Scalable Load Balancing in Networked Systems: A Survey of Recent Advances*

Mark Van der Boor[†] Sem C. Borst[†] Johan S. H. Van Leeuwaarden[‡] Debankur Mukherjee[§]

Abstract. In this survey we provide an overview of recent advances on scalable load balancing schemes which provide favorable delay performance and yet require minimal implementation overhead. The basic load balancing scenario involves a single dispatcher where tasks arrive that must immediately be forwarded to one of N single-server queues. The join-the-shortest-queue (JSQ) policy yields vanishing delays as N grows large, as in a centralized queuing arrangement, but involves a prohibitive communication burden. In contrast, JSQ(d) schemes that assign an incoming task to a server with the shortest queue among d servers selected uniformly at random require little communication, but lead to constant delays. In order to examine this fundamental trade-off between delay performance and implementation overhead, we discuss a body of recent research on JSQ(d(N)) schemes in which the diversity parameter d(N) depends on N and investigate the growth rate of d(N) required to match the optimal JSQ performance on fluid and diffusion scales.

Stochastic coupling techniques and scaling limits play an instrumental role in establishing this asymptotic optimality. We demonstrate how this methodology carries over to infinite-server settings, finite buffers, multiple dispatchers, servers arranged on graph topologies, and token-based load balancing schemes such as join-the-idle-queue (JIQ), thus providing a broad overview of the main trends in the field.

Key words. load balancing, stochastic coupling, limit theorems, Markov process, queuing theory

AMS subject classifications. 68M20, 60K25, 60K30, 68M14, 90B15, 90B22

DOI. 10.1137/20M1323746

Contents

I Introduction 555

2 Scalability Spectrum

559

^{*}Received by the editors March 9, 2020; accepted for publication (in revised form) October 22, 2021; published electronically August 4, 2022. This survey extends the short review presented at ICM 2018 [151], and section 6 provides a synopsis of a SIGMETRICS 2018 conference paper published in the Proceedings of the ACM on Measurement and Analysis of Computing Systems [116].

https://doi.org/10.1137/20M1323746

Funding: The work was partially supported by the National Science Foundation under grant CCF-2113027 and by The Netherlands Organization for Scientific Research (NWO) (Gravitation Grant NETWORKS 024.002.003).

[†]Department of Mathematics and Computer Science, Eindhoven University of Technology, Eindhoven, The Netherlands (m.v.d.boor@tue.nl, s.c.borst@tue.nl).

[‡]Department of Mathematics and Computer Science, Eindhoven University of Technology, Eindhoven, The Netherlands, and Tilburg University, Tilburg 5037 AB, The Netherlands (J.S.H. vanLeeuwaarden@tilburguniversity.edu).

[§]School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, GA 30318 USA (debankur.mukherjee@isye.gatech.edu).

		SCALABLE LOAD BALANCING IN NETWORKED SYSTEMS	555				
	2.1 2.2 2.3 2.4	Basic Model	559 561				
3	Preliminaries, JSQ Policy, and JSQ(d) Policies						
	3.1	Definitions, Limit Sequences, and Convergence Issues					
	3.2	Fluid Limit for JSQ (d) Policies					
	3.3 3.4	Fluid Limit for JSQ Policy					
	3.4	JSQ (d) Policies in Heavy-Traffic Regime					
4	Uni	versality of JSQ (d) Policies	57 3				
-	4.1	High-Level Outline of Proof Approach					
	4.2	Extension to Batch Arrivals					
5	Blocking and Infinite-Server Dynamics						
	5.1	Fluid Limit for JSQ Policy					
	5.2	Diffusion Limit for JSQ Policy					
	5.3	Universality of JSQ (d) Policies in Infinite-Server Dynamics	581				
6		Load Balancing in Graph Topologies					
	6.1	Asymptotic Optimality Criteria for Deterministic Graph Sequences					
	6.2	Asymptotic Optimality of Random Graph Sequences	593				
7		en-Based Load Balancing	594				
	7.1	Fluid-Level Optimality of JIQ Scheme					
	7.2 7.3	Diffusion-Level Optimality of JIQ Scheme					
	7.4	Joint Load Balancing and Autoscaling					
8	Red	undancy Policies and Alternative Scaling	605				
_	8.1	Redundancy-d Policies					
	8.2	Conventional Heavy Traffic					
	8.3	Nondegenerate Slowdown					
	8.4	Sparse-Feedback Regime					
	8.5	Scaling of Maximum Queue Length	609				
9		Extensions and Future Research Directions					
	9.1	JSQ(d) Policies with General Service Requirement Distributions					
	$9.2 \\ 9.3$	Heterogeneous Servers and Knowledge of Service Requirements Open Problems and Emerging Research Directions					
R,	efere		613				
	בוכו כו	ICC3	013				

I. Introduction. In this survey we review scalable load balancing algorithms (LBAs) which achieve excellent delay performance in large-scale systems and yet have a low implementation overhead. LBAs play a critical role in distributing service requests or tasks (e.g., computing jobs, database look-ups, file transfers) among servers

or distributed resources in parallel processing systems. The analysis and design of LBAs has attracted significant attention in recent years, mainly spurred on by crucial scalability challenges arising in cloud networks and data centers with massive numbers of servers. LBAs can be broadly categorized as static, dynamic, or some intermediate blend, depending on the amount of feedback or state information (e.g., congestion levels) that is used in allocating tasks. The use of state information naturally allows dynamic policies to achieve better delay performance, but also involves higher implementation complexity and a substantial communication burden. The latter issue is particularly pertinent in cloud networks and data centers with immense numbers of servers handling a huge influx of service requests. In order to understand the large-system characteristics, we examine scalability properties through the prism of asymptotic scalings where the system size grows large and identify LBAs that strike a balance between delay performance and implementation overhead.

The most basic load balancing scenario consists of N identical parallel servers and a dispatcher where tasks arrive sequentially. Arriving tasks must immediately be forwarded to one of the servers. Tasks are assumed to have unit-mean exponentially distributed service requirements, and the service discipline at each server is supposed to be oblivious to the actual service requirements. These assumptions, in conjunction with a Poisson arrival process, permit a Markovian state description for the evolution of the queue length process. Moreover, the symmetry arising from the homogeneity of tasks and exchangeability of the servers provides a particularly convenient basis for stochastic coupling arguments and scaling limits. In the early parts of this survey we will focus on this basic setup, which has been prevalent in the literature, but in later sections of the paper we will also discuss graph-based versions where the servers are no longer statistically identical. In addition, we will touch on scenarios with heterogeneous tasks, extensions to general service requirement distributions, and situations where advance knowledge of the service requirements is available.

In the above-described basic setup, the celebrated join-the-shortest-queue (JSQ) policy has several important stochastic optimality properties. In particular, the JSQ policy achieves the minimum mean overall delay among all nonanticipating policies that do not have any advance knowledge of the service requirements [34, 172]. In order to implement the JSQ policy, however, a dispatcher requires instantaneous knowledge of all the queue lengths, which may involve a prohibitive communication burden for a large number of servers N. This poor scalability has motivated the consideration of JSQ(d) policies, where an incoming task is assigned to a server with the shortest queue among $d \geq 2$ servers selected uniformly at random. Note that this involves an exchange of 2d messages per task, irrespective of the number of servers N. Seminal results in [114, 163] imply that even sampling as few as d=2 servers yields significant performance enhancements over purely random assignment (d=1) as N grows large, which is commonly referred to as the power-of-two or power-of-choice effect. In particular, when tasks arrive at rate λN , the queue length distribution at each individual server exhibits superexponential decay for any fixed $\lambda < 1$ as N grows large, a considerable improvement compared to exponential decay for purely random assignment.

The diversity parameter d thus induces a fundamental trade-off between the amount of communication overhead and the delay performance. Specifically, a random assignment policy does not entail any communication burden, but the mean waiting time remains *constant* as N grows large for any fixed $\lambda > 0$. In contrast, a nominal implementation of the JSQ policy (without maintaining state information at

the dispatcher) involves 2N messages per task, but the mean waiting time vanishes as N grows large for any fixed $\lambda < 1$. Although JSQ(d) policies with $d \geq 2$ yield major performance improvements over purely random assignment while reducing the communication burden by a factor O(N) compared to the JSQ policy, the mean waiting time $does\ not\ vanish$ in the limit. Hence, no fixed value of d will provide asymptotically optimal delay performance. This is evidenced by powerful results $[50,\ 51,\ 52]$ indicating that in the absence of any memory at the dispatcher, the communication overhead per task $must\ increase$ with N in order for any scheme to achieve a zero mean waiting time in the limit.

In the context of JSQ(d) policies, scalability specifically pertains to the intrinsic trade-off between delay performance and communication overhead as governed by the diversity parameter d, in conjunction with the relative load λ . In this survey we review scaling results which offer detailed insight into the latter trade-off in a regime where not only is the overall task arrival rate assumed to grow with N, but also the diversity parameter is allowed to depend on N. We write $\lambda(N)$ and d(N) to explicitly reflect that, and we provide a sketch of the analysis in [119] which identifies the growth rate of d(N) required in order to achieve a zero mean waiting time in the limit, depending on the scaling of $\lambda(N)$. This involves both fluid-scaled and diffusion-scaled versions of the queue length process in regimes where $\lambda(N)/N \to \lambda < 1$ and $(N - \lambda(N))/\sqrt{N} \to \infty$ $\beta > 0$ as $N \to \infty$, respectively; see section 3.1 for definitions of these objects. As will be discussed in detail there, the limiting processes are insensitive to the exact growth rate of d(N), as long as the latter is sufficiently fast, and in particular they coincide with the limiting processes for the JSQ policy. This demonstrates that the optimality of the JSQ policy can be preserved asymptotically while dramatically lowering the communication overhead.

As mentioned above, we will also consider network scenarios where the N servers are assumed to be interconnected by some underlying graph topology G_N . Tasks arrive at the various servers as independent Poisson processes of rate λ , and each incoming task is assigned to whichever server has the shortest queue among the one where it appears and its neighbors in G_N . Such network scenarios are not only of theoretical interest, but also of major practical relevance since they emerge in the modeling of so-called affinity relations and compatibility constraints between tasks and servers. Such features are increasingly common in data centers and cloud networks due to heterogeneity and data locality issues, and they also relate to scalability considerations; see section 6 for a further discussion and specific literature pointers. In the case that G_N is a clique (fully connected graph), each incoming task is assigned to the server with the shortest queue across the entire system, and the behavior is equivalent to that under the JSQ policy. The stochastic optimality properties of the JSQ policy thus imply that the queue length process in a clique will be "better" than in an arbitrary graph G_N . We will present sufficient conditions formulated in [116] for the fluid-scaled and diffusion-scaled versions of the queue length process in an arbitrary graph to be equivalent to the limiting processes in a clique as $N \to \infty$. The conditions demonstrate that the optimality of a clique can be asymptotically preserved while dramatically reducing the number of connections, provided the graph G_N is "suitably random"; see section 6 for a more formal statement.

While a zero waiting time can be achieved in the limit by sampling only $d(N) \ll N$ servers, the amount of communication overhead in terms of d(N) must still grow with N. This can be explained using the fact that a large number of servers need to be sampled for each incoming task to ensure that at least one of them is found

idle with high probability. This can be avoided by introducing memory at the dispatcher, in particular maintaining a record of vacant servers, and assigning tasks to idle servers, if there are any. This so-called join-the-idle-queue (JIQ) scheme [11, 101] has gained huge popularity recently and can be implemented through a simple tokenbased mechanism generating at most one message per task. The JIQ scheme is thus quite appealing from a scalability perspective, which raises the question of what the corresponding delay performance is in large-scale systems. We will therefore also review results implying not only that the fluid-scaled queue length process under the JIQ scheme asymptotically coincides with that under the JSQ policy as shown in [141], but also that this equivalence property extends to the diffusion-scaled queue length process as established in [118]. Thus, the use of memory allows the JIQ scheme to achieve asymptotically optimal delay performance with minimal communication overhead (at least in the idealized setting with statistically identical servers and homogeneous tasks). In particular, ensuring that tasks are assigned to idle servers whenever available is sufficient to achieve asymptotic optimality, and using any additional queue length information yields no meaningful performance benefits on the fluid or diffusion levels. It is worth pointing out though that the JIQ scheme is not optimal in certain asymptotic regimes such as the nondegenerate slow-down regime; see section 2.2 for a formal definition. In [73] it was shown that a minor modification of the JIQ scheme, called idle-one-first, which besides idle servers also keeps track of queues of length one, is asymptotically optimal; see section 8.3 for a detailed discussion.

On a methodological note, it is worth observing that a direct derivation of the fluid and diffusion limits in the above scenarios is quite challenging. Instead, the asymptotic equivalence results in [116, 118, 119] are derived by relating the relevant system occupancy processes to the corresponding processes under a JSQ policy, and showing that the deviation among these processes is asymptotically negligible on either fluid scale or diffusion scale under suitable assumptions on d(N) or G_N . The known fluid and diffusion limits for the JSQ policy thus yield the corresponding limit process for the JSQ(d(N)) policy, a load balancing graph G_N , and the JIQ scheme as by-products.

In this survey we highlight the stochastic coupling techniques that played an instrumental role in proving the asymptotic equivalence results in [116, 118, 119]. Although the stochastic coupling concepts provide an effective and overarching approach, they defy a systematic recipe and involve some degree of ingenuity and customization. Indeed, the specific coupling arguments that were developed in [116, 118, 119] are different from those that were originally used in establishing the stochastic optimality properties of the JSQ policy. Moreover, the specific coupling approaches differ in sometimes subtle but critical ways between a JSQ(d(N)) policy, a load balancing graph G_N , and the JIQ scheme, which all require the arguments to be suitably tailored. We also review further stochastic coupling constructions that were devised in [120] for scenarios with infinite-server dynamics.

While the results for load balancing graphs illustrate that the stochastic coupling techniques can be applied in "asymmetric" situations, it is fair to say that this approach is at its strongest in scenarios where all the servers are exchangeable and the evolution of the system occupancy can be represented in terms of a density-dependent Markov process. In these cases, the approach is particularly powerful in analyzing the system occupancy process on fluid or diffusion scale, where for many policies the behavior can be shown to asymptotically coincide with that of JSQ, for which fairly explicit characterizations are known.

Stochastic coupling does not seem to provide a directly useful approach for other functionals of the system occupancy process, such as the maximum queue length, where asymptotic equivalence with JSQ on fluid or diffusion scale does not provide any information, and in fact even asymptotically the behavior for many schemes is different. Applying stochastic coupling techniques in highly heterogeneous settings is also difficult since the lack of symmetry tends to break its underpinnings, and establishing scaling results for such scenarios remains as a particularly challenging subject for further research, as is further discussed in section 9.

A final caveat is in order. Load balancing is a broad subject which has been actively pursued for decades and has been investigated from a variety of perspectives in several communities (algorithm design, applied probability, complexity theory, performance evaluation). While this survey aims to touch on many of these aspects, reflecting historical developments and connecting various threads, it is impossible to exhaustively cover the load balancing literature in full detail. Rather than provide an encyclopedic overview, we therefore focus on scalability in terms of delay performance and implementation overhead in large-scale systems as the overarching theme, and highlight the combined power of stochastic coupling methods and scaling limits.

The survey is organized as follows. In section 2 we discuss various LBAs and evaluate their scalability properties. In section 3 we introduce some useful preliminary concepts, and then review fluid and diffusion limits for the JSQ policy as well as JSQ(d) policies with a fixed value of d. In section 4 we discuss the trade-off between delay performance and communication overhead as a function of the diversity parameter d, in conjunction with the relative load. In particular, we formulate asymptotic universality properties for JSQ(d) policies, which are extended to systems with server pools and network scenarios in sections 5 and 6, respectively. Section 7 is devoted to asymptotic optimality properties for the JIQ scheme. We discuss somewhat related redundancy policies and alternative scaling regimes and performance metrics in section 8. The survey is concluded in section 9 with a discussion of yet further extensions and several open problems and emerging research directions.

- **2. Scalability Spectrum.** In this section we review a wide spectrum of LBAs and examine scalability properties in terms of their delay performance vis-à-vis their associated implementation overhead in large-scale systems.
- **2.1. Basic Model.** Throughout this section and most of the paper, we focus on a basic scenario with N parallel single-server infinite-buffer queues and a single dispatcher where tasks arrive as a Poisson process with rate $\lambda(N)$, as depicted in Figure 1. Arriving tasks cannot be queued at the dispatcher and must immediately be forwarded to one of the servers. This canonical setup is commonly dubbed the *supermarket model*, in loose analogy with the everyday situation of choosing between parallel checkout lanes in supermarkets. Tasks are assumed to have unit-mean exponentially distributed service requirements, and the service discipline at each server is supposed to be oblivious to the actual service requirements.

When tasks do not get served and never depart, but simply accumulate, the above setup corresponds to a so-called balls-and-bins model, and we will further elaborate on the connections and differences with work in that domain in section 8.5.

2.2. Asymptotic Scaling Regimes. An exact analysis of the delay performance is quite involved, if not intractable, for all but the simplest LBAs. Numerical evaluation or simulation is not straightforward either, especially for high load levels and large system sizes. A common approach is therefore to consider various limit regimes, which

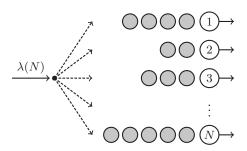


Fig. 1 Tasks arrive at the dispatcher as a Poisson process of rate $\lambda(N)$ and are forwarded to one of the N servers according to some specific LBA.

not only provide mathematical tractability and illuminate the fundamental properties, but are also natural in view of the typical conditions in which cloud networks and data centers operate. One can distinguish several asymptotic scalings that have been used for these purposes:

- (i) In the classical heavy-traffic regime, $\lambda(N) = \lambda N$ with a fixed number of servers N and a relative load λ that tends to one in the limit.
- (ii) In the conventional large-capacity or many-server regime, the relative load $\lambda(N)/N$ approaches a constant $\lambda < 1$ as the number of servers N grows large.
- (iii) The popular Halfin–Whitt regime, named after the authors of the seminal paper [74] where it was introduced and first analyzed, combines heavy traffic with a large capacity with

$$\frac{N - \lambda(N)}{\sqrt{N}} \to \beta > 0 \text{ as } N \to \infty,$$

so the relative capacity slack behaves as β/\sqrt{N} as the number of servers N grows large.

(iv) The so-called nondegenerate slow-down regime [8, 73] involves $N - \lambda(N) \to \gamma > 0$, so the relative capacity slack shrinks as γ/N as the number of servers N grows large.

The term nondegenerate slow-down refers to the fact that in the context of a centralized multiserver queue (where load balancing between servers occurs implicitly), the mean waiting time in regime (iv) tends to a strictly positive constant as $N \to \infty$, and is thus of similar magnitude to the mean service requirement. In contrast, in regimes (ii) and (iii), the mean waiting time in a multiserver queue decays exponentially fast in N or is of the order $1/\sqrt{N}$, respectively, as $N \to \infty$, while in regime (i) the mean waiting time grows arbitrarily large relative to the mean service requirement.

In the context of a centralized M/M/N queue, scalings (ii), (iii), and (iv) are commonly referred to as quality-driven (QD), quality-and-efficiency-driven (QED), and efficiency-driven (ED) regimes. These terms reflect that (ii) offers excellent service quality (vanishing waiting time), (iv) provides high resource efficiency (utilization approaching one), and (iii) achieves a combination of these two, providing the best of both worlds.

In the remainder of the paper we will focus on scalings (ii) and (iii) and refer to them as fluid and diffusion scalings, since it is natural to analyze the relevant system occupancy processes on fluid scale (1/N) and diffusion scale $(1/\sqrt{N})$ in these regimes,

respectively. In line with the central theme of this survey, we will not provide a detailed account of scalings (i) and (iv), which do not capture the large-scale perspective and do not allow for low delays, respectively. However, we will briefly mention some results for these regimes in sections 8.2 and 8.3.

2.3. Basic Load Balancing Algorithms.

2.3.1. Random Assignment: N Independent M/M/I Queues. One of the most basic LBAs is to assign each arriving task to a server selected uniformly at random. In that case, the various queues collectively behave as N independent M/M/I queues, each with arrival rate $\lambda(N)/N$ and unit service rate. In particular, at each of the queues, the total number of tasks in stationarity has a geometric distribution with parameter $\lambda(N)/N$. By virtue of the Poisson Arrivals See Time Averages (PASTA) property, the probability that an arriving task incurs a nonzero waiting time is $\lambda(N)/N$. The mean number of waiting tasks (excluding the possible task in service) at each of the queues is $\frac{\lambda(N)^2}{N(N-\lambda(N))}$, so the total mean number of waiting tasks is $\frac{\lambda(N)}{N-\lambda(N)}$, which by Little's law implies that the mean waiting time is $\frac{\lambda(N)}{N-\lambda(N)}$. In particular, when $\lambda(N) = N\lambda$, the probability that a task incurs a nonzero waiting time is λ and the mean waiting time of a task is $\frac{\lambda}{1-\lambda}$, independent of N, reflecting the independence of the various queues.

As we will see later, a broad range of queue-aware LBAs can deliver a probability of a nonzero waiting time and a mean waiting time that vanish asymptotically. While a random assignment policy is evidently not competitive with such queue-aware LBAs, it still plays a relevant role due to its strong degree of mathematical tractability. For example, the queue process under purely random assignment can be shown to provide an upper bound (in a stochastic majorization sense) for various more involved queue-aware LBAs for which even stability may be difficult to establish directly, yielding conservative performance bounds and stability guarantees.

A slightly better LBA is to assign tasks to the servers in a round-robin manner, dispatching every Nth task to the same server. In the fluid regime (ii), the interarrival time of tasks at each given queue will then converge to a constant $1/\lambda$ as $N \to \infty$. Thus each of the queues will behave as a D/M/1 queue in the limit, and the probability of a nonzero waiting time and the mean waiting time will be somewhat lower than under purely random assignment. However, both the probability of a nonzero waiting time and the mean waiting time will still tend to strictly positive values and not vanish as $N \to \infty$.

- **2.3.2.** Join-the-Shortest Queue (JSQ). Under the JSQ policy, each arriving task is assigned to the server with the currently shortest queue (ties are broken arbitrarily). In the basic model described above, the JSQ policy has several stochastic optimality properties and yields the "most balanced and smallest" queue process among all nonanticipating policies that do not have any advance knowledge of the service requirements [34, 172].
- **2.3.3.** Join-the-Smallest-Workload (JSW): Centralized M/M/N Queue. Under the JSW policy, each arriving task is assigned to the server with the currently smallest workload. Note that this is an anticipating policy, since it requires advance knowledge of the service requirements of all the tasks in the system. Further observe that this policy (myopically) minimizes the waiting time for each incoming task and mimics the operation of a centralized N-server queue with an FCFS (first come, first served) discipline. The equivalence with a centralized N-server queue with an FCFS disci-

pline yields an additional optimality property of the JSW policy: The vector of joint workloads at the various servers observed by each incoming task is smaller in the Schur convex sense than under any alternative admissible policy [43].

It is worth observing that the above optimality properties do not in fact rely on Poisson arrival processes or exponential service requirement distributions. At the same time, these optimality properties do not imply that the JSW policy minimizes the mean stationary waiting time. In our setting with Poisson arrivals and exponential service requirements, however, it can be shown through direct means that the total number of tasks under the JSW policy is stochastically smaller than under the JSQ policy. Indeed, in view of the equivalence with a centralized M/M/N queue, the total service completion rate under the JSW policy is given by $\min\{L, N\}$ when there are L tasks in total in the system, while under the JSQ policy the total service completion rate is at most equal to $\min\{L, N\}$, and may be lower than that when some servers are idle while tasks are queued up at other servers. Even though the JSW policy requires similar excessive communication overhead, aside from its anticipating nature, the above-mentioned equivalence in fact means that the total number of tasks behaves as a birth-death process, which renders it far more tractable than the JSQ policy. Specifically, it follows from textbook results for the centralized M/M/N queue that, given that all the servers are busy, the total number of waiting tasks is geometrically distributed with parameter $\lambda(N)/N$. The total mean number of waiting tasks is then $\Pi_W(N,\lambda(N))\frac{\lambda(N)}{N-\lambda(N)}$, and the mean waiting time is $\Pi_W(N,\lambda(N))\frac{1}{N-\lambda(N)}$, with $\Pi_W(N,\lambda(N))$ denoting the probability of the total occupancy in an M/M/N queue being N or larger, i.e., the probability of all servers being occupied and a task incurring a nonzero waiting time. The probability $\Pi_W(N,\lambda(N))$ can be obtained from the stationary distribution of the birth-death process representing the system occupancy and is described by the so-called Erlang-C formula as a function of the load and number of servers. The latter function can be expressed in semiexplicit wellapproximated "closed form" in terms of a normalizing constant which is the sum of an explicit infinite series. Standard results for the M/M/1 queue imply that the mean waiting time is $\frac{\lambda(N)}{N-\lambda(N)}$ for the random assignment policy considered in section 2.3.1. Thus it can immediately be concluded that the mean waiting time under the JSW policy is smaller by at least a factor of $\lambda(N)$.

In the fluid regime $\lambda(N) = N\lambda$, it can be shown that the probability $\Pi_W(N, \lambda(N))$ of a nonzero waiting time decays exponentially fast in N (see, for instance, [74]), and hence so does the mean waiting time. The pivotal results in [74] further demonstrate that in the diffusion regime (2.1), the probability $\Pi_W(N, \lambda(N))$ of a nonzero waiting time converges to a finite constant $\Pi_W^*(\beta)$. This implies that the mean waiting time of is of the order $1/\sqrt{N}$ and hence vanishes as $N \to \infty$.

2.3.4. Power-of-*d* **Load Balancing (JSQ**(d)). We have seen that the Achilles' heel of the JSQ policy is its excessive communication overhead in large-scale systems. This poor scalability has motivated consideration of so-called JSQ(d) policies, where an incoming task is assigned to a server with the shortest queue among d servers selected uniformly at random. The seminal results in [114, 163] demonstrate that in the fluid regime (ii), the stationary probability that there are i or more tasks at a given queue is proportional to $\lambda^{(d^i-1)/(d-1)}$ as $N \to \infty$, and thus it exhibits super-exponential decay as opposed to exponential decay for the random assignment policy considered in section 2.3.1.

As alluded to in section 1, the diversity parameter d thus induces a fundamental

trade-off between the amount of communication overhead and the performance in terms of queue lengths and delays. A rudimentary implementation of the JSQ policy (d=N, without replacement) involves O(N) communication overhead per task, but it can be shown that the probability of a nonzero waiting time and the mean waiting vanish as $N\to\infty$ in both the fluid and the diffusion regimes; see sections 3.3 and 3.4. Although JSQ(d) policies with a fixed parameter $d\geq 2$ yield major performance improvements over purely random assignment, as implied by the results in [114, 163], these results at the same time show that even in the fluid regime, the probability of a nonzero waiting time and the mean waiting time do not vanish as $N\to\infty$.

2.3.5. Token-Based Mechanisms: Join-the-Idle-Queue (JIQ). While a zero waiting time can be achieved in the limit by sampling only $d(N) \ll N$ servers, the amount of communication overhead in terms of d(N) must still grow with N. This can be countered by introducing memory at the dispatcher, in particular, maintaining a record of vacant servers, and assigning tasks to idle servers as long as there are any, or to a uniformly at random selected server otherwise. This so-called join-the-idle-queue (JIQ) scheme [11, 101] has received keen interest recently and can be implemented through a simple token-based mechanism. Specifically, idle servers send tokens to the dispatcher to advertise their availability, and when a task arrives and the dispatcher has tokens available, it assigns the task to one of the corresponding servers (and disposes of the token). Note that a server only issues a token when a task completion leaves its queue empty, thus generating at most one message per task. Surprisingly, the mean waiting time and the probability of a nonzero waiting time vanish under the JIQ scheme in both the fluid and the diffusion regimes, as we will further discuss in section 7. Thus, the use of memory allows the JIQ scheme to achieve asymptotically optimal delay performance with minimal communication overhead.

2.4. Performance Comparison. We now present some simulation results to compare the above-described LBAs in terms of delay performance. Specifically, we evaluate the mean waiting time and the probability of a nonzero waiting time in both a fluid regime $(\lambda(N)=0.9N)$ and a diffusion regime $(\lambda(N)=N-\sqrt{N})$. The simulations are conducted for $N=10,20,\ldots,200$ servers and run for 10000 time units. Each simulation starts with an empty system, but only jobs that leave after 2500 time units are counted in order to avoid transient effects. The probability of waiting and mean waiting time are computed using the empirical averages over all jobs that leave after 2500 time units. This procedure is repeated 20 times, and the results in Figure 2 show the mean waiting time and probability of waiting averaged across these 20 runs. An overview of the asymptotic delay performance and overhead associated with various LBAs is provided in Table 1.

We are specifically interested in distinguishing two classes of LBAs—those delivering a mean waiting time and probability of a nonzero waiting time that vanish asymptotically, and those that fail to do so—and relating that dichotomy to the associated communication overhead and memory requirement at the dispatcher. We give these classifications for both the fluid and the diffusion regimes.

JSQ, JIQ, and JSW. As mentioned earlier, JSQ, JIQ, and JSW have vanishing mean waiting times in both the fluid and the diffusion regimes, and this is supported by the figures, which further reflect the optimality of JSW in terms of mean waiting time. We can also observe a crucial difference, however, between JSW and JSQ/JIQ. Somewhat surprisingly, the probability of a nonzero waiting time does not vanish for JSW in the diffusion regime, while it does vanish for JSQ/JIQ. Since the mean

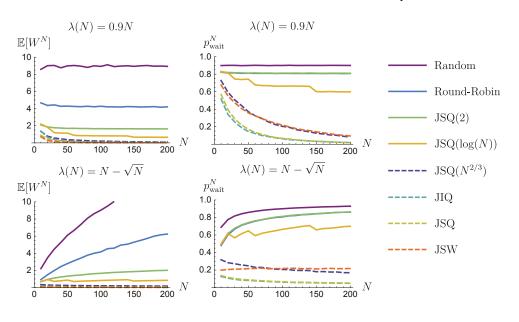


Fig. 2 Simulation results for mean waiting time $\mathbb{E}[W^N]$ and probability of a nonzero waiting time p_{wait}^N for both a fluid and a diffusion regime.

waiting time for JSW is smaller than for JSQ/JIQ, this implies that the mean of all nonzero waiting times (i.e., the mean waiting time conditional on having to wait) is an order of magnitude larger in JSQ/JIQ compared to JSW. This difference can be explained by the fact that JSW uses knowledge of the service requirements, whereas JSQ/JIQ do not. Indeed, when a task is placed in a queue under JSQ/JIQ, it will need to wait for a "normal" residual service time, whereas JSW exploits knowledge of that residual service time being relatively short among all N queues. Alternatively, taking the equivalent view of JSW as a centralized M/M/N queue, a task that needs to wait may find several tasks ahead of it in the queue, but this queue is served by N servers combined, whereas in JSQ/JIQ each queue is handled by just a single server. Conversely, when there are N or more tasks in the system in total, an arriving task will need to wait under JSW, while in JSQ/JIQ some of the servers may have several tasks in the queue, and the arriving task may still find an idle server with high probability. We will revisit the comparison between JSQ and a centralized M/M/N queue in section 3.4.

Random and Round-Robin. The mean waiting time does not vanish for random and round-robin in the fluid regime, as already mentioned in section 2.3.1. Moreover, the mean waiting time grows without bound in the diffusion regime for these two schemes. This is because the system can still be decomposed into single-server queues, and the loads of the individual M/M/1 and D/M/1 queues tend to 1.

