

Queueing System Topologies with Limited Flexibility

John N. Tsitsiklis †
MIT, LIDS
Cambridge, MA 02139
jnt@mit.edu

Kuang Xu †
MIT, LIDS
Cambridge, MA 02139
kuangxu@mit.edu

ABSTRACT

We study a multi-server model with n flexible servers and rn queues, connected through a fixed bipartite graph, where the level of flexibility is captured by the average degree, $d(n)$, of the queues. Applications in content replication in data centers, skill-based routing in call centers, and flexible supply chains are among our main motivations.

We focus on the scaling regime where the system size n tends to infinity, while the overall traffic intensity stays fixed. We show that a large capacity region (robustness) and diminishing queueing delay (performance) are jointly achievable even under very limited flexibility ($d(n) \ll n$). In particular, when $d(n) \gg \ln n$, a family of random-graph-based interconnection topologies is (with high probability) capable of stabilizing all admissible arrival rate vectors (under a bounded support assumption), while simultaneously ensuring a diminishing queueing delay, of order $\ln n/d(n)$, as $n \rightarrow \infty$. Our analysis is centered around a new class of virtual-queue-based scheduling policies that rely on dynamically constructed partial matchings on the connectivity graph.

Categories and Subject Descriptors

G.3 [Probability and Statistics]: Queuing theory; C.2.1 [Network Architecture and Design]: Distributed networks

General Terms

Performance, Theory

Keywords

queueing system, flexibility, partial resource pooling, random graph, expander graph, asymptotics

1. INTRODUCTION

At the heart of many modern queueing networks lies the problem of allocating processing resources (e.g., manufacturing plants,

web servers, or call-center staff) to meet multiple types of demands that arrive dynamically over time (e.g., orders, data queries, or customer inquiries). It is often the case that a *fully flexible* or *completely resource-pooled* system, where every unit of processing resource is capable of serving all types of demands, delivers the best possible performance. Our inquiry is, however, motivated by the unfortunate reality that such full flexibility is often infeasible due to overwhelming implementation costs (in the case of a data center) or human skill limitations (in the case of a skill-based call center).

What are the key benefits of flexibility and resource pooling in such queueing networks? Can we harness the same benefits even when the degree of flexibility is *limited*, and how should the network be designed and operated? These are the main questions that we address in this paper. While these questions can be approached from several different angles, we will focus on the metrics of *expected queueing delay* and *capacity region*; the former is a direct reflection of *performance*, while the latter measures the system's *robustness* against *demand uncertainties*, when the arrival rates for different demand types are unknown or likely to fluctuate over time. Our main message is positive: in the regime where the system size is large, improvements in both capacity region and delay are *jointly achievable* even under very limited flexibility, given a proper choice of the interconnection topology.

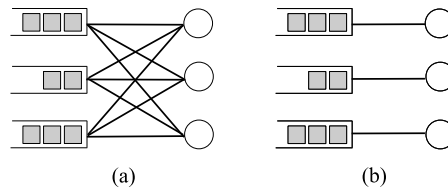


Figure 1: Extreme (in)flexibility: $d(n) = n$ vs. $d(n) = 1$.

Benefits of Full Flexibility. We begin by illustrating the benefits of flexibility and resource pooling using two simple examples.¹ Consider a system of n servers, each running at rate 1, and n queues, where each queue stores jobs of a particular demand type. For each $i \in \{1, \dots, n\}$, queue i receives a Poisson arrival stream of rate $\lambda_i = \rho \in (0, 1)$, independently from all other queues. The job sizes are independent and exponentially distributed with mean 1.

For the remainder of this paper, we will use as a measure of flexibility the average number of servers that a demand type can

† Research supported in part by the National Science Foundation, under grant CMMI-1234062, and an MIT-Xerox Fellowship. The authors are grateful for the feedback from the anonymous reviewers.

This document is an extended technical report for a conference paper appearing at SIGMETRICS'13, June, 2013, [26]. Updated February 22, 2016 for corrections in Figure 3.

¹The lack of mathematical rigor in this section is intended to make the results easier to state. Formal definitions will be given in later sections.

receive service from, denoted by $d(n)$. Let us consider the two extreme cases: a fully flexible system, with $d(n) = n$ (Figure 1(a)), and an inflexible system, with $d(n) = 1$ (Figure 1(b)). Fixing the traffic intensity ρ , and letting the system size, n , tend to infinity, we observe the following qualitative benefits of full flexibility.

1. Diminishing Delay: Let W be the steady-state expected time in queue. In the fully flexible case and under any work-conserving scheduling policy², the *collection* of all jobs in the system evolves as an $M/M/n$ queue with arrival rate ρn . It is not difficult to verify that the expected total number of jobs in queue is *bounded* by a constant independent of n , and hence the expected waiting time in queue (the time from entering the queue to the initiation of service) satisfies $\mathbb{E}(W) \rightarrow 0$, as $n \rightarrow \infty$.³ In contrast, the inflexible system is simply a collection of n *unrelated* $M/M/1$ queues, and hence the expected waiting time is $\mathbb{E}(W) = \frac{\rho}{1-\rho} > 0$, for all n . In other words, expected delay *diminishes* in a fully flexible system, as the system size increases, but stays bounded away from zero in the inflexible case.

2. Large Capacity Region: Suppose that we now allow the arrival rate λ_i to queue i to vary with i . For the fully flexible case, and treating it again as an $M/M/n$ queue, it is easy to see that the system is stable for all arrival rate vectors that satisfy $\sum_{i=1}^n \lambda_i < n$, whereas in the inflexible system, since all $M/M/1$ queues operate independently, we must have $\lambda_i < 1$, for all i , in order to achieve stability. Comparing the two, we see that the fully flexible system attains a much larger capacity region, and is hence more robust to uncertainties or changes in the arrival rates.

3. Joint Achievability: Finally, it is remarkable, though perhaps obvious, that a fully flexible system can achieve both benefits (diminishing delay and large capacity region) *simultaneously*, without sacrificing one for the other. In particular, for any $\rho \in (0, 1)$, the condition $\sum_{i=1}^n \lambda_i < \rho n$ for all n implies that $\mathbb{E}(W) \rightarrow 0$, as $n \rightarrow \infty$.

Preview of Main Result. Will the above benefits of flexibility continue to be present if the system is no longer fully flexible (i.e., if $d(n) \ll n$)? The main result of the paper (Theorem 1) shows that the objectives of diminishing delay and a large capacity region can still be jointly achieved, even when the amount of flexibility in the system is limited ($\ln n \ll d(n) \ll n$), as long as the arrival rates are appropriately bounded. However, when flexibility is limited, the interconnection topology and scheduling policy need to be chosen with care: our solution is based on connectivity graphs generated by the Erdős-Rényi random bipartite graph construction, combined with a new class of scheduling policies that rely on dynamically constructed partial matchings. Furthermore, the scheduling policies are completely oblivious to the exact values of λ_i , and adapt to them automatically.

1.1 Motivating Applications

We describe here some motivating applications for our model, which share a common overall architecture illustrated in Figure 2. **Content replication** is commonly used in data centers for bandwidth intensive operations such as database queries [2] or video streaming [9], by hosting the same piece of content on multiple servers. Here, a server corresponds to a physical machine in a data center, and each queue stores incoming demands for a particular piece of content (e.g., a video clip). A server j is connected to

²Under a work-conserving policy, a server is always busy whenever there is at least one job in the queues to which it is connected.

³The diminishing expected waiting time follows from the bounded expected total number of jobs in steady-state, Little's Law, and the fact that the total arrival rate is ρn , which goes to infinity as $n \rightarrow \infty$.

queue i if there is a copy of content i on server j , and $d(n)$ corresponds to the average number of replicas per content across the network. Similar structures also arise in **skill-based routing (SBR) in call centers**, where agents (servers) are assigned to answer calls from different categories (queues) based on their domains of expertise [15], and in **process-flexible supply chains** [3]-[8], where each plant (server) is capable of producing multiple product types (queues). In many of these applications, demand rates can be unpredictable and may change significantly over time (for instance, unexpected "spikes" in demand traffic are common in modern data centers [1]). The demand uncertainties make *robustness* an important criterion for system design. These practical concerns have been our primary motivation for studying the *trade-off* between robustness, performance, and the level of flexibility.

