \in \mathcal{S}^+} \frac{(2(q_s - \mu_s \text{IWL}) + 1)^2}{4\mu_s(a-1)}}_{v_2}.$$

Proof By combining Eq. (14), which expresses the positive optimal probabilities as a function of Λ_0 , into the objective function of Eq. (10) we obtain

$$\begin{aligned} f(P(\Lambda_0)) &= (a-1) \sum_{s \in \mathcal{S}^+} \frac{1}{\mu_s} \left(\frac{-2(q_s - \mu_s \text{IWL}) - 1 - \mu_s \Lambda_0}{2(a-1)} \right)^2 \\ &\quad + \sum_{s \in \mathcal{S}^+} \frac{2(q_s - \mu_s \text{IWL}) + 1}{\mu_s} \cdot \frac{-2(q_s - \mu_s \text{IWL}) - 1 - \mu_s \Lambda_0}{2(a-1)} \end{aligned}$$

$$\begin{aligned} &= \sum_{s \in \mathcal{S}^+} \frac{1}{4\mu_s(a-1)} \cdot \left((2(q_s - \mu_s \text{IWL}) + 1)^2 + (\mu_s \Lambda_0)^2 \right. \\ &\quad \left. + 2(2(q_s - \mu_s \text{IWL}) + 1)\mu_s \Lambda_0 \right) \\ &\quad + \sum_{s \in \mathcal{S}^+} \frac{-2(q_s - \mu_s \text{IWL}) + 1}{2\mu_s(a-1)} \cdot \frac{-2(q_s - \mu_s \text{IWL}) + 1 - \mu_s \Lambda_0}{2(a-1)} \\ &= \sum_{s \in \mathcal{S}^+} \frac{(2(q_s - \mu_s \text{IWL}) + 1)^2}{4\mu_s(a-1)} + \Lambda_0^2 \sum_{s \in \mathcal{S}^+} \frac{\mu_s}{4(a-1)} \\ &\quad + \Lambda_0 \sum_{s \in \mathcal{S}^+} \frac{2(2(q_s - \mu_s \text{IWL}) + 1)}{4(a-1)} \\ &\quad - \sum_{s \in \mathcal{S}^+} \frac{(2(q_s - \mu_s \text{IWL}) + 1)^2}{2\mu_s(a-1)} \\ &\quad - \Lambda_0 \sum_{s \in \mathcal{S}^+} \frac{(2(q_s - \mu_s \text{IWL}) + 1)}{2(a-1)} \\ &= \underbrace{\Lambda_0^2 \sum_{s \in \mathcal{S}^+} \frac{\mu_s}{4(a-1)}}_{v_1} - \underbrace{\sum_{s \in \mathcal{S}^+} \frac{(2(q_s - \mu_s \text{IWL}) + 1)^2}{4\mu_s(a-1)}}_{v_2}. \end{aligned}$$

This concludes the proof. \square

Lemma 2 enables us to compute v_1 and v_2 in $O(1)$ operations per iteration, and reduce the $\Omega(n)$ cost of line 12 in Algorithm 2 to $O(1)$. It follows that Algorithm 3 runs in $O(n \log n)$.

4.3 Linear complexity

The complexity of Algorithm 3 is dominated by sorting. For reducing the solution’s complexity further, we must remove the need for sorting both in Algorithms 1 and 3. To do so we use the median-of-medians approach [8].⁶ The correctness of the approach for computing the IWL is straight forward and it only remains to adapt the algorithm. Roughly, we seek the first server whose load would be equal or above the IWL (or that none exists). By using the median-of-medians algorithm to divide each sub-problem by ≈ 2 , and using previous calculations when moving to a more loaded server, we can compute the IWL in linear time. The pseudocode is presented in Algorithm 4.

Besides the IWL calculation, Algorithm 3 also relies on the servers being sorted to efficiently identify the probable set \mathcal{S}^+ . Thus, to prove the correctness of the median-of-medians approach here, we need to show that there is a sense

⁶ In essence, the median-of-medians approach relies on a sense of monotonicity. That is, for a given problem, the existence of a relation “ \leq ” between possible solutions such that for a possible solution x we can identify (in linear time) if the optimal solution x^* satisfies $x^* \leq x$ or $x \leq x^*$. Moreover, the approach is iterative and the identification problem must monotonically decrease in size by a constant factor at each iteration.

Algorithm 4: Computing the ideal workload.

```

1 Function ComputeIdealWorkLoad( $\mathcal{S}, \{q_s\}_{s \in \mathcal{S}}, \{\mu_s\}_{s \in \mathcal{S}}$ ,
a):
2    $\mathcal{T} \leftarrow \mathcal{S}; \mu_{tot}^* \leftarrow 0; l^* \leftarrow 0$ 
3    $IWL \leftarrow 0$ 
4   while  $\mathcal{T} \neq \emptyset$  do
5      $r \leftarrow \arg \text{median}_{s \in \mathcal{T}} \frac{q_s}{\mu_s}$ 
6      $\mathcal{T}_{upto} \leftarrow \{s \in \mathcal{T} \mid \frac{q_s}{\mu_s} \leq \frac{q_r}{\mu_r}\}$ 
7      $l \leftarrow l^* + \left(\frac{q_r}{\mu_r} - IWL\right) \mu_{tot}^* + \sum_{s \in \mathcal{T}_{upto}} \left(\frac{q_r}{\mu_r} - \frac{q_s}{\mu_s}\right) \mu_s$ 
8      $\mu_{tot} \leftarrow \mu_{tot}^* + \sum_{s \in \mathcal{T}_{upto}} \mu_s$ 
9     if  $l = a$  then
10      return  $\frac{q_r}{\mu_r}$ 
11     else if  $l > a$  then
12       $\mathcal{T} \leftarrow \mathcal{T}_{upto} \setminus \{r\}$ 
13     else
14       $IWL \leftarrow \frac{q_r}{\mu_r}$ 
15       $\mathcal{T} \leftarrow \mathcal{T} \setminus \mathcal{T}_{upto}$ 
16       $l^* \leftarrow l$ 
17       $\mu_{tot}^* \leftarrow \mu_{tot}$ 
18   return  $IWL + \frac{a - l^*}{\mu_{tot}^*}$ 

```

of monotonicity in the choice of \mathcal{S}^+ . More precisely, using the same notation as in Corollary 1, we additionally denote by $val[j]$ the value in line 16 of Algorithm 3 (line 13 of Algorithm 2) during the j^{th} iteration. Similarly, we denote by $\Lambda[j]$ the value in line 11 of Algorithm 3 during the j^{th} iteration. The monotonicity requirement formally translates to $\mathcal{S}_j, \mathcal{S}_{j'}$ with $j < j'$ satisfying the following two conditions: (1) $val[j] \geq val[j']$, and (2) that if $\Lambda_0[j] > 2IWL - \frac{2q_{i_j} + 1}{\mu_{i_j}}$ then $\Lambda_0[j'] > 2IWL - \frac{2q_{i_{j'}} + 1}{\mu_{i_{j'}}}$. The first part trivially holds since adding degrees of freedom to the possible assignment can only benefit us. I.e., an assignment to $\mathcal{S}_{j'}$ where only the servers in \mathcal{S}_j have positive probabilities while the rest are fixed to 0 is still possible, hence, an optimal solution is at least as good as that. It remains to show that the second part also holds. To that end we will use the following lemma.

Lemma 3 *If $\frac{a}{b} \geq \frac{c}{d}$ and $b, d > 0$, then $\frac{a}{b} \geq \frac{a+c}{b+d} \geq \frac{c}{d}$.*

Proof Simple arithmetic shows that

$$\begin{aligned} \frac{a}{b} &= \frac{a(b+d)}{b(b+d)} = \frac{a}{b+d} \cdot \frac{b+d}{b} = \frac{a}{b+d} \cdot \left(1 + \frac{d}{b}\right) \\ &= \frac{a}{b+d} + \frac{\frac{ad}{b}}{b+d} \geq \frac{a}{b+d} + \frac{c}{b+d} = \frac{a+c}{b+d}. \end{aligned}$$

Similarly,

$$\frac{c}{d} = \frac{c(b+d)}{d(b+d)} = \frac{c}{b+d} \cdot \frac{b+d}{d} = \frac{c}{b+d} \cdot \left(1 + \frac{b}{d}\right)$$

$$= \frac{c}{b+d} + \frac{\frac{cb}{d}}{b+d} \leq \frac{c}{b+d} + \frac{a}{b+d} = \frac{a+c}{b+d}.$$

The claim follows. □

Lemma 4 *Let $j, j' \in \{1, \dots, n\}$ denote the j^{th} and j'^{th} servers in the listing of \mathcal{S} according to the non-decreasing order of $\frac{2q_s + 1}{\mu_s}$. If $\Lambda_0[j] > 2IWL - \frac{2q_j + 1}{\mu_j}$, then for all $j' > j$ it holds that $\Lambda_0[j'] > 2IWL - \frac{2q_{j'} + 1}{\mu_{j'}}$.*

Proof In accordance with the previous notations, we denote by $\Lambda_n[j]$ and $\Lambda_d[j]$ the values of $\Lambda_{0,n}$ and $\Lambda_{0,d}$ in the j^{th} iteration of Algorithm 3 at lines 9 and 10 respectively. Let j be the first index at which $\Lambda_0[j] > 2IWL - \frac{2q_j + 1}{\mu_j}$. That is, we get an infeasible solution at line 12 of Algorithm 3. For $j + 1$ we have that

$$\begin{aligned} \Lambda_0[j + 1] &= \frac{\Lambda_n[j + 1]}{\Lambda_d[j + 1]} \\ &= \frac{\overbrace{\Lambda_n[j]}^a + \overbrace{2\mu_{j+1}IWL - (2q_{j+1} + 1)}^c}{\underbrace{\Lambda_d[j]}_b + \underbrace{\mu_{j+1}}_d}. \end{aligned}$$

Case 1: Suppose $\frac{c}{d} \geq \frac{a}{b}$. Then from Lemma 3 it follows that $\Lambda_d[j + 1] \geq \Lambda_d[j]$. Moreover, the non-decreasing order satisfies $2IWL - \frac{2q_{j+1} + 1}{\mu_{j+1}} \geq 2IWL - \frac{2q_j + 1}{\mu_j}$. Consequently,

$$\begin{aligned} \Lambda_0[j + 1] &\geq \Lambda_d[j] > 2IWL - \frac{2q_j + 1}{\mu_j} \\ &\geq 2IWL - \frac{2q_{j+1} + 1}{\mu_{j+1}}, \end{aligned}$$

and $\Lambda_0[j + 1]$ results in an infeasible solution as well.