JSQ(d) Policies. Three versions of JSQ(d) are included in Figure 2; d(N) = 2, $d(N) = \lfloor \log(N) \rfloor \to \infty$, and $d(N) = N^{2/3}$ for which $\frac{d(N)}{\sqrt{N} \log(N)} \to \infty$. Note that the graph for JSQ(log(N)), where the diversity parameter grows logarithmically in N, shows kneepoints due to the slow growth rate of log(N) and the fact that the actual integer value $d(N) = |\log(N)|$ occasionally jumps by 1. As can be seen in Figure 2,

Scheme	Queue length	Waiting time (fixed $\lambda < 1$)	Waiting time $(1 - \lambda \sim 1/\sqrt{N})$	Overhead per task
Random	$q_i^{\star} = \lambda^i$	$\frac{\lambda}{1-\lambda}$	$\Theta(\sqrt{N})$	0
JSQ(d)	$q_i^{\star} = \lambda^{(d^i - 1)/(d - 1)}$	$\Theta(1)$	$\Omega(\log N)$	2d
$d(N) \to \infty$	same as JSQ	same as JSQ	??	2d(N)
$\frac{d(N)}{\sqrt{N}\log(N)} \to \infty$	same as JSQ	same as JSQ	same as JSQ	2d(N)
JSQ	$q_1^{\star} = \lambda, q_2^{\star} = \mathrm{o}(1)$	o(1)	$\Theta(1/\sqrt{N})$	2N
JIQ	same as JSQ	same as JSQ	same as JSQ	≤ 1

Table I Queue length distribution, waiting times, and communication overhead for various LBAs.

the choices for which $d(N) \to \infty$ have vanishing wait in the fluid regime, while d=2 has not. Overall, we see that JSQ(d) policies clearly outperform random and roundrobin dispatching, while JSQ, JIQ, and JSW are better in terms of mean wait.

3. Preliminaries, JSQ Policy, and JSQ(d**) Policies.** In this section we first introduce some notation and preliminary concepts, and then we review fluid and diffusion limits for the JSQ policy as well as JSQ(d) policies with a fixed value of d.

3.1. Definitions, Limit Sequences, and Convergence Issues. We continue to focus on the basic scenario where all the servers are homogeneous, the service requirements are exponentially distributed, and the service discipline at each server is oblivious to the actual service requirements. Moreover, most of the LBAs under consideration do not distinguish between servers with equal queue lengths. Consequently, the queue length process is Markov on an enlarged filtration, allowing for random draws to resolve ties. In order to obtain a Markovian state description, it therefore suffices to only track the number of tasks, and in fact we do not need to keep a record of the number of tasks at each individual server, but only count the number of servers with a given number of tasks. Specifically, we represent the state of the system by a vector

$$\mathbf{Q}(t) := (Q_1(t), Q_2(t), \dots),$$

with $Q_i(t)$ denoting the number of servers with i or more tasks at time t, including the possible task in service, $i = 1, 2, \ldots$ Note that if we represent the queues at the various servers as (vertical) stacks, and arrange these from left to right in ascending order, then the value of Q_i corresponds to the width of the ith (horizontal) row, as depicted in the schematic diagram in Figure 3.

In order to examine the fluid and diffusion limits in regimes where the number of servers N grows large, we consider a sequence of systems indexed by N and attach a superscript N to the associated state variables. The fluid-scaled occupancy state is denoted by $\mathbf{q}^N(t) := (q_1^N(t), q_2^N(t), \ldots)$, with $q_i^N(t) = Q_i^N(t)/N$ representing the fraction of servers in the Nth system with i or more tasks as time $t, i = 1, 2, \ldots$ Let

$$S = \left\{ \mathbf{q} \in [0, 1]^{\infty} : q_i \le q_{i-1} \ \forall i = 2, 3, \dots, \text{ and } \sum_{i=1}^{\infty} q_i < \infty \right\}$$

be the set of all possible fluid-scaled states equipped with the ℓ_1 topology. Any (weak) limit $\mathbf{q}(\cdot)$ of the sequence of processes $\{\mathbf{q}^N(t)\}_{t\geq 0}$ in the conventional large-capacity

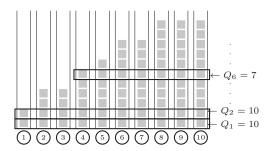


Fig. 3 The value of Q_i represents the width of the ith row, when the servers are arranged in nondescending order of their queue lengths.

regime (ii) as $N \to \infty$ (in a suitable topology on the space of functions on [0,T] taking values in \mathcal{S}) is called a *fluid limit*. In some frameworks in the literature this is also commonly referred to as a *mean-field limit* when the occupancy process is viewed as the (density-dependent) state evolution of a population of randomly interacting nodes or particles [15, 32, 91, 92]. Whenever we consider fluid limits, we assume the sequence of initial states is such that $\mathbf{q}^N(0) \to \mathbf{q}^\infty \in \mathcal{S}$ as $N \to \infty$.

The diffusion-scaled occupancy state is defined as $\bar{\mathbf{Q}}^N(t) = (\bar{Q}_1^N(t), \bar{Q}_2^N(t), \dots)$, with

$$(3.1) \qquad \quad \bar{Q}_1^N(t) = -\frac{N - Q_1^N(t)}{\sqrt{N}}, \qquad \bar{Q}_i^N(t) = \frac{Q_i^N(t)}{\sqrt{N}}, \quad i = 2, 3, \dots,$$

where we include a minus sign in the definition of $\bar{Q}_1^N(t)$ so as to adhere to the notation adopted in [36] which is the basis for the results that will be presented in section 3.4. Any (weak) limit $\mathbf{Q}(\cdot)$ of the sequence of processes $\{\mathbf{Q}^N(t)\}_{t\geq 0}$ in the Halfin–Whitt heavy-traffic regime (iii) as $N\to\infty$, once again in a suitable topology, is called a diffusion limit. Note that $-\bar{Q}_1^N(t)$ corresponds to the number of vacant servers, normalized by \sqrt{N} . The reason why $Q_1^N(t)$ is centered around N, while $Q_i^N(t)$, $i=2,3,\ldots$, are not, is that for the scalable LBAs that we consider the fraction of servers with exactly one task tends to one, whereas the fraction of servers with two or more tasks tends to zero as $N\to\infty$. For convenience, we will assume that each server has an infinite-capacity buffer, but all the results extend to the finite-buffer case; see, for instance, [36, 117, 118, 119, 120].

We conclude this subsection with a discussion of two important convergence issues associated with the above-defined scaling limits.

Accuracy of Asymptotic Approximations. A critical issue in the context of scaling limits is the rate of convergence and accuracy for finite-size systems. Some interesting results for the accuracy of mean-field approximations for interacting particle systems including load balancing models may be found in [62, 175, 176]. These results can be leveraged to develop refined approximations and improve the accuracy by adding expansion terms as demonstrated in [63, 64, 65].

Global Asymptotic Stability, Stationary Distributions, and Interchange of Limits. A further crucial issue in the context of scaling limits is whether limit processes that arise as $N \to \infty$ themselves have (unique) subsequential limits or limiting distributions as $t \to \infty$, and if so, how the stationary distributions of the prelimit processes (assuming they exist) relate to those limits. For fluid limits, which are usually

described in terms of a system of differential equations, the first question translates to the existence of a unique invariant point (fixed point) of these equations. While in most cases of practical interest such a unique invariant point tends to exist, this may be nontrivial to prove, and the existence of multiple invariant points cannot a priori be ruled out in general. In fact, existence of multiple invariant points has been shown in specific scenarios and is an indication of oscillatory behavior and so-called bistability issues in the original stochastic process for large N [66, 109]. Even when it can be established that a unique invariant point exists, the next question pertains to global attraction or global asymptotic stability. Specifically, the invariant point is said to be a global attractor, or globally asymptotically stable, if the fluid-limit process converges to this point for any initial condition. Global asymptotic stability has been established for various particular model instances, including the supermarket model with JSQ(d) load balancing strategies [114, 154, 163]. Common proof methodologies involve Lyapunov constructions [26, 47, 61], monotonicity properties [113, 145, 154, 163], and reversibility concepts [93], but there is no systematic recipe available, and the specific proof arguments tend to be highly tailored to the particular system under consideration. If global asymptotic stability of the invariant point can be established, then along with tightness this ensures that the sequence of stationary distributions of the prelimit process (assuming these exist) converge to this point; see, for instance, [16], with some of the key ideas and results dating back to much earlier work [79, 170]. This provides justification for an interchange of the large-scale $(N \to \infty)$ and stationary $(t \to \infty)$ limits, indicating that the invariant point provides a suitable approximation for the stationary distribution of the original stochastic process for sufficiently large values of N. In addition, the interchange of limits tends to furnish asymptotic independence among any finite subset of the queues [69]. Related results, convergence rates, and error probabilities are established in [104, 108]. Somewhat similar issues and observations apply for diffusion limits [53, 88].

3.2. Fluid Limit for JSQ(d) **Policies.** We first consider the fluid limit for JSQ(d) policies with an arbitrary but fixed value of d as characterized by the seminal results in [113, 163]. The result below is paraphrased from [113, 163].

Fluid Limit for JSQ(d). The sequence of processes $\{\mathbf{q}^N(t)\}_{t\geq 0}$ has a weak limit $\{\mathbf{q}(t)\}_{t\geq 0}$ that satisfies the system of differential equations

(3.2)
$$\frac{\mathrm{d}q_i(t)}{\mathrm{d}t} = \lambda(q_{i-1}^d(t) - q_i^d(t)) - (q_i(t) - q_{i+1}(t)), \quad i = 1, 2, \dots,$$

with $q_0(t) \equiv 1$ for all $t \geq 0$. The fluid-limit equations may be interpreted as follows. The first term represents the rate of increase in the fraction of servers with i or more tasks due to arriving tasks that are assigned to a server with exactly i-1 tasks. Note that the latter occurs in fluid state $\mathbf{q} \in \mathcal{S}$ with probability $q_{i-1}^d - q_i^d$, i.e., the probability that all d sampled servers have i-1 or more tasks, but not all of them have i or more tasks. The second term corresponds to the rate of decrease in the fraction of servers with i or more tasks due to service completions from servers with exactly i tasks, and the latter rate is given by $q_i - q_{i+1}$. The system in (3.2) characterizes the functional law of large numbers (FLLN) behavior of systems in regime (ii) under the JSQ(d) scheme. Weak convergence of the diffusion-scaled variation around the fluid-limit path to a certain Ornstein-Uhlenbeck process in the same load regime (both the transient behavior and in steady state) was shown in [70], establishing a functional central limit theorem (FCLT) result. Strong approximations for systems under the JSQ(d) scheme on any finite time interval by the deterministic system

in (3.2), a certain infinite-dimensional jump process, and a diffusion approximation were established in [107].

Now, assume $\lambda \in (0,1)$ for ergodicity of the queue length process. When the derivatives in (3.2) are set equal to zero for all i, the unique fixed point for any $d \geq 2$ is obtained as [113, 163]

(3.3)
$$q_i^* = \lambda^{\frac{d^i - 1}{d - 1}}, \quad i = 1, 2, \dots$$

It can be shown that the fixed point is asymptotically stable in the sense that $\mathbf{q}(t) \to \mathbf{q}^*$ as $t \to \infty$ for any initial fluid state \mathbf{q}^∞ with $\sum_{i=1}^\infty q_i^\infty < \infty$. As mentioned earlier, the fixed point reveals that the stationary queue length distribution at each individual server exhibits superexponential decay as $N \to \infty$, as opposed to exponential decay for a random assignment policy. As described above, this involves an interchange of the many-server $(N \to \infty)$ and stationary $(t \to \infty)$ limits. The justification is provided by the asymptotic stability of the fixed point along with a few further technical conditions.

3.3. Fluid Limit for JSQ Policy. We now turn to the fluid limit for the ordinary JSQ policy, which rather surprisingly was not rigorously established until fairly recently in [119], leveraging martingale functional limit theorems and time-scale separation arguments [80].

In order to state the fluid limit starting from an arbitrary fluid-scaled occupancy state, we first introduce some additional notation. For any fluid state $\mathbf{q} \in \mathcal{S}$, denote by $m(\mathbf{q}) = \min\{i \geq 0 : q_{i+1} < 1\}$ the minimum queue length among all servers. Now if $m(\mathbf{q}) = 0$, then define $p_0(\mathbf{q}) = 1$ and $p_i(\mathbf{q}) = 0$ for all $i = 1, 2, \ldots$ Otherwise, in the case that $m(\mathbf{q}) > 0$, define

(3.4)
$$p_i(\mathbf{q}) = \begin{cases} \min\left\{ (1 - q_{m(\mathbf{q})+1})/\lambda, 1 \right\} & \text{for } i = m(\mathbf{q}) - 1, \\ 1 - p_{m(\mathbf{q})-1}(\mathbf{q}) & \text{for } i = m(\mathbf{q}), \\ 0 & \text{otherwise.} \end{cases}$$

The fluid-limit result below is paraphrased from [119].

Fluid Limit of JSQ. For $\lambda \in (0,1)$, the weak limit of the sequence of processes $\{\mathbf{q}^N(t)\}_{t\geq 0}$ is given by a deterministic system $\{\mathbf{q}(t)\}_{t\geq 0}$ that satisfies the system of differential equations

(3.5)
$$\frac{d^+q_i(t)}{dt} = \lambda p_{i-1}(\mathbf{q}(t)) - (q_i(t) - q_{i+1}(t)), \quad i = 1, 2, \dots,$$

where d^+/dt denotes the right-derivative. The reason we have used the derivative in (3.2) and the right-derivative in (3.5) is that the limiting trajectory for the JSQ policy may not be differentiable at all time points. In fact, one of the major technical challenges in proving the fluid limit for the JSQ policy is that the drift of the process is not continuous, which leads to nonsmooth limiting trajectories; see [119] for further details. The uniqueness of the above weak limit was not established in [119] but follows from the recent result in [19, Theorem 2.1].

The fluid-limit trajectory in (3.5) can be interpreted as follows. The coefficient $p_i(\mathbf{q})$ represents the instantaneous fraction of incoming tasks assigned to servers with a queue length of exactly i in the fluid state $\mathbf{q} \in \mathcal{S}$. Note that a strictly positive fraction $1 - q_{m(\mathbf{q})+1}$ of the servers has a queue length of exactly $m(\mathbf{q})$. Clearly the

fraction of incoming tasks that gets assigned to servers with a queue length of $m(\mathbf{q})+1$ or larger is zero: $p_i(\mathbf{q}) = 0$ for all $i = m(\mathbf{q}) + 1, \dots$ Also, tasks at servers with a queue length of exactly i are completed at (normalized) rate $q_i - q_{i+1}$, which is zero for all $i = 0, \dots, m(\mathbf{q}) - 1$, and hence the fraction of incoming tasks that get assigned to servers with a queue length of $m(\mathbf{q}) - 2$ or less is zero as well: $p_i(\mathbf{q}) = 0$ for all $i=0,\ldots,m(\mathbf{q})-2$. This only leaves the fractions $p_{m(\mathbf{q})-1}(\mathbf{q})$ and $p_{m(\mathbf{q})}(\mathbf{q})$ to be determined. Now observe that the fraction of servers with a queue length of exactly $m(\mathbf{q})-1$ is zero. If $m(\mathbf{q})=0$, then clearly the incoming tasks will join an empty queue, and thus $p_{m(\mathbf{q})} = 1$ and $p_i(\mathbf{q}) = 0$ for all $i \neq m(\mathbf{q})$. Furthermore, if $m(\mathbf{q}) \geq 1$, since tasks at servers with a queue length of exactly $m(\mathbf{q})$ are completed at (normalized) rate $1 - q_{m(\mathbf{q})+1} > 0$, incoming tasks can be assigned to servers with a queue length of exactly $m(\mathbf{q}) - 1$ at that rate. We thus need to distinguish between two cases, depending on whether or not the normalized arrival rate λ is larger than $1 - q_{m(\mathbf{q})+1}$. If $\lambda < 1 - q_{m(\mathbf{q})+1}$, then all the incoming tasks can be assigned to a server with a queue length of exactly $m(\mathbf{q}) - 1$, so that $p_{m(\mathbf{q})-1}(\mathbf{q}) = 1$ and $p_{m(\mathbf{q})}(\mathbf{q}) = 0$. On the other hand, if $\lambda > 1 - q_{m(\mathbf{q})+1}$, then not all incoming tasks can be assigned to servers with a queue length of exactly $m(\mathbf{q}) - 1$ active tasks, and a positive fraction will be assigned to servers with a queue length of exactly $m(\mathbf{q})$: $p_{m(\mathbf{q})-1}(\mathbf{q}) = (1-q_{m(\mathbf{q})+1})/\lambda$ and $p_{m(\mathbf{q})}(\mathbf{q}) = 1 - p_{m(\mathbf{q})-1}(\mathbf{q}).$

In the case that $\lambda \in (0,1)$, the unique fixed point $\mathbf{q}^* = (q_1^*, q_2^*, \ldots)$ of the dynamical system in (3.5) is given by

(3.6)
$$q_i^* = \begin{cases} \lambda, & i = 1, \\ 0, & i = 2, 3, \dots \end{cases}$$

Note that the fixed point naturally emerges when $d \to \infty$ in the fixed point expression (3.3) for fixed d. However, the process-level results in [114, 163] for fixed d cannot be readily used to handle joint scalings of d and N and do not yield the entire fluid-scaled sample path for arbitrary initial states as given by (3.5). The fixed point in (3.6), in conjunction with an interchange of limits argument, indicates that in stationarity the fraction of servers with a queue length of two or more under the JSQ policy is negligible as $N \to \infty$.

3.4. Diffusion Limit for JSQ Policy. We next describe the diffusion limit for the JSQ policy in the Halfin–Whitt heavy-traffic regime (2.1), as derived in [36]. The statement below is paraphrased from [36]. Recall the centered and diffusion-scaled processes in (3.1).

Diffusion Limit for JSQ. For suitable initial conditions, the sequence of processes $\{\bar{\mathbf{Q}}^N(t)\}_{t\geq 0}$ converges weakly to the limit $\{\bar{\mathbf{Q}}(t)\}_{t\geq 0}$, where $(\bar{Q}_1(t), \bar{Q}_2(t), \ldots)$ is the unique solution to the system of SDEs

(3.7)
$$d\bar{Q}_1(t) = \sqrt{2}dW(t) - \beta dt - \bar{Q}_1(t) + \bar{Q}_2(t) - dU_1(t),$$

$$d\bar{Q}_2(t) = dU_1(t) - \bar{Q}_2(t),$$

and $\bar{Q}_i(t)=0$, $i\geq 3$, for $t\geq 0$, where W is standard Brownian motion and U_1 is the unique continuous nondecreasing nonnegative process satisfying $\int_0^\infty \mathbb{1}_{[\bar{Q}_1(t)<0]} dU_1(t) = 0$ and $U_1(0)=0$.

The diffusion-limit characterization in (3.7) may be interpreted as follows. First, recall that $-\bar{Q}_1$ corresponds to the number of vacant servers (normalized by \sqrt{N}), and observe that this number is governed by the number of arriving tasks on the one

hand (as long as the number of vacant servers is nonzero), with associated exponential rate $\lambda(N)$, and on the other hand the number of service completions at servers with exactly one task, with associated exponential rate $Q_1^N - Q_2^N$. Noting that $(N - \lambda(N))/\sqrt{N} \to \beta$, $\bar{Q}_1^N = -(N - Q_1^N)/\sqrt{N}$, and $\bar{Q}_2^N = Q_2^N/\sqrt{N}$, we recognize that these dynamics are reflected in the equation for $d\bar{Q}_1(t)$, with $\sqrt{2}dW(t)$ an additional diffusion term corresponding to the variation in the number of arrivals and service completions around the drift terms and $dU_1(t)$ a reflection term accounting for the fact that the number of vacant servers cannot be negative. More specifically, the term $dU_1(t)$ tracks the number of arriving tasks assigned to busy servers when there are no vacant servers, which explains why the derivative can only be positive when $\bar{Q}_1 < 0$. Now observe that, for suitable initial conditions, since $\beta < 0$, it is highly unlikely that all servers have two or more tasks, and the number of servers with three or more tasks is negligible on diffusion scale, as reflected in the fact that $\bar{Q}_i = 0, i \geq 3$. Also, the dynamics of the number of servers with two or more tasks are governed by the assignment of tasks to busy servers captured by the term $dU_1(t)$ and the service completions at servers with exactly two tasks, which is equal to Q_2 on diffusion scale since the number of servers with three or more tasks is negligible, explaining the equation for $d\bar{Q}_2(t)$.

The above convergence of the scaled occupancy measure was established in [36] for any finite time interval. The tightness of the sequence of diffusion-scaled steady-state occupancy measures $\{(\bar{Q}_1^N(\infty), \bar{Q}_2^N(\infty))\}_{N\geq 1}$ and the ergodicity of the limiting diffusion process (3.7), and hence the interchange of limits, were open problems until [24] further established that the weak-convergence result extends to the steady state as well, i.e., $\bar{\mathbf{Q}}^N(\infty)$ converges weakly to the random variable $(\bar{Q}_1(\infty), \bar{Q}_2(\infty), 0, 0, \ldots)$ as $N \to \infty$, where $(\bar{Q}_1(\infty), \bar{Q}_2(\infty))$ has the stationary distribution of the process (\bar{Q}_1, \bar{Q}_2) . Thus, the steady state of the diffusion process in (3.7) is proved to capture the asymptotic behavior of large-scale systems under the JSQ policy.

In [24] a Lyapunov function is obtained via a generator expansion framework using Stein's method, which establishes exponential ergodicity of (\bar{Q}_1, \bar{Q}_2) . Although this approach gives a good handle on the rate of convergence to stationarity, it sheds little light on the form of the stationary distribution of the limiting diffusion process (3.7) itself. In two companion papers [13, 14] the authors perform a detailed analysis of the steady state of this diffusion process. Using a classical regenerative process construction of the diffusion process in (3.7), [13] establishes that $\bar{Q}_1(\infty)$ has a Gaussian tail and the tail exponent is uniformly bounded by constants which do not depend on β , whereas $\bar{Q}_2(\infty)$ has an exponentially decaying tail and the coefficient in the exponent is linear in β . More precisely, for any $\beta > 0$ there exist positive constants C_1, C_2, D_1, D_2 not depending on β and positive constants $C^l(\beta)$, $C^u(\beta)$, $D^l(\beta)$, $D^u(\beta)$, $D_R(\beta)$ depending only on β such that

(3.8)
$$C^{l}(\beta)e^{-C_{1}x^{2}} \leq \mathbb{P}(\bar{Q}_{1}(\infty) < -x) \leq C^{u}(\beta)e^{-C_{2}x^{2}}, \quad x \geq C_{R}(\beta),$$
$$D^{l}(\beta)e^{-D_{1}\beta y} \leq \mathbb{P}(\bar{Q}_{2}(\infty) > y) \leq D^{u}(\beta)e^{-D_{2}\beta y}, \quad y \geq D_{R}(\beta).$$

It was further shown in [13] that there exists a positive constant C^* not depending on β such that almost surely along any sample path,

$$-2\sqrt{2} \le \liminf_{t \to \infty} \frac{\bar{Q}_1(t)}{\sqrt{\log t}} \le -1,$$
$$\frac{1}{\beta} \le \limsup_{t \to \infty} \frac{\bar{Q}_2(t)}{\log t} \le \frac{2}{\mathcal{C}^*\beta}.$$

Notice that the width of fluctuation of \bar{Q}_1 does not depend on the value of β , whereas that of \bar{Q}_2 is linear in β^{-1} .

Since the Nth system is ergodic and its arrival rate is $N - \beta \sqrt{N}$, it is straightforward to see that $\mathbb{E}(\bar{Q}_1^N(\infty)) = -\beta$ for all N, and hence it can also be derived from the evolution of the limiting diffusion process that $\mathbb{E}(\bar{Q}_1(\infty)) = -\beta$. Thus, intuitively, for large enough β , the system has mostly many idle servers, and the number of servers with queue length at least two diminishes. But the manner in which $\bar{Q}_2(\infty)$ scales as β becomes large is highly nontrivial. Specifically, it was shown in [14] that there exist $\beta_0 \geq 1$ and positive constants C_1, C_2, D_1, D_2 such that, for all $\beta \geq \beta_0$,

(3.9)
$$e^{-C_1\beta^2} \le \mathbb{E}\left(\bar{Q}_2(\infty)\right) \le e^{-C_2\beta^2},$$

$$\mathbb{P}\left(\bar{Q}_2(\infty) \ge e^{-e^{D_1\beta^2}}\right) \le e^{-D_2\beta^2},$$

i.e., the steady-state mean is of order $e^{-C\beta^2}$, but most of the steady-state mass concentrates at a much smaller scale $e^{-e^{D_1\beta^2}}$. This suggests intermittency in the behavior of the \bar{Q}_2 process, namely, \bar{Q}_2 is typically of order $e^{-e^{D_1\beta^2}}$, but during rare events when it achieves higher values, it takes a long time to decay. However, for small enough β , the behavior is qualitatively different. Since $\mathbb{E}(\bar{Q}_1(\infty)) = -\beta$, the system is expected to become more congested as β becomes smaller. As a result, intuitively, \bar{Q}_2 should increase in the distributional sense. In this regime as well, \bar{Q}_2 exhibits some striking behavior. Specifically, it was shown in [14] that there exist positive constants β^* , M_1 , and M_2 such that, for all $\beta \leq \beta^*$,

$$(3.10) \frac{M_1}{\beta} \le \mathbb{E}(\bar{Q}_2(\infty)) \le \frac{M_2}{\beta}.$$

Comparison with M/M/N Queue. The M/M/N queue in the Halfin–Whitt heavy-traffic regime has been studied quite extensively (see [46, 48, 49, 74, 155, 156, 157] and the references therein). In this case, the centered and scaled total number of tasks in the system $(\bar{S}^N(t)-N)/\sqrt{N}$ converges weakly to a diffusion process $\{\bar{S}(t)\}_{t\geq 0}$ [74, Theorem 2] with

(3.11)
$$\mathrm{d}\bar{S}(t) = \sqrt{2}\mathrm{d}W(t) - \beta\mathrm{d}t - \mathrm{d}\bar{S}(t)\mathbb{1}_{\left[\bar{S}(t) \le 0\right]},$$

where W is standard Brownian motion. As reflected in (3.7) and (3.11), the JSQ policy and the M/M/N system share some striking similarities in terms of the qualitative behavior of the total number of tasks in the system. In particular, both the number of idle servers and the number of waiting tasks are of the order $\Theta(\sqrt{N})$. This shows that, despite the distributed queuing operation, a suitable load balancing policy can deliver a similar combination of excellent service quality and high resource utilization efficiency in the QED regime (cf. section 2.2) as in a centralized queuing arrangement. Moreover, the interchange of limits result in [24] implies that for systems under the JSQ policy, $\bar{Q}_{tot}^N(\infty) := \sum_{i=1}^{\infty} Q_i^N(\infty)$ converges weakly to $\bar{Q}_1(\infty) + \bar{Q}_2(\infty)$, which has an exponential upper tail (large positive deviation) and a Gaussian lower tail (large negative deviation); see (3.8). This is again reminiscent of the corresponding tail asymptotics for the M/M/N queue. Note that since $\bar{S}(\cdot)$ is a simple combination of a Brownian motion with a negative drift (when all servers are fully occupied) and an Ornstein–Uhlenbeck process (when there are idle servers), the steady-state distribution $\bar{S}(\infty)$ can be computed explicitly and is indeed a combination of an exponential distribution and a Gaussian distribution.

There are, however, some clear differences between the diffusions in (3.7) and (3.11):

- (i) Observe that in the case of M/M/N systems, whenever there are waiting tasks (equivalent to Q_2 being positive in our case), the queue length has a constant negative drift towards zero. This leads to the exponential upper tail of $\bar{S}(\infty)$, by comparison with the stationary distribution of a reflected Brownian motion with constant negative drift. In the JSQ case, however, the rate of decrease of Q_2 is always proportional to itself, which makes it somewhat counterintuitive that its stationary distribution has an exponential tail.
- (ii) In the M/M/N system, the number of idle servers can be nonzero only when the number of waiting tasks is zero. Thus, the dynamics of both the number of idle servers and the number of waiting tasks are completely captured by the one-dimensional process \bar{S}^N and by the one-dimensional diffusion \bar{S} in the limit. But in the JSQ case, \bar{Q}_2 is never zero, and the dynamics of (\bar{Q}_1, \bar{Q}_2) are truly two-dimensional (although the diffusion is nonelliptic) with \bar{Q}_1 and \bar{Q}_2 interacting with each other in an intricate manner.
- (iii) From (3.7) we see that \bar{Q}_2 never hits zero. Thus, in steady state, there is no mass at $\bar{Q}_2 = 0$, and the system always has waiting tasks. This is in sharp contrast to the M/M/N case, where the system has no waiting tasks in steady state with positive probability.
- (iv) In the M/M/N system, a positive fraction of the tasks incurs a nonzero waiting time as $N \to \infty$, but a nonzero waiting time is only of length $1/(\beta\sqrt{N})$ in expectation. In contrast, in the JSQ case, it is easy to see that \bar{Q}_1 (the limit of the scaled number of idle servers) spends zero time at the origin, i.e., in steady state the fraction of arriving tasks that finds all servers busy vanishes in the large-N limit (in fact, this is of order $1/\sqrt{N}$; see [24]). However, such tasks will have to wait for the duration of a residual service time, implying that a nonzero waiting time is of order O(1) and does not vanish.
- (v) As $\beta \to 0$, [74, Proposition 2] implies that $\beta \bar{S}(\infty)$ for the M/M/N queue converges weakly to a unit-mean exponential distribution. In contrast, results in [14] show that $\beta(\bar{Q}_1(\infty) + \bar{Q}_2(\infty))$ converges weakly to a Gamma(2) random variable. This indicates that despite similar order of performance, due to the distributed operation, in terms of the number of waiting tasks JSQ is a factor 2 worse in expectation than the corresponding centralized system.
- **3.5.** JSQ(d) Policies in Heavy-Traffic Regime. Finally, we briefly discuss the behavior of JSQ(d) policies with a fixed value of d in the Halfin-Whitt heavy-traffic regime (2.1). While a complete characterization of the occupancy process for fixed d has remained elusive so far, significant partial results were obtained in [35]. In order to describe the transient asymptotics, introduce the following rescaled processes:

(3.12)
$$\bar{Q}_i^N(t) := \frac{N - Q_i^N(t)}{\sqrt{N}}, \quad i = 1, 2, \dots$$

Note that in contrast to (3.1), in (3.12) all components are centered by N. Also note that the sign of the first coordinate in (3.12) is the opposite of that in (3.1). The statement below is paraphrased from [35].

Process-Level Limit of JSQ(d) Policy in Halfin-Whitt Regime. Assuming that the initial states converge with respect to the product topology under the above scaling, [35, Theorem 2] establishes that on any finite time interval, $\bar{\mathbf{Q}}^N(\cdot)$ converges weakly

to a deterministic system $\bar{\mathbf{Q}}(\cdot)$ that satisfies the system of ODEs

$$d\bar{Q}_i(t) = -d(\bar{Q}_i(t) - \bar{Q}_{i-1}(t)) + \bar{Q}_{i+1}(t) - \bar{Q}_i(t), \quad i = 1, 2, \dots,$$

with the convention that $\bar{Q}_0(t) \equiv 0$. It is noteworthy that the scaled occupancy process loses its diffusive behavior for fixed d. It is further shown in [35] that with high probability the steady-state fraction of queues with length at least $\log_d(\sqrt{N}/\beta) - o(1)$ tasks approaches unity, which in turn implies that with high probability the steady-state delay is at least $\log_d(\sqrt{N}/\beta) - O(1)$ as $N \to \infty$. The diffusion approximation of the JSQ(d) policy in the Halfin–Whitt regime (2.1), starting from a different initial state, has been studied in [27].