1.2 Related Work

Bipartite graphs serve as a natural model of the relationships between demand types and service resources. It is well known in the supply chain literature that limited flexibility, corresponding to a sparse bipartite graph, can be surprisingly effective in resource allocation even when compared to a fully flexible system [3, 4, 5, 6, 7]. The use of sparse random graphs or expanders as flexibility structures that improve robustness has recently been studied in [8] in the context of supply chains, and in [9] for content replication. Similar to the robustness results reported in this paper, these works show that expanders or random graphs can accommodate a large set of demand rate vectors. However, in contrast to our work, nearly all analytical results in this literature focus on static allocation problems, where one tries to match supply with demand in a single slot, as opposed to the queueing context, where resource allocation decisions need to be made dynamically over time.

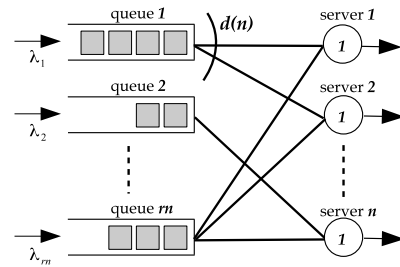


Figure 2: A processing network with rn queues and n servers.

In the queueing theory literature, the models we consider fall under the umbrella of so-called multi-class multi-server systems, where multiple queues and servers are connected through a bipartite graph. Under these (and similar) settings, complete resource pooling (full flexibility) is known to improve system performance [12, 13, 14], but much less is known when only limited flexibility is available, because systems with a non-trivial connectivity graph are very difficult to analyze, even under seemingly simple scheduling policies (e.g. first-come-first-serve) [10, 11]. Simulations in [15] show empirically that limited cross-training can be highly effective in a large call center under a skill-based routing algorithm. Using a very different set of modeling assumptions, [16] proposes a specific chaining structure with limited flexibility, which is shown to perform well under heavy traffic. Closer to the spirit of the current work is [17], which studies a partially flexible system where a fraction $p > 0$ of all processing resources are fully flexible, while the remaining fraction, $1 - p$, is dedicated to specific demand types, and shows an exponential improvement in de-

lay scaling under heavy-traffic. However, both [16] and [17] focus on the heavy-traffic regime, which is different from the current setting where traffic intensity is assumed to be fixed while the system size tends to infinity, and provide analytical results for the special case of uniform demand rates. Furthermore, with a constant fraction of fully flexible resources, the average degree in [17] scales linearly with the system size n , whereas we are interested in the case of a much slower degree scaling. Finally, the expected delay in the architecture considered in [17] does not tend to zero as the system size increases.

At a higher level, our work is focused on the joint trade-off between robustness, delay, and the degree of flexibility in a queueing network, which is little understood in the existing literature, especially for networks with a non-trivial interconnection topology.

On the technical end, we build on several existing ideas. The techniques of batching (cf. [18, 19]) and the use of virtual queues (cf. [20, 21]) have appeared in many contexts in queueing theory, but the specific models considered in the literature bear little resemblance to ours. The study of perfect matchings on a random bipartite graph dates back to the seminal work in [22]; while it has become a rich topic in combinatorics, we will refrain from giving a thorough literature survey because only some elementary and standard properties of random graphs are used in the current paper.

Organization of the Paper. We describe the model in Section 2 along with the notation to be used throughout. The main result is stated in Section 3, with a discussion of potential improvements postponed till Section 7. Section 3.2 studies a specific flexibility structure and compares it against the architecture proposed in this paper. The rest of the paper is devoted to the proof of our main result (Theorem 1), with an overview of the proof strategy given in Section 3.3.

2. MODEL AND NOTATION

Notation. To avoid clutter, we will minimize the use of floor and ceiling notation throughout the paper (e.g., writing T_{rn} instead of $T_{\lfloor rn \rfloor}$). We will assume that all values of interest are appropriately rounded up or down to an integer, whenever doing so does not cause ambiguity or confusion. We denote by \mathcal{G}_n the set of all $rn \times n$ bipartite graphs. For $g_n \in \mathcal{G}_n$, let $\deg(g_n)$ be the average degree among the rn left vertices. An $m \times n$ Erdős-Rényi random bipartite graph $G = (E, I \cup J)$,⁴ where each of the mn edges is present with probability p , is referred to as an (m, n, p) random graph. We will use $\mathbb{P}_{n,p}(\cdot)$ to denote the probability measure on \mathcal{G}_n induced by an (rn, n, p) random graph,

$$\mathbb{P}_{n,p}(g) = p^{|E|} (1-p)^{rn^2-|E|}, \quad \forall g \in \mathcal{G}_n, \quad (1)$$

where $|E|$ is the cardinality of the set of edges, E . For a subset of vertices, $M \subset I \cup J$, we will denote by $g|_M$ the graph induced by g on the vertices in M . Throughout, we will use the letter K to denote a generic constant.

The following short-hand notation for asymptotic comparisons will be used; here f and g are positive functions:

1. $f(x) \lesssim g(x)$ for $f(x) = \mathcal{O}(g(x))$, and $f(x) \gtrsim g(x)$ for $f(x) = \Omega(g(x))$.
2. $f(x) \gg g(x)$ for $\liminf_{x \rightarrow \infty} \frac{f(x)}{g(x)} = \infty$, and $f(x) \ll g(x)$ for $\limsup_{x \rightarrow \infty} \frac{f(x)}{g(x)} = 0$.
3. $f(x) \sim g(x)$ for $\lim_{x \rightarrow \infty} \frac{f(x)}{g(x)} = 1$.

⁴Throughout, we denote by E , I , and J the set of edges and left and right vertices, respectively.

When labeling a sequence, we will use the notation $a(n)$ if the value of a has a meaningful dependence on n , or a_n if n serves merely as an index. We will use $\text{Expo}(\lambda)$, $\text{Geo}(p)$, $\text{Bino}(n, p)$ as shorthands for the exponential, geometric and binomial distributions with the standard parameters, respectively. The expression $X \stackrel{d}{=} Y$ means that X and Y have the same distribution. Whenever possible, we will use upper-case letters for random variables, and lower-case letters for deterministic values.

The Model. We consider a sequence of systems operating in *continuous time*, indexed by an integer n , where the n th system consists of n servers and rn queues, where r is a positive constant that is fixed as n varies (Figure 2). The connectivity of the system is encoded by an $rn \times n$ undirected bipartite graph $g_n = (E, I \cup J)$, where I and J represent the set of queues and servers, respectively, and E the set of edges between them.⁵ A server $j \in J$ is *capable* of serving a queue $i \in I$ if and only if $(i, j) \in E$. We will denote by $\mathcal{N}(i)$ the set of servers in J connected to queue i , and similarly, by $\mathcal{N}(j)$ the set of queues in I connected to server j .

In the n th system, each queue i receives a stream of incoming jobs according to a Poisson process of rate $\lambda_{n,i}$, independent of all other streams. We will denote by λ_n the arrival rate vector, i.e., $\lambda_n = (\lambda_{n,1}, \lambda_{n,2}, \dots, \lambda_{n,rn})$. The sizes of the jobs are exponentially distributed with mean 1, independent from each other and from the arrival processes. All servers are assumed to run at a constant rate of 1. The system is assumed to be empty at time $t = 0$.

Jobs arriving at queue i can be assigned to any server $j \in \mathcal{N}(i)$ to receive service. The assignment is *binding*, in the sense that once the assignment is made, the job cannot be transferred to, or simultaneously receive service from, any other server. Moreover, service is *non-preemptive*, in the sense that once service is initiated for a job, the assigned server has to dedicate its full capacity to that job until its completion.⁶ Formally, if a server j just completed the service of a previous job at time t , its available actions are: **1. Serve a new job:** Server j can choose to fetch a job from any queue in $\mathcal{N}(j)$ and immediately start service. The server will remain occupied and take no other actions until the current job is completed. **2. Remain idle:** Server j can choose to remain idle. While in the idling state, it will be allowed to initiate a service (Action 1) at any point in time.

The performance of the system is fully determined by a *scheduling policy*, π , which specifies for each server $j \in J$, when to remain idle, when to serve a new job, and from which queue in $\mathcal{N}(j)$ to fetch a job when initiating a new service.

We only allow policies that are causal, in the sense that the decision at time t depends only on the history of the system (arrivals and service completions) up to t . We allow the scheduling policy to be *centralized* (i.e., to have full control over all servers' actions), and to be based on the knowledge of the graph g_n and the past history of all queues and servers. On the other hand, a policy does *not* observe the actual sizes of jobs before they are served, and **does not know the arrival rate vector λ_n** .