Case 2: Suppose $\frac{a}{b} > \frac{c}{d}$. Then by Lemma 3 we get that

$$\begin{aligned} \Lambda_0[j + 1] &> \frac{2\mu_{j+1}IWL - (2q_{j+1} + 1)}{\mu_{j+1}} \\ &= 2IWL - \frac{2q_{j+1} + 1}{\mu_{j+1}}, \end{aligned}$$

and $\Lambda_0[j + 1]$ results in an infeasible solution as well.

We have shown that if $\Lambda_0[j]$ results in an infeasible solution than $\Lambda_0[j + 1]$ also results in an infeasible solution. The claim follows by induction. □

With Lemma 4 proven, we can employ the median-of-medians approach to design Algorithm 5.

We have shown an algorithm that computes the dispatching probabilities in linear time complexity. Linear complexity is also theoretically optimal, since each server must be assigned a probability and it is possible for probabilities to

Algorithm 5: Find probabilities in $O(n)$ time

```

1 Function ComputeProbabilities( $\mathcal{S}, \{q_s\}_{s \in \mathcal{S}}, \{\mu_s\}_{s \in \mathcal{S}}, a,$ 
  IWL):
2    $\mathcal{T} \leftarrow \mathcal{S}$   $\Lambda_{0,n}^* \leftarrow -2(a-1); \Lambda_{0,d}^* \leftarrow 0$ 
3   while  $\mathcal{T} \neq \emptyset$  do
4      $r \leftarrow \arg \text{median}_{s \in \mathcal{T}} \frac{2q_s+1}{\mu_s}$ 
5      $\mathcal{T}_{upto} \leftarrow \left\{ s \in \mathcal{T} \mid \frac{2q_s+1}{\mu_s} \leq \frac{2q_r+1}{\mu_r} \right\}$ 
6      $\Lambda_{0,n} \leftarrow \Lambda_{0,n}^* + \sum_{s \in \mathcal{T}_{upto}} (2(\mu_s \text{IWL} - q_s) - 1)$ 
7      $\Lambda_{0,d} \leftarrow \Lambda_{0,d}^* + \sum_{s \in \mathcal{T}_{upto}} \mu_s$ 
8      $\Lambda_0 \leftarrow \frac{\Lambda_{0,n}}{\Lambda_{0,d}}$ 
9     if  $2\text{IWL} - \frac{2q_r+1}{\mu_r} < \Lambda_0$  then
10      /* infeasible; try smaller */
11       $\mathcal{T} \leftarrow \mathcal{T}_{upto} \setminus r$ 
12    else /* feasible; try larger */
13       $\mathcal{T} \leftarrow \mathcal{T} \setminus \mathcal{T}_{upto}$ 
14       $\Lambda_{0,n}^* \leftarrow \Lambda_{0,n}$ 
15       $\Lambda_{0,d}^* \leftarrow \Lambda_{0,d}$ 
16     $\Lambda_0^* \leftarrow \frac{\Lambda_{0,n}^*}{\Lambda_{0,d}^*}$ 
17     $P^* \leftarrow \left\{ \max\left\{0, \frac{-2(q_s - \mu_s \text{IWL}) - 1 - \mu_s \Lambda_0^*}{2(a-1)}\right\} \right\}_{s \in \mathcal{S}}$  return  $P^*$ 

```

be unique. However, our simulations in Sect. 7.3 show that for the system sizes we are interested in, Algorithm 3 runs faster. This is due to both constant factors, as well as due to sorting being $\Theta(n \log n)$ only in the worst case. In fact, since the data structures containing the servers according to the sorted orders are already sorted according to the previous round, the worst case complexity may be very rare. Clearly, in larger systems the asymptotic behavior will prevail and Algorithm 5 will be faster. Also, it is possible to accelerate Algorithm 5's average running time by computing an approximate median via sub-sampling instead of computing the exact one. Intuitively, a run of this variant is faster with high probability, since its expected time complexity is linear with a smaller constant and a strong concentration around its mean. But, for the context of this work we continue with the approach based on sorting.

5 Putting it all together

We now show how the algorithms from Sections 3 and 4, which compute the IWL and the optimal probabilities, can be employed in the complete dispatching procedure given in Algorithm 6. The complexity of an individual dispatcher's computation in any given round is $O(n \log n)$, and it is dominated by the sorting of n values in lines 2 and 3. If the sorted order is available to both COMPUTEIDEALWORKLOAD and COMPUTEPROBABILITIES, their running time is reduced

to $O(n)$. This is useful since (1) a designer may implement a sorted data structure in various ways according to what benefits her specific system's characteristics, and (2) if queue-length information is available before the job arrivals, the server ordering can be precomputed, further improving the online complexity.

Algorithm 6: STOCHASTICALLYCOORDINATEDDISPATCHING (SCD) Procedure.

```

input:  $\mathcal{S}, \{\mu_s\}_{s \in \mathcal{S}}$ 
1 for each round  $t = 1, 2, \dots$  do
2   Update queue-lengths  $\{q_s\}_{s \in \mathcal{S}}$ 
3    $\mathcal{S}_1 \leftarrow$  Sort  $\mathcal{S}$  according to  $\frac{q_s}{\mu_s}$ 
4    $\mathcal{S}_2 \leftarrow$  Sort  $\mathcal{S}$  according to  $\frac{2q_s+1}{\mu_s}$ 
5   upon receiving jobs  $a^{(d)}$  do
6     Estimate  $a$  by  $a_{\text{est},d}$ 
7     IWL  $\leftarrow$  COMPUTEIDEALWORKLOAD( $\mathcal{S}_1, \{q_s\}_{s \in \mathcal{S}}, \{\mu_s\}_{s \in \mathcal{S}}, a_{\text{est},d}$ )
8      $P \leftarrow$  COMPUTEPROBABILITIES( $\mathcal{S}_2, \{q_s\}_{s \in \mathcal{S}}, \{\mu_s\}_{s \in \mathcal{S}}, a_{\text{est},d}, \text{IWL}$ )
9     for each job in  $a^{(d)}$  do
10      draw a server  $s$  according to  $P$ 
11      send the job to server  $s$ 

```

5.1 Estimating the arrivals

While the derivation of the optimal probabilities assumed knowledge of the arrivals $\{a^{(1)}, \dots, a^{(m)}\}$, a dispatcher d only knows its own arrivals, i.e., $a^{(d)}$. Nevertheless, we note that the optimal probabilities depend only on the total number $a \triangleq \sum_{d \in \mathcal{D}} a^{(d)}$ of jobs that arrive, and not on the individual values $a^{(d)}$. We are thus left with estimating a . There are numerous optional methods for estimating a (e.g., assuming that it is the maximal capacity of the system, ML estimators, etc.). Following a simple and elegant approach taken in [24], we have dispatcher d estimate a by assuming that everyone else is receiving the same number of jobs as d :

$$a_{\text{est},d} = m \cdot a^{(d)}. \tag{18}$$

With this, Algorithm 6 is fully defined and can be implemented.

One reason why such an estimation scheme is effective is because the average estimation of the dispatchers exactly equals the total arrivals. That is,

$$\frac{1}{m} \sum_{d \in \mathcal{D}} a_{\text{est},d} = \frac{1}{m} \sum_{d \in \mathcal{D}} m a^{(d)} = \sum_{d \in \mathcal{D}} a^{(d)} = a. \tag{19}$$

Consequently, if some dispatchers overestimate the IWL and therefore assign smaller probabilities to less loaded servers, then other dispatchers underestimate the IWL and increase these probabilities. Roughly speaking, these deviations compensate for one another. In Sect. 7 we show how this simplistic estimation technique results in consistent state-of-the-art performance across many systems and metrics.

5.2 Stability

A formal guarantee for dispatching algorithms that is often considered desirable in the literature is called *stability* (see Sect. 6). Under mild assumptions on the stochastic nature of the arrival and service processes, stability ensures that the servers' queue-lengths will not grow unboundedly so long as the arrivals do not surpass all servers' total processing capacity.

Accordingly, in Sect. 6, we make the appropriate formal definitions and prove that *SCD* is stable. Moreover, our stability proof holds not only when applying the estimation technique in Eq. (18), but also applies for any estimation technique in which $1 \leq a_{est,d} < \infty$. Intuitively, this is because for $a_{est,d} = 1$, our policy behaves similarly to *SED*, and as $a_{est,d} \rightarrow \infty$ it approaches weighted-random. However, with any reasonable estimation (and in particular the one in Eq. (18) that we employ), the *SCD* procedure finds the best of both worlds. It eliminates the herding phenomenon incurred by *SED*, while also sending sufficient work to the less loaded servers, unlike the load oblivious weighted-random.

6 Detailed proof of strong stability

This section could be skipped at a first reading.

In line with standard practice [24,58,64], we make assumptions on the arrival and departure processes that make the system dynamics amenable to formal analysis. In particular, we assume, for all dispatchers $d \in \mathcal{D}$,

$$\begin{aligned} \{a^{(d)}(t)\}_{t=0}^\infty \text{ is an } i.i.d. \text{ process and,} \\ \mathbb{E}[a^{(d)}(0)] = \lambda_d, \quad \mathbb{E}[(a^{(d)}(0))^2] = \sigma_d. \end{aligned} \tag{20}$$

Likewise, for all servers $s \in \mathcal{S}$

$$\begin{aligned} \{c_s(t)\}_{t=0}^\infty \text{ is an } i.i.d. \text{ process,} \\ \mathbb{E}[c_s(0)] = \mu_s \text{ and } \mathbb{E}[(c_s(0))^2] = \varphi_s. \end{aligned} \tag{21}$$

Namely, for both the arrival and departure processes, we make the standard assumption that they are *i.i.d.* and have a

finite variance. We remark that the assumption that the arrival processes at the dispatchers are independent is made for ease of exposition and can be dropped by one skilled at the art at the cost of a more involved presentation (e.g., [58]).

Intuitively, our goal is to prove that if, on average, the total arrivals at the system are below the total processing capacity of all servers, then the expected queue-lengths at the servers are bounded by a constant. We next present the required terms to formalize this intuition.

Admissibility. We assume the total expected arrival rate to the system is admissible. Formally, it means that we assume that there exists an $\epsilon > 0$ such that $\sum_s \mu_s - \sum_d \lambda_d = \epsilon$.

We prove that for any multi-dispatcher heterogeneous system with admissible arrivals our dispatching policy is strongly stable. Our proof follows similar lines to the strong stability proof in [24]. The key difference is that we account for server heterogeneity. This makes the proof somewhat more involved. We next formally define the well-established *strong stability* criterion of interest. Strong stability is a strong form of stability for discrete-time queuing systems. Similarly to [24,58,64], since we assume that the arrival and departure processes has a finite variance, this criterion also implies throughput optimality and other strong theoretical guarantees that may be of interest (see [22,44,45] for details).