In [176] a broad framework involving Stein's method was introduced to analyze the rate of convergence of the stationary distribution under the JSQ(2) policy in a heavy-traffic regime, where $(N-\lambda(N))/\eta(N) \to \beta > 0$ as $N \to \infty$, with $\eta(N)$ a positive function diverging to infinity as $N \to \infty$. Note that the case $\eta(N) = \sqrt{N}$ corresponds to the Halfin–Whitt heavy-traffic regime (2.1). Using this framework, it was proved that when $\eta(N) = N^{\alpha}$ with some $4/5 < \alpha \le 1$,

$$(3.13) \quad \mathbb{E}\Big(\sum_{i=1}^{\infty} \left| q_i^N(\infty) - q_i^{N,\star} \right| \Big) \leq \frac{1}{N^{2\alpha - 1 - \xi}}, \quad \text{where} \quad q_i^{N,\star} = \left(\frac{\lambda(N)}{N}\right)^{2^i - 1},$$

and $\xi > 0$ is an arbitrarily small constant. Equation (3.13) not only shows that asymptotically the stationary occupancy measure concentrates at $\mathbf{q}^{N,\star}$, but also provides the rate of convergence.

4. Universality of JSQ(d) **Policies.** In this section we will further explore the trade-off between delay performance and communication overhead as a function of the diversity parameter d, in conjunction with the relative load. The latter trade-off will be examined in an asymptotic regime where not only does the total task arrival rate $\lambda(N)$ grow with N, but also the diversity parameter depends on N, and we write d(N) to explicitly reflect this dependence. We will specifically investigate what growth rate of d(N), depending on the scaling behavior of $\lambda(N)$, is required in order to asymptotically match the optimal performance of the JSQ policy and achieve a zero mean waiting time in the limit. The results presented in the remainder of the section are based on [119], where full proofs are also provided, unless specified otherwise.

Theorem 4.1 (universality of fluid limit for JSQ(d(N))). If $d(N) \to \infty$ as $N \to \infty$, then any fluid limit of the JSQ(d(N)) scheme coincides with that of the ordinary JSQ policy and, in particular, satisfies the system of differential equations in (3.5). Consequently, the stationary occupancy states converge to the unique fixed point as in (3.6).

THEOREM 4.2 (universality of diffusion limit for JSQ(d(N))). If $d(N)/(\sqrt{N}\log N) \to \infty$ as $N \to \infty$, then for suitable initial conditions the weak limit of the sequence of processes $\{\bar{\mathbf{Q}}^N(t)\}_{t\geq 0}$, under the JSQ(d(N)) policy, coincides with that of the ordinary JSQ policy and, in particular, is given by the system of SDEs in (3.7).

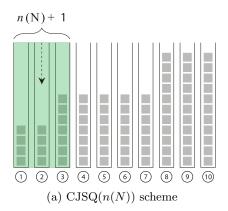
The above universality properties indicate that the JSQ overhead can be lowered by almost factors of O(N) and $O(\sqrt{N}/\log N)$ while retaining fluid- and diffusion-level optimality, respectively. In other words, Theorem 4.1 or 4.2 reveals that it is sufficient for d(N) to grow at any rate, or faster than $\sqrt{N}\log N$, in order to observe scaling benefits similar to those of a pooled system with N parallel single-server queues on fluid scale and diffusion scale, respectively. The stated conditions are in fact close to

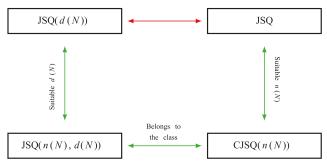
necessary, in the sense that if d(N) is uniformly bounded or $d(N)/(\sqrt{N}\log N) \to 0$ as $N \to \infty$, then, respectively, the fluid-limit and diffusion-limit paths under the JSQ(d(N)) scheme differ from those under the ordinary JSQ policy. In particular, if d(N) is uniformly bounded, the mean steady-state waiting time does not vanish as $N \to \infty$.

Remark 4.3. One implication of Theorem 4.1 is that in the subcritical regime any growth rate of d(N) is sufficient to achieve asymptotically vanishing steady-state probability of wait. This result is complemented by the results in [25, 97], where the steady-state analysis is extended in the heavy-traffic regime with $N^{\alpha}(1-\lambda(N)/N) \rightarrow$ $\beta > 0$ as $N \to \infty$ with $\alpha \in (0, 1/2)$. Note that the system approaches heavy traffic as the number of servers N grows large, but that the load is lighter than that in the Halfin-Whitt regime, which corresponds to $\alpha = 1/2$. Specifically, it is established in [97] that the steady-state probability of wait for the JSQ(d(N)) policy with $d(N) \geq$ $\frac{1}{\beta}N^{\alpha}\log N$ vanishes as $N\to\infty$. The results of [25] imply that when $\beta=1$ and d(N)=1 N^{γ} with $\alpha, \gamma \in (0, 1], k = \lceil (1 - \alpha)/\gamma \rceil$, and $2\alpha + \gamma(k - 1) > 1$, with probability tending to 1 as $N \to \infty$, the proportion of queues with queue length equal to k is at least $1-2N^{-1+\alpha+(k-1)\gamma}$ and there are no longer queues. A crucial distinction between the result stated in Theorem 4.2 and the results in [25, 97] is that the former analyzes the system on diffusion scale (and describes its behavior in terms of a limiting diffusion process), whereas [25, 97] analyze the system on fluid scale (and characterize its behavior in terms of a limiting fluid-scaled occupancy state). Much less is known when the asymptotic load is higher than the Halfin-Whitt regime, that is, when $N^{\alpha}(1 \lambda(N)/N \to \beta > 0$ as $N \to \infty$ with $\alpha \in (1/2,1)$. This is also known as the super-Halfin-Whitt regime. In this regime, when the system has a finite buffer capacity, [96] identifies a broad class of load balancing policies including the JSQ policy, idleone-first (I1F) policy, and the JSQ(d(N)) policy with $d(N) \geq N^{\alpha} \log^2 N$, for which, in steady state, $\mathbb{E}(Q_2^N(\infty))$ is $O(N^{\alpha} \log N)$ and $\mathbb{E}(Q_3^N(\infty))$ is $O(N^{-r(1-\alpha)-1})$, where r > 0 can be any constant independent of N. Further, [179] analyzes the process-level and steady-state limits of the occupancy process under the JSQ policy in the super-Halfin-Whitt regime and, in particular, shows that $Q_2^N(\infty)/N^{\alpha}$ converges weakly to a $Gamma(2, \beta)$ distribution (sum of two independent Exponential(β) distributions). Results in [19] allow for arbitrary growth rate of d(N) in the analysis of the JSQ(d(N))policy in the heavy-traffic regime. In that paper, the authors establish a process-level diffusion limit of the occupancy process under the JSQ(d(N)) policy for certain ranges of $\lambda(N)$ that depend on d(N). In particular, they include an alternative proof of the universality result in Theorem 4.2.

4.1. High-Level Outline of Proof Approach. The proofs of Theorems 4.1 and 4.2 rely on a stochastic coupling construction to bound the difference in the queue length processes between the JSQ policy and a scheme with an arbitrary value of d(N). This coupling is then exploited to obtain the fluid and diffusion limits of the JSQ(d(N)) policy, along with the associated fixed point, under the conditions stated in Theorems 4.1 and 4.2. Moreover, we will also allow the possibility that the servers have a finite buffer capacity B. In that case, whenever a task is assigned to a server that has B tasks in the queue (including the one currently in service), that task is lost forever. For an LBA Π , we will denote the total number of tasks lost up to time t by $L^{\Pi}(t)$.

A direct comparison between the JSQ(d(N)) scheme and the ordinary JSQ policy is not straightforward, which is why the CJSQ(n(N)) class of schemes is introduced as an intermediate scenario to establish the universality results. Just like the JSQ(d(N)) scheme, the schemes in the class CJSQ(n(N)) may be thought of as "sloppy" versions





(b) Asymptotic equivalence relations

Fig. 4 (a) High-level view of the CJSQ(n(N)) class of schemes, where, as in Figure 3, the servers are arranged in nondecreasing order of their queue lengths, and the arrival must be assigned through the green shaded region on the left. (b) The equivalence structure is depicted for various intermediate load balancing schemes to facilitate the comparison between the JSQ(d(N)) scheme and the ordinary JSQ policy.

of the JSQ policy, in the sense that tasks are not necessarily assigned to a server with the shortest queue length but to one of the n(N) + 1 lowest ordered servers, as is illustrated in Figure 4(a). In particular, for n(N) = 0, the class only includes the ordinary JSQ policy. Note that the JSQ(d(N)) scheme is guaranteed to identify the lowest-ordered server, but only among a randomly sampled subset of d(N) servers. In contrast, a scheme in the CJSQ(n(N)) class only guarantees that one of the n(N) + 1 lowest-ordered servers is selected, but across the entire pool of N servers. It is worthwhile to note that CJSQ(n(N)) is a class of policies, and that any policy which ensures that tasks are always assigned to one of the n(N) + 1 lowest ordered servers, no matter what the exact mechanism of the policy is, belongs to this class. The proof of the universality results in Theorems 4.1 and 4.2 has two parts, as further described below. The proof strategy is schematically represented in Figure 4(b).

Step I. Performance of Schemes in CJSQ(n(N)) **Class.** The first step is to show that for sufficiently small n(N), any scheme from the class CJSQ(n(N)) is still "close" to the ordinary JSQ policy. To achieve this, another type of sloppiness will be

introduced. Let MJSQ(n(N)) be a particular scheme that always assigns incoming tasks to precisely the (n(N)+1)th ordered server. Notice that this scheme is effectively the JSQ policy when the system always maintains n(N) idle servers or, equivalently, uses only N - n(N) servers, and $MJSQ(n(N)) \in CJSQ(n(N))$. For brevity, we will often suppress n(N) in the notation where it is clear from the context. We call any two systems S-coupled if they have synchronized arrival clocks and departure clocks of the kth longest queue for $1 \leq k \leq N$ ("S" in the name of the coupling stands for "server"). Note that the S-coupling between two systems with identical arrival and service rates always exists. Indeed, since tasks have identically and exponentially distributed service time requirements, synchronizing the departure clocks of the kth longest queue for k = 1, ..., N preserves the marginal dynamics of each system. Consider three S-coupled systems following, respectively, the JSQ policy, any scheme from the class CJSQ, and the MJSQ scheme. Recall that $Q_i^{\Pi}(t)$ is the number of servers with at least i tasks at time t and $L^{\Pi}(t)$ is the total number of lost tasks up to time t, for the schemes $\Pi = JSQ$, CJSQ, MJSQ. The following proposition provides a stochastic ordering for any scheme in the class CJSQ with respect to the ordinary JSQ policy and the MJSQ scheme.

PROPOSITION 4.4. Fix any $N \ge 1$, $1 \le B \le \infty$, and $0 \le n(N) \le N - 1$. Then, in the joint probability space constructed by the S-coupling of the three systems under, respectively, JSQ, MJSQ, and any scheme from the class CJSQ, the following ordering

- $\begin{array}{l} \text{is preserved almost surely throughout the sample path: for all } 1 \leq m \leq B \text{ and } t \geq 0, \\ \text{(i) } \left\{ \sum_{i=m}^{B} Q_i^{\text{JSQ}}(t) + L^{\text{JSQ}}(t) \right\}_{t \geq 0} \leq \left\{ \sum_{i=m}^{B} Q_i^{\text{CJSQ}}(t) + L^{\text{CJSQ}}(t) \right\}_{t \geq 0}, \\ \text{(ii) } \left\{ \sum_{i=m}^{B} Q_i^{\text{CJSQ}}(t) + L^{\text{CJSQ}}(t) \right\}_{t \geq 0} \leq \left\{ \sum_{i=m}^{B} Q_i^{\text{MJSQ}}(t) + L^{\text{MJSQ}}(t) \right\}_{t \geq 0}, \\ \end{array}$ provided the inequalities hold at time t = 0.

COROLLARY 4.5. Under the conditions of Proposition 4.4, for all $1 \le m \le B$ and

- $\begin{array}{l} \geq 0, \\ (\mathrm{i}) \ \ Q_m^{\mathrm{CJSQ}}(t) \geq \sum_{i=m}^B Q_i^{\mathrm{JSQ}}(t) \sum_{i=m+1}^B Q_i^{\mathrm{MJSQ}}(t) + L^{\mathrm{JSQ}}(t) L^{\mathrm{MJSQ}}(t), \\ (\mathrm{ii}) \ \ Q_m^{\mathrm{CJSQ}}(t) \leq \sum_{i=m}^B Q_i^{\mathrm{MJSQ}}(t) \sum_{i=m+1}^B Q_i^{\mathrm{JSQ}}(t) + L^{\mathrm{MJSQ}}(t) L^{\mathrm{JSQ}}(t), \\ \end{array}$ provided the inequalities hold at time t = 0.

It can be shown that if $n(N)/N \to 0$ as $N \to \infty$, then the MJSQ(n(N)) scheme has the same fluid limit along any subsequence as the ordinary JSQ policy, whenever the latter exists. Corollary 4.5 then implies that as long as $n(N)/N \to 0$, any scheme from the class CJSQ(n(N)) has the same fluid limit along any subsequence as the ordinary JSQ policy, whenever the latter exists.

Step 2. JSQ(d(N)) Has Same Limit as a Particular Scheme in CJSQ(n(N)).

The next step is to prove that for sufficiently large d(N) relative to n(N), one can construct a scheme belonging to the CJSQ(n(N)) class that differs "negligibly" from the JSQ(d(N)) scheme. Specifically, consider the JSQ(n(N), d(N)) scheme with $n(N), d(N) \leq N$, which is an intermediate blend between the CJSQ(n(N)) schemes and the JSQ(d(N)) scheme. At its first step, just as in the JSQ(d(N)) scheme, the JSQ(d(N), n(N)) scheme first chooses the shortest of d(N) random candidates but only sends the arriving task to that server's queue if it is one of the n(N) + 1 shortest queues. If it is not, then at the second step it picks any of the n(N) + 1 shortest queues uniformly at random and then sends the task to that server's queue. Note that by construction, JSQ(d(N), n(N)) is a scheme in CJSQ(n(N)). Consider two S-coupled systems with a JSQ(d(N)) and a JSQ(n(N), d(N)) scheme. Assume that at some specific arrival epoch, the incoming task is dispatched to the kth ordered server in the system under the JSQ(d(N)) scheme. If $k \in \{1, 2, ..., n(N) + 1\}$, then the system under the JSQ(n(N), d(N)) scheme also assigns the arriving task to the kth ordered server. Otherwise, it dispatches the arriving task uniformly at random among the first (n(N) + 1) ordered servers.

Next, it is established that if $d(N) \to \infty$, then for some n(N) with $n(N)/N \to 0$, the JSQ(d(N)) scheme and the JSQ(n(N), d(N)) scheme have the same fluid limit. Theorem 4.1 then follows by Step 1 and observing that the JSQ(n(N), d(N)) scheme belongs to the class CJSQ(n(N)).

The proof of Theorem 4.2 follows the same arguments as the proof of Theorem 4.1, but uses the candidate $n(N)/\sqrt{N} \to 0$ (instead of $n(N)/N \to 0$) in Step 1, and the candidate $d(N)/(\sqrt{N}\log(N)) \to \infty$ (instead of $d(N) \to \infty$) in Step 2.

4.2. Extension to Batch Arrivals. Consider an extension of the model in which tasks arrive in batches. We assume that the batches arrive as a Poisson process of rate $\lambda(N)/\ell(N)$ and have fixed size $\ell(N)>0$, so that the effective total task arrival rate remains $\lambda(N)$. We will show that for any growing batch size, fluid-level optimality can be achieved with O(1) communication overhead per task. For that, we define the JSQ(d(N)) scheme adapted for batch arrivals: When a batch arrives, the dispatcher samples $d(N) \geq \ell(N)$ servers without replacement and assigns the tasks to the $\ell(N)$ servers with the smallest queue lengths among the sampled servers.

THEOREM 4.6 (batch arrivals). Consider the batch arrival scenario with growing batch size $\ell(N) \to \infty$ and $\lambda(N)/N \to \lambda < 1$ as $N \to \infty$. For the JSQ(d(N)) scheme with $d(N) \ge \ell(N)/(1-\lambda-\varepsilon)$ for any fixed $\varepsilon > 0$, if $q_1^N(0) \to q_1(0) \le \lambda$ and $q_i^N(0) \to 0$ for all $i \ge 2$, then any (subsequent) weak limit of the sequence of processes $\{\mathbf{q}^N(t)\}_{t\ge 0}$ coincides with that of the ordinary JSQ policy and, in particular, is given by the system in (3.5).

Observe that for a fixed $\varepsilon > 0$, the communication overhead per task is on average given by $(1 - \lambda - \varepsilon)^{-1}$, which is O(1). Thus Theorem 4.6 ensures that in the case of batch arrivals with growing batch size, fluid-level optimality can be achieved with O(1) communication overhead per task. The result for the fluid-level optimality in stationarity can also be obtained indirectly by exploiting the fluid-limit result in [177]. Specifically, it can be deduced from the result in [177] that for batch arrivals with growing batch size, the JSQ(d(N)) scheme with suitably growing d(N) yields the same fixed point of the fluid limit as described in (3.6).

5. Blocking and Infinite-Server Dynamics. The basic scenario that we have focused on so far involves single-server queues. In this section we turn our attention to a system with parallel server pools, each with a fixed number B servers, where B can be either finite or infinite. As before, tasks arrive at a single dispatcher and must immediately be forwarded to one of the server pools, but they must also directly start execution or be discarded otherwise. As before, under the JSQ(d) policy, at each task arrival, the dispatcher selects d random server pools and assigns the task to the one with the lowest number of active tasks. When B is finite, a task that happens to land on a server pool with B active tasks is lost forever. In that case, the maximum total rate at which tasks can be processed in the system is BN, which we assume to be higher than the total arrival rate $\lambda(N)$. In other words, when $\lambda(N)/N = \lambda \in \mathbb{R}_+$, we assume $\lambda < B$. The execution times are assumed to be exponentially distributed and do not depend on the number of other tasks receiving service simultaneously. In order to distinguish it from the single-server queuing dynamics considered earlier, the

current scenario will henceforth be referred to as the "infinite-server dynamics."

As it turns out, the JSQ policy has similar stochastic optimality properties as in the case of single-server queues and, in particular, stochastically minimizes the cumulative number of discarded tasks [82, 111, 112, 139]. However, the JSQ policy also suffers from a similar scalability issue due to the excessive communication overhead in large-scale systems, which can be mitigated through JSQ(d) policies. Results in [147] and the more recent papers [89, 124, 127, 173] indicate that JSQ(d) policies provide similar power-of-choice gains for loss probabilities. It can be shown though that the optimal performance of the JSQ policy cannot be matched for any fixed value of d.

Motivated by these observations, we explore the trade-off between performance and communication overhead for infinite-server dynamics. We will demonstrate that the optimal performance of the JSQ policy can be asymptotically retained while drastically reducing the communication burden, mirroring the universality properties described in section 4 for single-server queues. The results presented in the remainder of the section are extracted from [120], where complete proofs are also provided, unless indicated otherwise.

5.1. Fluid Limit for JSQ Policy. Analogous to the single-server case, we represent the state of the Nth system by the vector $\mathbf{Q}^N(t) := (Q_1^N(t), Q_2^N(t), \ldots)$, with $Q_i^N(t)$ denoting the number of server pools with i or more active tasks at time t; the fluid-scaled occupancy state is denoted by $\mathbf{q}^N(t) := (q_1^N(t), q_2^N(t), \ldots)$, with $q_i^N(t) = Q_i^N(t)/N$ for $i \geq 1$. Also, as in subsection 3.3, for any fluid state $\mathbf{q} \in \mathcal{S}$, denote by $m(\mathbf{q}) = \min\{i \geq 0 : q_{i+1} < 1\}$ the minimum number of active tasks among all server pools with the convention that $q_{B+1} = 0$ if $B < \infty$. Now if $m(\mathbf{q}) = 0$, then define $p_0(\mathbf{q}) = 1$ and $p_i(\mathbf{q}) = 0$ for all $i = 1, 2, \ldots$. Otherwise, in the case $m(\mathbf{q}) > 0$, define

(5.1)
$$p_i(\mathbf{q}) = \begin{cases} \min \left\{ m(\mathbf{q})(1 - q_{m(\mathbf{q})+1})/\lambda, 1 \right\} & \text{for } i = m(\mathbf{q}) - 1, \\ 1 - p_{m(\mathbf{q})-1}(\mathbf{q}) & \text{for } i = m(\mathbf{q}), \\ 0 & \text{otherwise.} \end{cases}$$

Any weak limit of the sequence of processes $\{\mathbf{q}^N(t)\}_{t\geq 0}$ is given by a deterministic system $\{\mathbf{q}(t)\}_{t\geq 0}$ satisfying the differential equations

(5.2)
$$\frac{d^+q_i(t)}{dt} = \lambda p_{i-1}(\mathbf{q}(t)) - i(q_i(t) - q_{i+1}(t)), \quad i = 1, 2, \dots, B,$$

where d^+/dt denotes the right-derivative.

Equations (5.1) and (5.2) are to be contrasted with (3.4) and (3.5). While the form of the evolution equations (5.2) of the limiting dynamical system remains similar to (3.5), the rate of decrease of q_i is now $i(q_i - q_{i+1})$, reflecting the infinite-server dynamics.

Let $K := \lfloor \lambda \rfloor$ and $f := \lambda - K$ denote the integral and fractional parts of λ , respectively. Assuming $\lambda < B$, the unique fixed point of the dynamical system in (5.2) is given by

(5.3)
$$q_i^* = \begin{cases} 1, & i = 1, \dots, K, \\ f, & i = K+1, \\ 0, & i = K+2, \dots, B, \end{cases}$$

and thus $\sum_{i=1}^{B} q_i^* = \lambda$. This is consistent with the results in [124, 127, 173] for fixed d, where taking $d \to \infty$ yields the same fixed point. However, the results in

[124, 127, 173] cannot be used directly to handle joint scalings and do not yield the universality of the entire fluid-scaled sample path for arbitrary initial states. The fixed point in (5.3), in conjunction with an interchange of limits argument, indicates that in stationarity the fraction of server pools with at least K+2 and at most K-1 active tasks is negligible as $N \to \infty$.

- **5.2. Diffusion Limit for JSQ Policy.** As it turns out, the diffusion-limit results may be qualitatively different, depending on whether f=0 or f>0, and we will distinguish between these two cases accordingly. Observe that for any assignment scheme, in the absence of overflow events, the total number of active tasks evolves as the number of jobs in an $M/M/\infty$ system with arrival rate $\lambda(N)$ and unit service rate, for which the diffusion limit is well known [135]. For the JSQ policy we can establish, for suitable initial conditions, that the total number of server pools with K-2 or fewer and K+2 or more tasks is negligible on the diffusion scale. If f>0, the number of server pools with K-1 tasks is negligible as well, and the dynamics of the number of server pools with K or K+1 tasks can then be derived from the known diffusion limit of the total number of tasks mentioned above. In contrast, if f=0, the number of server pools with K-1 tasks is not negligible on the diffusion scale, and the limiting behavior is qualitatively different, but can still be characterized.
- **5.2.1. Diffusion-Limit Results for Nonintegral** λ **.** We first consider the case f > 0 and define $f(N) := \lambda(N) KN$. Based on the above observations, we define the following centered and scaled processes:

$$\bar{Q}_{i}^{N}(t) = N - Q_{i}^{N}(t) \ge 0 \quad \text{for} \quad i \le K - 1,$$

$$\bar{Q}_{K}^{N}(t) := \frac{N - Q_{K}^{N}(t)}{\log(N)} \ge 0,$$

$$\bar{Q}_{K+1}^{N}(t) := \frac{Q_{K+1}^{N}(t) - f(N)}{\sqrt{N}} \in \mathbb{R},$$

$$\bar{Q}_{i}^{N}(t) := Q_{i}^{N}(t) \ge 0 \quad \text{for} \quad i \ge K + 2.$$

Theorem 5.1 (diffusion limit for JSQ policy; f > 0). Assume $\bar{Q}_i^N(0)$ converges to $\bar{Q}_i(0)$ in \mathbb{R} , $i \geq 1$, and $\lambda(N)/N \to \lambda > 0$ as $N \to \infty$. Then

- (i) $\lim_{N\to\infty} \mathbb{P}\left[\sup_{t\in[0,T]} \bar{Q}_{K-1}^{N}(t) \leq 1\right) = 1$ and $\left\{\bar{Q}_{i}^{N}(t)\right\}_{t\geq0}$ converges weakly to $\left\{\bar{Q}_{i}(t)\right\}_{t\geq0}$, where $\bar{Q}_{i}(t)\equiv0$, provided $\lim_{N\to\infty}\mathbb{P}\left(\bar{Q}_{K-1}^{N}(0)\leq1\right)=1$, and $\bar{Q}_{i}^{N}(0)\stackrel{\mathbb{P}}{\to}0$ for $i\leq K-2$;
- (ii) $\{\bar{Q}_K^N(t)\}_{t\geq 0}$ is a stochastically bounded sequence of processes;
- (iii) $\{\bar{Q}_{K+1}^N(t)\}_{t\geq 0}$ converges weakly to $\{\bar{Q}_{K+1}(t)\}_{t\geq 0}$, where $\bar{Q}_{K+1}(t)$ is given by the Ornstein–Uhlenbeck process satisfying the $S\overline{D}E$

$$d\bar{Q}_{K+1}(t) = -\bar{Q}_{K+1}(t)dt + \sqrt{2\lambda}dW(t),$$

where W(t) is standard Brownian motion, provided $\bar{Q}_{K+1}^N(0)$ converges to $\bar{Q}_{K+1}(0)$ in \mathbb{R} :

(iv) for $i \geq K+2$, $\{\bar{Q}_i^N(t)\}_{t\geq 0}$ converges weakly to $\{\bar{Q}_i(t)\}_{t\geq 0}$, where $\bar{Q}_i(t) \equiv 0$, provided $\bar{Q}_i^N(0)$ converges to 0 in \mathbb{R} .

Theorem 5.1 implies that for suitable initial states, for large N, there will be almost no server pool with K-2 or fewer tasks and K+2 or more tasks on any finite time interval. Also, the number of server pools having fewer than K tasks is of order

 $\log(N)$, and there are $fN + O_P(\sqrt{N})$ server pools with precisely K+1 active tasks. Below we present some high-level intuition behind the scaling limits in Theorem 5.1.

High-Level Proof Idea. Observe that $\sum_{i=1}^K (N-Q_i^N(\cdot))$ increases by one at rate

$$\sum_{i=1}^{K} i(Q_i(t) - Q_{i+1}(t)) = \sum_{i=1}^{K} (Q_i(t) - Q_{K+1}(t)) \approx K(1 - f)N,$$

which is when there is a departure from some server pool with at most K active tasks, and if positive, it decreases by one at constant rate $\lambda(N) = (K+f)N + \mathrm{o}(N)$, which is whenever there is an arrival. Thus, $\sum_{i=1}^K (N-Q_i^N(\cdot))$ roughly behaves as a birth-and-death process with birth rate K(1-f)N and death rate (K+f)N. Since f>0, we have K+f>K(1-f), and on any finite time interval the maximum of such a birth-and-death process scales as $\log(N)$.

Similar to the argument above, the process $\sum_{i=1}^{K-1} \bar{Q}_i^N(\cdot)$ increases by one at rate

$$\begin{split} \sum_{i=1}^{K-1} i(Q_i^N(t) - Q_{i+1}^N(t)) &= \sum_{i=1}^{K-1} Q_i^N(t) - (K-1)Q_K^N(t) \\ &\leq (K-1)(N-Q_K^N(t)) = \mathcal{O}(\log(N)), \end{split}$$

which is when there is a departure from some server pool with at most K-1 active tasks, and if positive, it decreases by one at rate $\lambda(N)$, which is whenever there is an arrival. Thus, $\sum_{i=1}^{K-1} \bar{Q}_i^N(\cdot)$ roughly behaves as a birth-and-death process with birth rate $O(\log(N))$ and death rate O(N). This leads to the asymptotic result for $\sum_{i=1}^{K-1} \bar{Q}_i^N(\cdot)$ and, in particular, for $\bar{Q}_{K-1}^N(\cdot)$. This completes the proof of parts (i) and (ii) of Theorem 5.1.

Furthermore, since $\lambda < K+1$, the number of tasks that are assigned to server pools with at least K+1 tasks converges to zero in probability, and this completes the proof of part (iv) of Theorem 5.1.

Finally, all the above combined also means that on any finite time interval the total number of tasks in the system behaves with high probability as the total number of jobs in an $M/M/\infty$ system. Therefore, with the help of the diffusion-limit result for the $M/M/\infty$ system in [135, Theorem 6.14], we conclude the proof of part (iii) of Theorem 5.1.

5.2.2. Diffusion-Limit Results for Integral λ **.** We now turn to the case f=0 and assume that

(5.5)
$$\frac{KN - \lambda(N)}{\sqrt{N}} \to \beta \in \mathbb{R} \quad \text{as} \quad N \to \infty,$$

which can be thought of as an analogue of the Halfin–Whitt regime in (2.1). We now consider the following scaled quantities:

(5.6)
$$\zeta_1^N(t) := \frac{1}{\sqrt{N}} \sum_{i=1}^K (N - Q_i^N(t)), \qquad \zeta_2^N(t) := \frac{Q_{K+1}^N(t)}{\sqrt{N}}.$$

Theorem 5.2. Assuming the convergence of initial states, the process $\left\{ (\zeta_1^N(t),\zeta_2^N(t)) \right\}_{t\geq 0}$ converges weakly to the process $\left\{ (\zeta_1(t),\zeta_2(t)) \right\}_{t\geq 0}$ governed by

the system of SDEs

$$d\zeta_1(t) = \sqrt{2K} dW(t) - (\zeta_1(t) + K\zeta_2(t)) + \beta dt + dV_1(t),$$

$$d\zeta_2(t) = dV_1(t) - (K+1)\zeta_2(t),$$

where W is standard Brownian motion, and $V_1(t)$ is the unique continuous nondecreasing process satisfying $\int_0^t \mathbb{1}_{[\zeta_1(s)>0]} dV_1(s) = 0$ and $V_1(0) = 0$.

Unlike the f>0 case, the theorem above says that if f=0, then over any finite time horizon, there will be $\mathcal{O}_P(\sqrt{N})$ server pools with fewer than K or more than K active tasks, and hence most of the server pools have precisely K active tasks. The proof of Theorem 5.2 uses the reflection argument developed in [36]. Indeed, the proof follows by observing that the dynamics of $\left\{\left(\zeta_1^N(t),\zeta_2^N(t)\right)\right\}_{t\geq 0}$ resembles the dynamics of the JSQ policy in the Halfin–Whitt regime.

5.3. Universality of JSQ(d) Policies in Infinite-Server Dynamics. As in section 4, we now further explore the trade-off between performance and communication overhead as a function of the diversity parameter d(N), in conjunction with the load. We will specifically investigate what growth rate of d(N), depending on the scaling behavior of $\lambda(N)$, is required in order to asymptotically match the optimal performance of the JSQ policy.

THEOREM 5.3 (universality of fluid limit for JSQ(d(N)) and infinite-server dynamics). If $d(N) \to \infty$ as $N \to \infty$, then any (subsequent) fluid limit of the JSQ(d(N)) scheme coincides with that of the ordinary JSQ policy and, in particular, satisfies the system of differential equations in (5.2). Consequently, the stationary occupancy states converge to the unique fixed point as in (5.3).

In order to state the universality result on diffusion scale, define in the case f > 0

$$\bar{Q}_{i}^{d(N)}(t) := \frac{N - Q_{i}^{d(N)}(t)}{\sqrt{N}} \ge 0, \quad i \le K,$$

$$\bar{Q}_{K+1}^{d(N)}(t) := \frac{Q_{K+1}^{d(N)}(t) - f(N)}{\sqrt{N}} \in \mathbb{R},$$

$$\bar{Q}_{i}^{d(N)}(t) := \frac{Q_{i}^{d(N)}(t)}{\sqrt{N}} \ge 0 \quad \text{for} \quad i \ge K + 2,$$

and, otherwise, if f = 0,

$$\hat{Q}_{K-1}^{d(N)}(t) := \sum_{i=1}^{K-1} \frac{N - Q_i^{d(N)}(t)}{\sqrt{N}} \ge 0,$$

$$\hat{Q}_K^{d(N)}(t) := \frac{N - Q_K^{d(N)}(t)}{\sqrt{N}} \ge 0,$$

$$\hat{Q}_i^{d(N)}(t) := \frac{Q_i^{d(N)}(t)}{\sqrt{N}} \ge 0 \quad \text{for} \quad i \ge K + 1.$$

The scaling in (5.7) and (5.8) should be contrasted with (5.4) and (5.6), respectively.