⁵For notational simplicity, we omit the dependence of E , I , and J on n .

⁶While we restrict ourselves to binding and non-preemptive scheduling policies in this paper, other common architectures where (a) a server can serve multiple jobs concurrently (processor sharing), (b) a job can be served by multiple servers concurrently, or (c) job sizes are revealed upon entering the system, are clearly more powerful than the current setting, and are therefore capable of implementing the scheduling policy we consider here. Hence the performance upper bounds developed in this paper also apply to these more powerful variants.

Performance metric: Let W_k be the waiting time in queue for the k th job, defined as the time from the job's arrival to a queue until when it starts to receive service from some server. With a slight abuse of notation, we define the *expected waiting time*, $\mathbb{E}(W)$, as

$$\mathbb{E}(W) \triangleq \limsup_{k \rightarrow \infty} \mathbb{E}(W_k). \quad (2)$$

Note that $\mathbb{E}(W)$ captures the *worst-case* expected waiting time across all jobs in the long run, and is always well defined, even under scheduling policies that do not induce a steady-state distribution.

We are primarily interested in the scaling of $\mathbb{E}(W)$ in the regime where the total traffic intensity (i.e., the ratio between the total arrival rate and the total service capacity) stays fixed, while the size of the system, n , tends to infinity.

3. MAIN THEOREM

Before stating our main theorem, we first motivate a condition on the arrival rate vector, λ_n . Since we allow the arrival rates $\lambda_{n,i}$ to vary with i , and since on average each queue is connected to $d(n)$ servers, the range of fluctuations of the $\lambda_{n,i}$ with respect to i should not be too large compared to $d(n)$, or else the system would become unstable. We will therefore let the amount of *rate fluctuation* or *uncertainty* be bounded by some $u(n)$, an upper bound on $\lambda_{n,i}$ for all i . The following condition summarizes these points.

Condition 1. (Rate Condition) We fix a constant $\rho \in (0, 1)$ (referred to as the traffic intensity) and a sequence $\{u(n)\}_{n \geq 1}$ of positive reals with $\inf_n u(n) > 0$. For any $n \geq 1$, the arrival rate vector λ_n satisfies:

1. $\max_{1 \leq i \leq rn} \lambda_{n,i} \leq u(n)$.
2. $\sum_{i=1}^{rn} \lambda_{n,i} \leq \rho n$.

We denote by $\Lambda_n \subset \mathbb{R}_+^n$ the set of all arrival rate vectors that satisfy the above conditions.

We denote by $\mathbb{E}_\pi(W | g_n, \lambda_n)$ the expected waiting time (cf. Eq. (2)) under a scheduling policy π , a graph g_n , and an arrival rate vector λ_n . The following is the main result of the paper.

Theorem 1. We assume that $d(n) \gg \ln n$ and $u(n) \ll \sqrt{\frac{d(n)}{\ln n}}$, and fix $\rho \in (0, 1)$, $\gamma > 0$, and $n \geq 1$. Then, there exists a policy π_n such that for every arrival vector λ_n that satisfies Condition 1, the following hold.

1. There exists a bipartite graph, $g_n \in \mathcal{G}_n$ (possibly depending on λ_n) such that $\deg(g_n) \leq (1 + \gamma)d(n)$, and

$$\mathbb{E}_{\pi_n}(W | g_n, \lambda_n) \leq \frac{K u^2(n) \ln n}{d(n)} = o(1), \quad (3)$$

where $K > 0$ is a constant independent of n , g_n and λ_n .

2. Let G_n be a random graph in \mathcal{G}_n , chosen at random according to the Erdős-Rényi measure $\mathbb{P}_{n, \frac{d(n)}{n}}$ (cf. Eq. (1)). Then G_n has the property in Part 1, with high probability, uniformly over all $\lambda_n \in \Lambda_n$. Formally,

$$\mathbb{P}_{n, \frac{d(n)}{n}}(G_n \text{ has the property in Part 1}) \geq 1 - \delta_n, \quad \forall \lambda_n \in \Lambda_n,$$

where $\{\delta_n\}_{n \geq 1}$ is a sequence with $\lim_{n \rightarrow \infty} \delta_n = 0$.

The reader is referred to Section 7 for a discussion on potential improvements of the result.

3.1 Remarks on Theorem 1

Note that Part 1 of the theorem is a special case of Part 2, which states that for large values of n , and for a graph chosen according to the Erdős-Rényi model, the scheduling policy will have a large probability (with respect to the random graph generation) of achieving diminishing delay, as in Eq. (3). At a higher level, Eq. (3) relates three key characteristics of the system: *delay* $\mathbb{E}(W)$, *level of robustness* $u(n)$, and *degree of flexibility* $d(n)$. Furthermore, the scheduling policy only requires knowledge of ρ , not of the arrival rate vector; this is important in practice because this vector need not be known *a priori* or may change significantly over time.

The fact that a large capacity region (as given by Condition 1) is achievable when $d \ll n$ should not be too surprising: expander graphs (such as those obtained through a random graph construction) are known to have excellent max-flow performance. The more difficult portion of the argument is to show that a diminishing delay (a temporal property) can also be achieved without significantly sacrificing the range of admissible arrival rates, through appropriately designed scheduling policies. In our proof, the achievability of a large capacity region is not shown explicitly, but comes as a by-product of the delay analysis; in fact, the scheduling policy automatically finds, over time, a feasible flow decomposition over the connectivity graph (see Section 6).

Dependence on ρ .

It is possible to incorporate the traffic intensity ρ in the leading constant in Eq. (3), to obtain a bound of the form

$$\mathbb{E}_{\pi_n}(W | g_n, \lambda_n) \leq \frac{K'}{1 - \rho} \cdot \frac{u^2(n) \ln n}{d(n)}, \quad (4)$$

where K' is a constant independent of ρ . This captures a trade-off between the traffic intensity and the degree of flexibility in the system. A proof of Eq. (4) is given in Appendix A.7.

Adversarial Interpretation of Rate Robustness.

To better understand the power and limitations of the ‘‘rate robustness’’ entailed by Theorem 1, it is useful to view the system as a game between two players: an *Adversary* who chooses the rate vector λ_n , and a *Designer* who chooses the interconnection topology g_n and the scheduling policy π_n . The goal of the Designer is to achieve small average waiting time (which would also imply stability of all queues), while the goal of the Adversary is the opposite. The following definition will facilitate our discussion.

Definition 1. (Good Graphs) Let us fix n , $d(n)$, $u(n)$, ρ , γ , and K , as in Theorem 1. We define the set $\text{Good}(\lambda_n)$ of good graphs for λ_n as the subset of \mathcal{G}_n for which $\deg(g_n) \leq (1 + \gamma)d(n)$ and the inequality in Eq. (3) holds, for some policy π_n .

We now examine the robustness implications of Theorem 1 under different *orderings of the actions* in the game, listed according to increasing levels of difficulties for the Designer.

Level 1: Adversary acts first (weak adversary). The Designer gets to pick g_n and π_n *after* observing the realization of λ_n . Note that this weak form of adversary fails to capture the rate uncertainty arising in many applications, where the arrival rates are unknown at the time when the system is designed. For this setting, it is not hard to come up with a simple deterministic construction of a good graph and a corresponding policy. Formally, the set $\text{Good}(\lambda_n)$ is nonempty for every $\lambda_n \in \Lambda_n$. However, Part 1 of Theorem 1 actually makes a stronger statement. While the graph is chosen after observing λ_n , the policy is designed without knowledge of λ_n , albeit at the expense of a much more complex policy.

Level 2: Adversary and Designer act independently (intermediate adversary). Suppose that both the Adversary and the Designer make their decisions independently, *without* knowing each other’s choice. This case models most practical applications, where λ_n is generated by some exogenous mechanism, unrelated to the design process, and is not revealed to the Designer ahead of time. Here, our randomization in the construction of the graph offers protection against the choice of the Adversary. Part 2 of the theorem can be rephrased into the statement that there exists a policy for which

$$\inf_{\lambda_n \in \Lambda_n} \mathbb{P}_{n, \frac{d(n)}{n}} (G_n \in \text{Good}(\lambda_n)) \rightarrow 1. \quad (5)$$

As an extension, we may consider the case where the Adversary chooses λ_n at random, according to some probability measure μ_n on Λ_n , but still independently from G_n . However, this additional freedom to the Adversary does not make a difference, and it can be shown, through an easy application of Fubini’s Theorem, that

$$\inf_{\mu_n} \left(\mathbb{P}_{n, \frac{d(n)}{n}} \times \mu_n \right) (G_n \in \text{Good}(\lambda_n)) \rightarrow 1, \quad (6)$$

and where \times is used to denote product measure. The proof is given in Appendix A.8.