Definition 1 (Strong stability) A load balancing system is said to be strongly stable if for any admissible arrival rate it holds that

$$\limsup_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} \sum_{s \in \mathcal{S}} \mathbb{E}[q_s(t)] < \infty .$$

Now, we are ready to formalize the queue dynamics. Let $\bar{a}_s(t) = \sum_{d \in \mathcal{D}} \bar{a}_s^{(d)}(t)$ be the total number of arrivals at server s and round t . Then, the recursion describing the queue dynamics of server s over rounds is given by

$$q_s(t + 1) = \max\{0, q_s(t) + \bar{a}_s(t) - c_s(t)\} . \tag{22}$$

Squaring both sides of Eq. (22), rearranging, dividing by the server's processing capacity and omitting terms yields,

$$\begin{aligned} \frac{1}{\mu_s} (q_s(t + 1))^2 - \frac{1}{\mu_s} (q_s(t))^2 \\ \leq \frac{1}{\mu_s} (\bar{a}_s(t))^2 + \frac{1}{\mu_s} (c_s(t))^2 \\ - \frac{2q_s(t)}{\mu_s} (c_s(t) - \bar{a}_s(t)) . \end{aligned} \tag{23}$$

Summing Eq. (23) over the servers yields

$$\begin{aligned} & \sum_{s \in \mathcal{S}} \frac{1}{\mu_s} (q_s(t+1))^2 - \sum_{s \in \mathcal{S}} \frac{1}{\mu_s} (q_s(t))^2 \\ & \leq \sum_{s \in \mathcal{S}} \frac{1}{\mu_s} (\bar{a}_s(t))^2 + \sum_{s \in \mathcal{S}} \frac{1}{\mu_s} (c_s(t))^2 \\ & \quad - 2 \sum_{s \in \mathcal{S}} \frac{q_s(t)}{\mu_s} (c_s(t) - \bar{a}_s(t)). \end{aligned} \tag{24}$$

Denote $\mu_{tot} = \sum_{s \in \mathcal{S}} \mu_s$. Now, we define the following useful quantity for each server s : $w_s = \frac{\mu_s}{\mu_{tot}}$. Next, for each (s, d, k) , let $I_s^{d,k}(t)$ be an indicator function that takes the value of 1 with probability w_s and 0 otherwise such that

$$\sum_{s \in \mathcal{S}} I_s^{d,k}(t) = 1 \quad \forall d \in \mathcal{D}, k \in [1, \dots, a^{(d)}(t)].$$

We rewrite Eq. (24) and add and subtract the term $2 \sum_{s \in \mathcal{S}} \sum_{d \in \mathcal{D}} \sum_{k=1}^{a^{(d)}(t)} I_s^{d,k}(t) q_s(t)$ from the right hand side of the equation. This yields,

$$\begin{aligned} & \sum_{s \in \mathcal{S}} \frac{1}{\mu_s} (q_s(t+1))^2 - \sum_{s \in \mathcal{S}} \frac{1}{\mu_s} (q_s(t))^2 \\ & \leq \underbrace{\sum_{s \in \mathcal{S}} \frac{1}{\mu_s} (\bar{a}_s(t))^2 + \sum_{s \in \mathcal{S}} \frac{1}{\mu_s} (c_s(t))^2}_{(a)} \\ & \quad - 2 \underbrace{\sum_{s \in \mathcal{S}} \frac{q_s(t)}{\mu_s} (c_s(t) - \sum_{d \in \mathcal{D}} \sum_{k=1}^{a^{(d)}(t)} I_s^{d,k}(t))}_{(b)} \\ & \quad + 2 \underbrace{\sum_{s \in \mathcal{S}} \frac{q_s(t)}{\mu_s} \left(\sum_{d \in \mathcal{D}} \bar{a}_s^{(d)}(t) - \sum_{d \in \mathcal{D}} \sum_{k=1}^{a^{(d)}(t)} I_s^{d,k}(t) \right)}_{(c)}, \end{aligned} \tag{25}$$

where we also used $\bar{a}_s(t) = \sum_{d \in \mathcal{D}} \bar{a}_s^{(d)}(t)$ in (c). Our goal now is to take the expectation of Eq. (25). We start with analyzing Term (a) in Eq. (25). Denote $\mu_{\min} = \min_{s \in \mathcal{S}} \mu_s$. Taking expectation and using Eq. (20) and Eq. (21) we obtain

$$\begin{aligned} & \mathbb{E} \left[\sum_{s \in \mathcal{S}} \frac{1}{\mu_s} (\bar{a}_s(t))^2 + \sum_{s \in \mathcal{S}} \frac{1}{\mu_s} (c_s(t))^2 \right] \\ & \leq \frac{1}{\mu_{\min}} \mathbb{E} \left[\left(\sum_{d \in \mathcal{D}} a^{(d)}(t) \right)^2 \right] + \sum_{s \in \mathcal{S}} \frac{\varphi_s}{\mu_s} = \frac{1}{\mu_{\min}} \\ & \quad \cdot \left(\sum_{d \in \mathcal{D}} \sigma_d + \sum_{d \in \mathcal{D}} \sum_{\substack{d' \in \mathcal{D}, \\ d' \neq d}} \lambda_d \lambda_{d'} \right) + \sum_{s \in \mathcal{S}} \frac{\varphi_s}{\mu_s} \triangleq C. \end{aligned} \tag{26}$$

We now turn to analyze Term (b) in Eq. (25). Using the law of total expectation we obtain

$$\begin{aligned} & \mathbb{E} \left[\sum_{s \in \mathcal{S}} \frac{q_s(t)}{\mu_s} (c_s(t) - \sum_{d \in \mathcal{D}} \sum_{k=1}^{a^{(d)}(t)} I_s^{d,k}(t)) \right] \\ & = \mathbb{E} \left[\mathbb{E} \left[\sum_{s \in \mathcal{S}} \frac{q_s(t)}{\mu_s} (c_s(t) - \sum_{d \in \mathcal{D}} \sum_{k=1}^{a^{(d)}(t)} I_s^{d,k}(t)) \right] \right] \\ & \left[\{q_s(t)\}_{s \in \mathcal{S}} \right] = \sum_{s \in \mathcal{S}} \mathbb{E} \left[\frac{q_s(t)}{\mu_s} \left(\mu_s - w_s \sum_{d \in \mathcal{D}} \lambda_d \right) \right] \\ & = \sum_{s \in \mathcal{S}} \mathbb{E} \left[\frac{q_s(t)}{\mu_{tot}} \left(\mu_{tot} - \sum_{d \in \mathcal{D}} \lambda_d \right) \right] \\ & = \frac{\epsilon}{\mu_{tot}} \sum_{s \in \mathcal{S}} \mathbb{E} [q_s(t)], \end{aligned} \tag{27}$$

where we used Eq. (20), Eq. (21), the definition of $I_s^{d,k}(t)$ and the admissibility of the system. We also used the fact that $a^{(d)}(t)$ and $I_s^{d,k}(t)$ are independent which allowed us to employ Wald’s identity. We turn to analyze term (c) in Eq. (25).

Observe that

$$\begin{aligned} & \sum_{s \in \mathcal{S}} \frac{q_s(t)}{\mu_s} \left(\sum_{d \in \mathcal{D}} \bar{a}_s^{(d)}(t) - \sum_{d \in \mathcal{D}} \sum_{k=1}^{a^{(d)}(t)} I_s^{d,k}(t) \right) \\ & = \sum_{d \in \mathcal{D}} \left(\sum_{s \in \mathcal{S}} \frac{q_s(t)}{\mu_s} \left(\bar{a}_s^{(d)}(t) - \sum_{k=1}^{a^{(d)}(t)} I_s^{d,k}(t) \right) \right). \end{aligned} \tag{28}$$

This means that we can consider each dispatcher separately. In particular, for each dispatcher, we will apply the following Lemma.

Lemma 5 Consider dispatcher $d \in \mathcal{D}$. For any $a^{(d)}(t) > 0$, let $\{p_s^{(d)}(t)\}_{s \in \mathcal{S}}$ be the optimal probabilities computed by dispatcher d for round t (i.e., probabilities that respect Eq. (10) for the computed IWL). Then, for any $p_s^{(d)}(t), p_{s'}^{(d)}(t) > 0$, it holds that

$$\begin{aligned} & \frac{p_s^{(d)}(t)}{\mu_s} \leq \frac{p_{s'}^{(d)}(t)}{\mu_{s'}} \implies \\ & \frac{q_s(t)}{\mu_s} + \frac{a^{(d)}(t)}{\mu_{\min}} \geq \frac{q_s(t) + a^{(d)}(t)}{\mu_s} \geq \frac{q_{s'}(t)}{\mu_{s'}}. \end{aligned} \tag{29}$$

Proof See Sect. 6.1 □

With this result at hand, we produce the following construction. Let $p_{s_{(1)}}^{(d)}(t), p_{s_{(2)}}^{(d)}(t), \dots, p_{s_{(n)}}^{(d)}(t)$ be the optimal probabilities computed by dispatcher d ordered in a decreasing order of their value divided by the corresponding server

rate, i.e., $\frac{p_s^{(d)}(t)}{\mu_s}$. Then, according to Lemma 5, for any two servers $s_{(i)}, s_{(i+1)}$ with positive probabilities, we have that

$$\frac{q_{s_{(i+1)}}(t)}{\mu_{s_{(i+1)}}} + i \frac{a^{(d)}(t)}{\mu_{min}} \geq \frac{q_{s_{(i)}}(t)}{\mu_{s_{(i)}}} + (i-1) \frac{a^{(d)}(t)}{\mu_{min}}. \tag{30}$$

Now, for dispatcher $d \in \mathcal{D}$ consider

$$\begin{aligned} & \sum_{s_{(i)} \in \mathcal{S}} \frac{q_{s_{(i)}}(t)}{\mu_{s_{(i)}}} \left(\bar{a}_{s_{(i)}}^{(d)}(t) - \sum_{k=1}^{a^{(d)}(t)} I_{s_{(i)}}^{d,k}(t) \right) \\ &= \sum_{s_{(i)} \in \mathcal{S}} \left(\frac{q_{s_{(i)}}(t)}{\mu_{s_{(i)}}} + (i-1) \frac{a^{(d)}(t)}{\mu_{min}} \right) \cdot \underbrace{\left(\bar{a}_{s_{(i)}}^{(d)}(t) - \sum_{k=1}^{a^{(d)}(t)} I_{s_{(i)}}^{d,k}(t) \right)}_{(i)} \\ & - \sum_{s_{(i)} \in \mathcal{S}} \left((i-1) \frac{a^{(d)}(t)}{\mu_{min}} \right) \cdot \underbrace{\left(\bar{a}_{s_{(i)}}^{(d)}(t) - \sum_{k=1}^{a^{(d)}(t)} I_{s_{(i)}}^{d,k}(t) \right)}_{(ii)}. \end{aligned} \tag{31}$$

For (i), we use Lemma 6 which relies on a *majorization* argument similarly to [58,64].