THEOREM 5.4 (universality of diffusion limit for JSQ(d(N)) and infinite-server dynamics). Assume $d(N)/(\sqrt{N}\log N) \to \infty$. Under suitable initial conditions,

(i) if f > 0, then $\bar{Q}_i^{d(N)}(\cdot)$ converges to the zero process for $i \neq K+1$ and $\bar{Q}_{K+1}^{d(N)}(\cdot)$ converges weakly to the Ornstein-Uhlenbeck process satisfying the SDE

$$d\bar{Q}_{K+1}(t) = -\bar{Q}_{K+1}(t)dt + \sqrt{2\lambda}dW(t),$$

where W(t) is standard Brownian motion;

(ii) if f = 0, then $\hat{Q}_{K-1}^{d(N)}(\cdot)$ converges weakly to the zero process and $(\hat{Q}_{K}^{d(N)}(\cdot), \hat{Q}_{K+1}^{d(N)}(\cdot))$ converges weakly to $(\hat{Q}_{K}(\cdot), \hat{Q}_{K+1}(\cdot))$, described by the unique solution of the system of SDEs

$$d\hat{Q}_K(t) = \sqrt{2K} dW(t) - (\hat{Q}_K(t) + K\hat{Q}_{K+1}(t)) + \beta dt + dV_1(t),$$

$$d\hat{Q}_{K+1}(t) = dV_1(t) - (K+1)\hat{Q}_{K+1}(t),$$

where W is standard Brownian motion, and $V_1(t)$ is the unique continuous nondecreasing process satisfying $\int_0^t \mathbb{1}_{[\hat{Q}_K(s) \geq 0]} dV_1(s) = 0$ and $V_1(0) = 0$.

Having established the asymptotic results for the JSQ policy in sections 5.1 and 5.2, the proofs of the asymptotic results for the JSQ(d(N)) scheme in Theorems 5.3 and 5.4 involve establishing a universality result that shows that the limiting processes for the JSQ(d(N)) scheme are "asymptotically equivalent" to those for the ordinary JSQ policy for suitably large values of d(N). The notion of asymptotic equivalence between different schemes is formalized in the next definition.

DEFINITION 5.5. Let Π_1 and Π_2 be two schemes parameterized by the number of server pools N. For any positive function $g: \mathbb{N} \to \mathbb{R}_+$, we say that Π_1 and Π_2 are "g(N)-alike" if there exists a common probability space such that for any fixed $T \geq 0$, for all $i \geq 1$,

$$\sup_{t \in [0,T]} (g(N))^{-1} |Q_i^{\Pi_1}(t) - Q_i^{\Pi_2}(t)| \xrightarrow{\mathbb{P}} 0 \quad \text{as} \quad N \to \infty.$$

Intuitively speaking, if two schemes are g(N)-alike, then in some sense, the associated system occupancy states are indistinguishable on g(N)-scale. For brevity, for two schemes Π_1 and Π_2 that are g(N)-alike, we will often say that Π_1 and Π_2 have the same process-level limits on g(N)-scale. The next theorem states a sufficient criterion for the JSQ(d(N)) scheme and the ordinary JSQ policy to be g(N)-alike and, thus, it provides the key vehicle in establishing the universality result.

THEOREM 5.6. Let $g: \mathbb{N} \to \mathbb{R}_+$ be a function diverging to infinity. Then the JSQ policy and the JSQ(d(N)) scheme are g(N)-alike, with $g(N) \leq N$, if

$$(5.9) \hspace{1cm} (\mathrm{i}) \hspace{0.3cm} d(N) \to \infty \hspace{0.3cm} for \hspace{0.3cm} g(N) = \mathrm{O}(N),$$

$$(5.10) \qquad \qquad (ii) \quad d(N) \left(\frac{N}{g(N)} \log \left(\frac{N}{g(N)}\right)\right)^{-1} \to \infty \quad \textit{for} \quad g(N) = \mathrm{o}(N).$$

Theorem 5.6 yields the following two immediate corollaries.

COROLLARY 5.7. If $d(N) \to \infty$ as $N \to \infty$, then the JSQ(d(N)) scheme and the ordinary JSQ policy are N-alike.

COROLLARY 5.8. If $d(N)/(\sqrt{N}\log(N)) \to \infty$ as $N \to \infty$, then the JSQ(d(N)) scheme and the ordinary JSQ policy are \sqrt{N} -alike.

Observe that Corollaries 5.7 and 5.8 together with the asymptotic results for the JSQ policy in sections 5.1 and 5.2 imply Theorems 5.3 and 5.4. The rest of the section will be devoted to the proof of Theorem 5.6. The proof relies crucially on a novel coupling construction, which will be used to (lower and upper) bound the difference in occupancy states of two arbitrary schemes.

The Coupling Construction. Throughout the description of the coupling, we fix N and suppress the superscript N in the notation. Let $Q_i^{\Pi_1}(t)$ and $Q_i^{\Pi_2}(t)$ denote the number of server pools with at least i active tasks at time t in two systems following schemes Π_1 and Π_2 , respectively. With a slight abuse of terminology, we occasionally use Π_1 and Π_2 to refer to systems following schemes Π_1 and Π_2 , respectively. To couple the two systems, we synchronize the arrival epochs and maintain a single exponential departure clock with instantaneous rate at time t given by $M(t) := \max \left\{ \sum_{i=1}^B Q_i^{\Pi_1}(t), \sum_{i=1}^B Q_i^{\Pi_2}(t) \right\}$. We couple the arrivals and departures in the various server pools as follows:

Arrival: At each arrival epoch, assign the incoming task in each system to one of the server pools according to the respective schemes.

Departure: Define

$$H(t) := \sum_{i=1}^{B} \min \left\{ Q_i^{\Pi_1}(t), Q_i^{\Pi_2}(t) \right\}$$

and

$$p(t) := \begin{cases} \frac{H(t)}{M(t)} & \text{if } M(t) > 0, \\ 0 & \text{otherwise.} \end{cases}$$

At each departure epoch t_k (say), draw a uniform[0,1] random variable $U(t_k)$. The departures occur in a coupled way based upon the value of $U(t_k)$. In either of the systems, assign an active task index (i,j) if it is the jth task (in order of arrival) of the ith ordered server pool. Let $A_1(t)$ and $A_2(t)$ denote the set of all task indices present at time t in systems Π_1 and Π_2 , respectively. Color the indices (or tasks) in $A_1 \cap A_2$, $A_1 \setminus A_2$, and $A_2 \setminus A_1$ to be green, blue, and red, respectively, and note that $|A_1 \cap A_2| = H(t)$. Define a total order on the set of indices as follows: $(i_1, j_1) < (i_2, j_2)$ if $i_1 < i_2$, or $i_1 = i_2$ and $j_1 < j_2$. Now, if $U(t_k) \le p(t_k-)$, then select one green index uniformly at random and remove the corresponding tasks from both systems. Otherwise, if $U(t_k) > p(t_k-)$, then choose one integer m, uniformly at random from all the integers between 1 and M(t) - H(t) = M(t)(1 - p(t)), and remove the tasks corresponding to the mth smallest (according to the order defined above) red and blue indices in the corresponding systems. If the number of red (or blue) tasks is less than m, then do nothing in the corresponding system.

The above coupling has been schematically represented in Figure 5(a) and will henceforth be referred to as T-coupling, where T stands for "task-based." Now we need to show that, under the T-coupling, the two systems, considered independently, evolve according to their respective marginal statistical laws. This can be seen in several steps. Indeed, the T-coupling basically makes the departure rate uniform by the maximum number of tasks present in either of the two systems. Then, informally speaking, the green region signifies the common portion of tasks, and the red and blue regions represent the separate contributions. Without loss of generality, we assume that $|\mathcal{A}_1| \geq |\mathcal{A}_2|$. We make the following observations:

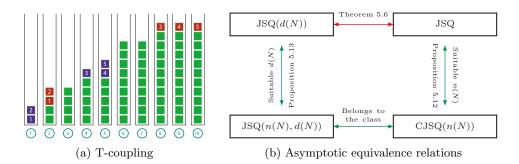


Fig. 5 (a) Superposition of the occupancy states at some particular time instant of schemes Π1 and Π2 when the server pools in both systems are arranged in nondecreasing order of the number of active tasks. The Π1 system is the union of the green and blue tasks, and the Π2 system is the union of the green and red tasks. (b) The equivalence structure is depicted for various intermediate load balancing schemes to facilitate the comparison between the JSQ(d(N)) scheme and the ordinary JSQ policy.

(i) The total departure rate from Π_i is

$$M(t)\left[p(t) + (1-p(t))\frac{|\mathcal{A}_i \setminus \mathcal{A}_{3-i}|}{M(t) - H(t)}\right] = |\mathcal{A}_1 \cap \mathcal{A}_2| + |\mathcal{A}_i \setminus \mathcal{A}_{3-i}| = |\mathcal{A}_i|, \ i = 1, 2.$$

- (ii) Since $|\mathcal{A}_1| \geq |\mathcal{A}_2|$, each task in Π_1 is equally likely to depart.
- (iii) Each task in Π_2 within $\mathcal{A}_1 \cap \mathcal{A}_2$ and each task within $\mathcal{A}_2 \setminus \mathcal{A}_1$ is equally likely to depart, and the probabilities of departures are proportional to $|\mathcal{A}_1 \cap \mathcal{A}_2|$ and $|\mathcal{A}_2 \setminus \mathcal{A}_1|$, respectively.

The T-coupling can be used to derive several stochastic inequality results that will play an instrumental role in proving Theorem 5.6. Recall the CJSQ(n(N)) class of schemes from section 4.1. From a high-level perspective, the proof follows a somewhat similar structure as in section 4.1.

Step I. Condition for g(N)-Alikeness of Schemes in CJSQ(n(N)) Class. The next lemma uses T-coupling to compare the occupancy processes of the JSQ policy with any scheme from the CJSQ(n(N)) class.

LEMMA 5.9. Let $Q_i^{\Pi_1}(t)$ and $Q_i^{\Pi_2}(t)$ denote the number of server pools with at least i tasks in two T-coupled systems under the JSQ policy and a scheme in the CJSQ(n(N)) class, respectively. Then, for any $k \in \{1, 2, ..., B\}$,

$$(5.11) \qquad \left\{ \sum_{i=1}^k Q_i^{\Pi_1}(t) - kn(N) \right\}_{t \ge 0} \le \left\{ \sum_{i=1}^k Q_i^{\Pi_2}(t) \right\}_{t \ge 0} \le \left\{ \sum_{i=1}^k Q_i^{\Pi_1}(t) \right\}_{t \ge 0},$$

provided the two systems start from the same occupancy states at t = 0. In particular, for all $k \ge 1$,

(5.12)
$$\sup_{t \ge 0} |Q_k^{\Pi_2}(t) - Q_k^{\Pi_1}(t)| \le kn(N).$$

Remark 5.10. The stochastic ordering in Lemma 5.9 is to be contrasted with the weak majorization results in [140, 143, 144, 167, 172] in the context of the ordinary JSQ policy in the single-server queuing scenario, and in [82, 111, 112, 139] in the

scenario of state-dependent service rates, nondecreasing with the number of active tasks. In the current infinite-server scenario, the results in [82, 111, 112, 139] imply that for any nonanticipating scheme Π taking assignment decisions based on the number of active tasks only, for all $t \geq 0$,

(5.13)
$$\sum_{m=1}^{\ell} X_{(m)}^{\text{JSQ}}(t) \leq_{st} \sum_{m=1}^{\ell} X_{(m)}^{\Pi}(t) \text{ for } \ell = 1, 2, \dots, N,$$

(5.14)
$$\left\{L^{\text{JSQ}}(t)\right\}_{t>0} \leq_{st} \left\{L^{\Pi}(t)\right\}_{t>0},$$

where $X_{(m)}^{\Pi}(t)$ is the number of tasks in the *m*th ordered server pool at time *t* in the system following scheme Π and $L^{\Pi}(t)$ is the total number of overflow events under policy Π up to time *t*. Observe that $X_{(m)}^{\Pi}$ can be visualized as the *m*th largest (rightmost) vertical bar (or stack) in Figure 3. Thus (5.13) says that the sum of the lengths of the ℓ largest *vertical* stacks in a system following any scheme Π is stochastically larger than or equal to that following the ordinary JSQ policy for any $\ell = 1, 2, \ldots, N$. Mathematically, this ordering can be equivalently written as

(5.15)
$$\sum_{i=1}^{B} \min \{\ell, Q_i^{\text{JSQ}}(t)\} \leq_{st} \sum_{i=1}^{B} \min \{\ell, Q_i^{\Pi}(t)\}$$

for all $\ell = 1, ..., N$. In contrast, in order to show asymptotic equivalence on various scales, we need to both upper and lower bound the occupancy states of the CJSQ(n(N)) schemes in terms of the JSQ policy, and therefore need a much stronger hold on the departure process. The T-coupling provides us just that, and has several useful properties that are crucial for our proof technique. For example, the T-coupling has an important feature that if two systems are T-coupled, then departures cannot increase the sum of the absolute differences of the Q_i -values, which is not true for the coupling considered in the above-mentioned literature. The left stochastic ordering in (5.11) also does not remain valid in those cases. Furthermore, observe that the right inequality in (5.11) (i.e., Q_i 's) implies the stochastic inequality is reversed in (5.15), which is counterintuitive in view of the well-established optimality properties of the ordinary JSQ policy. In the current infinite-server scenario where there is no queuing, this can be understood from the intuition that a better LBA has more tasks in service instead of fewer tasks in the queue. The fundamental distinction between the two coupling techniques is also reflected by the fact that the T-coupling does not allow for arbitrary nondecreasing state-dependent departure rate functions, unlike the couplings in [82, 111, 112, 139].

Remark 5.11 (comparison of T-coupling and S-coupling). As briefly mentioned earlier, in the current infinite-server scenario, the departures of the ordered server pools cannot be coupled, mainly since the departure rate at the mth ordered server pool, for some $m=1,2,\ldots,N$, depends on its number of active tasks. It is worthwhile to mention that the T-coupling in the current section is stronger than the S-coupling used in section 4 in the single-server queuing scenario. Observe that due to Lemma 5.9, the absolute difference of the occupancy states of the JSQ policy and any scheme from the CJSQ class at any time point can be bounded deterministically (without any terms involving the cumulative number of lost tasks). It is worth emphasizing that the universality result on some specific scale, stated in Theorem 5.6, does not depend on the behavior of the JSQ policy on that scale, whereas in the single-server queuing

scenario it does, mainly because the upper and lower bounds in Corollary 4.5 involve tail sums of two different policies. More specifically, in the single-server queuing scenario the fluid and diffusion limit results of $\mathrm{CJSQ}(n(N))$ class crucially use those of the $\mathrm{MJSQ}(n(N))$ scheme, while in the current scenario it does not—the results for the $\mathrm{MJSQ}(n(N))$ scheme come as a consequence of those for the $\mathrm{CJSQ}(n(N))$ class of schemes. Also, the bounds in Lemma 5.9 do not depend on t, and hence apply in the steady state as well. Moreover, the S-coupling compares the k highest horizontal bars, whereas the T-coupling in the current section compares the k lowest horizontal bars. As a result, the bounds on the occupancy states established in Corollary 4.5 involve tail sums of the occupancy states of the ordinary JSQ policy, which necessitates proving the convergence of tail sums of the occupancy states of the ordinary JSQ policy. In contrast, the bound in the infinite-server scenario involves only a single component (see (5.12)), thus proving convergence of each component suffices.

The goal in the first step is to show that for a suitable choice of n(N), the schemes in the $\mathrm{CJSQ}(n(N))$ class are indistinguishable on suitable scales. This is formalized in Proposition 5.12 below, which follows immediately from Lemma 5.9.

PROPOSITION 5.12. For any function $g: \mathbb{N} \to \mathbb{R}_+$ diverging to infinity, if $n(N)/g(N) \to 0$ as $N \to \infty$, then the JSQ policy and the CJSQ(n(N)) schemes are g(N)-alike.

Step 2. g(N)-Alikeness of JSQ(d(N)) and a Scheme in CJSQ(n(N)). Next we compare the CJSQ(n(N)) schemes with the JSQ(d(N)) scheme. The comparison follows a somewhat similar line of argument as in section 4.1 and involves a JSQ(n(N), d(N)) scheme which is an intermediate blend between the CJSQ(n(N)) schemes and the JSQ(d(N)) scheme. Specifically, the JSQ(d(N)) scheme selects a candidate server pool in the same way as the JSQ(d(N)) scheme. However, it only assigns the task to that server pool if it belongs to the n(N) + 1 lowest-ordered ones, and to a randomly selected server pool among them otherwise. By construction, the JSQ(n(N), d(N)) scheme belongs to the class CJSQ(n(N)).

The next proposition establishes a sufficient criterion on d(N) in order for the JSQ(d(N)) scheme and the JSQ(n(N),d(N)) scheme to be close in terms of g(N)-alikeness.

PROPOSITION 5.13. Assume $n(N)/g(N) \to 0$ as $N \to \infty$ for some function $g: \mathbb{N} \to \mathbb{R}_+$ diverging to infinity. The JSQ(n(N), d(N)) scheme and the JSQ(d(N)) scheme are g(N)-alike if the following condition holds:

(5.16)
$$\frac{n(N)}{N}d(N) - \log \frac{N}{q(N)} \to \infty \quad as \quad N \to \infty.$$

Finally, Proposition 5.13 in conjunction with Proposition 5.12 yields Theorem 5.6. The overall proof strategy as described above is schematically represented in Figure 5(b).

6. Load Balancing in Graph Topologies. In this section we return to the single-server queuing dynamics and extend the universality properties to network scenarios, where the N servers are assumed to be interconnected by some underlying graph topology G_N . Tasks arrive at the various servers as independent Poisson processes of rate λ , and each incoming task is assigned to whichever server has the smallest number of tasks among the one where it arrives and its neighbors in G_N . Ties are broken arbitrarily. Thus, in the case that G_N is a clique, each incoming task is assigned

to the server with the shortest queue across the entire system, and the behavior is equivalent to that under the JSQ policy. The stochastic optimality properties of the JSQ policy thus imply that the queue length process in a clique will be better balanced and smaller (in a majorization sense) than in an arbitrary graph G_N .

As stated in the introduction, network scenarios are not only of mathematical interest, but also of major relevance from an application perspective. For example, they emerge in modeling connectivity properties, geographic restrictions, and proximity relations in spatial network settings. Besides capturing such physical concepts in infrastructure networks, network scenarios also arise due to "logical relationships," in particular, so-called affinity notions and compatibility constraints between tasks and servers. Such features are increasingly common in data centers and cloud networks due to heterogeneity and data locality issues; see, for instance, [136, 169]. They also relate to the scalability considerations that are important in load balancing, as further explained below.

Sparse Graph Topologies. Besides the prohibitive communication overhead discussed earlier, a further scalability issue of the JSQ policy arises when executing a task that involves the use of some data. Storing such data for all possible tasks on all servers will typically require an excessive amount of storage capacity. These two burdens can be effectively mitigated in sparser graph topologies where tasks that arrive at a specific server i are only allowed to be forwarded to a subset of the servers \mathcal{N}_i . For the tasks that arrive at server i, queue length information then only needs to be obtained from servers in \mathcal{N}_i , and it suffices to store replicas of the required data on the servers in \mathcal{N}_i . The subset \mathcal{N}_i containing the peers of server i can be naturally viewed as its neighbors in some graph topology G_N . Here we consider the case of undirected graphs, but most of the analysis can be extended to directed graphs.

While sparser graph topologies relieve the scalability issues associated with a clique, the queue length process will be worse (in the majorization sense) because of the limited connectivity. Surprisingly, however, even quite sparse graphs can asymptotically match the optimal performance of a clique, provided they are suitably random, as we will further describe below.

The above model has been studied in [61, 147], focusing on certain fixed-degree graphs and, in particular, the ring topologies for which [113] has already presented simulation results. The results demonstrate that the flexibility to forward tasks to a few neighbors, or even just one, with possibly shorter queues significantly improves the performance in terms of the waiting time and tail distribution of the queue length. This resembles the *power-of-choice* gains observed for JSQ(d) policies in complete graphs.

However, the results in [61, 147] also establish that the performance sensitively depends on the underlying graph topology, and that selecting from a fixed set of d-1 neighbors typically does not match the performance of resampling d-1 alternate servers for each incoming task from the entire population, as in the power-of-d scheme in a complete graph. Further interesting results for the performance of LBAs in a network context, with a focus on tail asymptotics, may be found in [38, 110].

Supermarket Model on Graphs. For the case when each arriving task is routed to the shortest of $d \geq 2$ randomly selected neighboring queues, the process-level convergence over any finite time interval has been established recently in [28]. In this work, the authors analyze the evolution of the queue length process at an arbitrary tagged server as the system size becomes large. The main ingredient is a careful analysis of local occupancy measures associated with the neighborhood of each server

and the argument that under suitable conditions their asymptotic behavior is the same for all servers. Under mild conditions on the graph topology G_N (diverging minimum degree and the ratio between minimum degree and maximum degree in each connected component converges to 1), for a suitable initial occupancy measure, [28, Theorem 2.1] establishes that for any fixed $d \geq 2$, the global occupancy state process for the JSQ(d) scheme on G_N has the same weak limit in (3.2) as that on a clique, as the number of vertices N becomes large. Further, a propagation of chaos property was shown to hold for this system, in the sense that the queue lengths at any finite collection of tagged servers are asymptotically independent, and the queue length process for each server converges in distribution (in the path space) to a certain McKean-Vlasov process [28, Theorem 2.2]. Furthermore, when the graph sequence is random, with the Nth graph given as an Erdős-Rényi random graph (ERRG) on N vertices with average degree d(N), note that there are two types of randomness that drive the dynamics of the process, one being the randomness of the underlying graph and the other being the randomness of the arrival/departure processes given the graph. This setup falls under the framework of random processes in a random environment. Here one is typically interested in two types of convergence results: (1) annealed convergence, where one looks at the dynamics of the sequence of occupancy processes averaged over the randomness of the underlying graph, and (2) quenched convergence, where one samples a sequence of random graphs with increasing N and, given that sequence, considers the dynamics of the sequence of occupancy process. In [28] annealed convergence is established under the condition $d(N) \to \infty$, and the quenched convergence is shown under a stronger condition $d(N)/\log N \to \infty$.

Asymptotic Optimality on Graphs. We return to the case when each incoming task is assigned to whichever server has the smallest number of tasks among the one where it arrives and its neighbors in G_N . The results presented in the remainder of the section are based on [116], where full proofs are also provided, unless indicated otherwise. As mentioned earlier, the queue length process in a clique will be better balanced and smaller (in a majorization sense) than in an arbitrary graph G_N . Accordingly, a graph G_N is said to be N-optimal or \sqrt{N} -optimal when the queue length process on G_N is equivalent to that on a clique on an N-scale or \sqrt{N} -scale, respectively. Roughly speaking, a graph is N-optimal if the fraction of nodes with i tasks for $i=0,1,\ldots$ behaves as in a clique as $N\to\infty$. The fluid-limit results for the JSQ policy discussed in section 3.3 imply that the latter fraction is zero in the limit for all i > 2 in a clique in stationarity, i.e., the fraction of servers with two or more tasks vanishes in any graph that is N-optimal, and consequently the mean waiting time vanishes as well as $N \to \infty$. Furthermore, the diffusion-limit results of [36] for the JSQ policy discussed in section 3.4 imply that the number of nodes with zero tasks and the number with two tasks both scale as \sqrt{N} as $N \to \infty$. Again, loosely speaking, a graph is \sqrt{N} -optimal if in the heavy-traffic regime the number of nodes with zero tasks and the number with two tasks when scaled by \sqrt{N} both evolve as in a clique as $N \to \infty$. Formal definitions of asymptotic optimality on an N-scale or \sqrt{N} -scale will be introduced in Definition 6.1 below.

As one of the main results, we will demonstrate that, remarkably, asymptotic optimality can be achieved in quite sparse ERRGs. We prove that a sequence of ERRGs indexed by the number of vertices N with $d(N) \to \infty$ as $N \to \infty$ is N-optimal. We further establish that the latter growth condition for the average degree is in fact necessary in the sense that any graph sequence that contains $\Theta(N)$ bounded-degree vertices cannot be N-optimal. This implies that a sequence of ERRGs with finite

average degree cannot be N-optimal. The growth rate condition is more stringent for optimality on \sqrt{N} -scale in the heavy-traffic regime. Specifically, we prove that a sequence of ERRGs indexed by the number of vertices N with $d(N)/(\sqrt{N}\log(N)) \to \infty$ as $N \to \infty$ is \sqrt{N} -optimal.

The above results demonstrate that the asymptotic optimality of cliques on an N-scale and a \sqrt{N} -scale can be achieved in far sparser graphs, where the number of connections is reduced by nearly a factor N and $\sqrt{N}/\log(N)$, respectively, provided the topologies are suitably random in the ERRG sense. This translates into equally significant reductions in communication overhead and storage capacity, since both are roughly proportional to the number of connections.

Arbitrary Graph Topologies. The key challenge in the analysis of load balancing on arbitrary graph topologies is that one needs to keep track of the evolution of the number of tasks at each vertex along with their corresponding neighborhood relationship. This creates a major problem in constructing a tractable Markovian state descriptor and renders a direct analysis of such processes highly intractable, as already alluded to in [113]. Consequently, even asymptotic results for load balancing processes on an arbitrary graph have remained scarce so far. We take a radically different approach and aim to compare the load balancing process on an arbitrary graph with that on a clique. Specifically, rather than analyze the behavior for a given class of graphs or degree value, we explore the types of topologies and degree properties for which the performance is asymptotically similar to that in a clique.

Stochastic Coupling for Graphs. Our proof arguments build on the stochastic coupling constructions developed in section 4 for JSQ(d) policies. Specifically, we view the load balancing process on an arbitrary graph as a "sloppy" version of that on a clique, and thus construct several other intermediate sloppy versions. By constructing novel couplings, we develop a method of comparing the load balancing process on an arbitrary graph with that on a clique. In particular, we bound the difference between the fraction of vertices with i or more tasks in the two systems for $i = 1, 2, \ldots$, to obtain asymptotic optimality results. From a high-level viewpoint, conceptually related graph conditions for asymptotic optimality were examined using quite different techniques in [145, 146] in a dynamic scheduling framework (as opposed to the load balancing context).

Notation. For $k=1,\ldots,N$, denote by $X_k(G_N,t)$ the queue length at the kth server at time t (including the task possibly in service), and by $X_{(k)}(G_N,t)$ the queue length at the kth ordered server at time t when the servers are arranged in nondecreasing order of their queue lengths (ties can be broken in some way that will be evident from the context). Let $Q_i(G_N,t)$ denote the number of servers with queue length at least i at time t and $q_i(G_N,t) = Q_i(G_N,t)/N$, $i=1,2,\ldots$ It is important to note that $(q_i(G_N,t))_{i\geq 1}$ is itself not a Markov process. Given the graph G_N , the queue length process $(X_k(G_N,t))_{k=1}^N$ is Markovian under the model assumptions and $(q_i(G_N,t)_{i\geq 1})$ is a function of $(X_k(G_N,t))_{k=1}^N$. Also, in the Halfin–Whitt heavy-traffic regime (2.1), define the centered and scaled processes

(6.1)
$$\bar{Q}_1(G_N, t) = -\frac{N - Q_1(G_N, t)}{\sqrt{N}}, \quad \bar{Q}_i(G_N, t) = \frac{Q_i(G_N, t)}{\sqrt{N}},$$

analogously to (3.1).

Asymptotic Optimality. In general, the optimality of the clique topology is strict, but it turns out that near-optimality can be achieved asymptotically in a broad

class of other graph topologies. Therefore, we now introduce two notions of asymptotic optimality that will be useful to characterize the performance in large-scale systems.

DEFINITION 6.1 (asymptotic optimality). A graph sequence $\mathbf{G} = \{G_N\}_{N\geq 1}$ is called "asymptotically optimal on N-scale" or "N-optimal" if, for any $\lambda < 1$, the process $(q_1(G_N, \cdot), q_2(G_N, \cdot), \ldots)$ converges weakly, on any finite time interval, to a process $(q_1(\cdot), q_2(\cdot), \ldots)$ satisfying (3.5).

Moreover, a graph sequence $\mathbf{G} = \{G_N\}_{N\geq 1}$ is called "asymptotically optimal on \sqrt{N} -scale" or " \sqrt{N} -optimal" if, in the Halfin-Whitt heavy-traffic regime (2.1), on any finite time interval, the process $(\bar{Q}_1(G_N,\cdot),\bar{Q}_2(G_N,\cdot),\ldots)$ as in (6.1) converges weakly to the process $(\bar{Q}_1(\cdot),\bar{Q}_2(\cdot),\ldots)$ given by (3.7).

Intuitively speaking, if a graph sequence is N-optimal or \sqrt{N} -optimal, then in some sense, the associated occupancy processes are indistinguishable from those of the sequence of cliques on N-scale or \sqrt{N} -scale. In other words, on any finite time interval their occupancy processes can differ from those in cliques by at most o(N) or $o(\sqrt{N})$, respectively. We will interchangeably use the terms fluid scale and diffusion scale to refer to N-scale and \sqrt{N} -scale, respectively. In particular, exploiting interchange of the stationary $(t \to \infty)$ and many-server $(N \to \infty)$ limits, we obtain that for any N-optimal graph sequence $\{G_N\}_{N\geq 1}$, as $N\to\infty$,

$$q_1(G_N, \infty) \to \lambda$$
 and $q_i(G_N, \infty) \to 0$ for all $i = 2, \ldots, B$,

implying that the stationary fraction of servers with queue length two or larger and the mean waiting time vanish. It is worthwhile to point out that the above interchange of limits requires the ergodicity of the queue length process for each fixed N, a certain tightness of the sequence $\{(q_1(G_N, \infty), q_2(G_N, \infty), \ldots)\}_{N\geq 1}$, and the global stability of the fluid limits.

6.1. Asymptotic Optimality Criteria for Deterministic Graph Sequences. We now proceed to develop a criterion for asymptotic optimality of an arbitrary deterministic graph sequence on different scales. Next this criterion will be leveraged to establish optimality of a sequence of random graphs. We start by introducing some useful notation and two measures of well-connectedness. Let G = (V, E) be any graph. For a subset $U \subseteq V$, define $COM(U) := |V \setminus N[U]|$ to be the cardinality of the set of all vertices that are disjoint from U and its immediate neighbors, where $N[U] := U \cup \{v \in V : \exists u \in U \text{ with } (u, v) \in E\}$. For any fixed $\varepsilon > 0$ define

$$(6.2) \quad \operatorname{dis}_1(G,\varepsilon) := \sup_{U \subseteq V, |U| \ge \varepsilon |V|} \operatorname{com}(U), \qquad \operatorname{dis}_2(G,\varepsilon) := \sup_{U \subseteq V, |U| \ge \varepsilon \sqrt{|V|}} \operatorname{com}(U).$$

The next theorem provides sufficient conditions for asymptotic optimality on N-scale and \sqrt{N} -scale in terms of the above two well-connectedness measures.

THEOREM 6.2. For any graph sequence $G = \{G_N\}_{N \geq 1}$,

- (i) **G** is N-optimal if, for any $\varepsilon > 0$, DIS₁ $(G_N, \varepsilon)/N \to 0$ as $N \to \infty$;
- (ii) **G** is \sqrt{N} -optimal if, for any $\varepsilon > 0$, $DIS_2(G_N, \varepsilon)/\sqrt{N} \to 0$ as $N \to \infty$.