Level 3: Designer acts first (strong adversary). In this case, the Adversary chooses λ_n *after* observing the g_n and π_n picked by the Designer. While such an adversary may be too strong for most practical applications, it is still interesting to ask whether there exist *fixed* g_n and π_n that will work well for *any* λ_n that satisfies Condition 1 (or even weaker conditions). We conjecture that this is the case.

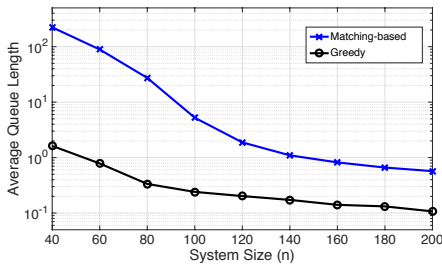


Figure 3: Diminishing Delay. Simulations with $n \times n$ networks ($r = 1$), $\lambda_{n,i} = \rho = 0.95$, and $d(n) = \sqrt{n}$.

On Practical Scheduling Policies.

The scheduling policy that we will use in the proof of Theorem 1 was mainly designed for the purpose of analytical tractability, rather than practical efficiency. For instance, simulations suggest that a seemingly naive greedy heuristic, where any free server chooses a job from a longest queue to which it is connected, can achieve a superior delay scaling (see Figure 3).⁷ Unfortunately, deriving an explicit delay upper bound for the greedy policy or other similar heuristics appears to be challenging. See also the discussion in Section 7.

3.2 A Comparison with Modular Architectures

Assume for simplicity that there is an equal number of queues and servers (i.e., $r = 1$). In a *modular architecture*, the designer partitions the system into $n/d(n)$ *separate* sub-networks. Each

⁷The scheduling policy being simulated is a discrete-time version of the continuous-time policy analyzed in this paper.

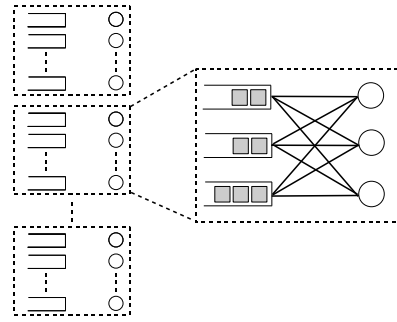


Figure 4: A modular architecture, consisting of $n/d(n)$ disjointed clusters, each of size $d(n)$. Within each cluster, all servers are connected to all queues.

sub-network consists of $d(n)$ queues and servers that are fully connected within the sub-network (Figure 4). Note that this architecture guarantees a degree of exactly $d(n)$, by construction. Assume that all queues have the same arrival rate $\rho \in (0, 1)$ and that each server uses an arbitrary work-conserving service policy. Since each sub-network is *fully connected*, it is not difficult to see that the modular architecture achieves a diminishing queuing delay as $n \rightarrow \infty$, as long as $d(n) \rightarrow \infty$. This seems even better than the random-graph-based architecture, which requires $d(n)$ to scale at least as fast as $\ln n$. However, this architecture fares much worse in terms of robustness, as we now proceed to discuss.

Since the sub-networks are completely disconnected from each other, the set of arrival rates that are admissible is severely reduced due to the requirement that the total arrival rate in *each* sub-network be less than $d(n)$, which is much more restrictive than the rate constraint given in Condition 1. Thus, the modular design may achieve a better delay scaling when the arrival rates are (almost) uniform, but at the risk of instability when the arrival rates are random or chosen adversarially, because the processing resources are not shared across different sub-networks.

3.3 Proof Overview

Sections 4 through 6 contain the proof of Theorem 1. We will provide an explicit construction of scheduling policies π_n , and then show that they achieve the delay scaling in Eq. (3). At the core of the construction is the use of *virtual queues*, whose operation is described in detail in Section 4.

The proof of Theorem 1 is completed in two steps. In Section 5, we first prove the result in the special case where all λ_i are equal to some $\lambda < 1/r$. The symmetry brought about by the uniform arrival rates will significantly simplify the notation, while almost all intuition developed here will carry over to the proof for the general case. We then show in Section 6 that, perhaps surprisingly, the original scheduling policies designed for the uniform arrival rate case can be applied directly to achieve the delay scaling in Eq. (3) for *any* λ_n satisfying Condition 1.

4. VIRTUAL QUEUE AND THE SCHEDULING POLICY

This section provides a detailed construction of a virtual-queue-based scheduling policy that achieves the delay scaling in Theorem 1. We begin by describing some high-level ideas behind our design.

Regularity vs. Discrepancies. Setting aside computational issues, an efficient scheduling policy is difficult to design because the future is unpredictable and random: one does not know *a pri-*

ori which part of the network will become more loaded, and hence current resource allocation decisions must take into account all possibilities of future arrivals and job sizes.

However, as the size of the system, n , becomes large, certain *regularities* in the arrival processes begin to emerge. To see this, consider the case where $r = 1$, $\lambda_{n,i} = \lambda < 1$ for all n and i , and assume that at time $t > 0$, all servers are busy serving some job. Now, during the interval $[t, t + \gamma_n)$, “roughly” $\lambda n \gamma_n$ new jobs will arrive, and $n \gamma_n$ servers will become available. For this $[t, t + \gamma_n)$ interval, denote by Γ the set of queues that received a job, and by Δ the set of roughly $n \gamma_n$ servers that became available. If γ_n is chosen so that $\lambda n \gamma_n \ll n$, these incoming jobs are likely to spread out across the queues, so that most queues receive at most one job. Assuming that this is indeed the case, we see that the connectivity graph g_n restricted to $\Gamma \cup \Delta$, denoted $g_n|_{\Gamma \cup \Delta}$, is a subgraph sampled uniformly at random among all $(\lambda n \gamma_n \times n \gamma_n)$ -sized subgraphs of g_n . When $n \gamma_n$ is sufficiently large, and g_n is *well connected* (as in an Erdős-Rényi random graph with appropriate edge probability), we may expect that, with high probability, $g_n|_{\Gamma \cup \Delta}$ admits a matching (Definition 4) that includes the entire set Γ , in which case all $\lambda n \gamma_n$ jobs can start receiving service by the end of the interval.

Note that when n is sufficiently large, despite the randomness in the arrivals, the symmetry in the system makes delay performance at a short time scale *insensitive* to the exact locations of the arrivals. Treated collectively, the structure of the set of arrivals and available servers in a small interval becomes less random and more “regular,” as $n \rightarrow \infty$. Of course, for any finite n , the presence of randomness means that *discrepancies* (events that deviate from the expected regularity) will be present. For instance, the following two types of events will occur with small, but non-negligible, probability.

1. Arrivals may be located in a poorly connected subset of g_n .
2. Arrivals may concentrate on a small number of queues.

We need to take care of these discrepancies, and show that their negative impact on performance is insignificant.

Following this line of thought, our scheduling policy aims to use most of the resources to dynamically target the *regular* portion of the traffic (via matchings on subgraphs), while ensuring that the impact of the *discrepancies* is well contained. In particular, we will create two classes of *virtual queues* to serve these two objectives:

1. A single *Matching queue* that targets *regularity* in arrival and service times, and *discrepancies* of the first type.
2. A collection of rn *Residual queues*, one for each (physical) queue, which targets the discrepancies in arrival and service times of the second type.

The queues are “virtual,” as opposed to “physical,” in the sense that they merely serve to conceptually simplify the description of the scheduling policy.

Good Graphs: The management of the virtual queues must comply with the underlying connectivity graph, g_n , which is fixed over time. We informally describe here some desired properties of g_n ; more detailed definitions and performance implications will be given in subsequent sections, as a part of the queueing analysis. We are interested in graphs that belong to a set \mathcal{H}_n , defined as the intersection of the following three subsets of \mathcal{G}_n :

1. $\hat{\mathcal{G}}_n$ (cf. Lemma 1): For the case $r = 1$, $g_n \in \hat{\mathcal{G}}_n$ if it admits a full matching. For the more general case, a suitable generalization is provided in the context of Lemma 1. This property will be used in both Matching and Residual queues to handle *discrepancies*.
2. $\tilde{\mathcal{G}}_n$ (cf. Lemmas 5-7) We have $g_n \in \tilde{\mathcal{G}}_n$ if a randomly sampled sublinear-sized subgraph admits a full matching, with

high probability. This property will be used in the Matching queue to take advantage of *regularity*.