Lemma 6

$$\mathbb{E} \left[\sum_{s_{(i)} \in \mathcal{S}} \left(\frac{q_{s_{(i)}}(t)}{\mu_{s_{(i)}}} + (i-1) \frac{a^{(d)}(t)}{\mu_{min}} \right) \cdot \left(\bar{a}_{s_{(i)}}^{(d)}(t) - \sum_{k=1}^{a^{(d)}(t)} I_{s_{(i)}}^{d,k}(t) \right) \right] \leq 0. \tag{32}$$

Proof See Sect. 6.2. □

We continue to analyze (ii). By Eq. (20), it holds that

$$\begin{aligned} & \mathbb{E} \left[\sum_{s_{(i)} \in \mathcal{S}} \left((i-1) \frac{a^{(d)}(t)}{\mu_{min}} \right) \cdot \left(\bar{a}_{s_{(i)}}^{(d)}(t) - \sum_{k=1}^{a^{(d)}(t)} I_{s_{(i)}}^{d,k}(t) \right) \right] \\ & \leq \mathbb{E} \left[\sum_{s_{(i)} \in \mathcal{S}} \left((i-1) \frac{(a^{(d)}(t))^2}{\mu_{min}} \right) \right] \\ & = \frac{\sigma_d(n^2 - n)}{2\mu_{min}}. \end{aligned} \tag{33}$$

By taking the expectation of Term (c) and using Eq. (28), Lemma 6 and Eq. (33), we obtain

$$\begin{aligned} & \mathbb{E} \left[\sum_{s \in \mathcal{S}} \frac{q_s(t)}{\mu_s} \left(\sum_{d \in \mathcal{D}} \bar{a}_s^{(d)}(t) - \sum_{d \in \mathcal{D}} \sum_{k=1}^{a^{(d)}(t)} I_s^{d,k}(t) \right) \right] \\ & \leq \sum_{d \in \mathcal{D}} \frac{\sigma_d(n^2 - n)}{2\mu_{min}} \triangleq D. \end{aligned} \tag{34}$$

Finally, by applying Eq. (26), Eq. (27) and Eq. (34), we can take the expectation of the right-hand side of Eq. (25). This yields

$$\begin{aligned} & \mathbb{E} \left[\sum_{s \in \mathcal{S}} \frac{1}{\mu_s} (q_s(t+1))^2 \right] - \mathbb{E} \left[\sum_{s \in \mathcal{S}} \frac{1}{\mu_s} (q_s(t))^2 \right] \\ & \leq C - \frac{2\epsilon}{\mu_{tot}} \sum_{s \in \mathcal{S}} \mathbb{E}[q_s(t)] + 2D. \end{aligned} \tag{35}$$

Summing Eq. (35) over rounds $0, \dots, T-1$, multiplying by $\frac{\mu_{tot}}{2\epsilon T}$ and rearranging yields,

$$\begin{aligned} & \frac{1}{T} \sum_{t=0}^{T-1} \sum_{s \in \mathcal{S}} \mathbb{E}[q_s(t)] \\ & \leq \frac{(C + 2D)\mu_{tot}}{2\epsilon} + \frac{\mu_{tot}}{2\epsilon T} \mathbb{E} \sum_{s \in \mathcal{S}} \frac{1}{\mu_s} (q_s(0))^2, \end{aligned} \tag{36}$$

where we omitted the non-positive term $\mathbb{E} \left[- \sum_{s \in \mathcal{S}} \frac{1}{\mu_s} (q_s(T))^2 \right]$ due to the telescopic series on the left hand side of Eq. (35). Taking limits of Eq. (36) and making the standard assumption that the system is initialized with bounded queue-lengths, i.e., $\mathbb{E} \left[\sum_{s \in \mathcal{S}} (q_s(0))^2 \right] < \infty$ yields:

$$\limsup_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} \sum_{s \in \mathcal{S}} \mathbb{E}[q_s(t)] \leq \frac{(C+2D)\mu_{tot}}{2\epsilon} < \infty. \tag{37}$$

This concludes the proof.

Note that the proof does not rely on the specific method used to estimate the total arrivals (which, in turn, affects the calculated dispatching probabilities). We only rely on the fact that the estimation respects $1 \leq a_{est,d} \leq \infty$ and on the way the probabilities are calculated based on the estimation. This allows the freedom to design even better estimation techniques than Eq. (18) as we briefly mentioned in the main text of the paper.

6.1 Proof of lemma 5

Assume by the way of contradiction that $\frac{p_s^{(d)}(t)}{\mu_s} \leq \frac{p_{s'}^{(d)}(t)}{\mu_{s'}}$ but $\frac{q_s(t)+a^{(d)}(t)}{\mu_s} < \frac{q_{s'}(t)}{\mu_{s'}}$. For ease of exposition, we next omit the dispatcher superscript d and the round index t . Consider an alternative solution where all other probabilities are identical except that we set $\tilde{p}_s = p_s + p_{s'}$ and $\tilde{p}_{s'} = 0$. By the definition of the error function given in Eq. (10), we have that the difference in the error is

$$\begin{aligned} \text{diff} &= (a-1) \frac{(p_s+p_{s'})^2}{\mu_s} + 2(p_s+p_{s'}) \left(\frac{q_s}{\mu_s} - \text{IWL} + \frac{0.5}{\mu_s} \right) \\ &\quad - \left((a-1) \frac{p_s^2}{\mu_s} + 2p_s \left(\frac{q_s}{\mu_s} - \text{IWL} + \frac{0.5}{\mu_s} \right) + (a-1) \frac{p_{s'}^2}{\mu_{s'}} \right. \\ &\quad \left. + 2p_{s'} \left(\frac{q_{s'}}{\mu_{s'}} - \text{IWL} + \frac{0.5}{\mu_{s'}} \right) \right). \end{aligned}$$

Simplifying yields,

$$\begin{aligned} \text{diff} &= (a-1) \left(\frac{p_{s'}^2 + 2p_s p_{s'}}{\mu_s} - \frac{p_{s'}^2}{\mu_{s'}} \right) \\ &\quad + 2p_{s'} \left(\frac{q_s}{\mu_s} + \frac{0.5}{\mu_s} \right) - 2p_{s'} \left(\frac{q_{s'}}{\mu_{s'}} + \frac{0.5}{\mu_{s'}} \right). \end{aligned}$$

We now split into two cases.

Case 1: $\mu_s \geq \mu_{s'}$. In this case

$$\begin{aligned} \text{diff} &\leq (a-1) \left(\frac{2p_s p_{s'}}{\mu_s} \right) + 2p_{s'} \left(\frac{q_s}{\mu_s} \right) - 2p_{s'} \left(\frac{q_{s'}}{\mu_{s'}} \right) \\ &< (a-1) \left(\frac{2p_s p_{s'}}{\mu_s} \right) - a \frac{2p_{s'}}{\mu_s} < 0. \end{aligned}$$

Case 2: $\mu_s < \mu_{s'}$. In this case

$$\begin{aligned} \text{diff} &\leq (a-1) \left(\frac{2p_s p_{s'}}{\mu_{s'}} \right) + 2p_{s'} \left(\frac{q_s}{\mu_s} \right) - 2p_{s'} \left(\frac{q_{s'}}{\mu_{s'}} \right) + \frac{p_{s'}}{\mu_s} \\ &< (a-1) \left(\frac{2p_s p_{s'}}{\mu_{s'}} \right) - a \frac{2p_{s'}}{\mu_s} + \frac{p_{s'}}{\mu_s} < 0. \end{aligned}$$

In both cases the new solution has a lower error. This is a contradiction to the assumption that $\frac{q_s(t)+a^{(d)}(t)}{\mu_s} < \frac{q_{s'}(t)}{\mu_{s'}}$. This concludes the proof.

6.2 Proof of lemma 6

For any $s \in \mathcal{S}$, we have that

$$\mathbb{E} \left[\bar{a}_{s(i)}^{(d)}(t) \mid a^{(d)}(t), \{q_s(t)\}_{s \in \mathcal{S}} \right] = p_{s(i)}^{(d)}(t) \cdot a^{(d)}(t),$$

and,

$$\begin{aligned} &\mathbb{E} \left[\sum_{k=1}^{a^{(d)}(t)} I_{s(i)}^{d,k}(t) \mid a^{(d)}(t), \{q_s(t)\}_{s \in \mathcal{S}} \right] \\ &= \frac{\mu_{s(i)}}{\mu_{tot}} \cdot a^{(d)}(t). \end{aligned}$$

This means that

- (i) $p_{s(i)}^{(d)}(t) \cdot a^{(d)}(t) = \mu_{s(i)} \cdot a^{(d)}(t) \cdot \frac{p_{s(i)}^{(d)}(t)}{\mu_{s(i)}}$,
- (ii) $\frac{\mu_{s(i)}}{\mu_{tot}} \cdot a^{(d)}(t) = \mu_{s(i)} \cdot a^{(d)}(t) \cdot \frac{1}{\mu_{tot}}$.

Observe that the term $\mu_{s(i)} \cdot a^{(d)}(t)$ is identical in both (i) and (ii), and recall that $\frac{p_{s(i)}^{(d)}(t)}{\mu_{s(i)}}$ is monotonically *non-increasing* in i . This allows us to obtain the following *majorization* [62] Lemma.

Lemma 7 Let $k \in \{0, 1, \dots, n-1\}$. Then,

$$\sum_{i=n}^{n-k} p_{s(i)}^{(d)}(t) \cdot a^{(d)}(t) \leq \sum_{i=n}^{n-k} \frac{\mu_{s(i)}}{\mu_{tot}} \cdot a^{(d)}(t) \quad \forall k,$$

where an equality is obtained for $k = n-1$.

Proof See Sect. 6.3. □

Finally, since $\left(\frac{q_{s(i)}(t)}{\mu_{s(i)}} + (i-1) \cdot \frac{a^{(d)}(t)}{\mu_{min}} \right)$ is a monotonically *non-decreasing* in i , we obtain the following Lemma.