From a high-level perspective, the conditions in Theorem 6.2 (i) and (ii) require that neighborhoods of any $\Theta(N)$ and $\Theta(\sqrt{N})$ vertices contain at least $N-\mathrm{o}(N)$ and $N-\mathrm{o}(\sqrt{N})$ vertices, respectively. As we will see below in Theorems 6.8 and 6.10, the conditions in Theorem 6.2 impose suitable levels of connectivity in the graph topology in order for it to be asymptotically optimal on fluid and diffusion scales, while significantly reducing the total number of connections. The next corollary is an immediate consequence of Theorem 6.2.

COROLLARY 6.3. Let $\mathbf{G} = \{G_N\}_{N \geq 1}$ be any graph sequence. Then (i) if the minimum degree in G_N equals $N - \mathrm{o}(N)$, then \mathbf{G} is N-optimal, and (ii) if the minimum degree in G_N equals $N - \mathrm{o}(\sqrt{N})$, then \mathbf{G} is \sqrt{N} -optimal.

The rest of the subsection is devoted to a discussion of the main proof arguments for Theorem 6.2, focusing on the proof of N-optimality. The proof of \sqrt{N} -optimality follows along similar lines. We establish in Proposition 6.4 that if a system is able to assign each task to a server in the set, denoted by $\mathcal{S}^N(n(N))$, of the n(N)+1 nodes with shortest queues, where n(N) is o(N), then it is N-optimal. Since the underlying graph is not a clique, however (otherwise there is nothing to prove), for any n(N) not every arriving task can be assigned to a server in $\mathcal{S}^N(n(N))$. Hence we further prove in Proposition 6.5 a stochastic comparison property implying that if on any finite time interval of length t, the number of tasks $\Delta^N(t)$ that are not assigned to a server in $\mathcal{S}^N(n(N))$ is $o_P(N)$, then the system is N-optimal as well. The N-optimality can then be concluded when $\Delta^N(t)$ is $o_P(N)$, which we establish in Proposition 6.6 under the condition that $\mathrm{DIS}_1(G_N,\varepsilon)/N\to 0$ as $N\to\infty$ as stated in Theorem 6.2.

To further explain the idea described in the above proof outline, it is useful to adopt a slightly different point of view toward load balancing processes on graphs. From a high-level viewpoint, a load balancing process can be thought of as follows: there are N servers, which are assigned incoming tasks by some scheme. The assignment scheme can arise from some topological structure, in which case we will call it topological load balancing, or it can arise from some other property of the occupancy process, in which case we will call it nontopological load balancing. As mentioned earlier, the JSQ policy on the clique is optimal among the set of all nonanticipating schemes, irrespective of being topological or nontopological. Also, load balancing on graph topologies other than a clique can be thought of as a "sloppy" version of that on a clique, when each server only has access to partial information on the occupancy state. Below we first introduce a different type of sloppiness in the task assignment scheme and show that under a limited amount of sloppiness, optimality is retained on a suitable scale. Next we will construct a scheme, which is a hybrid of topological and nontopological schemes, whose behavior is simultaneously close to both the load balancing process on a suitable graph and that on a clique.

A Class of Sloppy Load Balancing Schemes. Fix some function $n: \mathbb{N} \to \mathbb{N}$ and recall the set $\mathcal{S}^N(n(N))$ from before as well as the class $\mathrm{CJSQ}(n(N))$ from section 4.1, where each arriving task is assigned to one of the servers in $\mathcal{S}^N(n(N))$. It should be emphasized that for any scheme in $\mathrm{CJSQ}(n(N))$, we are not imposing any restrictions on how the incoming task should be assigned to a server in $\mathcal{S}^N(n(N))$. The scheme only needs to ensure that the arriving task is assigned to some server in $\mathcal{S}^N(n(N))$ with respect to some tie-breaking mechanism. Observe that using Corollary 4.5 and following arguments as in the proofs of Theorems 4.1 and 4.2, we obtain the next proposition, which provides a sufficient criterion for asymptotic optimality of any scheme in $\mathrm{CJSQ}(n(N))$.

PROPOSITION 6.4. For $0 \le n(N) < N$, let $\Pi \in \mathrm{CJSQ}(n(N))$ be any scheme. (i) If $n(N)/N \to 0$ as $N \to \infty$, then Π is N-optimal. (ii) If $n(N)/\sqrt{N} \to 0$ as $N \to \infty$, then Π is \sqrt{N} -optimal.

A Bridge between Topological and Nontopological Load Balancing. For any graph G_N and $n \leq N$, we first construct a scheme called $I(G_N, n)$, which is an intermediate blend between the topological load balancing process on G_N and some

kind of nontopological load balancing on N servers. The choice of n = n(N) will be clear from the context.

To describe the scheme $I(G_N, n)$, first synchronize the arrival epochs at server v in both systems, v = 1, 2, ..., N. Further, synchronize the departure epochs at the kth ordered server with the kth smallest number of tasks in the two systems, k = 1, 2, ..., N. When a task arrives at server v at time t, say, it is assigned in the graph G_N to a server $v' \in N[v]$ according to its own statistical law. For the assignment under the scheme $I(G_N, n)$, first observe that if

(6.3)
$$\min_{u \in N[v]} X_u(G_N, t) \le \max_{u \in S(n)} X_u(G_N, t),$$

then there exists some tie-breaking mechanism for which $v' \in N[v]$ belongs to $\mathcal{S}(n)$ under G_N . Pick such an ordering of the servers, and assume that v' is the kth ordered server in that ordering for some $k \leq n+1$. Under $I(G_N, n)$ assign the arriving task to the kth ordered server (breaking ties arbitrarily in this case). Otherwise, if (6.3) does not hold, then the task is assigned to one of the n+1 servers with minimum queue lengths under G_N uniformly at random.

Denote by $\Delta^N(I(G_N, n), T)$ the cumulative number of arriving tasks up to time $T \geq 0$ for which (6.3) is violated under the above coupling. The next proposition shows that the load balancing process under the scheme $I(G_N, n)$ is close to that on the graph G_N in terms of the random variable $\Delta^N(I(G_N, n), T)$.

Proposition 6.5. The inequality

(6.4)
$$\sum_{i=1}^{B} |Q_i(G_N, t) - Q_i(I(G_N, n), t)| \le 2\Delta^N(I(G_N, n), t) \qquad \forall \ t \ge 0$$

is preserved almost surely, provided the two systems start from the same occupancy state at t = 0.

In order to conclude optimality on N-scale or \sqrt{N} -scale, it remains to be shown that the term $\Delta^N(I(G_N,n),T)$ is sufficiently small. The next proposition provides suitable asymptotic bounds for $\Delta^N(I(G_N,n),T)$ under the conditions on $\mathrm{DIS}_1(G_N,\varepsilon)$ and $\mathrm{DIS}_2(G_N,\varepsilon)$ stated in Theorem 6.2. For N-optimality, the idea is that since for all $\varepsilon>0$, $\mathrm{DIS}_1(G_N,\varepsilon)$ is $\mathrm{o}(N)$, one can show that there is a number $n_\varepsilon(N)=\mathrm{o}(N)$ such that $\mathrm{COM}(U)=\mathrm{o}(N)$ uniformly over all $U\subseteq V_N$ with $|U|\geq n_\varepsilon(N)$. Consequently, this can be used to show that on any finite time interval, "most of the tasks" will be assigned to one of the $n_\varepsilon(N)$ servers with smallest queue lengths. This enables us to couple the system with a scheme from the class $\mathrm{CJSQ}(n_\varepsilon(N))$. The idea is similar when we consider \sqrt{N} -optimality.

Proposition 6.6.

(i) For any $\varepsilon > 0$, there exist $\varepsilon' > 0$ and $n_{\varepsilon'}(N)$ with $n_{\varepsilon'}(N)/N \to 0$ as $N \to \infty$, such that if $\text{DIS}_1(G_N, \varepsilon')/N \to 0$ as $N \to \infty$, then for all T > 0,

$$\mathbb{P}\left(\Delta^N(I(G_N,n_{\varepsilon'}),T)/N>\varepsilon\right)\to 0.$$

(ii) For any $\varepsilon > 0$, there exist $\varepsilon' > 0$ and $m_{\varepsilon'}(N)$ with $m_{\varepsilon'}(N)/\sqrt{N} \to 0$ as $N \to \infty$, such that if $DIS_2(G_N, \varepsilon')/\sqrt{N} \to 0$ as $N \to \infty$, then for all T > 0,

$$\mathbb{P}\left(\Delta^N(I(G_N, m_{\varepsilon'}), T) / \sqrt{N} > \varepsilon\right) \to 0.$$

The proof of Theorem 6.2 then readily follows by combining Propositions 6.4–6.6 and observing that the scheme $I(G_N, n)$ belongs to the class CJSQ(n) by construction.

From the conditions of Theorem 6.2, it follows that if for all $\varepsilon > 0$, $DIS_1(G_N, \varepsilon)$ and $\mathrm{DIS}_2(G_N,\varepsilon)$ are $\mathrm{o}(N)$ and $\mathrm{o}(\sqrt{N})$, respectively, then the total number of edges in G_N must be $\omega(N)$ and $\omega(N\sqrt{N})$, respectively. Theorem 6.7 states that the superlinear growth rate of the total number of edges is not only sufficient, but also necessary in the sense that any graph with O(N) edges is asymptotically suboptimal on N-scale.

THEOREM 6.7. Let $G = \{G_N\}_{N\geq 1}$ be any graph sequence, such that there exists a fixed integer $M < \infty$ with

(6.5)
$$\limsup_{N \to \infty} \frac{\#\{v \in V_N : d_v \le M\}}{N} > 0,$$

where d_v is the degree of the vertex v. Then G is suboptimal on N-scale.

To prove Theorem 6.7, we show that starting from an all-empty state, in finite time, a positive fraction of servers in G_N will have at least two tasks. This establishes that the occupancy processes when scaled by N cannot agree with those in the sequence of cliques, and hence $\{G_N\}_{N\geq 1}$ cannot be N-optimal. The idea of the proof can be explained as follows: If a system contains $\Theta(N)$ bounded-degree vertices, then starting from an all-empty state, in any finite time interval there will be $\Theta(N)$ servers u, say, for which all the servers in N[u] have at least one task. For all such servers an arrival at u must produce a server with queue length two. It follows that the instantaneous rate at which servers of queue length two are formed is bounded away from zero, and hence $\Theta(N)$ servers of queue length two are produced in finite time.

6.2. Asymptotic Optimality of Random Graph Sequences. Next we investigate how the load balancing process behaves on random graph topologies. Specifically, we aim to understand what types of graphs are asymptotically optimal in the presence of randomness (i.e., in an average-case sense). Theorem 6.8 establishes sufficient conditions for asymptotic optimality of a sequence of inhomogeneous random graphs. Recall that a graph G' = (V', E') is called a supergraph of G = (V, E) if V = V' and $E \subseteq E'$.

THEOREM 6.8. Let $G = \{G_N\}_{N\geq 1}$ be a graph sequence such that, for each N, $G_N = (V_N, E_N)$ is a supergraph of the inhomogeneous random graph G'_N where any two vertices $u, v \in V_N$ share an edge with probability p_{uv}^N , independently of each other.

(i) If $\inf \{p_{uv}^N : u, v \in V_N\}$ is $\omega(1/N)$, then \mathbf{G} is N-optimal.

(ii) If $\inf \{p_{uv}^N : u, v \in V_N\}$ is $\omega(\log(N)/\sqrt{N})$, then \mathbf{G} is \sqrt{N} -optimal.

The proof of Theorem 6.8 relies on Theorem 6.2. Specifically, if G_N satisfies conditions (i) and (ii) in Theorem 6.8, then the corresponding conditions (i) and (ii) in Theorem 6.2 hold.

As an immediate corollary of Theorem 6.8 we obtain an optimality result for the sequence of ERRGs. Let ERRG(N, p(N)) denote a graph on N vertices, such that any pair of vertices share an edge with probability p(N).

COROLLARY 6.9. Let $G = \{G_N\}_{N \geq 1}$ be a graph sequence such that, for each N, G_N is a supergraph of ERRG(N, p(N)) and d(N) = (N-1)p(N). Then (i) if $d(N) \to \infty$ ∞ as $N \to \infty$, then **G** is N-optimal; (ii) if $d(N)/(\sqrt{N}\log(N)) \to \infty$ as $N \to \infty$, then **G** is \sqrt{N} -optimal.

Theorem 6.2 can be further leveraged to establish the optimality of the following sequence of random graphs. For any $N \geq 1$ and $d(N) \leq N-1$ such that Nd(N) is even, construct the erased random regular graph on N vertices as follows: Initially, attach d(N) half-edges to each vertex. Call all such half-edges unpaired. At each step, pick one half-edge arbitrarily and pair it to another half-edge uniformly at random among all unpaired half-edges to form an edge, until all the half-edges have been paired. Thus, note that there can be more than one edge between two vertices (i.e., multiedge) or a half-edge of a vertex can be paired with another half-edge of the same vertex (self-loops). Such a graph is known as a regular multigraph. In fact, it is known [153, Proposition 7.7] that the above pairing procedure results in a random graph that has a uniform distribution over all regular multigraphs with degree d(N). Now the erased random regular graph is formed by erasing all the self-loops and collapsing the multiple edges to a single edge, which thus produces a simple graph.

Theorem 6.10. Let $\mathbf{G} = \{G_N\}_{N\geq 1}$ be a sequence of erased random regular graphs with degree d(N). Then (i) if $d(N) \to \infty$ as $N \to \infty$, then \mathbf{G} is N-optimal; (ii) if $d(N)/(\sqrt{N}\log(N)) \to \infty$ as $N \to \infty$, then \mathbf{G} is \sqrt{N} -optimal.

Note that, due to Theorem 6.7, we can conclude that the growth rate condition for N-optimality in Corollary 6.9 (i) and Theorem 6.10 (i) is not only sufficient, but necessary as well. Thus, informally speaking, N-optimality is achieved under the minimum condition required as long as the underlying topology is suitably random.

- 7. Token-Based Load Balancing. While a zero waiting time can be achieved in the limit by sampling only d(N) = o(N) servers, as sections 4 and 6 showed, even in network scenarios, the amount of communication overhead in terms of d(N) must still grow with N. As mentioned earlier, this can be avoided by introducing memory at the dispatcher, in particular, maintaining a record of only vacant servers and assigning tasks to idle servers, if there are any, or to a uniformly at random selected server otherwise. This so-called join-the-idle-queue (JIQ) scheme [11, 101] can be implemented through a simple token-based mechanism generating at most one message per task. Remarkably enough, even with such low communication overhead, the mean waiting time and the probability of a nonzero waiting time vanish under the JIQ scheme in both the fluid and diffusion regimes, as we will discuss in the next two subsections. It is worth emphasizing, though, that the JIQ scheme is *not* optimal in the nondegenerate slow-down regime, which was introduced in section 2.2 and will be further discussed in section 8.3.
- **7.1. Fluid-Level Optimality of JIQ Scheme.** We first consider the fluid limit of the JIQ policy. It is not hard to show that the number of busy servers under the JIQ scheme is stochastically larger (in the path space) than that for the JSQ(1) policy (tasks assigned uniformly at random). Consequently, the JIQ scheme is stable whenever $\lambda < 1$. Recall that $q_i^N(\infty)$ is a random variable denoting the process $q_i^N(\cdot)$ in steady state. Under significantly more general conditions (in the presence of finitely many heterogeneous server pools and for general service time distributions with decreasing hazard rate) it was proved in [141] that under the JIQ scheme,

$$(7.1) \qquad q_1^N(\infty) \to \lambda, \qquad q_i^N(\infty) \to 0 \quad \text{ for all } i \geq 2, \qquad \text{as} \quad N \to \infty.$$

The above equation in conjunction with the PASTA property yields that the steady-state probability of a nonzero wait vanishes as $N \to \infty$, thus exhibiting asymptotic optimality of the JIQ policy on fluid scale.

High-Level Outline of Proof Idea. Loosely speaking, the proof of (7.1) consists of three principal components:

(i) Starting from an all-empty state, the asymptotic rate of increase of q_1 is given by the arrival rate λ . Also, the rate of decrease is q_1 . Thus, on a small time interval dt, the rate of change of q_1 is given by

(7.2)
$$\frac{\mathrm{d}q_1(t)}{\mathrm{d}t} = \lambda - q_1(t).$$

Under the above dynamics, the system occupancy states converge to the unique fixed point of the above ODE, given by the point $(\lambda, 0, 0, ...)$.

- (ii) The occupancy process is monotone, in the sense that (a) starting from an all-empty state, the occupancy process is componentwise stochastically nondecreasing in time (in the sense of stochastic dominance), and (b) the occupancy process at any fixed time t, starting from an arbitrary state, componentwise stochastically dominates the occupancy process at time t, starting from an allempty state.
- (iii) Under the JIQ scheme, the system is stable, and hence the occupancy process is ergodic. Since $q_1(t)$ is the instantaneous rate of departure from the system, ergodicity implies that in steady state there can be $at \ most \ \lambda$ fraction of busy servers (containing at least one task). In fact, it further establishes that the steady-state fraction of servers with more than one tasks vanishes asymptotically.

Points (i) and (ii) above imply that starting from any state the system must have at least λ fraction of busy servers, and finally this along with point (iii) establishes that the steady-state occupancy process must converge to $(\lambda, 0, 0, \ldots)$.

7.2. Diffusion-Level Optimality of JIQ Scheme. We now turn to the diffusion limit of the JIQ scheme established in [118]. Recall the centered and scaled occupancy process as in (3.1) and the Halfin-Whitt heavy-traffic regime in (2.1).

Theorem 7.1 (diffusion limit for JIQ). Assume that $\lambda(N)$ satisfies (2.1). Under suitable initial conditions the weak limit of the sequence of centered and diffusion-scaled occupancy processes in (3.1) coincides with that of the ordinary JSQ policy and, in particular, is given by the system of SDEs in (3.7).

The above theorem implies that for suitable states, on any finite time interval, the occupancy process of a system under the JIQ policy is indistinguishable from that under the JSQ policy.

High-Level Outline of Proof Idea. The proof of Theorem 7.1 relies on a novel coupling construction introduced in [118] and described below in detail. The idea is to compare the occupancy processes of two systems following JIQ and JSQ policies, respectively. Comparing the JIQ and JSQ policies is facilitated when viewed as follows: (i) If there is an idle server in the system, both JIQ and JSQ perform similarly. (ii) Also, when there is no idle server and only $O(\sqrt{N})$ servers with queue length two or more, JSQ assigns the arriving task to a server with queue length one. In that case, since JIQ assigns at random, the probability that the task will land on a server with queue length two or more, and thus JIQ acts differently than JSQ, is $O(1/\sqrt{N})$. Since on any finite time interval the number of times an arrival finds all servers busy is at most $O(\sqrt{N})$, all the arrivals except O(1) of them are assigned in exactly the same manner in both JIQ and JSQ, which then leads to the same scaling limit for both policies.

The diffusion limit result in Theorem 7.1 is in fact true for an even broader class of load balancing schemes. As in section 4.1, let B denote the buffer capacity (possibly infinite) of each server, and in the case $B < \infty$, if a task is assigned to a server with B outstanding tasks, it is instantly discarded. For an LBA Π , we will denote the total number of tasks lost up to time t by $L^{\Pi}(t)$. Define the class of schemes

$$\Pi^{(N)} := \{ \Pi(d_0, d_1, \dots, d_{B-1}) : d_0 = N, 1 \le d_i \le N, 1 \le i \le B - 1, B \ge 2 \},$$

where in the scheme $\Pi(d_0, d_1, \ldots, d_{B-1})$ with buffer capacity B, the dispatcher assigns an incoming task to the server with the minimum queue length among d_k (possibly function of N) servers selected uniformly at random when the minimum queue length across the system is $k, k = 0, 1, \ldots, B-1$. The system analyzed in [36] (JSQ with B=2) can be written as $\Pi(N,N)$, JIQ can be expressed as $\Pi(N,1,1,\ldots)$, and JIQ with a buffer capacity B=2 is $\Pi(N,1)$.

The crux of the argument in proving diffusion-level optimality for any scheme in $\Pi^{(N)}$ goes as follows: First, the occupancy process under the scheme $\Pi(N, d_1, \ldots, d_{B-1})$ is sandwiched between those under $\Pi(N, 1)$ and $\Pi(N, d_1)$. More specifically, the ℓ_1 -distance between the occupancy processes under $\Pi(N, d_1, \ldots, d_{B-1})$ and $\Pi(N, 1)$ is bounded by the number of items lost due to full buffers. Next, this loss is bounded using the number of servers with queue length 2 in $\Pi(N, N)$. This allows the use of the results in [36] and yields that on any finite time interval with high probability an O(1) number of items is lost due to full buffers, which is negligible on \sqrt{N} scale. Specifically, this shows that for suitable initial states, the schemes $\Pi(N, 1)$ and $\Pi(N, d_1)$, along with any scheme in the class $\Pi^{(N)}$, have the same diffusion limit in the Halfin–Whitt heavy-traffic regime. We conclude this subsection by describing the coupling construction stating the stochastic inequalities and giving a brief proof sketch for Theorem 7.1.

The Coupling Construction. We now construct a coupling between two systems following any two schemes, say, $\Pi_1 = \Pi(l_0, l_1, \ldots, l_{B-1})$ and $\Pi_2 = \Pi(d_0, d_1, \ldots, d_{B'-1})$ in $\Pi^{(N)}$, respectively, to establish the desired stochastic ordering results. Note that Π_1 and Π_2 have (possibly different) buffer capacities B and B', respectively. With slight abuse of notation we will denote by Π_i the system following scheme Π_i , i = 1, 2.

For the arrival process we couple the two systems as follows. First, we synchronize the arrival epochs of the two systems. Now assume that in the systems Π_1 and Π_2 , the minimum queue lengths are k and m, respectively, for some $k \leq B-1$, $m \leq B'-1$. Therefore, when a task arrives, the dispatchers in Π_1 and Π_2 have to select l_k and d_m servers, respectively, and then have to send the task to the server with the minimum queue length among the respectively selected servers. Since the servers are being selected uniformly at random, we can assume without loss of generality, as in the stack construction, that the servers are arranged in nondecreasing order of their queue lengths and are indexed in increasing order. Hence, observe that when a few server indices are selected, the server having the minimum of those indices will be the server with the minimum queue length among them. Thus, in this case the dispatchers in Π_1 and Π_2 select l_k and d_m random numbers (without replacement) from $\{1, 2, \dots, N\}$ and then send the incoming task to the servers having indices that are the minimum of those selected numbers. Now, note that selecting l_k (or d_m) random servers is equivalent to selecting a random permutation of $\{1, 2, \ldots, N\}$, say, $(\sigma_1, \sigma_2, \ldots, \sigma_N)$, and selecting first l_k (or d_m) indices. To couple the assignment decisions of the two systems, at each arrival epoch a single random permutation of $\{1, 2, \dots, N\}$ is drawn,

denoted by $\Sigma^{(N)} := (\sigma_1, \sigma_2, \dots, \sigma_N)$. Define $\sigma_{(i)} := \min_{j \leq i} \sigma_j$. Then observe that system Π_1 sends the task to the server with the index $\sigma_{(l_k)}$ and system Π_2 sends the task to the server with the index $\sigma_{(d_m)}$. Since at each arrival epoch both systems use a common random permutation, they take decisions in a coupled manner.

For the potential departure process, couple the service completion times of the kth queue in both scenarios, $k = 1, 2, \dots, N$. More precisely, for the potential departure process assume that we have a single synchronized $\exp(N)$ clock independent of arrival epochs for both systems. Now when this clock alarm rings, a number k is uniformly selected from $\{1, 2, \dots, N\}$ and a potential departure occurs from the kth queue in both systems. If, at a potential departure epoch, an empty queue is selected, then we do nothing. Since the service time requirements are i.i.d. exponentially distributed, the memoryless property ensures that the two schemes, considered independently, still evolve according to their appropriate statistical laws under the above coupling.

Proposition 7.2. For two schemes $\Pi_1 = \Pi(l_0, l_1, \dots, l_{B-1})$ and $\Pi_2 = \Pi(d_0, d_1, \dots, l_{B-1})$ $\ldots, d_{B'-1}$) with $B \leq B'$, assume $l_0 = \cdots = l_{B-2} = d_0 = \cdots = d_{B-2} = d$, $l_{B-1} \leq \cdots \leq d_{B-2} = d$ $\begin{array}{l} d_{B-1}, \ and \ either \ d=N \ or \ d \leq d_{B-1}. \ Then \ the \ following \ hold: \\ (i) \ \{Q_i^{\Pi_1}(t)\}_{t \geq 0} \leq_{st} \{Q_i^{\Pi_2}(t)\}_{t \geq 0} \ for \ i=1,2,\ldots,B, \\ (ii) \ \{\sum_{i=1}^B Q_i^{\Pi_1}(t) + L^{\Pi_1}(t)\}_{t \geq 0} \geq_{st} \{\sum_{i=1}^{B'} Q_i^{\Pi_2}(t) + L^{\Pi_2}(t)\}_{t \geq 0}, \end{array}$

- (iii) $\{\Delta(t)\}_{t\geq 0} \geq \{\sum_{i=B+1}^{B'} Q_i^{\Pi_2}(t)\}_{t\geq 0}$ almost surely under the coupling defined above, for any fixed $N \in \mathbb{N}$ where $\Delta(t) := L^{\Pi_1}(t) - L^{\Pi_2}(t)$, provided that at time t = 0 the above ordering holds.

Proof of Theorem 7.1. Let $\Pi = \Pi(N, d_1, \dots, d_{B-1})$ be a load balancing scheme in the class $\Pi^{(N)}$. Denote by Π_1 the scheme $\Pi(N,d_1)$ with buffer size B=2 and let Π_2 denote the JIQ policy $\Pi(N,1)$ with buffer size B=2.

Observe that from Proposition 7.2 we have, under the coupling defined above,

$$\begin{aligned} |Q_i^\Pi(t) - Q_i^{\Pi_2}(t)| &\leq |Q_i^\Pi(t) - Q_i^{\Pi_1}(t)| + |Q_i^{\Pi_1}(t) - Q_i^{\Pi_2}(t)| \\ &\leq |L^{\Pi_1}(t) - L^\Pi(t)| + |L^{\Pi_2}(t) - L^{\Pi_1}(t)| \leq 2L^{\Pi_2}(t) \end{aligned}$$

for all $i \geq 1$ and $t \geq 0$ with the understanding that $Q_i(t) = 0$ for all i > B, for a scheme with buffer capacity B. The third inequality above is due to Proposition 7.2 (iii), which in particular says that $\{L^{\Pi_2}(t)\}_{t\geq 0} \geq \{L^{\Pi_1}(t)\}_{t\geq 0} \geq \{L^{\Pi}(t)\}_{t\geq 0}$ almost surely under the coupling. Now we have the following lemma.

LEMMA 7.3. For all $t \geq 0$, under the assumptions of Theorem 7.1, $\{L^{\Pi_2}(t)\}_{N\geq 1}$ forms a tight sequence.

Since $L^{\Pi_2}(t)$ is nondecreasing in t, the above lemma in particular implies that

(7.4)
$$\sup_{t \in [0,T]} \frac{L^{\Pi_2}(t)}{\sqrt{N}} \xrightarrow{\mathbb{P}} 0.$$

For any scheme $\Pi \in \Pi^{(N)}$, from (7.3) we know that

$$\{Q_i^{\Pi_2}(t) - 2L^{\Pi_2}(t)\}_{t \ge 0} \le \{Q_i^{\Pi}(t)\}_{t \ge 0} \le \{Q_i^{\Pi_2}(t) + 2L^{\Pi_2}(t)\}_{t \ge 0}.$$

Combining (7.3) and (7.4) shows that if the weak limits under the \sqrt{N} scaling exist, they must be the same for all the schemes in the class $\Pi^{(N)}$. Also, as described in section 3, the weak limit for $\Pi(N,N)$ exists and the common weak limit can be described by the unique solution of the SDEs in (3.7). Hence, the proof of Theorem 7.1 is complete.

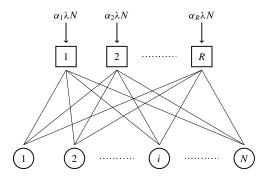


Fig. 6 Schematic view of the model with R dispatchers and N servers.

7.3. Multiple Dispatchers. So far we have focused on a basic scenario with a single dispatcher, but it is not uncommon for LBAs to operate across multiple dispatchers. While the presence of multiple dispatchers does not affect the queuing dynamics of JSQ(d) policies, it does matter for the JIQ scheme, which uses memory at the dispatcher. In order to examine the impact, we consider in this subsection a scenario with N parallel identical servers as before and $R \geq 1$ dispatchers, as depicted in Figure 6. Tasks arrive at dispatcher r as a Poisson process of rate $\alpha_r \lambda N$, with $\alpha_r > 0$, $r = 1, \ldots, R$, $\sum_{r=1}^R \alpha_r = 1$, and λ denoting the task arrival rate per server. For conciseness, we denote $\alpha = (\alpha_1, \ldots, \alpha_R)$, and without loss of generality we assume that the dispatchers are indexed such that $\alpha_1 \geq \alpha_2 \geq \cdots \geq \alpha_R$.

When a server becomes idle, it sends a token to one of the dispatchers selected uniformly at random, advertising its availability. When a task arrives at a dispatcher which has tokens available, one of the tokens is selected, and the task is immediately forwarded to the corresponding server.

We distinguish two scenarios when a task arrives at a dispatcher which has no tokens available, referred to as the *blocking* and *queuing* scenarios, respectively. In the blocking scenario, the incoming task is blocked and instantly discarded. In the queuing scenario, the arriving task is forwarded to one of the servers selected uniformly at random. If the selected server happens to be idle, then the outstanding token at one of the other dispatchers is revoked.

In the queuing scenario we assume $\lambda < 1$, which is not only necessary but also sufficient for stability. It is not difficult to show that the joint queue length process is stochastically majorized by a scheme that assigns each task to a server chosen uniformly at random. In the latter case, the system decomposes into N independent M/M/1 queues, each of which has load $\lambda < 1$ and is stable.

Scenarios with multiple dispatchers have received limited attention in the literature, and the scant papers that exist [101, 115, 142] almost exclusively assume that the loads at the various dispatchers are strictly equal, i.e., $\alpha_1 = \cdots = \alpha_R = 1/R$. In these cases the fluid limit, for suitable initial states, is the same as in (7.2) for a single dispatcher and, in particular, the fixed point is the same, hence the JIQ scheme continues to achieve asymptotically optimal delay performance with minimal communication overhead. The results in [142] in fact show that the JIQ scheme remains asymptotically optimal even when the servers are heterogeneous, while it is readily seen that JSQ(d) policies cannot even provide maximum stability (i.e., achieve stability whenever feasible) in that case for any fixed value of d. As one of the few

exceptions, [149] allows the loads at the various dispatchers to be different. It is not uncommon for such skewed load patterns to arise, for example, when the various dispatchers receive tasks from external sources, making it difficult to perfectly balance the task arrival rates.

Results for Blocking Scenario. For the blocking scenario, denote by $B(R, N, \lambda, \alpha)$ the steady-state blocking probability of an arbitrary task. It is established in [149] that

$$B(R, N, \lambda, \alpha) \to \max\{1 - R\alpha_R, 1 - 1/\lambda\} \text{ as } N \to \infty.$$

This result shows that in the many-server limit the system performance in terms of blocking is determined either by the relative load of the least-loaded dispatcher or by the aggregate load. This may be explained informally as follows. Let \bar{x}_0 be the expected fraction of busy servers in steady state, so that each dispatcher receives tokens on average at a rate $\bar{x}_0 N/R$. We distinguish two cases, depending on whether or not a positive fraction of the tokens reside at the least-loaded dispatcher R in the limit. If that is the case, then the task arrival rate $\alpha_R \lambda N$ at dispatcher R must equal the rate $\bar{x}_0 N/R$ at which it receives tokens, i.e., $\bar{x}_0/R = \alpha_R \lambda$. Otherwise, the task arrival rate $\alpha_R \lambda N$ at dispatcher R must be no less than the rate $\bar{x}_0 N/R$ at which it receives tokens, i.e., $\bar{x}_0/R \leq \alpha_R \lambda$. Since dispatcher R is the least loaded, it then follows that $\bar{x}_0/R \leq \alpha_r \lambda$ for all $r = 1, \ldots, R$, which means that the task arrival rate at all the dispatchers is higher that the rate at which tokens are received. Thus the fraction of tokens at each dispatcher is zero in the limit, i.e., the fraction of idle servers is zero, implying $\bar{x}_0 = 1$. Combining the two cases, and observing that $\bar{x}_0 \leq 1$, we conclude that $\bar{x}_0 = \min\{R\alpha_R\lambda, 1\}$. Because of Little's law, \bar{x}_0 is related to the blocking probability B as $\bar{x}_0 = \lambda(1 - B)$. This yields $1 - B = \min\{R\alpha_R\lambda, 1/\lambda\}$ or, equivalently, $B = \max\{1 - R\alpha_R, 1 - 1/\lambda\}.$

The above explanation also reveals that, somewhat counterintuitively, it is the least-loaded dispatcher that throttles tokens and leaves idle servers stranded, thus acting as a bottleneck. Specifically, in the limit, dispatcher R (or the set of least-loaded dispatchers in the case of ties) inevitably ends up with all the available tokens, if any. The accumulation of tokens hampers the visibility of idle servers to the heavier-loaded dispatchers, and it leaves idle servers stranded while tasks queue up at other servers.