3. \mathcal{L}_n (cf. Section 5.3) : We have $g_n \in \mathcal{L}_n$ if g_n has average degree approximately equal to $d(n)$. This property is to comply with our degree constraint.

Once the description of the policy is completed, the proof will consist of establishing the following:

- (i) If $g_n \in \mathcal{H}_n$, then the claimed delay bound holds.
- (ii) The probability that a random bipartite graph belongs to \mathcal{H}_n tends to 1, as $n \rightarrow \infty$.

Inputs to the Scheduling Policy: Besides n and r , the scheduling policy uses the following inputs:

1. ρ , the traffic intensity as defined in Condition 1,
2. ϵ , a constant in $(0, 1 - \rho)$,
3. $b(n)$, a batch size function,
4. g_n , the interconnection topology.

4.1 Arrivals to Virtual Queues

The arrivals to the Matching and Residual queues are arranged in *batches*. Roughly speaking, a batch is a set of jobs that are treated collectively as a single entity that arrives to a virtual queue. We define a sequence of random times $\{T_B(k)\}_{k \in \mathbb{Z}_+}$, by letting $T_B(0) = 0$, and for all $k \geq 1$,

$$T_B(k) \triangleq \text{time of the } k \frac{\rho}{r} b(n) \text{th arrival to the system,}$$

where $b(n) \in \mathbb{Z}_+$ is a design parameter, and will be referred to as the *batch parameter*.⁸ We will refer to the time period $(T_B(k-1), T_B(k)]$ as the *kth batch period*.

Definition 2. (Arrival Times to Virtual Queues) *The time of arrival of the kth batch to all virtual queues is $T_B(k)$, and the corresponding interarrival time is $A(k) \triangleq T_B(k+1) - T_B(k)$.*

While all virtual queues share the same arrival times, the contents of their respective batches are very different. We will refer to them as the *Matching batch* and *Residual batch*, respectively.

The *kth Matching batch* is the set of jobs that *first* arrive to their respective queues during the *kth* batch period. Since for each queue only the first job belongs to the Matching batch, it is convenient to represent the Matching batch as the set of queues that receive at least one job during the batch period.

The *kth* batch arriving to the *i*th Residual queue is the set of all jobs, except for the first one, that arrive to queue *i* during the *kth* batch.

Definition 3. (Size of a Residual Batch) *Let $H_i(k)$ be the total number of jobs arriving to queue i during the kth batch period. The size of the i-th Residual batch is*

$$R_i(k) \triangleq (H_i(k) - 1)^+. \quad (7)$$

A graphical illustration of the Matching and Residual batches is given in Figure 5.

⁸We chose not to absorb the constant ρ/r into $b(n)$ at this point because this will allow for simpler notation in subsequent sections.

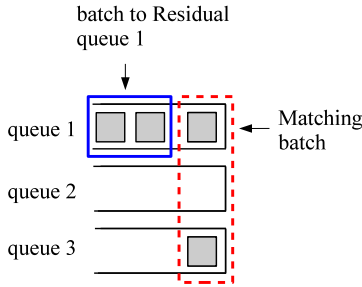


Figure 5: An illustration of the arrivals during a single batch period.

4.2 State Transitions and Service Rules

This section describes how the batches will be served by the physical servers. Before getting into the details, we first describe the general ideas.

1. The sizes of Residual batches are typically quite small, and the physical servers will process them using a (non-adaptive) randomized scheduling rule.
2. For each Matching batch, we will first “collect” a number of available servers, equal to the size of the batch:
 - (a) With high probability, all jobs in the Matching batch can be simultaneously processed by these servers through a matching over g_n .
 - (b) With small probability, some jobs in the Matching batch are located in a poorly connected subset of g_n and cannot be matched with the available servers. In this case, all jobs in the batch will be served, one at a time, according to a *fixed* server-to-queue assignment (a “clear” phase).

To implement the above queueing dynamics, we will specify the evolution of *states* and *actions* of all physical servers and virtual queues, to be described in detail in the remainder of this section.

4.2.1 States and Actions of (Physical) Servers and Residual Queues

We first introduce the notion of an assignment function, which will be used to schedule the physical servers. An assignment function L maps each queue $i \in I$ to a server $L(i) \in J$. We say that L is an *efficient assignment function* for the connectivity graph g_n if $(i, L(i)) \in E$ for all $i \in I$, and

$$|L^{-1}(j)| \leq r + 1, \quad (8)$$

for all $j \in J$. As will become clear in the sequel, our scheduling policy will use an assignment function L to ensure that every Residual queue receives at least a “minimum service rate” from some server. Let $\hat{\mathcal{G}}_n$ be the set of all $g_n \in \mathcal{G}_n$ such that g_n has an efficient assignment function. With our random graph construction, an efficient assignment function exists with high probability. The proof can be found in Appendix A.1.

Lemma 1. *Let $p(n) = d(n)/n$ and assume that $d(n) \gg \ln n$. Then,*

$$\lim_{n \rightarrow \infty} \mathbb{P}_{n,p(n)}(\hat{\mathcal{G}}_n) = 1.$$

A physical server can be, at any time, in one of two states: “BUSY” and “STANDBY;” cf. Fig. 4.2.1. The *end* of a period in the BUSY state will be referred to as an *event point*, and the time interval between two consecutive event points an *event period*. At each event

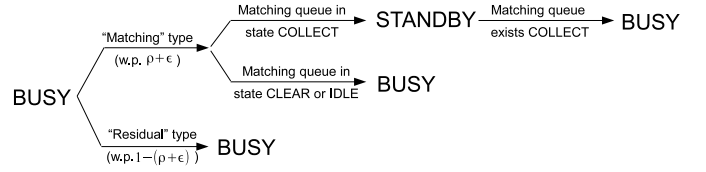


Figure 6: State evolution of a physical server.

point, a new decision is to be made as to which virtual queue the server will choose to serve, during the next event period. All servers are initialized in a BUSY state, with the time until the first event point distributed as $\text{Expo}(1)$, independently across all servers.

Recall that ϵ be a parameter in $(0, 1 - \rho)$. The state transitions defined below also involve the current state of the Matching queue, whose evolution will be given in Section 4.2.2. When a server $j \in J$ is at the k th event point:

1. With probability $\rho + \epsilon$, the k th event period is of type “matching”:
 - (a) If the Matching queue is in state COLLECT, server j enters state STANDBY.
 - (b) If the Matching queue is in state CLEAR, let B' be the Matching batch at the front of the Matching queue, and let $M' \subset I$ be the set of queues that still contain an unserved job from batch B' . Let $i^* = \min\{i : i \in M'\}$.
 - i) If $L(i^*) = j$, then server j starts processing the job in queue i^* that belongs to B' , entering state BUSY.
 - ii) Else, server j goes on a vacation of length $\text{Expo}(1)$, entering state BUSY.
 - (c) If the Matching queue is in state IDLE, server j goes on a vacation of length $\text{Expo}(1)$, entering state BUSY.
2. With probability $1 - (\rho + \epsilon)$, the k th event period is of type “residual.” Let i be an index drawn uniformly at random from the set $L^{-1}(j)$.
 - (a) If the i th Residual queue is non-empty, server j starts processing a job from the Residual batch that is currently at the front of the i th Residual queue, entering state BUSY.⁹
 - (b) Else, server j goes on a vacation of length $\text{Expo}(1)$, entering state BUSY.

The above procedure describes all the state transitions for a single server, except for one case: when in state STANDBY, any server can be ordered by the Matching queue to start processing a job, or initiate a vacation period. The transition out of the STANDBY state will be specified in the next subsection.

4.2.2 States and Actions of the Matching Queue

As mentioned earlier, the Matching queue is not a physical entity, but a mechanism that coordinates the actions of the physical servers. The name “Matching” reflects the fact that this virtual queue is primarily focused on using matchings to schedule batches, where a matching is defined as follows.

Definition 4. (Matching in a Bipartite Graph) *Let $F \subset E$ be a subset of the edges of $g = (E, I \cup J)$. We say that F is a matching, if $|\{j' : (i, j') \in F\}| \leq 1$, and $|\{i' : (i', j) \in F\}| \leq 1$, for all $i \in I, j \in J$. A matching F is said to contain $S \subset I \cup J$ if for all $s \in S$, there exists some edge in F that is incident on s , and is said to be full if it contains I or J .*

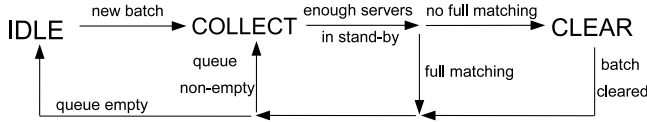


Figure 7: State evolutions of the Matching queue.