Lemma 8 For all k :

$$\begin{aligned} &\sum_{i=n}^{n-k} \left(\frac{q_{s(i)}(t)}{\mu_{s(i)}} + (i-1) \cdot \frac{a^{(d)}(t)}{\mu_{min}} \right) p_{s(i)}^{(d)}(t) \cdot a^{(d)}(t) \\ &\leq \sum_{i=n}^{n-k} \left(\frac{q_{s(i)}(t)}{\mu_{s(i)}} + (i-1) \cdot \frac{a^{(d)}(t)}{\mu_{min}} \right) \frac{\mu_{s(i)}}{\mu_{tot}} \cdot a^{(d)}(t). \end{aligned}$$

Proof See Sect. 6.4. □

In particular, Lemma 8 holds for $k = n-1$. Therefore, using the Law of total expectation and Lemma 8 we obtain

$$\begin{aligned} &\mathbb{E} \left[\mathbb{E} \left[\sum_{s(i) \in \mathcal{S}} \left(\frac{q_{s(i)}(t)}{\mu_{s(i)}} + (i-1) \cdot \frac{a^{(d)}(t)}{\mu_{min}} \right) \right. \right. \\ &\quad \cdot \left. \left. \left(\bar{a}_{s(i)}^{(d)}(t) - \sum_{k=1}^{a^{(d)}(t)} I_{s(i)}^{d,k}(t) \right) \mid a^{(d)}(t), \{q_s(t)\}_{s \in \mathcal{S}} \right] \right] \\ &= \mathbb{E} \left[\sum_{s(i) \in \mathcal{S}} \left(\frac{q_{s(i)}(t)}{\mu_{s(i)}} + (i-1) \cdot \frac{a^{(d)}(t)}{\mu_{min}} \right) \right. \\ &\quad \cdot \left. \left(p_{s(i)}^{(d)}(t) \cdot a^{(d)}(t) - \frac{\mu_{s(i)}}{\mu_{tot}} \cdot a^{(d)}(t) \right) \right] \leq 0. \end{aligned} \tag{38}$$

This concludes the proof.

6.3 Proof of lemma 7

By the way of contradiction, consider the smallest $k \in \{0, 1, \dots, n - 1\}$ such that

$$\sum_{i=n}^{n-k} \mu_{s(i)} \cdot a^{(d)}(t) \cdot \frac{p_{s(i)}^{(d)}(t)}{\mu_{s(i)}} > \sum_{i=n}^{n-k} \mu_{s(i)} \cdot a^{(d)}(t) \cdot \frac{1}{\mu_{tot}}.$$

Dividing both terms by $a^{(d)}(t)$ yields

$$\sum_{i=n}^{n-k} \mu_{s(i)} \frac{p_{s(i)}^{(d)}(t)}{\mu_{s(i)}} > \sum_{i=n}^{n-k} \mu_{s(i)} \cdot \frac{1}{\mu_{tot}}.$$

Since $\frac{p_{s(i)}^{(d)}(t)}{\mu_{s(i)}}$ is a monotonically *non-increasing* sequence in i , we must have that if $k' \geq k$ then $\frac{p_{s(n-k')}^{(d)}(t)}{\mu_{s(n-k')}} > \frac{1}{\mu_{tot}}$. This means that it must hold that

$$\sum_{i=n}^1 \mu_{s(i)} \frac{p_{s(i)}^{(d)}(t)}{\mu_{s(i)}} > \sum_{i=n}^1 \mu_{s(i)} \cdot \frac{1}{\mu_{tot}} = 1.$$

This is a contradiction to the fact that $\sum_{i=n}^1 \mu_{s(i)} \cdot \frac{1}{\mu_{tot}} = \sum_{i=1}^n \mu_{s(i)} \cdot \frac{p_{s(i)}^{(d)}(t)}{\mu_{s(i)}} = \sum_{i=1}^n p_{s(i)}^{(d)}(t) = 1$ (that is, both are probability vectors and thus their sum must be equal to 1).

6.4 Proof of lemma 8

Denote $\eta_{(i)} = \left(\frac{q_{s(i)}(t)}{\mu_{s(i)}} + (i - 1) \cdot \frac{a^{(d)}(t)}{\mu_{min}} \right)$ and recall that

$\eta_{(i)}$ is a monotonically *non-decreasing* in i and that $\frac{p_{s(i)}^{(d)}(t)}{\mu_{s(i)}}$ is monotonically *non-increasing* in i .

Consider the smallest index j such that $p_{s(j)}^{(d)}(t) \leq \frac{\mu_{s(j)}}{\mu_{tot}}$. If $j = 1$ than by Lemma 7 and since $\sum_{i=1}^n p_{s(i)}^{(d)}(t) = 1$, it must hold that $p_{s(j)}^{(d)}(t) = \frac{\mu_{s(j)}}{\mu_{tot}} \forall i$ and we are done. Now, assume $j > 1$.

Consider the sums

$$\begin{aligned} & \sum_{i=j-1}^1 \eta_{(i)} \cdot \left(p_{s(i)}^{(d)}(t) - \frac{\mu_{s(i)}}{\mu_{tot}} \right) \\ & \leq \max\{\eta_{(i)}\}_{i \in \{1, \dots, j-1\}} \sum_{i=1}^{j-1} \left(p_{s(i)}^{(d)}(t) - \frac{\mu_{s(i)}}{\mu_{tot}} \right), \end{aligned}$$

and,

$$\begin{aligned} & \sum_{i=n}^j \eta_{(i)} \cdot \left(p_{s(i)}^{(d)}(t) - \frac{\mu_{s(i)}}{\mu_{tot}} \right) \\ & \leq \min\{\eta_{(i)}\}_{i \in \{j, \dots, n\}} \sum_{i=n}^j \left(p_{s(i)}^{(d)}(t) - \frac{\mu_{s(i)}}{\mu_{tot}} \right). \end{aligned}$$

Notice that

$$\max\{\eta_{(i)}\}_{i \in \{1, \dots, j-1\}} \leq \min\{\eta_{(i)}\}_{i \in \{j, \dots, n\}},$$

and,

$$\sum_{i=1}^{j-1} \left(p_{s(i)}^{(d)}(t) - \frac{\mu_{s(i)}}{\mu_{tot}} \right) = - \sum_{i=n}^j \left(p_{s(i)}^{(d)}(t) - \frac{\mu_{s(i)}}{\mu_{tot}} \right).$$

Therefore, we obtain

$$\begin{aligned} & \sum_{i=n}^1 \eta_{(i)} \cdot \left(p_{s(i)}^{(d)}(t) - \frac{\mu_{s(i)}}{\mu_{tot}} \right) \\ & = \sum_{i=j-1}^1 \eta_{(i)} \cdot \left(p_{s(i)}^{(d)}(t) - \frac{\mu_{s(i)}}{\mu_{tot}} \right) \\ & \quad + \sum_{i=n}^j \eta_{(i)} \cdot \left(p_{s(i)}^{(d)}(t) - \frac{\mu_{s(i)}}{\mu_{tot}} \right) \\ & \leq \max\{\eta_{(i)}\}_{i \in \{1, \dots, j-1\}} \sum_{i=1}^{j-1} \left(p_{s(i)}^{(d)}(t) - \frac{\mu_{s(i)}}{\mu_{tot}} \right) \\ & \quad + \min\{\eta_{(i)}\}_{i \in \{j, \dots, n\}} \sum_{i=n}^j \left(p_{s(i)}^{(d)}(t) - \frac{\mu_{s(i)}}{\mu_{tot}} \right) \\ & = \sum_{i=1}^{j-1} \left(p_{s(i)}^{(d)}(t) - \frac{\mu_{s(i)}}{\mu_{tot}} \right) \\ & \quad \cdot \left(\max\{\eta_{(i)}\}_{i \in \{1, \dots, j-1\}} - \min\{\eta_{(i)}\}_{i \in \{j, \dots, n\}} \right) \leq 0 \end{aligned} \tag{39}$$

This concludes the proof.

7 Evaluation

In this section, we describe simulations and compare our solution both to well-established algorithms and to the most recent state-of-the-art algorithms for the heterogeneous multi-dispatcher setting. The performance of load balancing techniques is usually evaluated at high loads, so that the job arrival rates are sufficiently high to utilize the servers.

In such a setting, we are interested in two main performance criteria. The first criterion is the average response time

over all client requests (i.e., jobs). For each request, we measure the number of rounds it spent in the system (i.e., from its arrival to a dispatcher to its departure from the server that processed it). The second criterion is the response times' tail distribution. This metric is crucial since it often represents client experience due to two reasons: (1) clients may send multiple requests (e.g., browsing a website), and delaying a small subset of these may ruin client experience; (2) a client request may be broken into multiple smaller tasks and the request time will be determined by the last tasks to complete (e.g., search engines). In both cases, it is important to meet the desired tail latency for the 95th, 99th, or even the 99.9th percentile of the distribution [11,24,46].

7.1 Setup

In our simulations, every run of each algorithm lasts for 10^5 rounds. A round is composed of three phases as described in Sect. 2. Namely, first new requests arrive at the dispatchers. The dispatchers must then act immediately and independently from each other and dispatch each request to a server for processing. Finally, the servers perform work and possibly finish requests. Finished requests immediately depart from the system.

Each dispatcher has its own arrival process of requests. Specifically, the number of requests that arrive to each dispatcher d at each round t is drawn from a Poisson distribution with parameter λ . Formally, $a^{(d)}(t) \sim \text{Pois}(\lambda_d)$. Each server has its own processing rate. Specifically, the number of requests that server s can process in each round t is geometrically distributed with parameter μ_s . Formally, $c_s(t) \sim \text{Geom}(\frac{1}{1+\mu_s})$. Our modeling of the arrival and departure processes follow standard practice (e.g., [2,34,43,57,58]).

Clearly, when the average number of requests that arrive at the system surpasses the average processing rate of all servers combined, the system is considered *infeasible*, i.e., it cannot process all requests and must drop some of them. It is therefore a standard assumption that the system is *admissible* [20,24,34,40,42,52–54,58,60,64], i.e., that the total average arrival rate at the system is upper-bounded by the total processing capacity of all servers. Accordingly, this is a setting we are interested in examining. We define the *offered load* to be $\rho = \frac{\mathbb{E}[\sum_{d \in \mathcal{D}} a^{(d)}(0)]}{\mathbb{E}[\sum_{s \in \mathcal{S}} c_s(0)]} = \frac{\sum_{d \in \mathcal{D}} \lambda_d}{\sum_{s \in \mathcal{S}} \mu_s}$.

To be admissible, it must hold that $\rho < 1$. Therefore, in all our simulations we test the performance of the different dispatching algorithms for $\rho \in (0, 1)$.

We have implemented 10 different dispatching techniques in addition to ours. These include both well-established techniques and the most recent state-of-the-art techniques. In particular, we compare against *JSQ* [15,61,63], *SED* [18,19,28,50], *JSQ(2)* [36,41,59], *WR* (weighted random), *JIQ* [34,42,52,53,56], *LSQ* [58,64], as well as *hJSQ(2)*, *hJIQ*

and *hLSQ*. The last three policies, i.e., *hJSQ(2)*, *hJIQ* and *hLSQ* are the adaptations of *JSQ(2)*, *JIQ* and *LSQ* to account for server heterogeneity.⁷ We also compare against the recent *TWF* policy of [24] that achieves stochastic coordination for homogeneous systems. For a fair comparison, in all our experiments, we use the same random seed across all algorithms, resulting in identical arrival and departure processes.