Results for Queuing Scenario. For the queuing scenario, denote by $W(R, N, \lambda, \alpha)$ a random variable with the steady-state waiting-time distribution of an arbitrary task. It is shown in [149] that, for a fixed $\lambda < 1$ and $N \to \infty$,

$$\mathbb{E}[W(R, N, \lambda, \alpha)] \to \frac{\lambda_2(R, \lambda, \alpha)}{1 - \lambda_2(R, \lambda, \alpha)},$$

where

$$\lambda_2(R, \lambda, \alpha) = 1 - \frac{1 - \lambda \sum_{i=1}^{r^*} \alpha_i}{1 - \lambda r^* / R}$$

with

$$r^* = \sup \left\{ r \middle| \alpha_r > \frac{1}{R} \frac{1 - \lambda \sum_{i=1}^r \alpha_i}{1 - \lambda r/R} \right\}$$

and the convention that $r^* = 0$ if $\alpha_1 = \cdots = \alpha_R = 1/R$. In particular,

$$\lambda_2(2,\lambda,(1-\alpha_2,\alpha_2)) = \lambda \frac{1-2\alpha_2}{2-\lambda},$$

so that

$$\mathbb{E}[W(2, N, \lambda, (1 - \alpha_2, \alpha_2))] \to \frac{\lambda(1 - 2\alpha_2)}{2 - 2\lambda(1 - \alpha_2)}.$$

Here λ_2 can be interpreted as the rate at which tasks are forwarded to randomly selected servers. Furthermore, dispatchers $1, \ldots, r^*$ receive tokens at a lower rate than the incoming tasks and, in particular, $\lambda_2^* = 0$ if and only if $r^* = 0$.

When the arrival rates at all dispatchers are strictly equal, i.e., $\alpha_1 = \cdots = \alpha_R = 1/R$, the above results indicate that the stationary blocking probability and the mean waiting time asymptotically vanish as $N \to \infty$, which is in agreement with the observations in [142] mentioned above. However, when the arrival rates at the various dispatchers are not perfectly equal, so that $\alpha_R < 1/R$, the blocking probability and mean waiting time are strictly positive in the limit, even for arbitrarily low overall load and an arbitrarily small degree of skewness in the arrival rates. Thus, the ordinary JIQ scheme fails to achieve asymptotically optimal performance for heterogeneous dispatcher loads.

Enhancements. In order to counter the above-described performance degradation for asymmetric dispatcher loads, [149] proposes two enhancements.

Enhancement 1 (nonuniform token allotment). When a server becomes idle, it sends a token to dispatcher r with probability β_r .

Enhancement 2 (token exchange mechanism). Any token is transferred to a uniformly randomly selected dispatcher at rate ν .

Note that the token exchange mechanism only creates a constant communication overhead per task as long as the rate ν does not depend on the number of servers N, and thus preserves the scalability of the basic JIQ scheme. The above enhancements can achieve asymptotically optimal performance for suitable values of the β_r parameters and the exchange rate ν .

Large Number of Dispatchers. In the above setup we assumed the number of dispatchers to remain fixed as the number of servers grows large, but a further natural scenario would be for the number of dispatchers R(N) to scale with the number of servers, as considered in [115], where the author analyzes the case R(N) = rN for some constant r, so that the relative load of each dispatcher is λr . The term "I-queue" is used for the queue of (idle) servers that is known by one of the dispatchers. A server is added to an I-queue when it becomes idle. With fluid limits and fixed-point calculations, the analysis in [115] determines the fraction of I-queues with i queued servers and the fraction of servers with i tasks in queue that are in the jth position in one of the I-queues. The fixed point can be computed numerically.

Anticipation. In [115] it is also proposed that servers issue their availability tokens to the dispatchers before they are idle, e.g., when they have just one task remaining. This appears beneficial at very high load when there are (on average) fewer idle servers than dispatchers, and tasks would frequently be assigned to uniformly at random selected servers otherwise. Two variants are introduced. First, an LCFS (last come, first served) scheme in which the server that is in the I-queue the least amount of time is chosen for the incoming task. Second, a server that became idle may probe d I-queues, after which it chooses the least loaded among the d selected servers. Both variants lead to small performance improvements.

7.4. Joint Load Balancing and Autoscaling. Besides delay performance and implementation overhead, a further key attribute in the context of large-scale cloud

networks and data centers is energy consumption. So-called autoscaling algorithms have emerged as a popular mechanism for adjusting service capacity in response to varying demand levels, to minimize energy consumption while meeting performance targets, but have mostly been investigated in settings with a centralized queue, and queue-driven autoscaling techniques have been widely investigated in the literature [7, 54, 94, 95, 98, 99, 100, 132, 148, 171]. In systems with a centralized queue it is common to put servers to "sleep" while the demand is low, since servers in sleep mode consume much less energy than active servers. Under Markovian assumptions, the behavior of these mechanisms can be described in terms of various incarnations of M/M/N queues with setup times. There are several further recent papers that examine on-demand server addition/removal in a somewhat different vein [128, 130]. Unfortunately, data centers and cloud networks with massive numbers of servers are too complex to maintain any centralized queue, which involves a prohibitively high communication burden to obtain instantaneous state information.

Motivated by these observations, the authors of [121] propose a joint load balancing and autoscaling strategy, which retains the excellent delay performance and low implementation overhead of the ordinary JIQ scheme and at the same time minimizes the energy consumption. The strategy is referred to as TABS (token-based auto-balance scaling) and operates as follows:

- When a server becomes idle, it sends a "green" message to the dispatcher, waits for an exp(μ) time (standby period), and turns itself off by sending a "red" message to the dispatcher (the corresponding green message is destroyed).
- When a task arrives, the dispatcher selects a green message at random if there are any and assigns the task to the corresponding server (the corresponding green message is replaced by a "yellow" message). Otherwise, the task is assigned to an arbitrary busy server, and if at that arrival epoch there is a red message at the dispatcher, then it selects one at random and the setup procedure of the corresponding server is initiated, replacing its red message with an "orange" message. The setup procedure takes exp(ν) time after which the server becomes active.
- Any server which activates due to the latter event sends a green message to the dispatcher (the corresponding orange message is replaced), waits for an $\exp(\mu)$ time for a possible assignment of a task, and again turns itself off by sending a red message to the dispatcher.

The TABS scheme gives rise to a distributed operation in which servers are in one of four states (busy, idle-on, idle-off, or standby) and advertise their state to the dispatcher via exchange of tokens. Figure 7 illustrates this token-based exchange protocol. Note that setup procedures are never aborted and continue even when idle-on servers do become available. Recently, dynamic scaling and load balancing with variable service capacity and on-demand agents has been further examined in [67].

To describe systems under the TABS scheme, we use $\mathbf{Q}^N(t) := (Q_1^N(t), Q_2^N(t), \ldots, Q_B^N(t))$ to denote the system occupancy state at time t as before, where $B \geq 1$ is a finite buffer capacity. Also, let $\Delta_0^N(t)$ and $\Delta_1^N(t)$ denote the number of idle-off servers and servers in setup mode at time t, respectively. The fluid-scaled quantities are denoted by the respective small letters, $q_i^N(t) := Q_i^N(t)/N$, $\delta_0^N(t) = \Delta_0^N(t)/N$, and $\delta_1^N(t) = \Delta_1^N(t)/N$. For brevity of notation, we will write $\mathbf{q}^N(t) = (q_1^N(t), \ldots, q_B^N(t))$ and $\delta^N(t) = (\delta_0^N(t), \delta_1^N(t))$. The results presented in the remainder of the section are extracted from [121], unless indicated otherwise.

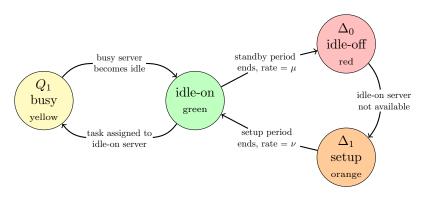


Fig. 7 Illustration of server on-off decision rules in the TABS scheme, along with message colors and state variables.

Fluid Limit. Under suitable initial conditions, on any finite time interval, with probability 1, any sequence in N has a further subsequence along which the sequence of processes $(\mathbf{q}^N(\cdot), \boldsymbol{\delta}^N(\cdot))$ converges to a deterministic limit $(\mathbf{q}(\cdot), \boldsymbol{\delta}(\cdot))$ that satisfies the system of ODEs

(7.5)
$$\frac{d^{+}q_{i}(t)}{dt} = \lambda(t)p_{i-1}(\mathbf{q}(t), \boldsymbol{\delta}(t), \lambda(t)) - (q_{i}(t) - q_{i+1}(t)), \ i = 1, \dots, B, \\ \frac{d^{+}\delta_{0}(t)}{dt} = u(t) - \frac{d^{+}\xi(t)}{dt}, \qquad \frac{d^{+}\delta_{1}(t)}{dt} = \frac{d^{+}\xi(t)}{dt} - \nu\delta_{1}(t),$$

where by convention $q_{B+1}(\cdot) \equiv 0$ and

$$u(t) = 1 - q_1(t) - \delta_0(t) - \delta_1(t), \quad \frac{\mathrm{d}^+ \xi(t)}{\mathrm{d}t} = \lambda(t) (1 - p_0(\mathbf{q}(t), \boldsymbol{\delta}(t), \lambda(t))) \mathbb{1}_{[\delta_0(t) > 0]}.$$

For any $(\mathbf{q}, \boldsymbol{\delta})$ and $\lambda > 0$, $(p_i(\mathbf{q}, \boldsymbol{\delta}, \lambda))_{i \geq 0}$ are given by

$$p_0(\mathbf{q}, \boldsymbol{\delta}, \lambda) = \begin{cases} 1 & \text{if } u = 1 - q_1 - \delta_0 - \delta_1 > 0, \\ \min\{\lambda^{-1}(\delta_1 \nu + q_1 - q_2), 1\} & \text{otherwise,} \end{cases}$$
$$p_i(\mathbf{q}, \boldsymbol{\delta}, \lambda) = (1 - p_0(\mathbf{q}, \boldsymbol{\delta}, \lambda))(q_i - q_{i+1})q_1^{-1}, i = 1, \dots, B.$$

We now provide an intuitive explanation of the fluid limit stated above. The term u(t) corresponds to the asymptotic fraction of idle-on servers in the system at time t, and $\xi(t)$ represents the asymptotic cumulative number of server setups (scaled by N) that have been initiated during [0,t]. The coefficient $p_i(\mathbf{q}, \boldsymbol{\delta}, \lambda)$ can be interpreted as the instantaneous fraction of incoming tasks that are assigned to some server with queue length i, when the fluid-scaled occupancy state is $(\mathbf{q}, \boldsymbol{\delta})$ and the scaled instantaneous arrival rate is λ . Observe that as long as u > 0, there are idle-on servers, and hence all the arriving tasks will join idle servers. This explains why if u > 0, $p_0(\mathbf{q}, \boldsymbol{\delta}, \lambda) = 1$ and $p_i(\mathbf{q}, \boldsymbol{\delta}, \lambda) = 0$ for $i = 1, \ldots, B - 1$. If u = 0, then observe that servers become idle at rate $q_1 - q_2$ and servers in setup mode turn on at rate $\delta_1 \nu$. Thus the idle-on servers are created at a total rate $\delta_1 \nu + q_1 - q_2$. If this rate is larger than the arrival rate λ , then almost all the arriving tasks can be assigned to idle servers. Otherwise, only a fraction $(\delta_1 \nu + q_1 - q_2)/\lambda$ of arriving tasks join idle servers. The rest of the tasks are distributed uniformly among busy servers,

so a proportion $(q_i - q_{i+1})q_1^{-1}$ are assigned to servers having queue length i. For any $i = 1, \ldots, B$, q_i increases when there is an arrival at some server with queue length i-1, which occurs at rate $\lambda p_{i-1}(\mathbf{q}, \boldsymbol{\delta}, \lambda)$, and it decreases when there is a departure from some server with queue length i, which occurs at rate $q_i - q_{i-1}$. Since each idle-on server turns off at rate μ , the fraction of servers in the off mode increases at rate μu . Observe that if $\delta_0 > 0$, for each task that cannot be assigned to an idle server, a setup procedure is initiated at one idle-off server. As noted above, $\xi(t)$ captures the (scaled) cumulative number of setup procedures initiated up to time t. Therefore, the fraction of idle-off servers and the fraction of servers in setup mode decreases and increases by $\xi(t)$, respectively, during [0, t]. Finally, since each server in setup mode becomes idle-on at rate ν , the fraction of servers in setup mode decreases at rate $\nu \delta_1$.

Fixed Point and Global Stability. In the case of a constant arrival rate $\lambda(t) \equiv \lambda < 1$, any fluid sample path in (7.5) has a unique fixed point,

(7.6)
$$\delta_0^* = 1 - \lambda, \quad \delta_1^* = 0, \quad q_1^* = \lambda, \quad \text{and} \quad q_i^* = 0$$

for $i=2,\ldots,B$. Indeed, it can be verified that $p_0(\mathbf{q}^*,\boldsymbol{\delta}^*,\lambda)=1$ and $u^*=0$ for $(\mathbf{q}^*,\boldsymbol{\delta}^*)$ given by (7.6) so that the derivatives of $q_i, i=1,\ldots,B, \,\delta_0$, and δ_1 become zero, and that these cannot be zero at any other fluid-scaled occupancy state. Note that, at the fixed point, a fraction λ of the servers have exactly one task, while the remaining fraction have zero tasks, independently of the values of the parameters μ and ν .

In order to establish the convergence of the sequence of steady states, we need the global stability of the fluid limit, i.e., starting from any fluid-scaled occupancy state, any fluid sample path described by (7.5) converges to the unique fixed point (7.6) as $t \to \infty$. More specifically, irrespective of the starting state,

(7.7)
$$(\mathbf{q}(t), \boldsymbol{\delta}(t)) \to (\mathbf{q}^*, \boldsymbol{\delta}^*) \quad \text{as} \quad t \to \infty,$$

where $(\mathbf{q}^*, \boldsymbol{\delta}^*)$ is as defined in (7.6).

Interchange of Limits. The global stability can be leveraged to show that the steady-state distribution of the Nth system, for large N, can be well approximated by the fixed point of the fluid limit in (7.6). Specifically, it justifies the interchange of the many-server $(N \to \infty)$ and stationary $(t \to \infty)$ limits. Since the buffer capacity B at each server is supposed to be finite, for every N, the Markov process $(\mathbf{Q}^N(t), \Delta_0^N(t), \Delta_1^N(t))$ is irreducible, has a finite state space, and thus has a unique steady-state distribution. Let π^N denote the steady-state distribution of the Nth system, i.e.,

$$\pi^{N}(\cdot) = \lim_{t \to \infty} \mathbb{P}\left(\mathbf{q}^{N}(t) = \cdot, \boldsymbol{\delta}^{N}(t) = \cdot\right).$$

The fluid-limit result and the global stability thus yield that π^N converges weakly to π as $N \to \infty$, where π is given by the Dirac mass concentrated upon $(\mathbf{q}^*, \boldsymbol{\delta}^*)$ defined in (7.6).

Remark 7.4. Note that the above interchange of limits result was obtained under the assumption that the queues have finite buffers, and analysis of the infinite-buffer scenario was left open. The key challenge in the latter case stems from the fact that the system stability under the usual subcritical load assumption is not automatic. In fact, as explained in [122], when the number of servers N is fixed, the stability may not hold even under a subcritical load assumption. In [122] the stability issue of the

TABS scheme was addressed and the convergence of the sequence of steady states was shown for the infinite-buffer scenario. In particular, it was established that for a fixed choice of parameters $\lambda < 1$, $\mu > 0$, and $\nu > 0$, the system with N servers under the TABS scheme is stable for large enough N. There the authors introduced an induction-based approach that uses both the conventional fluid limit (in the sense of a large starting state) and the mean-field fluid limit (when $N \to \infty$) in an intricate fashion to prove the large-N stability of the system.

Performance Metrics. As mentioned earlier, two key performance metrics are the expected waiting time of tasks $\mathbb{E}[W^N]$ and energy consumption $\mathbb{E}[P^N]$ for the Nth system in steady state. In order to quantify the energy consumption, we assume that the energy usage of a server is P_{full} when busy or in setup mode, P_{idle} when idle-on, and zero when turned off. Evidently, for any value of N, at least a fraction λ of the servers must be busy in order for the system to be stable, and hence λP_{full} is the minimum mean energy usage per server needed for stability. We will define $\mathbb{E}[Z^N] = \mathbb{E}[P^N] - \lambda P_{\text{full}}$ as the relative energy wastage accordingly. The interchange of limits result can be leveraged to obtain that asymptotically the expected waiting time and energy consumption for the TABS scheme vanish in the limit, for any strictly positive values of μ and ν . More specifically, for a constant arrival rate $\lambda(t) \equiv \lambda < 1$, for any $\mu > 0$, $\nu > 0$, as $N \to \infty$,

- (a) zero mean waiting time: $\mathbb{E}[W^N] \to 0$;
- (b) zero energy wastage: $\mathbb{E}[Z^N] \to 0$.

The key implication is that the TABS scheme, while only involving constant communication overhead per task, provides performance in a distributed setting that is as good at the fluid level as can possibly be achieved, even in a centralized queue, or with unlimited information exchange.

Comparison to Ordinary JIQ Policy. Consider again a constant arrival rate $\lambda(t) \equiv \lambda$. It is worthwhile to observe that the component \mathbf{q} of the fluid limit as in (7.5) coincides with that for the ordinary JIQ policy where servers always remain on, when the system following the TABS scheme starts with all the servers being idle-on and $\lambda + \mu < 1$. To see this, observe that the component \mathbf{q} depends on $\boldsymbol{\delta}$ only through $(p_{i-1}(\mathbf{q}, \boldsymbol{\delta}))_{i\geq 1}$. Now, $p_0 = 1$, $p_i = 0$ for all $i \geq 1$, whenever $q_1 + \delta_0 + \delta_1 < 1$, irrespective of the precise values of $(\mathbf{q}, \boldsymbol{\delta})$. Moreover, starting from the above initial state, δ_1 can increase only when $q_1 + \delta_0 = 1$. Therefore, the fluid limits of \mathbf{q} in (7.5) and the ordinary JIQ scheme are identical if the system parameters (λ, μ, ν) are such that $q_1(t) + \delta_0(t) < 1$ for all $t \geq 0$. Let $y(t) = 1 - q_1(t) - \delta_0(t)$. The solutions to the differential equations

$$\frac{\mathrm{d}q_1(t)}{\mathrm{d}t} = \lambda - q_1(t), \quad \frac{\mathrm{d}y(t)}{\mathrm{d}t} = q_1(t) - \lambda - \mu y(t),$$

 $y(0) = 1, q_1(0) = 0$, are given by

$$q_1(t) = \lambda(1 - e^{-t}), \quad y(t) = \frac{e^{-(1+\mu)t}}{\mu - 1} (e^t(\lambda + \mu - 1) - \lambda e^{\mu t}).$$

Notice that if $\lambda + \mu < 1$, then y(t) > 0 for all $t \ge 0$ and, thus, $q_1(t) + \delta_0(t) < 1$ for all $t \ge 0$. The fluid-level optimality of the JIQ scheme was described in section 7.1. This observation thus establishes the optimality of the fluid-limit trajectory under the TABS scheme for suitable parameter values in terms of response time performance.

From the energy usage perspective, under the ordinary JIQ policy, since the asymptotic steady-state fractions of busy servers (q_1^*) and idle-on servers are given by λ and $1 - \lambda$, respectively, the asymptotic steady-state (scaled) energy usage is given by

$$\mathbb{E}[P^{\text{JIQ}}] = \lambda P_{\text{full}} + (1 - \lambda)P_{\text{idle}} = \lambda P_{\text{full}}(1 + (\lambda^{-1} - 1)f),$$

where $f = P_{\text{idle}}/P_{\text{full}}$ is the relative energy consumption of an idle server. As described earlier, the asymptotic steady-state (scaled) energy usage under the TABS scheme is λP_{full} . Thus the TABS scheme reduces the asymptotic steady-state energy usage by $\lambda P_{\text{full}}(\lambda^{-1}-1)f = (1-\lambda)P_{\text{idle}}$, which amounts to a relative saving of $(\lambda^{-1}-1)f/(1+(\lambda^{-1}-1)f)$. In summary, the TABS scheme performs as well as the ordinary JIQ policy in terms of the waiting time and communication overhead while providing a significant energy saving.

- **8. Redundancy Policies and Alternative Scaling.** In this section we discuss somewhat related redundancy policies, alternative scaling regimes, and some additional performance metrics of interest.
- **8.1. Redundancy-**d **Policies.** So-called redundancy-d policies involve a somewhat similar operation to JSQ(d) policies and also share the primary objective of ensuring low delays [6, 162]. In a redundancy-d policy, $d \geq 2$ candidate servers are selected uniformly at random (with or without replacement) for each arriving task, just like in a JSQ(d) policy. Rather than forwarding the task to the server with the shortest queue, however, replicas are dispatched to all sampled servers. Note that the initial replication to d servers selected uniformly at random does not entail any communication burden, but the elimination of redundant copies at a later stage does involve a significant amount of information exchange and complexity.

Two common options can be distinguished for the removal of redundant clones. In the first variant, as soon as the first replica starts service, the other clones are abandoned. In this case, a task gets executed by the server which had the smallest workload at the time of arrival (and which may or may not have had the shortest queue length) among the sampled servers. This can be interpreted as a power-of-d version of the JSW policy discussed in section 2.3.3. The optimality properties of the JSW policy mentioned in that subsection suggest that redundancy-d policies should outperform JSQ(d) policies, which appears to be supported by simulation experiments.

In the second option the other clones of the task are not eliminated until the first replica has completed service (which may or may not have been the first replica to start service). While a task is only handled by one of the servers in the former case, it may be processed by several servers in the latter case. When the service times are exponentially distributed and independent for the various clones, the aggregate amount of time spent by all the servers until completion remains exponentially distributed with the same mean. An exact analysis of the delay distribution in systems with N=2 or N=3 servers is provided in [57, 58], and exact expressions for the mean delay with an arbitrary number of servers are established in [59]. The limiting delay distribution in the many-server regime (ii) is derived in [55, 60] based on an asymptotic independence assumption among the servers. In general, the mean aggregate amount of time devoted to a task and the resulting delay may be larger or smaller for fewer or more variable service time distributions, also depending on the number of replicas per task [133, 138, 164, 165]. In particular, for heavy-tailed service time distributions, the mean aggregate time spent on a task may be considerably reduced by virtue of the redundancy. Indeed, even if the first replica to start service has an

extremely long service time, that is not likely to be the case for the other clones as well. In spite of the extremely long service time of the first replica, it is therefore unlikely for the aggregate amount of time spent on the task or its waiting time to be large. This provides a significant performance benefit to redundancy-d policies over JSQ(d) policies, and has also motivated a strong interest in adaptive replication schemes [4, 84, 85].

A further closely related model is where k of the replicas need to complete service, $1 \le k \le d$, in order for the task to finish, which is relevant in the context of storage systems with coding and MapReduce tasks [86, 87]. The special case where k=d=N corresponds to a classical fork-join system. The authors of [78] present a unified approach for analyzing the stability and performance of a broad class of workload-dependent task assignment and replication policies based on considering the so-called cavity process in a many-server regime with $N \to \infty$. This class of policies includes both versions of the redundancy-d policy as well as the above-mentioned k-out-of-d service mechanism system.

8.2. Conventional Heavy Traffic. In this subsection we briefly discuss a few asymptotic results for LBAs in the classical heavy-traffic regime as described in section 2.2, where the number of servers N is fixed and the relative load tends to one in the limit.

The papers [39, 40, 134, 178] establish diffusion limits for the JSQ policy in a sequence of systems with Markovian characteristics as in our basic model setup, but where in the Kth system the arrival rate is $K\lambda + \hat{\lambda}\sqrt{K}$, while the service rate of the ith server is $K\mu_i + \hat{\mu}_i\sqrt{K}$, $i=1,\ldots,N$, with $\lambda = \sum_{i=1}^N \mu_i$, inducing critical load as $K\to\infty$. It is proved that for suitable initial conditions the queue lengths are of the order $O(\sqrt{K})$ over any finite time interval and exhibit a state-space collapse property. In particular, a properly scaled version of the joint queue length process lives in a one-dimensional rather than N-dimensional space, reflecting that the various queue lengths evolve in lock-step, with the relative proportions remaining virtually identical in the limit, while the aggregate queue length varies.

Atar, Keslassy, and Mendelson [9] investigate a similar scenario and establish diffusion limits for three policies: the JSQ(d) policy, the redundancy-d policy (where the redundant clones are abandoned as soon as the first replica starts service), and a combined policy called replicate-to-shortest-queues (RSQ), where d replicas are dispatched to the d-shortest queues. Note that the latter policy requires instantaneous knowledge of all the queue lengths and hence involves similarly excessive communication overhead as the ordinary JSQ policy, besides the substantial information exchange associated with the removal of redundant copies. Conditions are derived for the values of the relative service rates μ_i , $i = 1, \ldots, N$, in conjunction with the diversity parameter d, in order for the queue lengths under the JSQ(d) and redundancy-d policies to be of the order $O(\sqrt{K})$ over any finite time interval and exhibit statespace collapse. The conditions for the two policies are distinct, but in both cases they are weaker for larger values of d, as intuitively expected. While the conditions for the values of μ_i depend on d, whenever they are met, the actual diffusion-scaled queue length processes do not depend on the exact value of d in the limit, showing a certain resemblance to the universality property as identified in section 2.3.4 for the conventional large-capacity and Halfin-Whitt regimes.

The authors of [181] consider a slightly different model setup with a time-slotted operation and identify a class Π of LBAs that not only provide throughput optimality (or maximum stability, i.e., keep the queues stable in a suitable sense whenever feasible

to do so at all), but also achieve heavy-traffic delay optimality in the sense that the properly scaled aggregate queue length is the same as that in a centralized queue where all the resources are pooled as the load tends to one. As it turns out, the class Π includes JSQ(d) policies with $d \geq 2$, but it does not include the JIQ scheme, which tends to degenerate into a random assignment policy when idle servers are rarely available. The authors further propose a threshold-based policy which has low implementation complexity like the JIQ scheme, but does belong to the class Π and hence achieves heavy-traffic delay optimality. A later paper [180] establishes both necessary and sufficient conditions for threshold-based task assignment policies to achieve heavy-traffic optimality in terms of mean delay.

8.3. Nondegenerate Slowdown. In this subsection we briefly discuss a few of the scarce asymptotic results for LBAs in the so-called nondegenerate slow-down regime described in section 2.2, where $N-\lambda(N)\to\gamma>0$ as the number of servers N grows large. In a centralized queue the process tracking the evolution of the number of waiting tasks, suitably accelerated and normalized by N, converges in this regime to a Brownian motion with drift $-\gamma$ reflected at zero as $N\to\infty$, as demonstrated in [8]. In stationarity, the number of waiting tasks, normalized by N, converges in this regime to an exponentially distributed random variable with parameter γ as $N\to\infty$. Hence, the mean number of waiting tasks must be at least of the order N/γ , and the waiting time cannot vanish as $N\to\infty$ under any policy.

The authors of [73] characterize the diffusion-scaled queue length process under the JSQ policy in this asymptotic regime. They further compare the diffusion limit for the JSQ policy with that for a centralized queue as described above, as well as several LBAs such as the JIQ scheme and a refined version called idle-one-first (I1F), where a task is assigned to a server with exactly one task if no idle server is available and to a randomly selected server otherwise.

In [73], it is proved that the diffusion limit for the JIQ scheme is no longer asymptotically equivalent to that for the JSQ policy in this asymptotic regime, and the JIQ scheme fails to achieve asymptotic optimality in that respect, as opposed to the behavior in the conventional large-capacity and Halfin–Whitt heavy-traffic regimes discussed in section 2.3.5. In contrast, the I1F scheme does preserve the asymptotic equivalence with the JSQ policy in terms of the diffusion-scaled queue length process, and thus retains asymptotic optimality in that sense.

These results provide further indication that the amount and accuracy of queue length information needed to achieve asymptotic equivalence with the JSQ policy depend not only on the scale dimension (e.g., fluid or diffusion), but also on the load regime. Put differently, the finer the scale and the higher the load, the more strictly one can distinguish various LBAs in terms of the relative performance compared to the JSQ policy.

8.4. Sparse-Feedback Regime. As described in section 2.3.5, the JIQ scheme involves a communication overhead of at most one message per task, and yet it achieves optimal delay performance in the fluid and diffusion regimes. However, even just one message per task may still be prohibitive, especially when tasks do not involve big computational tasks, but instead small data packets which require little processing. In such situations the sheer message exchange volume in providing queue length information may be disproportionate to the actual amount of processing required.

Motivated by the above issues, [150] proposes and examines a novel class of LBAs which also leverage memory at the dispatcher, but allow the communication overhead to be seamlessly adapted and reduced to below that of the JIQ scheme. Specifically, in

the proposed schemes, the various servers provide occasional queue status notifications to the dispatcher, in either a synchronous or an asynchronous fashion. The dispatcher uses these reports to maintain queue estimates and forwards incoming tasks to the server with the lowest queue estimate. The queue estimate for a server is incremented for every task assigned and set to the true queue length at update moments, but is never lowered in between updates. Note that when the update frequency per server is δ , the number of messages per task is $d = \delta/\lambda$, with $\lambda < 1$ denoting the arrival rate per server.

The results in [150] demonstrate that the proposed schemes markedly outperform JSQ(d) policies with the same number of $d \geq 1$ messages per task and they can achieve a vanishing waiting time in the many-server limit when the update frequency δ exceeds $\lambda/(1-\lambda)$. In the case when servers only report zero queue lengths and suppress updates for nonzero queues, the update frequency required for a vanishing waiting time can in fact be lowered to just λ , matching the one message per task involved in the JIQ scheme.

From a scalability viewpoint, the most pertinent regime is d < 1 where only very sparse server feedback is required. It is shown in [150] that the proposed schemes then outperform the corresponding sparsified versions of the JIQ scheme where idle servers only provide notifications to the dispatcher with probability d. In order to further explore the performance for d < 1 in the many-server limit, [150] investigates fluid limits for the synchronous case as well as the asynchronous case with exponential update intervals. The fixed points of the fluid limit are leveraged to derive the stationary queue length distribution as a function of the update frequency.

Additionally, [150] examines the performance in the ultralow feedback regime where the update frequency δ goes to zero, and in particular establishes a somewhat counterintuitive dichotomy. In the synchronous case, the behavior of each of the individual queues approaches that of a single-server queue with a near-deterministic arrival process and exponential service times, with the mean waiting time tending to a finite constant. In contrast, in the asynchronous case, the individual queues experience saw-tooth behavior with oscillations and waiting times that grow without bound.