The Matching queue can be in one of three states: IDLE, COLLECT, and CLEAR. It is initialized to state IDLE.

1. If the Matching queue is in state IDLE, it remains so until the arrival of a Matching batch, upon which the Matching queue enters state COLLECT.
2. If the Matching queue is in state COLLECT, it remains so until there are exactly $(\rho/r)b(n)$ servers in state STANDBY.¹⁰ Denote by $\Gamma \subset I$ the Matching batch currently at the front of the Matching queue, and by $\Delta \subset J$ the set of servers in STANDBY at the end of the COLLECT period. The Matching queue makes the following decisions:
 - (a) If $g_n|_{\Gamma \cup \Delta}$ admits a full matching F :
 - i) Let all matched servers in Δ start processing a job in the current Matching batch according to the matching F , entering state BUSY. If there are unmatched servers in Δ (which will occur if $|\Gamma| < |\Delta| = (\rho/r)b(n)$), let all unmatched servers enter state BUSY by initiating a vacation, with a length independently distributed according to Expo (1).
 - ii) This completes the departure of a Matching batch from the Matching queue. If the Matching queue remains non-empty at this point, it returns to state COLLECT (to deal with the rest of the Matching batch); else, it enters state IDLE.
 - (b) If $g_n|_{\Gamma \cup \Delta}$ does not admit a full matching, the Matching queue enters state CLEAR. All servers in state STANDBY enter state BUSY by initiating a vacation, with a length independently distributed according to Expo (1).
3. If the Matching queue is in state CLEAR, it remains so until all jobs in the current Matching batch have started receiving processing from one of the servers, upon which it enters state COLLECT if the Matching queue remains non-empty, or state IDLE, otherwise.

By now, we have described how the batches are formed (Section 4.1), and how each physical server operates (Section 4.2). The scheduling policy is hence fully specified, provided that the underlying connectivity graph g admits an efficient assignment function.

5. DYNAMICS OF VIRTUAL QUEUES – UNIFORM ARRIVAL RATES

We now analyze the queueing dynamics induced by the virtual-queue-based scheduling policy introduced in Section 4. In this section, we will prove Theorem 1 for the special case of uniform arrival rates, where

$$\lambda_{n,i} = \lambda < 1/r, \quad (9)$$

$i \in I$. In particular, we have $\rho = \lambda r < 1$. Starting with this section, we will focus on a specific batch size function of the form

$$b(n) = K \frac{n \ln n}{d(n)} = K \frac{n}{y(n)}, \quad (10)$$

⁹The Residual batch at the front of the Residual queue is defined to be the oldest Residual batch that contains any as yet unserved job.

¹⁰This particular number is equal to the total number of jobs arriving in a single batch period.

for a suitably chosen constant K , where $y(n)$ is defined by

$$y(n) = \frac{d(n)}{\ln n}. \quad (11)$$

The specifics of the choice of K will be given later. We have assumed that $d(n) \gg \ln n$, and hence $y(n) \rightarrow \infty$ as $n \rightarrow \infty$. Note that this implies that $b(n) \ll n$. As will be seen soon, this guarantees that only a small fraction of the jobs will enter the Residual queues.

The main idea behind the delay analysis is rather simple: we will treat each virtual queue as a $GI/GI/1$ queue, and use Kingman's bound to derive an upper bound on the expected waiting time in queue. The combination of a batching policy with Kingman's bound is a fairly standard technique in queueing theory for deriving delay upper bounds (see, e.g., [19]). Our main effort will go into characterizing the various queueing primitives associated with the virtual queues (arrival rates, traffic intensity, and variances of inter-arrival and processing times), as well as in resolving the statistical dependence induced by the interactions among the physical servers. We begin by stating a simple fact on the inter-arrival-time distribution for the virtual queues. Even though in this section we assume that $r = 1$, we state the lemma for the general case.

Lemma 2. *The inter-arrival times of batches to the virtual queues, $\{A(k)\}_{k \geq 1}$, are i.i.d., with $\mathbb{E}(A(k)) = b(n)/(rn)$, and $\text{Var}(A(k)) \lesssim b(n)/n^2$.*

Proof. By definition, $A(k)$ is equal in distribution to the time until a Poisson process with rate λrn records $\lambda b(n)$ arrivals. Therefore, $\mathbb{E}(A(k)) = (\lambda b(n))/(\lambda rn) = b(n)/(rn)$, and $\text{Var}(A(k)) = \lambda b(n) \cdot \frac{1}{(\lambda rn)^2} \lesssim b(n)/n^2$. \square

Definition 5. (Service Times in Virtual Queues) *Consider the k th batch arriving at a virtual queue (Matching or Residual). Define the time of service initiation, E_k , to be when the batch first reaches the front of the queue, and the time of departure, D_k , to be when the last job in the batch starts receiving service from one of the physical servers. Let $D_k = E_k$ if the batch contains no jobs. The service time for the k batch is defined to be*

$$S(k) = D_k - E_k.$$

The interval $[E_k, D_k]$ is referred to as the service period of the k th batch.

We will denote by $S^M(k)$ and $S^R_j(k)$ the service time for the k th batch in the Matching queue and the j th Residual queue, respectively. Note that our scheduling policy (Section 4.2) induces interactions among the physical servers, and hence in general, the service times in a Residual queue are neither independent across batches or queues nor identically distributed.

5.1 Delays in Matching Queues

We first turn our attention to the service times $S^M(k)$ of the batches in the Matching queue. Based on our construction, it is not difficult to verify that the $S^M(k)$ are i.i.d. for any fixed graph, g_n . Furthermore, the value of $S^M(k)$ is equal to the length of a COLLECT period plus possibly the length of a CLEAR period, if the subgraph formed during the COLLECT period, $g_n|_{\Gamma \cup \Delta}$, fails to contain a matching that includes all queues in the current batch. We will therefore write

$$S^M(k) \stackrel{d}{=} S_{col} + X_{q(g_n)} \cdot S_{cle}, \quad (12)$$

where S_{col} and S_{cle} correspond to the length of a COLLECT and CLEAR period, respectively, and X_q is a Bernoulli random variable with $\mathbb{P}(X_q = 1) = q$. The value of $q(g_n)$ is defined by

$$q(g_n) = \mathbb{P}\{g_n|_{\Gamma \cup \Delta} \text{ does not admit a full matching}\}, \quad (13)$$

where the probability is taken over the randomness in Γ and Δ , but is conditional on $G_n = g_n$.

We begin by analyzing the properties of S_{cle} . Recall from Section 4.2.1 that, when the Matching queue is in state CLEAR, the jobs in the current Matching batch are to be served in a sequential fashion (Step 1-(b)) by a designated server. In particular, letting i^* be the smallest index of the queues that contain an unserved job from the current Matching batch, then the job in queue i^* , denoted by b , will be served by server $L(i^*)$ whenever that server enters an event period of type ‘‘matching.’’ No other unprocessed jobs in the Matching batch can be served until job b begins to receive service, and the amount of time it takes until this occurs is exponentially distributed with mean $1/(\rho + \epsilon)$, since the probability for an event period at a physical server to be of type ‘‘matching’’ is $(\rho + \epsilon)$ (Step 1). Since there are at most $\lambda b(n)$ jobs in a Matching batch, the length of a CLEAR period, S_{cle} , is no greater than the amount of time it takes for a Poisson process of rate $\rho + \epsilon$ to record $\lambda b(n)$ arrivals. Arguing similar to the proof of Lemma 2, we have the following lemma. Note that the distribution of S_{cle} does not depend on the structure of g_n , as long as g_n admits an efficient assignment function (cf. Lemma 1).

Lemma 3. $\mathbb{E}(S_{cle}) = \frac{\lambda}{\rho + \epsilon} b(n)$, and $\mathbb{E}(S_{cle}^2) \lesssim b^2(n)$.