7.2 Response time

We ran the evaluation over four different systems and two different server heterogeneity configurations to represent: (a) moderate heterogeneity that may appear by having different generations of CPUs and hardware configurations of servers and, (b) higher heterogeneity settings that may appear in the presence of accelerators (e.g., FPGA or ASIC). Namely, sub-figures denoted by (a) show the evaluation results for systems where $\mu_s \sim U[1, 10]$. That is, the service rate of each server in each system is randomly drawn from the uniform distribution over the real interval $[1,10]$. Similarly, sub-figures denoted by (b) show the evaluation results for systems where $\mu_s \sim U[1, 100]$. That is, the service rate of each server in each system is randomly drawn from the uniform distribution over the real interval $[1,100]$.

Figure 3 shows the average response time of requests as a function of the offered load at the system. It is evident that *SCD* consistently achieves the best results across all systems and offered loads. Figure 4 shows the response time tail distribution in systems with 100 servers and 10 dispatchers. Again, *SCD* achieves the best results with no clear second best. For example, for moderate heterogeneity, at the offered load of $\rho = 0.99$ and considering the 10^{-4} percentile, which is often of interest, *SCD* improves over the second-best algorithm (*hLSQ* in this specific case) by over $2.1 \times$. *TWF*, which is the second-best at the average-response-time metric, degrades here and is outperformed by *SED* and *hLSQ* since they account for server heterogeneity where *TWF* does not. For higher heterogeneity, the margin from the second-best (*LSQ*) increases to $2.3 \times$. Note that with this higher heterogeneity, the delay tail distribution of *TWF* and *JSQ* is significantly degraded since they do not account for server heterogeneity (by more than an order of magnitude even for a load of $\rho = 0.7$).

⁷ Similarly to *SED*, the servers are ranked by their expected delay, i.e., by $\frac{q_s}{\mu_s}$, instead of by q_s . Likewise, when random sampling of servers occurs, servers are sampled proportionally to their processing rates rather than uniformly. Specifically, the probability to sample server s is $\frac{\mu_s}{\sum_{s \in \mathcal{S}} \mu_s}$ instead of $\frac{1}{n}$.

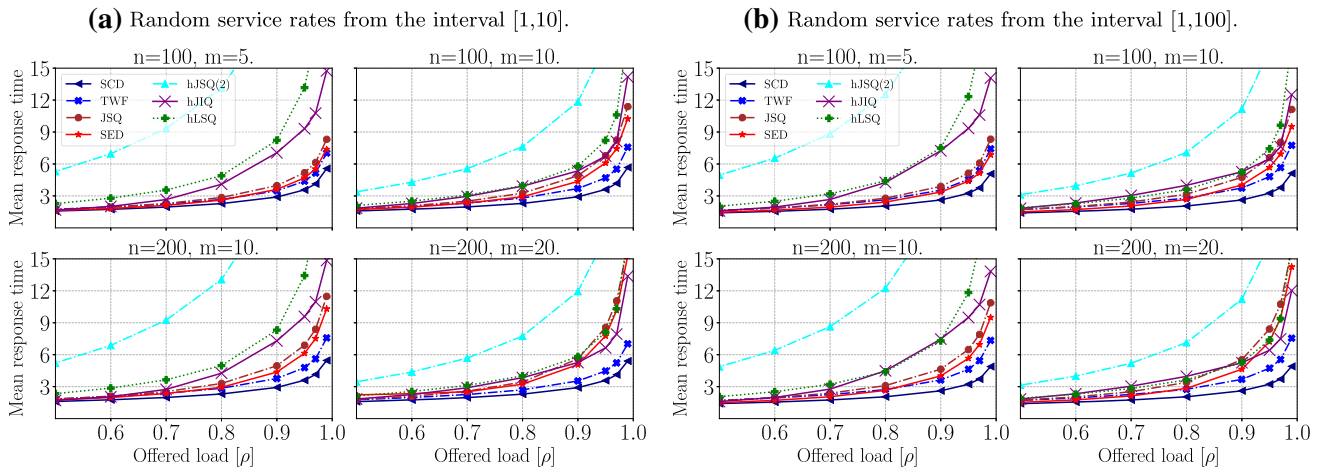
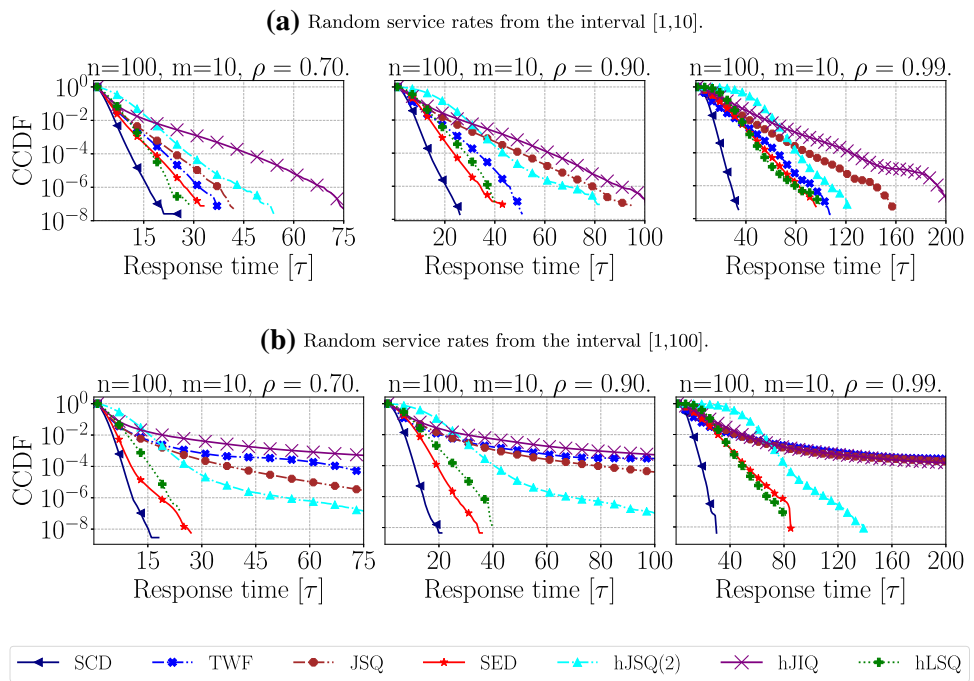


Fig. 3 The figures depict the average response time as a function of the offered load over four different systems. The x-axis represents the offered load ρ . The y-axis represents the average response time in number of rounds. Sub-figure **a** shows systems with moderate heterogeneity. Sub-figure **b** shows systems with higher heterogeneity

Fig. 4 The figures depict the response-time tail distribution over a system with 100 servers and 10 dispatchers over three different offered loads. The x-axis represents the response time in number of rounds (denoted by τ). The y-axis represents the complementary cumulative distribution function (CCDF). Sub-figure **a** shows systems with moderate heterogeneity. Sub-figure **b** shows systems with higher heterogeneity



For clarity, we discussed here only the 6 most competitive algorithms out of the 10 implemented. Below, in Sect. 7.2.1, we provide, for the interested reader, additional simulation results that compare between *SCD* and *JSQ(2)*, *JIQ*, *LSQ* and *WR* (weighted random). In a nutshell, the *JSQ(2)*, *JIQ* and *LSQ* algorithms are less competitive since they do not account for server heterogeneity, whereas *WR* ignores queue length information and fails to leverage less loaded servers in a timely manner.

7.2.1 Additional simulation results

We show complementary results to Figures 3 and 4 comparing *SCD* to *JSQ(2)*, *JIQ*, *LSQ* and weighted random (*WR*) in Figs. 5 and 6⁸. It is evident that *SCD* significantly outperforms all these techniques across all systems, metrics, and offered loads. Indeed, these techniques are less competitive than the six presented in the main text. This is because *JSQ(2)*, *JIQ*, and *LSQ* do not account for server heterogeneity whereas *WR* ignores queue length information.

⁸ In *WR*, any request is sent to server s with probability $\frac{\mu_s}{\mu_{tot}}$.

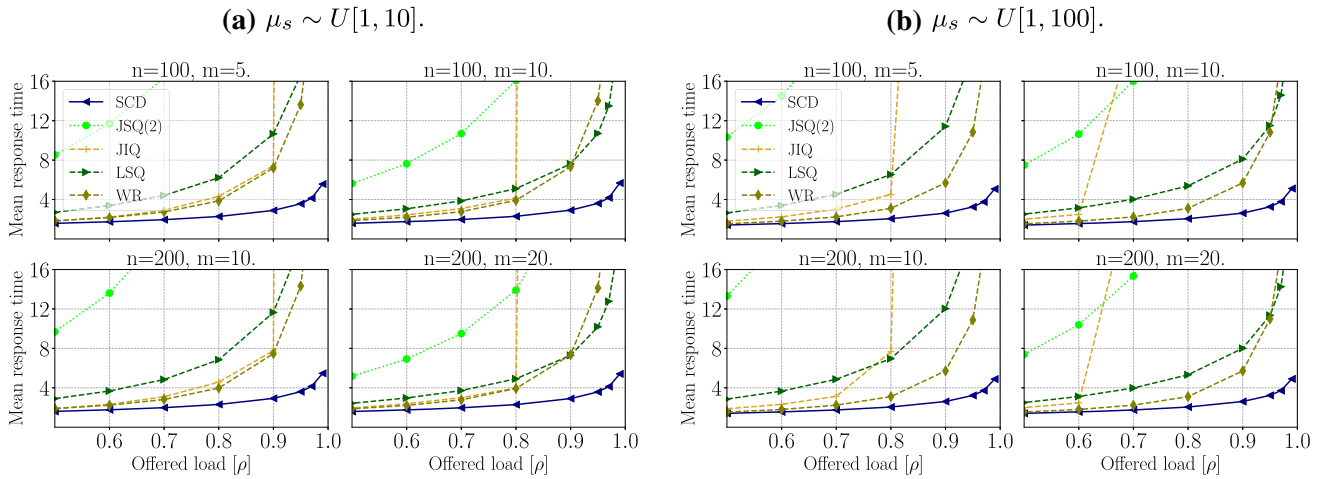
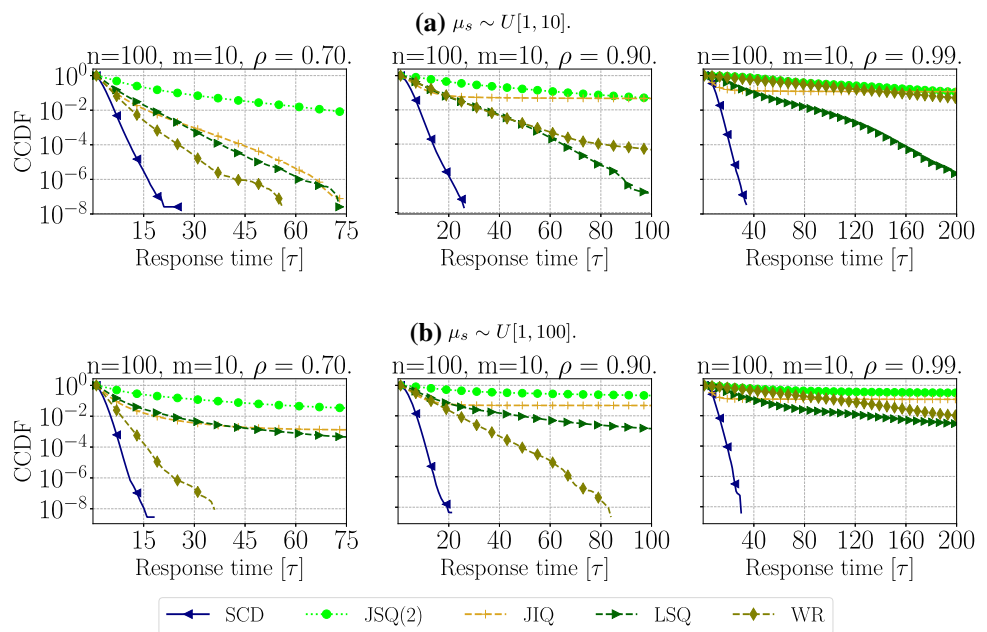