In order to achieve a vanishing waiting time, the dispatcher must assign each incoming task to an idle server with high probability, and thus be able to identify on average at least one idle server for every incoming task. When the amount of memory at the dispatcher is limited, the dispatcher may in fact have to identify more idle servers on average to ensure that at least one is available with high probability for each incoming task, as is also reflected in the results of [50, 51, 52]. These conditions, in conjunction with the fact that the fraction of idle servers in equilibrium is $1-\lambda$, translate into a minimum required communication overhead for various families of algorithms. For example, if the dispatcher samples a server at random, it will find that server idle with probability $1-\lambda$, so in the absence of any memory it will need to sample a number of servers that grows with N for each incoming task, while with unlimited memory, it will need to sample on average $1/(1-\lambda)$ servers per incoming task. Likewise, if servers report their queue status to the dispatcher, then an arbitrary server will report as idle with probability $1-\lambda$, so they all need to do that every $\lambda/(1-\lambda)$ time units on average, i.e., $1/(1-\lambda)$ times on average per incoming task. When only idle servers report their status to the dispatcher, as in the JIQ algorithm, they only need to do so at most once per incoming task. When servers report their status asynchronously rather than all simultaneously, or idle servers only after some delay, the associated memory requirement at the dispatcher can be reduced.

8.5. Scaling of Maximum Queue Length. So far we have focused on the asymptotic behavior of LBAs in terms of the number of servers with a certain queue length, either on fluid scale or diffusion scale, in various regimes as $N \to \infty$. A related but different performance metric is the maximum queue length M(N) among all servers as $N \to \infty$. The authors of [103] showed that for fixed $d \ge 2$ the stationary maximum queue length M(N) in a system under the JSQ(d) policy is concentrated on at most two adjacent values which are $\log(\log(N))/\log(d)+O(1)$, whereas for purely random assignment (d=1), it scales as $\log(N)/\log(1/\lambda)$ and does not concentrate on a bounded range of values. This is yet a further manifestation of the power-of-choice effect.

An earlier paper [102] had already shown a similar result for the maximum bin occupancy under a power-of-d policy in a balls-and-bins context where arriving items (balls) do not get served and never depart, but simply accumulate in bins, so that (stationary) queue lengths are not meaningful. The maximum bin occupancy under purely random assignment, however, scales as $\log(N)/\log(\log(N))$ and does concentrate on two adjacent values, in contrast with the queuing scenario mentioned above.

In fact, the very notion of randomized load balancing and power-of-d strategies was introduced in a balls-and-bins setting in the seminal paper [10]. Several further variations and extensions in that context have been considered in [1, 17, 18, 31, 33, 45, 68, 129, 131, 161]. One of the earliest papers on graph-based load balancing was also concerned with a balls-and-bins setting [90].

As alluded to above, there are natural parallels between the balls-and-bins setup and the queuing scenario that we have focused on so far. These commonalities are, for example, reflected in the fact that power-of-d strategies yield similar dramatic performance improvements over purely random assignment in both settings. However, there are also quite fundamental differences between the balls-and-bins setup and the queuing scenario, besides the obvious contrasts in the performance metrics. This distinction is already reflected in the different scaling behavior under purely random assignment of the maximum queue length in a queuing scenario and the maximum bin occupancy in a balls-and-bins setting, as mentioned above. A further manifestation of this is provided by the fact that a simple round-robin strategy produces a perfectly balanced allocation in a balls-and-bins setup but is far from optimal in a queuing scenario, as observed in section 2.3.1. In particular, the stationary fraction of servers with two or more tasks under a round-robin strategy remains positive in the limit as $N \to \infty$, whereas it vanishes under the JSQ policy. On a related account, since tasks get served and eventually depart in a queuing scenario, less balanced allocations with a large portion of vacant servers will generate fewer service completions and result in a larger total number of tasks. Thus different schemes yield not only various degrees of balance, but also variations in the aggregate number of tasks in the system, which is not the case in a balls-and-bins setup.

9. Extensions and Future Research Directions. Throughout most of the paper we have focused on the supermarket model as a canonical setup and adopted several common assumptions in that context: (i) all servers are identical; (ii) the service requirements are exponentially distributed; (iii) no advance knowledge of the service requirements is available; (iv) in particular, the service discipline at each server is oblivious to the actual service requirements. As mentioned earlier, the stochastic optimality of the JSQ policy, and hence its central role as an ideal performance benchmark, rely critically on these assumptions. The latter also broadly applies to the stochastic coupling techniques and asymptotic universality properties that we have considered in the previous sections.

In this section we turn to a brief overview of results for scenarios where some of the above assumptions are relaxed, in particular allowing for general service requirement distributions and possibly heterogeneous servers, along with some broader methodological issues. In section 9.1 we focus on the behavior of JSQ(d) policies in such scenarios, mainly in the large-N limit, while also briefly commenting on the JIQ policy. In section 9.2 we discuss strategies which specifically exploit knowledge of server speeds or service requirements of arriving tasks in making task assignment decisions, and may not necessarily use queue length information, mostly in a fixed-N regime. While nonexponential service requirement distributions and heterogeneous settings cover a major share of the extensions beyond the supermarket model, there is also a plethora of further model variations that have been considered in the literature. An exhaustive listing is simply out of reach, but some notable examples within the realm of scaling laws include [105, 106, 166, 168].

9.1. JSQ(d) Policies with General Service Requirement Distributions. The authors of [41, 42] use direct probabilistic methods and fluid limits to obtain stability conditions for finite-size systems with a renewal arrival process, an FCFS discipline at each server, various state-dependent routing policies, including JSQ, and general service requirement distributions, which may depend on the task type, the server, or both. Using fluid limits as well as Lyapunov functions, [20, 21] show that JSQ(d) policies achieve stability for any subcritical load in finite-size systems with a renewal arrival process, identical servers, nonidling local service disciplines, and general service requirement distributions. In addition, the author derives uniform bounds on the tails of the marginal queue length distributions and uses them to prove relative compactness of these distributions.

The authors of [22, 23] examine mean-field limits for JSQ(d) policies with generally distributed service requirements, leveraging the above-mentioned tail bounds and relative compactness. They establish that similar power-of-choice benefits occur as originally demonstrated for exponentially distributed service requirements in [113, 163], provided a certain "ansatz" holds asserting that finite subsets of queues become independent in the large-N limit. The latter "propagation of chaos" property is shown to hold in several settings, e.g., when the service requirement distribution has a decreasing hazard rate and the discipline at each server is FCFS, or when the service requirement distribution has a finite second moment and the load is sufficiently low. The ansatz also always holds for the power-of-d version of the JSW rather than JSQ policy; see Theorem 2.1 in [23].

It is further shown in [22, 23] that the arrival process at any given server tends to a state-dependent Poisson process, and that the queue length distribution becomes insensitive with respect to the service requirement distribution when the service discipline is either processor sharing or LCFS with preemptive resume. This may be explained based on the insensitivity property of queues with state-dependent Poisson arrivals and symmetric service disciplines.

There are strong plausibility arguments that a similar asymptotic insensitivity property should hold for the JIQ policy in a queuing scenario, even if the discipline at each server is not symmetric but FCFS, for example. So far, however, this has only been rigorously established for service requirement distributions with decreasing hazard rate in [141]. This result was in fact proved for systems with heterogeneous server pools and was further extended in [142] to systems with multiple symmetric dispatchers. As it turns out, general service requirement distributions with an increasing hazard rate give rise to major technical challenges due to the lack of certain

monotonicity properties. This has only allowed a proof of the asymptotic zero-wait property for the JIQ policy for load values strictly below 1/2 so far [44].

A fundamental technical issue associated with any general service requirement distribution is that the joint queue length no longer provides a suitable state description, and that the state space required for a Markovian description is no longer countable. The authors of [2,3] introduce a particle representation for the state of the system and describe the state dynamics for a JSQ(d) policy via a sequence of interacting measure-valued processes. They prove that as N grows large, a suitably scaled sequence of state processes converges to a hydrodynamic limit which is characterized as the unique solution of a countable system of coupled deterministic measure-valued equations, i.e., a system of PDEs rather than the usual ODEs. They also establish a "propagation of chaos" result, meaning that finite collections of queues are asymptotically independent.

The authors of [123, 125, 126] analyzed the performance and stability of static probabilistic routing strategies and power-of-d policies in the large-N limit in systems with exponential service requirement distributions, but heterogeneous server pools and a processor-sharing discipline at each server. They also considered variants of the JSQ(d) policy which account for the server speed in the selection criterion as well as hybrid combinations of the JSQ(d) policy with static probabilistic routing. Related results for heterogeneous loss systems rather than queuing scenarios are presented in [89, 124, 127]. As the results in [125, 126] reflect, ordinary JSQ(d) policies may fail to sample the faster servers sufficiently often in such scenarios, and therefore they can fail to achieve maximum stability, let alone asymptotic optimality. In [123] a weighted version of JSQ(d) policies is presented that does provide maximum stability, without requiring any specific knowledge of the underlying system parameters and server speeds in particular.

The authors of [158, 159, 160] examine mean-field limits for power-of-d policies in many-server loss systems as well processor-sharing queues with phase-type service requirement distributions. They observe that the fixed point suggests a similar insensitivity property of the stationary occupancy distribution, as mentioned above. In view of the insensitivity of loss systems with possibly state-dependent Poisson arrivals, this may be interpreted as an indirect indication that the arrival process at any given server pool tends to a state-dependent Poisson arrival process in the large-N limit. In a somewhat different strand of work, the authors of [83] investigate the behavior of blocking probabilities in various load regimes in systems with many single-server finite-buffer queues, a processor-sharing discipline at each server, and an insensitive routing policy.

9.2. Heterogeneous Servers and Knowledge of Service Requirements. The bulk of the literature has focused on systems with identical servers, and scenarios with nonidentical server speeds have received relatively limited attention. A natural extension of the JSQ policy is to assign jobs to the server with the normalized shortest queue length or, equivalently, assuming exponentially distributed service requirements, the shortest expected delay. While such a generalized JSQ (GJSQ) or shortest-expected-delay (SED) strategy tends to perform well [12], it is not strictly optimal in general [34], and the true optimal strategy may in fact have a highly complicated structure.

The authors of [137] present approximations for the performance of GJSQ policies in a fixed-N regime with generally distributed service requirements and a processor-sharing discipline at each server, extending the analysis in [72] for the ordinary JSQ

policy with homogeneous servers. In [81] necessary and sufficient conditions are established for JSQ(d) policies to be optimal in systems with heterogeneous server speeds in a classical heavy-traffic regime. The authors of [56] examine fluid limits for a system with both "fast" and "slow" servers and task assignment policies which may only have limited knowledge of the speeds of individual servers. In [29] and [174] fluid limits and heavy-traffic limits are investigated, respectively, for a somewhat related system where speeds are also heterogeneous but depend on the combination of the server and the task types due to affinity relations.

In a separate line of work, the authors of [37] consider static dispatching policies in a fixed-N regime with heterogeneous servers and an FCFS or processor-sharing discipline at each server. The assignment decision may depend on the service requirement of the arriving task, but not on the actual queue lengths or any other state information. In case of FCFS the optimal routing policy is shown to have a nested size interval structure, generalizing the strict size interval structure of the task assignment strategies in [76] which are optimal for homogeneous servers. In the case of processor sharing, the knowledge of the service requirements of arriving tasks is irrelevant, in the absence of any state information.

The authors of [5] consider static probabilistic routing policies in the somewhat similar setup of a fixed-N regime with multiple task types, servers with heterogeneous speeds, and a processor-sharing discipline at each server. The routing probabilities are selected so as to minimize either the global weighted holding cost or the expected holding cost for an individual task, and may depend on the type of the task and its service requirement, but not on any other state information.

When knowledge of the service requirements of arriving tasks is available, it is natural to exploit that for the purpose of local scheduling at the various servers and, for example, use size-based disciplines. The impact of the local scheduling discipline and server heterogeneity on the performance and degree of efficiency of load balancing strategies is examined in [30]. The authors of [71] show that any given dispatching policy can be augmented with a "guardrails" feature to ensure minimization of the mean delay in a classical heavy-traffic regime in systems where the local scheduling is governed by the shortest-remaining processing-time (SRPT) policy. An interesting broader issue concerns the relative benefits provided by exploiting knowledge of service requirements of arriving tasks versus using information on queue lengths or workloads at the various servers, which strongly depend on the service requirement distribution [77].

9.3. Open Problems and Emerging Research Directions. If we now return to scalable load balancing as the central theme of this survey, and consider the above-described extensions in that light, it is striking how scant the results are if any of the assumptions (i)–(iv) as stated at the beginning of section 2.2 are dropped. On further thought, the paucity of results from a scalability viewpoint is perhaps not so surprising since it is not even clear what the optimal achievable (delay) performance is in the absence of these assumptions, leaving aside any trade-off with communication overhead.

The graph-based load balancing scenario considered in section 6 moves beyond assumption (i) of all servers being identical as it entails that different incoming tasks can only be served by different subsets of the servers. Thus, it is not clear what the optimal assignment policy is, but since the server speeds are still homogeneous, it can be argued that the JSQ policy provides a bound on the achievable performance. The results obtained in [116] as reviewed in section 6 establish suitable conditions in terms

of the graph for that lower bound to be asymptotically achievable.

Important extensions of these results are presented in [136, 169], which allow for more general compatibility constraints between different task types and different servers represented in terms of a bipartite graph and examine conditions in terms of the latter graph for the achievable performance to be asymptotically equivalent to that in case of full compatibility. Informally speaking, both papers establish conditions in terms of the connectivity properties of the bipartite compatibility graph for similar performance to be achievable as in a fully flexible system. More specifically, the authors of [136] focus on scenarios with identical server speeds and uniform loads across the various job types and establish process-level limits indicating convergence of the system occupancy under JSQ policies to that in the supermarket model with full flexibility. The authors of [169] allow for heterogeneous server speeds and arbitrary load distributions and use drift methods to prove bounds and demonstrate that speedaware extensions of the JSQ and JIQ strategies achieve vanishing waiting times and minimum expected response times. Interestingly, the results in [136, 169] also entail a certain notion of universality as in [116], with similar achievable performance as in a fully flexible system under relatively sparse compatibility relations. An open question is what the associated communication overhead is with these policies, and whether it is close in any sense to the minimum communication overhead required for asymptotically optimal delay performance.

A further extension of the above two models is where the service rates can depend in an arbitrary way on the pairwise combination of the task and the server. In that case it is also open what the minimum required overhead is to achieve asymptotically optimal performance, and it even remains to be established what the asymptotically optimal performance is.

Both these questions are also largely open for nonexponential service requirement distributions, even in the absence of any compatibility constraints. It is evident that for nearly deterministic service requirements, a zero mean waiting time can be achieved without any communication overhead at arbitrarily high subcritical load (using open-loop policies such as round-robin). It might thus be natural to expect that for extremely variable service requirements, correspondingly high communication overhead is needed to achieve a zero mean waiting time. However, this is countered by the asymptotic insensitivity of the JIQ policy, which has been proven for service requirement distributions with decreasing hazard rate, as mentioned earlier. Also, the amount of communication overhead can in fact be reduced by *not* issuing messages when a server is busy at predefined time instants rather than sending messages when a server is idle [152]. All in all, it seems largely open exactly how the amount of communication overhead required for vanishing waiting time depends on the service requirement distribution in conjunction with the system load.

Finally, throughout we have tacitly assumed that each task involves a single processing operation that can be handled by a single server. In reality, however, tasks can have a highly complex structure and consist of several subtasks that can be executed by multiple servers subject to certain precedence constraints; see, for instance, [75] for references and further background. The above questions also seem to be totally open in these scenarios.

REFERENCES

 M. Adler, S. Chakrabarti, M. Mitzenmacher, and L. Rasmussen, Parallel randomized load balancing, in Proc. STOC '95, 1995, pp. 238–247, https://doi.org/10.1145/225058.

- 225131. (Cited on p. 609)
- [2] R. AGHAJANI, X. LI, AND K. RAMANAN, The PDE method for the analysis of randomized load balancing networks, Proc. ACM Meas. Anal. Comput. Syst., 1 (2017), art. 38, https: //doi.org/10.1145/3154497. (Cited on p. 611)
- [3] R. AGHAJANI AND K. RAMANAN, The hydrodynamic limit of a randomized load balancing network, Ann. Appl. Probab., 29 (2019), pp. 2114–2174, http://arxiv.org/abs/1707.02005.
 (Cited on p. 611)
- [4] M. AKTAS, P. PENG, AND E. SOLJANIN, Effective straggler mitigation: Which clones should attack and when?, ACM Perform. Eval. Rev., 45 (2017), pp. 12–14. (Cited on p. 606)
- [5] E. Altman, U. Ayesta, and B. J. Prabhu, Load balancing in processor sharing systems, Telecommun. Syst., 47 (2011), pp. 35–48, https://doi.org/10.1007/s11235-010-9300-8.
 (Cited on p. 612)
- [6] G. ANANTHANARAYANAN, A. GHODSI, S. SHENKER, AND I. STOICA, Effective straggler mitigation: Attack of the clones, in Proc. NSDI '13, 2013, pp. 185–198, https://www.usenix.org/ conference/nsdi13/technical-sessions/presentation/ananthanarayanan. (Cited on p. 605)
- [7] L. L. H. Andrew, M. Lin, and A. Wierman, Optimality, fairness, and robustness in speed scaling designs, ACM SIGMETRICS Perform. Eval. Rev., 38 (2010), pp. 37–48, https: //doi.org/10.1145/1811099.1811044. (Cited on p. 601)
- [8] R. Atar, A diffusion regime with nondegenerate slowdown, Oper. Res., 60 (2012), pp. 490–500, https://doi.org/10.1287/opre.1110.1030. (Cited on pp. 560, 607)
- R. Atar, I. Keslassy, and G. Mendelson, Subdiffusive load balancing in time-varying queueing systems, Oper. Res., 67 (2019), pp. 1678–1698, https://doi.org/10.1287/opre. 2019.1851. (Cited on p. 606)
- [10] Y. AZAR, A. Z. BRODER, A. R. KARLIN, AND E. UPFAL, Balanced allocations, in Proc. STOC '94, 1994, pp. 593–602. (Cited on p. 609)
- [11] R. BADONNEL AND M. BURGESS, Dynamic pull-based load balancing for autonomic servers, in Proc. IEEE NOMS, 2008, pp. 751–754, https://doi.org/10.1109/NOMS.2008.4575205. (Cited on pp. 558, 563, 594)
- [12] S. A. BANAWAN AND J. ZAHORJAN, Load sharing in heterogeneous queueing systems, in Proc. IEEE INFOCOM '89, Vol. 2, 1989, pp. 731–739, https://doi.org/10.1109/INFCOM.1989. 101521. (Cited on p. 611)
- [13] S. BANERJEE AND D. MUKHERJEE, Join-the-shortest queue diffusion limit in Halfin-Whitt regime: Tail asymptotics and scaling of extrema, Ann. Appl. Probab., 29 (2019), pp. 1262–1309, https://doi.org/10.1214/18-AAP1436, http://arxiv.org/abs/1803.03306. (Cited on p. 570)
- [14] S. BANERJEE AND D. MUKHERJEE, Join-the-shortest queue diffusion limit in Halfin-Whitt regime: Sensitivity on the heavy traffic parameter, Ann. Appl. Probab., 30 (2020), pp. 80– 144, https://doi.org/10.1214/19-AAP1496. (Cited on pp. 570, 571, 572)
- [15] M. Benaïm and J.-Y. Le Boudec, A class of mean field interaction models for computer and communication systems, Perform. Eval., 65 (2008), pp. 823–838, https://doi.org/10. 1016/j.peva.2008.03.005. (Cited on p. 566)
- [16] M. Benaïm and J.-Y. Le Boudec, On Mean Field Convergence and Stationary Regime, preprint, https://arxiv.org/abs/1111.5710, 2011. (Cited on p. 567)
- [17] P. BERENBRINK, A. CZUMAJ, A. STEGER, AND B. VÖCKING, Balanced allocaton: The heavily loaded case, in Proc. STOC '00, 2000, pp. 745–754. (Cited on p. 609)
- [18] P. BERENBRINK, A. CZUMAJ, A. STEGER, AND B. VÖCKING, Balanced allocations: The heavily loaded case, SIAM J. Comput., 35 (2006), pp. 1350–1385, https://doi.org/10.1137/ S009753970444435X. (Cited on p. 609)
- [19] S. BHAMIDI, A. BUDHIRAJA, AND M. DEWASKAR, Near Equilibrium Fluctuations for Supermarket Models with Growing Choices, preprint, https://arxiv.org/abs/2006.03621, 2020. (Cited on pp. 568, 574)
- [20] M. Bramson, Stability of two families of queueing networks and a discussion of fluid limits, Queueing Syst., 28 (1998), pp. 7–31, https://doi.org/10.1023/A:1019182619288. (Cited on p. 610)
- [21] M. Bramson, Stability of join the shortest queue networks, Ann. Appl. Probab., 21 (2011), pp. 1568–1625, https://doi.org/10.1214/10-AAP726. (Cited on p. 610)
- [22] M. BRAMSON, Y. LU, AND B. PRABHAKAR, Randomized load balancing with general service time distributions, in Proc. ACM SIGMETRICS 2010, 2010, pp. 275–286, https://doi. org/10.1145/1811039.1811071. (Cited on p. 610)
- [23] M. BRAMSON, Y. LU, AND B. PRABHAKAR, Asymptotic independence of queues under randomized load balancing, Queueing Syst., 71 (2012), pp. 247–292, https://doi.org/10.1007/ s11134-012-9311-0. (Cited on p. 610)

- [24] A. BRAVERMAN, Steady-state analysis of the join-the-shortest-queue model in the Halfin-Whitt regime, Math. Oper. Res., 45 (2020), pp. 1069–1103. (Cited on pp. 570, 571, 572)
- [25] G. BRIGHTWELL, M. FAIRTHORNE, AND M. J. LUCZAK, The supermarket model with bounded queue lengths in equilibrium, J. Stat. Phys., 173 (2018), pp. 1149–1194, https://doi.org/ 10.1007/s10955-018-2044-7. (Cited on p. 574)
- [26] A. BUDHIRAJA, P. DUPUIS, M. FISCHER, AND K. RAMANAN, Limits of relative entropies associated with weakly interacting particle systems, Electron. J. Probab., 20 (2015), art. 80. (Cited on p. 567)
- [27] A. BUDHIRAJA AND E. FRIEDLANDER, Diffusion approximations for load balancing mechanisms in cloud storage systems, Adv. Appl. Probab., 51 (2019), pp. 41–86, https://doi.org/10. 1017/apr.2019.3. (Cited on p. 573)
- [28] A. BUDHIRAJA, D. MUKHERJEE, AND R. WU, Supermarket model on graphs, Ann. Appl. Probab., 29 (2019), pp. 1740–1777, https://doi.org/10.1214/18-AAP1437, https://arxiv.org/abs/1712.07607. (Cited on pp. 587, 588)
- [29] E. CARDINAELS, S. C. BORST, AND J. S. H. VAN LEEUWAARDEN, Job assignment in large-scale service systems with affinity relations, Queueing Syst., 93 (2019), pp. 227–268, https://doi.org/10.1007/s11134-019-09633-y. (Cited on p. 612)
- [30] H.-L. CHEN, J. R. MARDEN, AND A. WIERMAN, On the impact of heterogeneity and back-end scheduling in load balancing designs, in Proc. IEEE INFOCOM '09, 2009, pp. 2267–2275, https://doi.org/10.1109/INFCOM.2009.5062152. (Cited on p. 612)
- [31] A. CZUMAJ, F. MEYER AUF DER HEIDE, AND V. STEMANN, Shared memory simulations with triple-logarithmic delay, in Algorithms-ESA '95, Lecture Notes in Comput. Sci. 979, Springer, Berlin, Heidelberg, 1995, pp. 46-59, https://doi.org/10.1007/3-540-60313-1_ 133. (Cited on p. 609)
- [32] R. W. R. Darling and J. R. Norris, Differential equation approximations for Markov chains, Probab. Surveys, 5 (2008), pp. 37–79, https://doi.org/10.1214/07-PS121. (Cited on p. 566)
- [33] M. DIETZFELBINGER AND F. MEYER AUF DER HEIDE, Simple, efficient shared memory simulations, in Proc. SPAA '93, 1993, pp. 110–119, https://doi.org/10.1145/165231.165246. (Cited on p. 609)
- [34] A. EPHREMIDES, P. VARAIYA, AND J. WALRAND, A simple dynamic routing problem, IEEE Trans. Automat. Control, 25 (1980), pp. 690–693. (Cited on pp. 556, 561, 611)
- [35] P. ESCHENFELDT AND D. GAMARNIK, Supermarket queueing system in the heavy traffic regime. Short queue dynamics, Ann. Appl. Probab., to appear. (Cited on pp. 572, 573)
- [36] P. ESCHENFELDT AND D. GAMARNIK, Join the shortest queue with many servers. The heavy traffic-asymptotics, Math. Oper. Res., 43 (2018), pp. 867–886, https://doi.org/moor.2017. 0887. (Cited on pp. 566, 569, 570, 581, 588, 596)
- [37] H. Feng, V. Misra, and D. Rubenstein, Optimal state-free, size-aware dispatching for heterogeneous M/G/-type systems, Perform. Eval., 62 (2005), pp. 475–492, https://doi. org/10.1016/j.peva.2005.07.031. (Cited on p. 612)
- [38] R. FOLEY AND R. MCDONALD, Join the shortest queue: Stability and exact asymptotics, Ann. Appl. Prob., 11 (2001), pp. 569-607. (Cited on p. 587)
- [39] G. FOSCHINI AND J. SALZ, A basic dynamic routing problem and diffusion, IEEE Trans. Commun., 26 (1978), pp. 320–327, https://doi.org/10.1109/TCOM.1978.1094075. (Cited on p. 606)
- [40] G. J. FOSCHINI, On heavy traffic diffusion analysis and dynamic routing in packet switched networks, Comp. Perf., (1977), pp. 499–513. (Cited on p. 606)
- [41] S. G. Foss and N. I. Chernova, Ergodicity of partially accessible multichannel communication systems, Probl. Peredachi Inf., 27 (1991), pp. 9–14, http://www.mathnet.ru/php/archive.phtml?wshow=paper{&}jrnid=ppi{&}paperid=552{&}option{_-}lang=eng. (Cited on p. 610)
- [42] S. G. FOSS AND N. I. CHERNOVA, On the stability of a partially accessible multi-station queue with state-dependent routing, Queueing Syst., 29 (1998), pp. 55–73, https://doi.org/10. 1023/A:1019175812444. (Cited on p. 610)
- [43] S. G. Foss and N. I. Chernova, On optimality of the FCFS discipline in multiserver queueing systems and networks, Siberian Math. J., 42 (2001), pp. 372–385. (Cited on p. 562)
- [44] S. G. Foss and A. L. Stolyar, Large-scale join-idle-queue system with general service times, J. Appl. Probab., 54 (2017), pp. 995–1007, https://doi.org/10.1017/jpr.2017.49. (Cited on p. 611)
- [45] D. FOTAKIS, R. PAGH, P. SANDERS, AND P. SPIRAKIS, Space efficient hash tables with worst case constant access time, Theory Comput. Syst., 38 (2005), pp. 229–248, https://doi. org/10.1007/s00224-004-1195-x. (Cited on p. 609)

- [46] B. Fralix, C. Knessl, and J. S. H. van Leeuwaarden, First passage times to congested states of many-server systems in the Halfin-Whitt regime, Stoch. Models, 30 (2014), pp. 162–186, https://doi.org/10.1080/15326349.2014.900385. (Cited on p. 571)
- [47] C. FRICKER AND N. GAST, Incentives and redistribution in homogeneous bike-sharing systems with stations of finite capacity, EURO J. Transp. Logist., 5 (2016), pp. 261–291, https: //doi.org/10.1007/s13676-014-0053-5. (Cited on p. 567)
- [48] D. Gamarnik and D. A. Goldberg, On the rate of convergence to stationarity of the M/M/N queue in the Halfin-Whitt regime, Ann. Appl. Probab., 23 (2013), pp. 1879– 1912, https://doi.org/10.1214/12-AAP889, https://arxiv.org/abs/1003.2004. (Cited on p. 571)
- [49] D. Gamarnik and D. A. Goldberg, Steady-state GI/GI/N queue in the Halfin-Whitt regime, Ann. Appl. Probab., 23 (2013), pp. 2382–2419, https://doi.org/10.1214/12-AAP905, https://arxiv.org/abs/1103.1709. (Cited on p. 571)
- [50] D. GAMARNIK, J. TSITSIKLIS, AND M. ZUBELDIA, Delay, memory, and messaging tradeoffs in distributed service systems, in Proc. ACM SIGMETRICS 2016/Performance 2016, 2016, pp. 1–12, https://doi.org/10.1145/2896377.2901478. (Cited on pp. 557, 608)
- [51] D. GAMARNIK, J. TSITSIKLIS, AND M. ZUBELDIA, Delay, memory, and messaging tradeoffs in distributed service systems, Stoch. Syst., 8 (2018), pp. 45–74. (Cited on pp. 557, 608)
- [52] D. GAMARNIK, J. TSITSIKLIS, AND M. ZUBELDIA, A lower bound on the queueing delay in resource constrained load balancing, Ann. Appl. Prob., 30 (2020), pp. 870–901. (Cited on pp. 557, 608)
- [53] D. Gamarnik and A. Zeevi, Validity of heavy traffic steady-state approximations in generalized Jackson networks, Ann. Appl. Probab., 16 (2006), pp. 56–90, https://doi.org/10.1214/105051605000000638. (Cited on p. 567)
- [54] A. GANDHI, S. DOROUDI, M. HARCHOL-BALTER, AND A. SCHELLER-WOLF, Exact analysis of the M/M/k/setup class of Markov chains via recursive renewal reward, in Proc. ACM SIGMETRICS 2013, 2013, pp. 153–166, https://doi.org/10.1145/2465529.2465760. (Cited on p. 601)
- [55] K. GARDNER, M. HARCHOL-BALTER, A. SCHELLER-WOLF, M. VELEDNITSKY, AND S. ZBARSKY, Redundancy-d: The power of d choices for redundancy, Oper. Res., 65 (2017), pp. 1078– 1094, https://doi.org/10.1287/opre.2016.1582. (Cited on p. 605)
- [56] K. GARDNER AND C. STEPHENS, Smart dispatching in heterogeneous systems, ACM SIG-METRICS Perform. Eval. Rev., 47 (2019), pp. 12–14, https://doi.org/10.1145/3374888. 3374894. (Cited on p. 612)
- [57] K. GARDNER, S. ZBARSKY, S. DOROUDI, M. HARCHOL-BALTER, E. HYYTIÄ, AND A. SCHELLER-WOLF, Reducing latency via redundant requests, in Proc. ACM SIGMETRICS 2015, 2015, pp. 347–360, https://doi.org/10.1145/2796314.2745873. (Cited on p. 605)
- [58] K. Gardner, S. Zbarsky, S. Doroudi, M. Harchol-Balter, E. Hyytiä, and A. Scheller-Wolf, Queueing with redundant requests: Exact analysis, Queueing Syst., 83 (2016), pp. 227–259, https://doi.org/10.1007/s11134-016-9485-y. (Cited on p. 605)
- [59] K. GARDNER, S. ZBARSKY, M. HARCHOL-BALTER, AND A. SCHELLER-WOLF, The power of d choices for redundancy, in Proc. ACM SIGMETRICS 2016/Performance 2016, 2016, pp. 409-410, https://doi.org/10.1145/2896377.2901497. (Cited on p. 605)
- [60] K. GARDNER, S. ZBARSKY, M. VELEDNITSKY, M. HARCHOL-BALTER, AND A. SCHELLER-WOLF, Understanding response time in the redundancy-d system, ACM SIGMETRICS Perform. Eval. Rev., 44 (2016), pp. 33–35, https://doi.org/10.1145/3003977.3003989. (Cited on p. 605)
- [61] N. Gast, The power of two choices on graphs: The pair-approximation is accurate, ACM SIG-METRICS Perform. Eval. Rev., 43 (2015), pp. 69–71, https://doi.org/10.1145/2825236. 2825263. (Cited on pp. 567, 587)
- [62] N. GAST, Expected values estimated via mean-field approximation are 1/N-accurate, Proc. ACM Meas. Anal. Comput. Syst., 1 (2017), art. 17, https://doi.org/10.1145/3084454. (Cited on p. 566)
- [63] N. GAST, L. BORTOLUSSI, AND M. TRIBASTONE, Size expansions of mean field approximation: Transient and steady-state analysis, Perform. Eval., 129 (2019), pp. 60–80, https://doi. org/10.1016/j.peva.2018.09.005. (Cited on p. 566)
- [64] N. GAST, D. LATELLA, AND M. MASSINK, A refined mean field approximation for synchronous population processes, ACM SIGMETRICS Perform. Eval. Rev., 46 (2018), pp. 30–32, https://doi.org/10.1145/3305218.3305230. (Cited on p. 566)
- [65] N. GAST AND B. VAN HOUDT, A refined mean field approximation, Proc. ACM Meas. Anal. Comput. Syst., 1 (2018), art. 33, https://doi.org/10.1145/3154491. (Cited on p. 566)
- [66] R. J. Gibbens, P. Hunt, and F. P. Kelly, Bistability in communication networks, in Dis-