To analyze S_{col} , we argue as follows. A COLLECT period consists of the time until a first server in Δ completes service, followed by the time until one of the remaining servers completes service, etc., until we have a total of $\lambda b(n)$ service completions. Thus the length of the period is the sum of $\lambda b(n)$ independent exponential random variables with parameters $\lambda(\rho + \epsilon)n$, $\lambda(\rho + \epsilon)(n - 1)$, \dots , $\lambda(\rho + \epsilon)(n - \lambda b(n) + 1)$. Using the fact that $b(n) \ll n$, this is essentially the same situation as in our analysis of S_{cle} , and we have the following result.

Lemma 4. $\mathbb{E}(S_{col}) \sim (\lambda/(\rho + \epsilon)) \cdot (b(n)/n)$ and $\mathbb{E}(S_{col}^2) \lesssim b^2(n)/n^2$.

Finally, we want to focus our attention to a set $\tilde{\mathcal{G}}_n \subset \mathcal{G}_n$ of graphs that possess the following two properties.

1. The value of $\mathbb{E}(X_{q(g_n)}) = q(g_n)$ is small, for all $g_n \in \tilde{\mathcal{G}}_n$. This property will help us upper bound the service time $S^M(k)$, using Eq. (12) and the moment bounds for S_{cle} and S_{col} developed earlier.
2. The set $\tilde{\mathcal{G}}_n$ has high probability under the Erdős-Rényi random graph model.

We start with some definitions. For $m \leq n$, let $\mathcal{M}_{m,n}$ be the family of all $m \times m$ subsets of $I \cup J$, that is, $h \in \mathcal{M}_{m,n}$ if and only if $|h \cap I| = |h \cap J| = m$. Let $m(n)$ be such that $m(n) \rightarrow \infty$ as $n \rightarrow \infty$. Let $\mathbf{P}_{\mathcal{M}_{m(n),n}}$ be a probability measure on $\mathcal{M}_{m(n),n}$. Let, for each $g \in \mathcal{G}_n$,

$$l(g) = \mathbf{P}_{\mathcal{M}_{m(n),n}}\left(\{h \in \mathcal{M}_{m(n),n} : g|_h \text{ does not admit a full matching}\}\right),$$

and $p(n) = d(n)/n$. We now define $\tilde{\mathcal{G}}_n$ as the set of graphs in \mathcal{G}_n that satisfy

$$l(g) \leq m^2(n) (1 - p(n))^{m(n)}. \quad (14)$$

Informally, this is a set of graphs which, for the given measure $\mathbf{P}_{\mathcal{M}_{m(n),n}}$ on subgraphs, have a high probability that a random subgraph will have a full matching. Consistent with the general outline of the proof given in Section 4, we will show that a) graphs in $\tilde{\mathcal{G}}_n$ have favorable delay guarantees (Proposition 1) and b) that random graphs are highly likely to belong to $\tilde{\mathcal{G}}_n$ (Lemma 5).

Lemma 5. (Probability of Full Matching on Random Subgraphs) With $p(n) = d(n)/n$, we have

$$\lim_{n \rightarrow \infty} \mathbb{P}_{n,p(n)}(\tilde{\mathcal{G}}_n) = 1, \quad (15)$$

Remark on Lemma 5: Eq. (15) states that graphs in $\tilde{\mathcal{G}}_n$ can be found using the Erdős-Rényi model, with high probability. This probabilistic statement is not to be confused with Eq. (14), which is a *deterministic* property that holds for all $g \in \tilde{\mathcal{G}}_n$. For the latter property, the randomness lies only in the sampling of subgraphs (via $\mathbf{P}_{\mathcal{M}_{m(n),n}}$). This distinction is important for our analysis, because the interconnection topology, g_n , stays fixed over time, while the random subgraph, $g_n|_{\Gamma \cup \Delta}$, is drawn independently for different batches.

Before proving Lemma 5, we state the following useful fact. The proof consists of a simple argument using Hall’s marriage theorem and a union bound (cf. Lemma 2.1 in [23] for a proof of the lemma).

Lemma 6. Let G be an (n, n, p) random bipartite graph. Then

$$\mathbb{P}_{n,p}(G \text{ does not admit a full matching}) \leq 3n(1 - p)^n. \quad (16)$$

Proof. (Lemma 5) Let H be a random element of $\mathcal{M}_{m(n),n}$, distributed according to $\mathbf{P}_{\mathcal{M}_{m(n),n}}(h)$. Let G be an $(rn, n, p(n))$ random bipartite graph over $I \cup J$, generated independently of H . Note that the distribution of G restricted to any $m(n)$ -by- $m(n)$ subset of $I \cup J$ is that of an $(m(n), m(n), p(n))$ random bipartite graph. Therefore, by Lemma 6, we have

$$\begin{aligned} \mathbb{E}(l(G)) &= \mathbb{P}(G|_H \text{ does not admit a full matching}) \\ &\leq 3m(n) (1 - p(n))^{m(n)}, \end{aligned} \quad (17)$$

where $\mathbb{P}(\cdot)$ represents the distributions of both G and H . Eq. (17), combined with Markov’s inequality, yields

$$\begin{aligned} \mathbb{P}_{n,p(n)}(\tilde{\mathcal{G}}_n) &= 1 - \mathbb{P}(l(G) > m^2(n) (1 - p(n))^{m(n)}) \\ &\geq 1 - \frac{\mathbb{E}(l(G))}{m^2(n) (1 - p(n))^{m(n)}} \geq 1 - \frac{3}{m(n)}, \end{aligned} \quad (18)$$

which converges to 1 as $n \rightarrow \infty$, because $m(n) \rightarrow \infty$. \square

Using Lemma 5, we can now obtain an upper bound on the value of $q(g_n)$. The proof is given in Appendix A.3.

Lemma 7. Let $b(n) = Kn/y(n)$, as in Eq. (10), for some constant $K > 8/\lambda$. If $g_n \in \tilde{\mathcal{G}}_n$, then

$$q(g_n) \leq \frac{1}{y^2(n)} \cdot n^{-(\lambda K - 4)/2}.$$

We are now ready to bound the mean and variance of the service time distribution for the Matching queue, using Eq. (12) and Lemmas 3, 4, and 7. The proof is given in Appendix A.4.

Lemma 8. (Service Time Statistics for the Matching Queue) Assume that $g_n \in \tilde{\mathcal{G}}_n$, and let $b(n) = Kn/y(n)$, for some constant $K > 8/\lambda$. The service times of the Matching batches, $S^M(k)$, are i.i.d., with

$$\mathbb{E}(S^M(k)) \sim \frac{\lambda}{\rho + \epsilon} \cdot \frac{1}{y(n)}, \text{ and } \text{Var}(S^M(k)) \lesssim \frac{1}{y^2(n)}.$$

Let W_M be the steady-state **waiting time of a batch in the Matching queue**, where waiting time is defined to be the time from when the batch is formed until when all jobs in the batch have started to receive service from a physical server. The following is the main result of this subsection.

Proposition 1. (Delays in the Matching Queue) Assume $g_n \in \hat{\mathcal{G}}_n \cap \tilde{\mathcal{G}}_n$, where $\hat{\mathcal{G}}_n$ and $\tilde{\mathcal{G}}_n$ were defined before Lemmas 1 and 5, respectively. We have

$$\mathbb{E}(W_M) \lesssim \frac{1}{y(n)}. \quad (19)$$

Proof. We use Kingman's bound, which states that the expected value of the waiting time in a $GI/GI/1$ queue, W , is bounded by $\mathbb{E}(W) \leq \tilde{\lambda} \frac{\sigma_a^2 + \sigma_s^2}{2(1-\tilde{\rho})}$, where $\tilde{\lambda}$ is the arrival rate, $\tilde{\rho}$ is the traffic intensity, and σ_a^2 and σ_s^2 are the variances for the interarrival times and service times, respectively. Using Lemmas 2 and 8, the claim follows by substituting

$$\begin{aligned} \tilde{\lambda} &= \frac{1}{\mathbb{E}(A(k))} \lesssim \frac{n}{b(n)} = \frac{y(n)}{K}, \\ \tilde{\rho} &= \frac{\mathbb{E}(S^M(k))}{\mathbb{E}(A(k))} \leq \frac{\frac{\lambda}{\rho+\epsilon} \frac{K}{y(n)}}{\frac{K}{ry(n)}} = \frac{\rho}{\rho+\epsilon} < 1, \\ \sigma_a^2 &= \text{Var}(A(k)) \sim \frac{1}{\lambda} \frac{b(n)}{n^2} \lesssim \frac{1}{ny(n)}, \\ \sigma_s^2 &= \text{Var}(S^M(k)) \lesssim \frac{1}{y(n)^2}, \end{aligned} \quad (20)$$

for all sufficiently large values of n . We obtain

$$\mathbb{E}(W_M) \lesssim y(n) \left(\frac{1}{ny(n)} + \frac{1}{y(n)^2} \right) \lesssim \frac{1}{y(n)}. \quad \square$$