Fig. 5 Complementary results to Figure 3. Comparing *SCD* versus *JSQ(2)*, *JIQ*, *LSQ* and *WR*

Fig. 6 Complementary results to Figure 4. Comparing *SCD* versus *JSQ(2)*, *JIQ*, *LSQ* and *WR*



7.3 Computation costs

We next test *SCD*'s execution running times. That is, given the system state and arrivals, how much time does it take for a dispatcher to calculate the dispatching probabilities for that round? To answer that question, we implemented all dispatching techniques, and in particular *SCD* using algorithm 2, 3 and 5(+4), as well as *JSQ* and *SED* in C++ and optimized them for run-time purposes. All times were measured using a single core setup on a machine with an Intel Core i7-7700 CPU @3.60GHz and 16GB DDR3 2133MHz RAM.

For each algorithm in each round, we measure the time it takes each dispatcher to calculate its requests assignment to the servers. While the asymptotic complexity of the algo-

rithms is fixed, the complexity of each particular instance may require a different computation, depending on the number of arrived requests and server queue-lengths. Therefore, we report the cumulative distribution function (CDF) of those times.

Figures 7 and 8 shows the results for the setting described with $\mu_s \sim U[1, 10]$ and with $\mu_s \sim U[1, 100]$, respectively. For both heterogeneity levels, the running time of *SCD* via the algorithm from Sect. 4.2 scales similarly to *JSQ* and *SED* as expected. I.e., all three have complexity $O(n \log n)$. Overall, the measured running time of *SCD* via Algorithm 5, which employs the MoM mechanism, is similar to that of Algorithm 3. *SCD* via Algorithm 2 is slower.

Although Algorithm 5 is asymptotically optimal, for the system sizes we are interested in, Algorithm 3 runs slightly

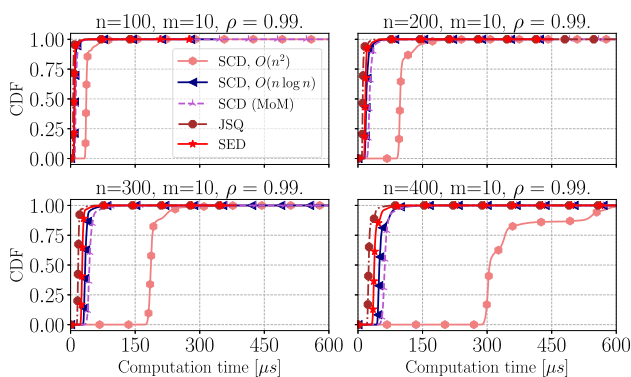


Fig. 7 Execution run times over systems with an increasing number of servers and $\mu_s \sim U[1, 10]$

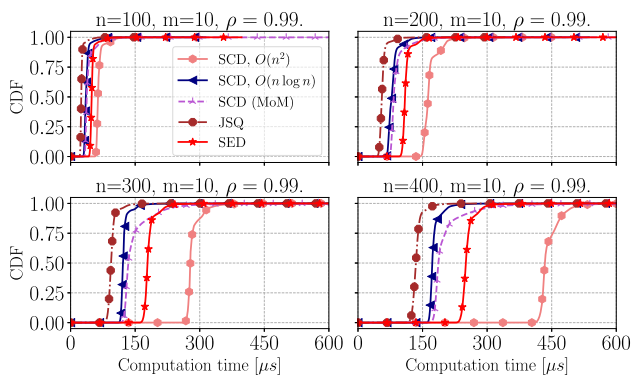


Fig. 8 Execution run times over systems with an increasing number of servers and $\mu_s \sim U[1, 100]$

faster. However, as the system grows larger the relative difference smallens (from circa 31% with 100 servers to about 25% with 400 servers, on average). For very large systems the linear behavior of Algorithm 5 will prevail, making it faster than Algorithm 3. (The conditions for which this occurs, are beyond the scope of this paper.)⁹ Unlike the linear *SCD* algorithm, the asymptotic effect of the quadratic behavior of Algorithm 2 is apparent in these system sizes, and it is clearly slower than the rest. Another interesting phenomenon is that with the increased heterogeneity, though constant, there is a larger gap among *SCD*, *JSQ* and *SED*. In particular, *SED* becomes slower than *SCD*. This is due to the following reason. All algorithms use dedicated and optimized-for-speed data structures.¹⁰ However, with the increased heterogeneity, there is an increased gap between the exact number of operations *JSQ* and *SED* require to update their data structures when assigning new requests to servers. For *JSQ*, when

⁹ Additionally, as discussed in Sect. 4.3, it is possible to accelerate Algorithm 5 using a non-deterministic variant with probabilistic run-time guarantees.

¹⁰ for *JSQ* and *SED* these are min-heaps that always keep the next best server at the top of the heap so we do not have to sort the server according to their queue length (*JSQ*) or their load (*SED*) after each update.

assigning a new request to a server s we need to update $q_s(t) \leftarrow q_s(t) + 1$ disregarding the server processing rate. This results in a predictable behavior of the data structure where only a few operations are needed to fix it (often a single operation). However, for *SED* the behavior is less predictable. The update in this case is $\frac{q_s(t)}{\mu_s} \leftarrow \frac{q_s(t)+1}{\mu_s}$ where the addition is proportionally inverse the server processing rate, i.e., proportional to $\frac{1}{\mu_s}$. Thus, the assignment of requests to different servers often requires additional operations to fix the data structure. These running time measurements prove to be consistent. In particular, we obtained similar results for other systems and server heterogeneity levels.

We conclude that from the algorithm running time perspective, *SCD*, when implemented via the algorithm described in Sect. 4.2, incurs an acceptable computational overhead, similar to that of *JSQ* and *SED*, which are in use in today’s high-performance load balancers [38,55].

8 Discussion

Large scale computing systems are ubiquitous today more than ever. Often, these systems have many distributed aspects which substantially affect their performance. The need to address the distributed nature of such systems, both algorithmically and fundamentally, is imperative. In this work, we presented *SCD*, a load balancing algorithm that addresses modern computer clusters’ distributed and heterogeneous nature in a principled manner. Extensive simulation results demonstrate that *SCD* outperforms the state-of-the-art load balancing algorithms across different systems and metrics. Regarding computation complexity, we designed *SCD* to run in $O(n \log n)$ time, as well as an algorithm that runs in asymptotically optimal $O(n)$ time. Therefore, it is no harder to employ in practical systems than traditional approaches such as *JSQ* and *SED*.

Our work leaves several open problems. For example: (1) The amount of work that a job requires may depend on specific features of the server processing it. Can information about the nature of jobs and features of servers be used to further improve the stochastic coordination among the dispatchers? (2) It may be the case that not all dispatchers are simultaneously connected to all servers, breaking the symmetry among the dispatchers. How should we incorporate such (possibly dynamic) connectivity information while maintaining stochastic coordination? (3) Can shared randomness (among the dispatchers) be used for better coordination? a possible challenge is that the analysis becomes even more involved (e.g., in some instances having shared randomness appear to have no advantage - e.g., (a) zero arrivals at all dispatchers except one; (b) probable sets that include one server).

Clearly, distributed load balancing presents many interesting challenges, and a broad set of issues for study that can impact the efficiency of practical systems.

Acknowledgements Guy Goren was partly supported by a grant from the Technion Hiroshi Fujiwara cyber security research center and the Israel cyber bureau, as well as by a Jacobs fellowship. Yoram Moses is the Israel Pollak academic chair at the Technion. His work was supported in part by the Israel Science Foundation under grant 2061/19. In addition, we would like to thank the reviewers for insightful suggestions that have contributed to the paper.