- order in Physical Systems, G. Grimmett and D. Welsh, eds., The Clarendon Press, 1990, pp. 113–128. (Cited on p. 567)
- [67] D. GOLDSZTAJN, A. FERRAGUT, F. PAGANINI, AND M. JONCKHEERE, Controlling the number of active instances in a cloud environment, in Proc. IFIP Performance 2017, 2017, pp. 15–20. (Cited on p. 601)
- [68] G. H. Gonnet, Expected length of the longest probe sequence in hash code searching, J. ACM, 28 (1981), pp. 289–304, https://doi.org/10.1145/322248.322254. (Cited on p. 609)
- [69] C. Graham, Chaoticity on path space for a queueing network with selection of the shortest queue among several, J. Appl. Probab., 37 (2000), pp. 198–211, https://doi.org/10.1239/ jap/1014842277. (Cited on p. 567)
- [70] C. GRAHAM, Functional central limit theorems for a large network in which customers join the shortest of several queues, Probab. Theory Related Fields, 131 (2005), pp. 97–120, https://doi.org/10.1007/s00440-004-0372-9. (Cited on p. 567)
- [71] I. GROSOF, Z. SCULLY, AND M. HARCHOL-BALTER, Load balancing guardrails: Keeping your heavy traffic on the road to low response times, Proc. ACM Meas. Anal. Comput. Syst., 3 (2019), art. 42, https://doi.org/10.1145/3341617.3326157. (Cited on p. 612)
- [72] V. GUPTA, M. HARCHOL-BALTER, K. SIGMAN, AND W. WHITT, Analysis of join-the-shortest-queue routing for web server farms, Perform. Eval., 64 (2007), pp. 1062–1081, https://doi.org/10.1016/j.peva.2007.06.012. (Cited on p. 611)
- [73] V. GUPTA AND N. WALTON, Load balancing in the non-degenerate slowdown regime, Oper. Res., 67 (2019), pp. 281–294, https://doi.org/10.1287/opre.2018.1768. (Cited on pp. 558, 560, 607)
- [74] S. Halfin and W. Whitt, Heavy-traffic limits for queues with many exponential servers, Oper. Res., 29 (1981), pp. 567–588, https://doi.org/10.1287/opre.29.3.567. (Cited on pp. 560, 562, 571, 572)
- [75] M. HARCHOL-BALTER, Open problems in queueing theory inspired by datacenter computing, Queueing Syst., 97 (2020), pp. 3–37. (Cited on p. 613)
- [76] M. HARCHOL-BALTER, M. E. CROVELLA, AND C. D. MURTA, On choosing a task assignment policy for a distributed server system, J. Parallel Distr. Comput., 59 (1999), pp. 204–228, https://doi.org/10.1006/JPDC.1999.1577. (Cited on p. 612)
- [77] M. HARCHOL-BALTER, A. SCHELLER-WOLF, AND A. R. YOUNG, Surprising results on task assignment in server farms with high-variability workloads, in Proc. ACM SIGMET-RICS '09/Performance '09, 2009, pp. 287–298, https://doi.org/10.1145/1555349.1555383. (Cited on p. 612)
- [78] T. HELLEMANS, T. BODAS, AND B. VAN HOUDT, Performance analysis of workload dependent load balancing policies, Proc. ACM Meas. Anal. Comput. Syst., 3 (2019), art. 35, https: //doi.org/10.1145/3326150. (Cited on p. 606)
- [79] P. Hunt, Loss networks under diverse routing: The symmetric star network, Adv. Appl. Probab., 27 (1995), pp. 255–272. (Cited on p. 567)
- [80] P. Hunt and T. Kurtz, Large loss networks, Stochastic Process. Appl., 53 (1994), pp. 363-378, https://doi.org/10.1016/0304-4149(94)90071-X. (Cited on p. 568)
- [81] D. Hurtado-Lange and S. T. Maguluri, Throughput and delay optimality of power-of-d choices in inhomogeneous load balancing systems, Oper. Res. Lett., 49 (2021), pp. 616– 622, https://doi.org/10.1016/j.orl.2021.06.010. (Cited on p. 612)
- [82] P. K. Johri, Optimality of the shortest line discipline with state-dependent service rates, European J. Oper. Res., 41 (1989), pp. 157–161, https://doi.org/10.1016/0377-2217(89) 90378-0. (Cited on pp. 578, 584, 585)
- [83] M. JONCKHEERE AND B. J. PRABHU, Asymptotics of insensitive load balancing and blocking phases, in ACM SIGMETRICS 2016/Performance 2016, 2016, pp. 311–322, https://doi. org/10.1145/2896377.2901454. (Cited on p. 611)
- [84] G. Joshi, Boosting service capacity via adaptive task replication, ACM SIGMETRICS Perform. Eval. Rev., 45 (2017), pp. 9–11, https://doi.org/10.1145/3152042.3152046. (Cited on p. 606)
- [85] G. Joshi, Synergy via redundancy: Boosting service capacity via adaptive task replication, in Proc. IFIP Performance 2017, 2017, pp. 21–28. (Cited on p. 606)
- [86] G. JOSHI, E. SOLJANIN, AND G. WORNELL, Efficient replication of queued tasks for latency reduction in cloud systems, in Proc. Allerton 2015, 2015, pp. 107–114, https://doi.org/ 10.1109/ALLERTON.2015.7446992. (Cited on p. 606)
- [87] G. JOSHI, E. SOLJANIN, AND G. WORNELL, Efficient redundancy techniques for latency reduction in cloud systems, ACM Trans. Model. Perform. Eval. Comput. Syst., 2 (2017), pp. 1–30, https://doi.org/10.1145/3055281. (Cited on p. 606)
- [88] W. Kang and K. Ramanan, Asymptotic approximations for stationary distributions of many-

- server queues with abandonment, Ann. Appl. Probab., 22 (2012), https://doi.org/10.1214/10-AAP738. (Cited on p. 567)
- [89] A. KARTHIK, A. MUKHOPADHYAY, AND R. R. MAZUMDAR, Choosing among heterogeneous server clouds, Queueing Syst., 85 (2017), pp. 1–29, https://doi.org/10.1007/s11134-016-9488-8. (Cited on pp. 578, 611)
- [90] K. Kenthapadi and R. Panigrahy, Balanced allocation on graphs, in Proc. SODA '06, SIAM, Philadelphia, 2006, pp. 434–443. (Cited on p. 609)
- [91] T. G. Kurtz, Strong approximation theorems for density dependent Markov chains, Stoch. Proc. Appl., 6 (1978), pp. 223–240, https://doi.org/10.1016/0304-4149(78)90020-0. (Cited on p. 566)
- [92] J. LE BOUDEC, D. MCDONALD, AND J. MUNDINGER, A generic mean field convergence result for systems of interacting objects, in Proc. QUEST '07, 2007, pp. 3–18, https://doi.org/ 10.1109/QEST.2007.8. (Cited on p. 566)
- [93] J.-Y. LE BOUDEC, The stationary behaviour of fluid limits of reversible processes is concentrated on stationary points, Netw. Het. Media, 8 (2013), pp. 529–540. (Cited on p. 567)
- [94] M. LIN, Z. LIU, A. WIERMAN, AND L. L. H. ANDREW, Online algorithms for geographical load balancing, in Proc. IGCC'12, 2012, pp. 1–10, https://doi.org/10.1109/IGCC.2012. 6322266. (Cited on p. 601)
- [95] M. Lin, A. Wierman, L. L. H. Andrew, and E. Thereska, Dynamic right-sizing for power-proportional data centers, IEEE/ACM Trans. Netw., 21 (2013), pp. 1378–1391, https://doi.org/10.1109/TNET.2012.2226216. (Cited on p. 601)
- [96] X. LIU AND L. YING, Universal scaling of distributed queues under load balancing in the super-Halfin-Whitt regime, IEEE/ACM Trans. Netw., 30 (2022), pp. 190–201, https:// arxiv.org/abs/1912.11904. (Cited on p. 574)
- [97] X. LIU AND L. YING, A simple steady-state analysis of load balancing algorithms in the sub-Halfin-Whitt regime, J. Appl. Probab., 57 (2020), pp. 578–596, https://doi.org/10.1017/ jpr.2020.13. (Cited on p. 574)
- [98] Z. LIU, Y. CHEN, C. BASH, A. WIERMAN, D. GMACH, Z. WANG, M. MARWAH, AND C. HYSER, Renewable and cooling aware workload management for sustainable data centers, ACM SIGMETRICS Perform. Eval. Rev., 40 (2012), pp. 175–186, https://doi.org/10.1145/ 2318857.2254779. (Cited on p. 601)
- [99] Z. LIU, M. LIN, A. WIERMAN, S. H. LOW, AND L. L. H. ANDREW, Geographical load balancing with renewables, ACM SIGMETRICS Perform. Eval. Rev., 39 (2011), pp. 62–66, https: //doi.org/10.1145/2160803.2160862. (Cited on p. 601)
- [100] Z. LIU, M. LIN, A. WIERMAN, S. H. LOW, AND L. L. H. ANDREW, Greening geographical load balancing, in Proc. ACM SIGMETRICS 2011, 2011, pp. 233-244, https://doi.org/ 10.1145/1993744.1993767. (Cited on p. 601)
- [101] Y. Lu, Q. Xie, G. Kliot, A. Geller, J. R. Larus, and A. Greenberg, Join-idle-queue: A novel load balancing algorithm for dynamically scalable web services, Perform. Eval., 68 (2011), pp. 1056–1071, https://doi.org/10.1016/j.peva.2011.07.015. (Cited on pp. 558, 563, 594, 598)
- [102] M. J. LUCZAK AND C. MCDIARMID, On the power of two choices: Balls and bins in continuous time, Ann. Appl. Probab., 15 (2005), pp. 1733–1764, https://doi.org/10.1214/105051605000000205. (Cited on p. 609)
- [103] M. J. LUCZAK AND C. McDIARMID, On the maximum queue length in the supermarket model, Ann. Probab., 34 (2006), pp. 493-527, https://doi.org/10.1214/00911790500000710. (Cited on p. 609)
- [104] M. J. LUCZAK AND C. McDIARMID, Asymptotic distributions and chaos for the supermarket model, Electron. J. Probab., 12 (2007), pp. 75–99. (Cited on p. 567)
- [105] M. J. LUCZAK AND C. MCDIARMID, Balanced routing of random calls, Ann. Appl. Probab., 25 (2015), pp. 1279–1324. (Cited on p. 610)
- [106] M. J. LUCZAK, C. MCDIARMID, AND E. UPFAL, On-line routing of random calls in networks, Probab. Theory Related Fields, 125 (2003), pp. 457–482. (Cited on p. 610)
- [107] M. J. LUCZAK AND J. R. NORRIS, Strong approximation for the supermarket model, Ann. Appl. Probab., 15 (2005), pp. 2038–2061, https://doi.org/10.1214/105051605000000368. (Cited on p. 568)
- [108] M. J. LUCZAK AND J. R. NORRIS, Averaging over fast variables in the fluid limit for Markov chains: Application to the supermarket model with memory, Ann. Appl. Probab., 23 (2013), pp. 957–986. (Cited on p. 567)
- [109] D. Martirosyan and P. Robert, The equilibrium states of large networks of Erlang queues, Adv. Appl. Probab., 52 (2020), pp. 617–654, https://arxiv.org/abs/1811.04763. (Cited on p. 567)

- [110] D. McDonald and S. Turner, Comparing load balancing algorithms for distributed queueing networks, in Analysis of Communication Networks: Call Centres, Traffic and Performance, Fields Inst. Commun. 28, AMS, 2000, pp. 109–133. (Cited on p. 587)
- [111] R. MENICH, Optimality of shortest queue routing for dependent service stations, in Proc. CDC '87, 1987, pp. 1069–1072, https://doi.org/10.1109/CDC.1987.272565. (Cited on pp. 578, 584, 585)
- [112] R. MENICH AND R. F. SERFOZO, Optimality of routing and servicing in dependent parallel processing systems, Queueing Syst., 9 (1991), pp. 403–418, https://doi.org/10.1007/ BF01159224. (Cited on pp. 578, 584, 585)
- [113] M. MITZENMACHER, The Power of Two Choices in Randomized Load Balancing, Ph.D. thesis, University of California, Berkeley, 1996. (Cited on pp. 567, 568, 587, 589, 610)
- [114] M. MITZENMACHER, The power of two choices in randomized load balancing, IEEE Trans. Parallel Distrib. Syst., 12 (2001), pp. 1094–1104, https://doi.org/10.1109/71.963420. (Cited on pp. 556, 562, 563, 567, 569)
- [115] M. MITZENMACHER, Analyzing distributed Join-Idle-Queue: A fluid limit approach, in Proc. Allerton 2016, 2016, pp. 312–318, https://doi.org/10.1109/ALLERTON.2016.7852246. (Cited on pp. 598, 600)
- [116] D. MUKHERJEE, S. BORST, AND J. VAN LEEUWAARDEN, Asymptotically optimal load balancing topologies, Proc. ACM Meas. Anal. Comput. Syst., 2 (2018), art. 14, https://doi.org/10. 1145/3179417. (Cited on pp. 554, 557, 558, 588, 612, 613)
- [117] D. MUKHERJEE, S. C. BORST, J. S. H. VAN LEEUWAARDEN, AND P. A. WHITING, Asymptotic optimality of power-of-d load balancing in large-scale systems, Math. Oper. Res., 45 (2020), pp. 1535–1571, https://doi.org/10.1287/moor.2019.1042. (Cited on p. 566)
- [118] D. MUKHERJEE, S. BORST, J. VAN LEEUWAARDEN, AND P. WHITING, Universality of load balancing schemes on the diffusion scale, J. Appl. Probab., 53 (2016), pp. 1111–1124, https://doi.org/10.1017/jpr.2016.68. (Cited on pp. 558, 566, 595)
- [119] D. Mukherjee, S. Borst, J. van Leeuwaarden, and P. Whiting, Universality of powerof-d load balancing in many-server systems, Stoch. Syst., 8 (2018), pp. 265–292, https: //doi.org/10.1287/stsy.2018.0016. (Cited on pp. 557, 558, 566, 568, 573)
- [120] D. Mukherjee, S. Borst, J. van Leeuwaarden, and P. Whiting, Asymptotic optimality of power-of-d load balancing in large-scale systems, Math. Oper. Res., 45 (2020), pp. 1535– 1571, https://doi.org/10.1287/moor.2019.1042. (Cited on pp. 558, 566, 578)
- [121] D. MUKHERJEE, S. DHARA, S. BORST, AND J. VAN LEEUWAARDEN, Optimal service elasticity in large-scale distributed systems, Proc. ACM Meas. Anal. Comput. Syst., 1 (2017), art. 25, https://doi.org/10.1145/3084463. (Cited on p. 601)
- [122] D. MUKHERJEE AND A. STOLYAR, Join-Idle-Queue with service elasticity: Large-scale asymptotics of a non-monotone system, Stoch. Syst., 9 (2019), pp. 338–358, http://arxiv.org/abs/1803.07689. (Cited on p. 603)
- [123] A. Mukhopadhyay, A. Karthik, and R. R. Mazumdar, Randomized assignment of jobs to servers in heterogeneous clusters of shared servers for low delay, Stoch. Syst., 6 (2016), pp. 90–131, https://doi.org/10.1214/15-SSY179. (Cited on p. 611)
- [124] A. MUKHOPADHYAY, A. KARTHIK, R. R. MAZUMDAR, AND F. GUILLEMIN, Mean field and propagation of chaos in multi-class heterogeneous loss models, Perform. Eval., 91 (2015), pp. 117–131, https://doi.org/10.1016/j.peva.2015.06.008. (Cited on pp. 578, 579, 611)
- [125] A. MUKHOPADHYAY AND R. R. MAZUMDAR, Rate-based randomized routing in large heterogeneous processor sharing systems, in Proc. ITC 26, 2014, pp. 1–9, https://doi.org/10. 1109/ITC.2014.6932934. (Cited on p. 611)
- [126] A. MUKHOPADHYAY AND R. R. MAZUMDAR, Analysis of randomized Join-the-Shortest-Queue (JSQ) schemes in large heterogeneous processor-sharing systems, IEEE Trans. Control Netw. Syst., 3 (2016), pp. 116–126, https://doi.org/10.1109/TCNS.2015.2428331. (Cited on p. 611)
- [127] A. MUKHOPADHYAY, R. R. MAZUMDAR, AND F. GUILLEMIN, The power of randomized routing in heterogeneous loss systems, in Proc. ITC '27, 2015, pp. 125–133, https://doi.org/10. 1109/ITC.2015.22. (Cited on pp. 578, 579, 611)
- [128] L. M. NGUYEN AND A. L. STOLYAR, A service system with randomly behaving on-demand agents, ACM SIGMETRICS Perform. Eval. Rev., 44 (2016), pp. 365–366, https://doi. org/10.1145/2964791.2901484. (Cited on p. 601)
- [129] R. PAGH AND F. F. RODLER, Cuckoo hashing, J. Algorithms, 51 (2004), pp. 122–144, https://doi.org/10.1016/j.jalgor.2003.12.002. (Cited on p. 609)
- [130] G. PANG AND A. L. STOLYAR, A service system with on-demand agent invitations, Queueing Syst., 82 (2016), pp. 259–283, https://doi.org/10.1007/s11134-015-9464-8. (Cited on p. 601)

- [131] R. PANIGRAHY, Efficient hashing with lookups in two memory accesses, in Proc. SODA '05, SIAM, Philadelphia, 2005, pp. 830–839. (Cited on p. 609)
- [132] J. PENDER AND T. PHUNG-DUC, A law of large numbers for M/M/c/delayoff-setup queues with nonstationary arrivals, in Proc. ASMTA '16, Cham, 2016, pp. 253–268, https://doi. org/10.1007/978-3-319-43904-4_18. (Cited on p. 601)
- [133] F. POLOCZEK AND F. CIUCU, Contrasting effects of replication in parallel systems: From overload to underload and back, in Proc. ACM SIGMETRICS 2016/Performance 2016, 2016, pp. 375–376, https://doi.org/10.1145/2964791.2901499. (Cited on p. 605)
- [134] M. I. REIMAN, Some diffusion approximations with state space collapse, in Modelling and Performance Evaluation Methodology, Springer, Berlin, Heidelberg, 1984, pp. 207–240. (Cited on p. 606)
- [135] P. ROBERT, Stochastic Networks and Queues, Springer, Berlin, Heidelberg, 2003, https://doi.org/10.1007/978-3-662-13052-0. (Cited on pp. 579, 580)
- [136] D. RUTTEN AND D. MUKHERJEE, Load balancing under strict compatibility constraints, in Proc. ACM SIGMETRICS 2021, 2021, pp. 51–52. (Cited on pp. 587, 613)
- [137] J. SELEN, I. ADAN, AND S. KAPODISTRIA, Approximate performance analysis of generalized join the shortest queue routing, in Proc. VALUETOOLS '16, 2016, pp. 103–110, https: //doi.org/10.4108/eai.14-12-2015.2262695. (Cited on p. 611)
- [138] N. B. SHAH, K. LEE, AND K. RAMCHANDRAN, When do redundant requests reduce latency?, IEEE Trans. Commun., 64 (2016), pp. 715–722, https://doi.org/10.1109/TCOMM.2015. 2506161. (Cited on p. 605)
- [139] P. D. SPARAGGIS, D. TOWSLEY, AND C. G. CASSANDRAS, Extremal properties of the shortest/longest non-full queue policies in finite-capacity systems with state-dependent service rates, J. Appl. Probab., 30 (1993), pp. 223–236. (Cited on pp. 578, 584, 585)
- [140] P. D. SPARAGGIS, D. TOWSLEY, AND C. G. CASSANDRAS, Sample path criteria for weak majorization, Adv. Appl. Probab., 26 (1994), pp. 155–171, https://doi.org/10.2307/1427584. (Cited on p. 584)
- [141] A. L. STOLYAR, Pull-based load distribution in large-scale heterogeneous service systems, Queueing Syst., 80 (2015), pp. 341–361, https://doi.org/10.1007/s11134-015-9448-8. (Cited on pp. 558, 594, 610)
- [142] A. L. STOLYAR, Pull-based load distribution among heterogeneous parallel servers: The case of multiple routers, Queueing Syst., 85 (2017), pp. 31–65, https://doi.org/10.1007/ s11134-016-9508-8. (Cited on pp. 598, 600, 610)
- [143] D. TOWSLEY, Application of majorization to control problems in queueing systems, in Scheduling Theory and Its Applications, P. Chrétienne, E. G. Coffman, J. K. Lenstra, and Z. Liu, eds., John Wiley & Sons, Chichester, UK, 1995, pp. 295–311. (Cited on p. 584)
- [144] D. TOWSLEY, P. SPARAGGIS, AND C. CASSANDRAS, Optimal routing and buffer allocation for a class of finite capacity queueing systems, IEEE Trans. Automat. Control, 37 (1992), pp. 1446–1451, https://doi.org/10.1109/9.159590. (Cited on p. 584)
- [145] J. N. TSITSIKLIS AND K. Xu, On the power of (even a little) resource pooling, Stoch. Syst., 2 (2012), pp. 1–66, https://doi.org/10.1287/11-SSY033. (Cited on pp. 567, 589)
- [146] J. N. TSITSIKLIS AND K. Xu, Flexible queueing architectures, Oper. Res., 65 (2017), pp. 1398–1413, https://doi.org/10.1287/opre.2017.1620. (Cited on p. 589)
- [147] S. R. TURNER, The effect of increasing routing choice on resource pooling, Probab. Eng. Inf. Sci., 12 (1998), pp. 109–124, https://doi.org/10.1017/S0269964800005088. (Cited on pp. 578, 587)
- [148] R. URGAONKAR, U. C. KOZAT, K. IGARASHI, AND M. J. NEELY, Dynamic resource allocation and power management in virtualized data centers, in Proc. IEEE NOMS 2010, 2010, pp. 479–486, https://doi.org/10.1109/NOMS.2010.5488484. (Cited on p. 601)
- [149] M. VAN DER BOOR, S. BORST, AND J. VAN LEEUWAARDEN, Load balancing in large-scale systems with multiple dispatchers, in Proc. IEEE INFOCOM 2017, 2017, pp. 1–9, https://doi.org/10.1109/INFOCOM.2017.8057012. (Cited on pp. 599, 600)
- [150] M. VAN DER BOOR, S. BORST, AND J. VAN LEEUWAARDEN, Hyper-scalable JSQ with sparse feedback, Proc. ACM Meas. Anal. Comput. Syst., 3 (2019), art 4. (Cited on pp. 607, 608)
- [151] M. VAN DER BOOR, S. BORST, J. VAN LEEUWAARDEN, AND D. MUKHERJEE, Scalable load balancing in networked systems: Universality properties and stochastic coupling methods, in Proc. Int. Cong. Math. 2018, B. Sirakov, P. N. de Souza, and M. Viana, eds., 2018, pp. 3881–3912. (Cited on p. 554)
- [152] M. VAN DER BOOR, M. ZUBELDIA AND S. BORST, Zero-wait load balancing with sparse messaging, Oper. Res. Lett., 48 (2020), pp. 368–375. (Cited on p. 613)
- [153] R. VAN DER HOFSTAD, Random Graphs and Complex Networks, Vol. 1, Cambridge University Press, Cambridge, 2017, https://doi.org/10.1017/9781316779422. (Cited on p. 594)

- [154] B. VAN HOUDT, Global attraction of ODE-based mean field models with hyperexponential job sizes, Proc. ACM Meas. Anal. Comput. Syst., 3 (2019), art. 23. (Cited on p. 567)
- [155] J. VAN LEEUWAARDEN AND C. KNESSL, Transient behavior of the Halfin-Whitt diffusion, Stoch. Proc. Appl., 121 (2011), pp. 1524–1545, https://doi.org/10.1016/j.spa.2011.03.007. (Cited on p. 571)
- [156] J. VAN LEEUWAARDEN AND C. KNESSL, Spectral gap of the Erlang A model in the Halfin-Whitt regime, Stoch. Syst., 2 (2012), pp. 149–207, https://doi.org/10.1214/10-SSY012. (Cited on p. 571)
- [157] J. VAN LEEUWAARDEN, B. MATHIJSEN, AND B. ZWART, Economies-of-scale in many-server queueing systems: Tutorial and partial review of the QED Halfin-Whitt heavy-traffic regime, SIAM Rev., 61 (2019), pp. 403-440, https://doi.org/10.1137/17M1133944. (Cited on p. 571)
- [158] T. VASANTAM, A. MUKHOPADHYAY, AND R. MAZUMDAR, Mean-field analysis of loss models with mixed-Erlang distributions under power-of-d routing, in ITC 29, 2017, pp. 250–258, https://doi.org/10.23919/ITC.2017.8064362. (Cited on p. 611)
- [159] T. VASANTAM, A. MUKHOPADHYAY, AND R. MAZUMDAR, The mean-field behavior of processorsharing systems with general job lengths under the SQ(d) policy, Perform. Eval., 127 (2018), pp. 120–153. (Cited on p. 611)
- [160] T. VASANTAM, A. MUKHOPADHYAY, AND R. MAZUMDAR, Insensitivity of the mean-field limit of loss systems under SQ(d) routing, Adv. Appl. Prob., 51 (2019), pp. 1027–1066, https://doi.org/10.1017/apr.2019.41. (Cited on p. 611)
- [161] B. VÖCKING, How asymmetry helps load balancing, in Proc. FOCS '99, 1999, pp. 131–140, https://dl.acm.org/doi/abs/10.5555/795665.796480. (Cited on p. 609)
- [162] A. VULIMIRI, P. B. GODFREY, R. MITTAL, J. SHERRY, S. RATNASAMY, AND S. SHENKER, Low latency via redundancy, in Proc. CoNEXT '13, 2013, pp. 283–294, https://doi.org/10.1145/2535372.2535392. (Cited on p. 605)
- [163] N. D. VVEDENSKAYA, R. L. DOBRUSHIN, AND F. I. KARPELEVICH, Queueing system with selection of the shortest of two queues: An asymptotic approach, Problemy Peredachi Informatsii, 32 (1996), pp. 20–34. (Cited on pp. 556, 562, 563, 567, 568, 569, 610)
- [164] D. WANG, G. JOSHI, AND G. WORNELL, Efficient task replication for fast response times in parallel computation, Proc. ACM SIGMETRICS Perform. Eval. Rev., 42 (2014), pp. 599– 600, https://doi.org/10.1145/2637364.2592042. (Cited on p. 605)
- [165] D. WANG, G. JOSHI, AND G. WORNELL, Using straggler replication to reduce latency in large-scale parallel computing, ACM SIGMETRICS Perform. Eval. Rev., 43 (2015), pp. 7–11, https://doi.org/10.1145/2847220.2847223. (Cited on p. 605)
- [166] W. Wang, Q. Xie, and M. Harchol-Balter, Zero queueing for multi-server jobs, Proc. ACM Meas. Anal. Comput. Syst., 5 (2021), art. 7. (Cited on p. 610)
- [167] R. R. Weber, On the optimal assignment of customers to parallel servers, J. Appl. Probab., 15 (1978), pp. 406–413, https://doi.org/10.2307/3213411. (Cited on p. 584)
- [168] W. Weng and W. Wang, Achieving zero asymptotic queueing delay for parallel jobs, Proc. ACM Meas. Anal. Comput. Syst., 4 (2020), art. 42. (Cited on p. 610)
- [169] W. Weng, X. Zhou, and R. Srikant, Optimal load balancing with locality constraints, Proc. ACM Meas. Anal. Comput. Syst., 4 (2020), art. 45. (Cited on pp. 587, 613)
- [170] W. WHITT, Blocking when service is required from several facilities simultaneously, AT&T Tech. J., 64 (1985), pp. 1807–1856. (Cited on p. 567)
- [171] A. WIERMAN, L. L. H. ANDREW, AND A. TANG, Power-aware speed scaling in processor sharing systems: Optimality and robustness, Perform. Eval., 69 (2012), pp. 601–622, https://doi.org/10.1016/j.peva.2012.07.002. (Cited on p. 601)
- [172] W. WINSTON, Optimality of the shortest line discipline, J. Appl. Probab., 14 (1977), pp. 181–189, http://www.jstor.org/stable/3213271. (Cited on pp. 556, 561, 584)
- [173] Q. XIE, X. DONG, Y. LU, AND R. SRIKANT, Power of d choices for large-scale bin packing: A loss model, Proc. ACM SIGMETRICS Perform. Eval. Rev., 43 (2015), pp. 321–334, https://doi.org/10.1145/2796314.2745849. (Cited on pp. 578, 579)
- [174] A. Yekkehkhany, A. Hojjati, and M. H. Hajiesmaili, GB-PANDAS: Throughput and heavy-traffic optimality analysis for affinity scheduling, ACM SIGMETRICS Perform. Eval. Rev., 45 (2017), pp. 2–14, https://doi.org/10.1145/3199524.3199528. (Cited on p. 612)
- [175] L. Ying, On the approximation error of mean-field models, in Proc. ACM SIGMETRICS 2016/Performance 2016, 2016, pp. 285–297, https://doi.org/10.1145/2896377.2901463. (Cited on p. 566)
- [176] L. Ying, Stein's method for mean field approximations in light and heavy traffic regimes, Proc. ACM Meas. Anal. Comput. Syst., 1 (2017), art. 12, https://doi.org/10.1145/3084449.

(Cited on pp. 566, 573)

- [177] L. Ying, R. Srikant, and X. Kang, The power of slightly more than one sample in randomized load balancing, in Proc. IEEE INFOCOM 2015, 2015, pp. 1131–1139, https://doi.org/10.1109/INFOCOM.2015.7218487. (Cited on p. 577)
- [178] H. ZHANG, G.-H. HSU, AND R. WANG, Heavy traffic limit theorems for a sequence of shortest queueing systems, Queueing Syst., 21 (1995), pp. 217–238. (Cited on p. 606)
- [179] Z. ZHAO, S. BANERJEE, AND D. MUKHERJEE, Many-Server Asymptotics for Join-the-Shortest Queue in the Super-Halfin-Whitt Scaling Window, preprint, https://arxiv.org/abs/2106. 00121, 2020. (Cited on p. 574)
- [180] X. Zhou, J. Tan, and N. Shroff, Heavy-traffic delay optimality in pull-based load balancing systems: Necessary and sufficient conditions, Proc. ACM Meas. Anal. Comput. Syst., 2 (2018), art. 41, https://doi.org/10.1145/3287323. (Cited on p. 607)
- [181] X. Zhou, F. Wu, J. Tan, Y. Sun, and N. Shroff, Designing low-complexity heavy-traffic delay-optimal load balancing schemes: Theory to algorithms, Proc. ACM Meas. Anal. Comput. Syst., 1 (2017), art. 39, https://doi.org/10.1145/3154498. (Cited on p. 606)