References

- Adler, M., Chakrabarti, S., Mitzenmacher, M., Rasmussen, L.: Parallel randomized load balancing. *Random Struct. Algorithms* **13**(2), 159–188 (1998)
- Anselmi, J., Dufour, F.: Power-of-d-choices with memory: fluid limit and optimality. *Math. Oper. Res.* (2020). <https://doi.org/10.1287/moor.2019.1014>
- Assadi, S., Bernstein, A., Langley, Z.: Improved bounds for distributed load balancing. arXiv preprint [arXiv:2008.04148](https://arxiv.org/abs/2008.04148), (2020)
- Atar, R., Keslassy, I., Orda, A., Vargaftik, S.: Persistent-idle load-distribution. *Stochastic Systems*, Gal Mendelson (2020)
- Azar, Y., Broder, A.Z., Karlin, A.R., Upfal, E.: Balanced allocations. In: *Proceedings of the twenty-sixth annual ACM symposium on theory of computing*, 593–602, (1994)
- Barbette, T., Tang, C., Yao, H., Kostić, D., Maguire Jr, G.Q., Papadimitratos, P., Chiesa, M.: A high-speed load-balancer design with guaranteed per-connection-consistency. In *17th Symposium on Networked Systems Design and Implementation (NSDI)*, 2020, 667–683, (2020)
- Berenbrink, P., Brinkmann, A., Friedetzky, T., Nagel, L.: Balls into non-uniform bins. *J. Parallel Distrib. Comput.* **74**(2), 2065–2076 (2014)
- Blum, M., Floyd, R.W., Pratt, V.R., Rivest, R.L., Tarjan, R.E., et al.: Time bounds for selection. *J. Comput. Syst. Sci.* **7**(4), 448–461 (1973)
- Boyd, S., Vandenberghe, L.: *Convex Optimization*. Cambridge University Press, Cambridge, England (2004)
- Czygrinow, A., Hanćkowiak, M., Szymańska, E., Wawrzyniak, W.: Distributed 2-approximation algorithm for the semi-matching problem. In: *International Symposium on Distributed Computing*, 210–222. Springer, (2012)
- Dean, J., Barroso, L.A.: The tail at scale. *Commun. ACM* **56**(2), 74–80 (2013)
- Delimitrou, C., Kozyrakis, C.: Paragon: Qos-aware scheduling for heterogeneous datacenters. *ACM SIGPLAN Notices* **48**(4), 77–88 (2013)
- Duato, J., Pena, A.J., Silla, F., Mayo, R., Quintana-Ortí, E.S.: rcuda: Reducing the number of gpu-based accelerators in high performance clusters. In: *2010 International Conference on High Performance Computing & Simulation*, 224–231. IEEE, (2010)
- Eisenbud, D.E., Yi, C., Contavalli, C., Smith, C., Kononov, R., Mann-Hielscher, E., Cilingiroglu, A., Cheyney, B., Shang, W., Hosein, J.D.: Maglev: a fast and reliable software network load balancer. In: *13th USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, 523–535, (2016)
- Eryilmaz, A., Srikant, R.: Asymptotically tight steady-state queue length bounds implied by drift conditions. *Queueing Syst.* **72**(3–4), 311–359 (2012)
- Foss, S., Chernova, N.: On the stability of a partially accessible multi-station queue with state-dependent routing. *Queueing Syst.* **29**(1), 55–73 (1998)
- Gandhi, R., Liu, H.H., Hu, Y.C., Lu, G., Padhye, J., Yuan, L., Zhang, M.: Duet: Cloud scale load balancing with hardware and software. *ACM SIGCOMM Comput. Commun. Rev.* **44**(4), 27–38 (2014)
- Gardner, K., Jaleel, J.A., Wickeham, A., Doroudi, S.: Scalable load balancing in the presence of heterogeneous servers. *Performance Evaluation*, (2021)
- Gardner, K., Stephens, Cole: Smart dispatching in heterogeneous systems. *ACM SIGMETRICS Perform. Eval. Rev.* **47**(2), 12–14 (2019)
- Garrett, O.: NGINX and the “Power of Two Choices” Load-Balancing Algorithm., 2018. <https://www.nginx.com/blog/nginx-power-of-two-choices-load-balancing-algorithm>, published on November 12, (2018)
- Garrett, O.: HTTP Load Balancing., (2021). <https://www.nginx.com/>
- Georgiadis, L., Neely, M.J., Tassiulas, L., et al.: Resource allocation and cross-layer control in wireless networks. *Found. Trends Netw.* **1**(1), 1–144 (2006)
- Goren, G., Vargaftik, S.: Implementation of the SCD algorithms., (2021). <https://github.com/guytechnion/stochastically-coordinated-dispatching>
- Goren, G., Vargaftik, S., Moses, Y.: Distributed dispatching in the parallel server model. In: *34th International Symposium on Distributed Computing (DISC 2020)*, (2020)
- Govindan, R., Minei, I., Kallahalla, M., Koley, B., Vahdat, A.: Evolve or die: High-availability design principles drawn from googles network infrastructure. In: *Proceedings of the 2016 ACM SIGCOMM Conference*, 58–72, (2016)
- Halldórsson, M., Köhler, S., Patt-Shamir, B., Rawitz, D.: Distributed backup placement in networks. *Distrib. Comput.* **31**(2), 83–98 (2018)
- Huang, M., Wu, D., Yu, C.H., Fang, Z., Interlandi, M., Condie, T., Cong, J.: Programming and runtime support to blaze fpga accelerator deployment at datacenter scale. In: *Proceedings of the Seventh ACM Symposium on Cloud Computing*, 456–469, (2016)
- Jaleel, J.A., Wickeham, A., Doroudi, S., Gardner, K.: A general “power-of-d” dispatching framework for heterogeneous systems. (2020)
- Karp, R.M., Luby, M., Meyer auf der Heide, F.: Efficient pram simulation on a distributed memory machine. *Algorithmica*, **16**(4):517–542, (1996)
- Karush, W.: Minima of functions of several variables with inequalities as side conditions. Master’s Thesis, Department of Mathematics, University of Chicago, (1939)
- Kleinberg, R., Piliouras, G., Tardos, É.: Load balancing without regret in the bulletin board model. *Distrib. Comput.* **24**(1), 21–29 (2011)
- Kuhn, H.W., Tucker, A.W.: Nonlinear programming. In: *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, 481–492, Berkeley, Calif., (1951). University of California Press
- Lenzen, C., Wattenhofer, R.: Tight bounds for parallel randomized load balancing. *Distrib. Comput.* **29**(2), 127–142 (2016)
- Lu, Y., Xie, Q., Kliot, G., Geller, A., Larus, J.R., Greenberg, A.: Join-idle-queue: a novel load balancing algorithm for dynamically scalable web services. *Perform. Eval.* **68**(11), 1056–1071 (2011)
- Lu, Y., Xie, Q., Kliot, G., Geller, A., Larus, J.R., Greenberg, A.: Join-Idle-Queue: a novel load balancing algorithm for dynamically scalable web services., (2011). <https://www.microsoft.com/en-us/research/publication/join-idle-queue-a-novel-load-balancing-algorithm-for-dynamically-scalable-web-services/>

36. Luczak, M.J., McDiarmid, Colin, et al.: On the maximum queue length in the supermarket model. *Annals Probab.* **34**(2), 493–527 (2006)
37. Mars, J., Tang, L., Hundt, R.: Heterogeneity in “homogeneous” warehouse-scale computers: A performance opportunity. *IEEE Comput. Archit. Lett.* **10**(2), 29–32 (2011)
38. Tony Mauro of F5. Choosing an NGINX Plus Load-Balancing Technique., October 29, (2015). <https://www.nginx.com/blog/choosing-nginx-plus-load-balancing-techniques/#:~:text=With>
39. McMullen, T.: Load balancing is impossible, (2016). <https://www.youtube.com/watch?v=kpvbOzHUakA>
40. Mitzenmacher, M.: How useful is old information? *IEEE Trans. Parallel Distrib. Syst.* **11**(1), 6–20 (2000)
41. Mitzenmacher, M.: The power of two choices in randomized load balancing. *IEEE Trans. Parallel Distrib. Syst.* **12**(10), 1094–1104 (2001)
42. Mitzenmacher, M.: Analyzing distributed join-idle-queue: a fluid limit approach. In: 2016 54th Annual Allerton Conference on Communication, Control, and Computing, 312–318. IEEE, (2016)
43. Mitzenmacher, M., Prabhakar, B., Shah, D.: Load balancing with memory. In: 43rd Annual IEEE Symposium on Foundations of Computer Science., 799–808. IEEE, (2002)
44. Neely, M.J.: Optimal energy and delay tradeoffs for multiuser wireless downlinks. *IEEE Trans. Inf. Theory* **53**(9), 3095–3113 (2007)
45. Neely, M.J.: Stability and capacity regions or discrete time queueing networks. arXiv preprint [arXiv:1003.3396](https://arxiv.org/abs/1003.3396), (2010)
46. Nishtala, R., Carpenter, P., Petrucci, V., Martorell, X.: Hipster: Hybrid task manager for latency-critical cloud workloads. In: IEEE International Symposium on High Performance Computer Architecture (HPCA), 409–420, (2017)
47. Nocedal, J., Wright, S.: *Numerical Optimization*. Springer Science & Business Media, Berlin/Heidelberg (2006)
48. Prekas, G., Kogias, M., Bugnion, E.: Zygus: Achieving low tail latency for microsecond-scale networked tasks. In: 26th Symposium on Operating Systems Principles (SOSP), pages 325–341, (2017)
49. Schurman, E., Brutlag, J.: The user and business impact of server delays, additional bytes, and http chunking in web search. In: Velocity Web Performance and Operations Conference. O’Reilly, (2009)
50. Selen, J., Adan, I., Kapodistria, S., van Leeuwen, J.: Steady-state analysis of shortest expected delay routing. *Queueing Syst.* **84**(3), 309–354 (2016)
51. Smith, M.: Netflix Technology Blog. Rethinking Netflix’s Edge Load Balancing. September 2018., (2018). <https://netflixtechblog.com/netflix-edge-load-balancing-695308b5548c>
52. Stolyar, A.L.: Pull-based load distribution in large-scale heterogeneous service systems. *Queueing Syst.* **80**(4), 341–361 (2015)
53. Stolyar, A.L.: Pull-based load distribution among heterogeneous parallel servers: the case of multiple routers. *Queueing Syst.* **85**(1–2), 31–65 (2017)
54. Tarreau, W.: HAProxy. Test Driving “Power of Two Random Choices” Load Balancing., 2019. <https://www.haproxy.com/blog/power-of-two-load-balancing/>, published on February 15, (2019)
55. Tarreau, W.: The Reliable, High Performance TCP/HTTP Load Balancer., (2021). <http://www.haproxy.org/>
56. van der Boor, M., Borst, S., van Leeuwen, J.: Load balancing in large-scale systems with multiple dispatchers. In: IEEE INFOCOM, (2017)
57. van der Boor, M., Borst, S., van Leeuwen, J.: Hyper-scalable jsq with sparse feedback. *Proc. ACM Meas. Anal. Comput. Syst.* **3**(1), 1–37 (2019)
58. Vargatik, S., Keslassy, I., Orda, A.: LSQ: load balancing in large-scale heterogeneous systems with multiple dispatchers. *IEEE/ACM Transactions on Networking*, (2020)
59. Vvedenskaya, N.D., Dobrushin, R.L., Karpelevich, F.I.: Queueing system with selection of the shortest of two queues: an asymptotic approach. *Problemy Peredachi Informatsii*, **32**(1):20–34, (1996)
60. Wang, C., Feng, C., Cheng, J.: Distributed join-the-idle-queue for low latency cloud services. *IEEE/ACM Trans. Netw.* **26**(5), 2309–2319 (2018)
61. Weber, R.R.: On the optimal assignment of customers to parallel servers. *J. Appl. Probab.* **15**(2), 406–413 (1978)
62. Wikipedia contributors. Majorization — Wikipedia, the free encyclopedia. <https://en.wikipedia.org/w/index.php?title=Majorization&oldid=993822990>, (2020). [Online; accessed 29-January-2021]
63. Winston, W.: Optimality of the shortest line discipline. *J. Appl. Probab.* **14**(1), 181–189 (1977)
64. Zhou, X., Shroff, N., Wierman, A.: Asymptotically optimal load balancing in large-scale heterogeneous systems with multiple dispatchers. *Perform. Eval.* **145**, 102146 (2021)
65. Zhou, X., Fei, W., Tan, J., Sun, Y., Shroff, N.: Designing low-complexity heavy-traffic delay-optimal load balancing schemes: Theory to algorithms. *ACM POMACS* **1**(2), 39 (2017)

Publisher’s Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.