5.2 Delays in a Residual Queue

The delay analysis for the Residual queues is conceptually simpler compared to that of the Matching queue, but requires more care on the technical end. The main difficulty comes from the fact that, due to the physical servers' interactions with the Matching queue, the service times in a Residual queue i , $\{S_i^R(k)\}_{k \geq 0}$, are neither identically distributed nor independent over k (Section 4.2.1). To overcome this difficulty, we will use a trick to restore the i.i.d. nature of the service times. In particular, we will create a per-sample-path coupling of the $S_i^R(k)$'s to an i.i.d. sequence $\{\tilde{S}_i^R(k)\}_{k \geq 0}$, such that $S_i^R(k) \leq \tilde{S}_i^R(k)$ for all k , almost surely. If we pretend that the $\tilde{S}_i^R(k)$'s are indeed the true service times, an upper bound on the expected waiting time can be established using Kingman's bound. We finish the argument by showing that this upper bound applies to the expected value of the original waiting time. The following proposition is the major technical result of this section.

Proposition 2. Assume $g_n \in \hat{\mathcal{G}}_n$, where $\hat{\mathcal{G}}_n$ was defined in Lemma 1. Fix any $i \in I$. Let k be the index for batches. Denote by $\{T_i^R(k)\}_{k \geq 0}$ and $\{S_i^R(k)\}_{k \geq 0}$ the sequences of interarrival and service times at the i th Residual queue, respectively, defined on a common probability space. There exists a sequence $\{\tilde{S}_i^R(k)\}_{k \geq 0}$, defined on the same probability space, that satisfies:

1. $\mathbb{P}(\tilde{S}_i^R(k) \geq S_i^R(k), \forall k \in \mathbb{N}) = 1$.
2. Elements of $\{\tilde{S}_i^R(k)\}$ are i.i.d. over k , and independent of $\{T_i^R(k)\}$.
3. $\mathbb{E}(\tilde{S}_i^R(1)) \lesssim \frac{b^2(n)}{n^2}$, and $\mathbb{E}(\tilde{S}_i^R(1)^2) \lesssim \frac{b^2(n)}{n^2}$.

Proof. The proof is given in Appendix A.5. It involves an explicit coupling construction of the sequence $\{\tilde{S}_i^R(k)\}$, which is then shown to satisfy all three claims. Technicalities aside, the proof makes use of the following simple observations: **(1)** the event periods can be "prolonged" via coupling to be made independent, without qualitatively changing the scaling of the service time, $S_i^R(k)$, and **(2)** the first and second moments of $S_i^R(k)$ are mainly influenced by the size of the Residual batch, $R(k)$, which is small ($\lesssim \frac{b(n)}{n}$) by design. \square

Denote by W_i^R the steady-state **waiting times in the i th Residual queue**, where waiting time is defined to be the time from when the batch is formed till when all jobs in the batch have started to receive service. The next proposition is the main result of this subsection, which is proved using the stochastic dominance result of Proposition 2, and the same Kingman's bound as in Proposition 1. The proof is given in A.6.

Proposition 3. (Delay in a Residual Queue) Assume $g_n \in \hat{\mathcal{G}}_n \cap \tilde{\mathcal{G}}_n$, where $\hat{\mathcal{G}}_n$ and $\tilde{\mathcal{G}}_n$ were defined before Lemmas 1 and 5, respectively. We have¹¹

$$\max_{i \in I} \mathbb{E}(W_i^R) \lesssim \frac{1}{y(n)}. \quad (21)$$

5.3 Proof of Theorem 1 Under Uniform Arrival Rates

Proof. Let $\mathcal{H}'_n = \hat{\mathcal{G}}_n \cap \tilde{\mathcal{G}}_n$ for all $n \geq 0$, where $\hat{\mathcal{G}}_n$ and $\tilde{\mathcal{G}}_n$ were defined before Lemmas 1 and 5, respectively. Since each job is served either through the Matching queue or a Residual queue, the total queueing delay is no more than that of the waiting time in a virtual queue plus the time to form a batch ($A(k)$). Hence, letting $b(n) = Kn/y(n)$, by Lemma 2, and Propositions 1 and 3, we have

$$\begin{aligned} \mathbb{E}_\pi(W|g_n, \lambda_n) &\leq \mathbb{E}(A(1)) + \max\left\{\mathbb{E}(W^M), \mathbb{E}(W_1^R)\right\} \\ &\leq \frac{K'}{y(n)} = K' \cdot \frac{\ln n}{d(n)}, \end{aligned} \quad (22)$$

if $g_n \in \mathcal{H}'_n$, where K' is a constant independent of n and g_n . It is not difficult to show, by the weak law of large numbers, that there exist $\epsilon_n \downarrow 0$, such that

$$\mathbb{P}\left(1 - \epsilon_n \leq \frac{\text{deg}(G_n)}{d(n)} \leq 1 + \epsilon_n\right) = 1, \quad (23)$$

for all n , where G_n is a $(rn, n, \frac{d(n)}{n})$ random graph. Let $\mathcal{L}_n = \left\{g_n \in \mathcal{G}_n : \frac{\text{deg}(g_n)}{d(n)} \in [1 - \epsilon_n, 1 + \epsilon_n]\right\}$, and $\mathcal{H}_n = \mathcal{H}'_n \cap \mathcal{L}_n$. By Eq. (23), and Lemmas 1 and 7, we have

$$\mathbb{P}_{n, \frac{d(n)}{n}}(\mathcal{H}_n) \geq 1 - \delta_n, \quad (24)$$

for all n , for some $\delta_n \downarrow 0$. Note that the definitions of $\hat{\mathcal{G}}_n, \tilde{\mathcal{G}}_n$ and \mathcal{L}_n do not depend on the arrival rates λ_n , and hence the value of δ_n also does not depend on λ_n .

Eqs. (22) and (24) combined prove the first two claims of the theorem. Since all λ_i are equal in this case, the scheduling policy is oblivious to the arrival rates by definition. This completes the proof of Theorem 1 when $\lambda_{n,i} = \lambda < 1/r$ for all n and i . \square

¹¹The expectation is defined in the sense of Eq. (2).

6. GENERAL CASE AND ARRIVAL-RATE OBLIVIOUS SCHEDULING

We now complete the proof of Theorem 1 for the general case, for λ_n satisfying Condition 1. In particular, we will show that the original virtual-queue-based scheduling policy given in Section 4.2 automatically achieves the $u^2(n) \ln n/d(n)$ delay scaling, while being fully oblivious to the values of the $\lambda_{n,i}$. We will use a definition of \mathcal{H}_n identical to that in the uniform case, so that the probability of this set will still converge to 1. To avoid replicating our earlier arguments, and since the delays in each of the virtual queues are completely characterized by the statistics of arrival times and job sizes, we will examine the changes in each of these queuing primitives as we move to the general case. Throughout the section, we will analyze a system where the arrival rate vector, λ_n , satisfies Condition 1, and the scheduling rules are the same as before, with $b(n) = Kn/y(n)$.

1. Inter-arrival times of batches. The inter-arrival times to all virtual queues are equal to the time it takes to collect enough jobs to form a single batch. By the merging property of Poisson processes, the aggregate stream of arrivals to the system is a Poisson process of rate $\sum_{i \in I} \lambda_{n,i}$, with

$$\sum_{i \in I} \lambda_{n,i} \leq \rho n. \quad (25)$$

We will make a further assumption, that there exists $\tilde{\rho}$, with $0 < \tilde{\rho} < \rho$, such that

$$\sum_{i \in I} \lambda_{n,i} \geq \tilde{\rho} n. \quad (26)$$

This assumption will be removed by the end of this section by a small modification to the scheduling policy. Given Eqs. (25) and (26), a result analogous to Lemma 2 holds:

$$\mathbb{E}(A(k)) = \frac{b(n)}{n}, \quad \text{and} \quad \text{Var}(A(k)) \lesssim \frac{b(n)}{n^2}. \quad (27)$$

2. Service times for Residual queues. The distributions of the times at which a server tries to fetch a job from a Residual queue can vary due to the different values of the $\lambda_{n,i}$'s. However, it is not difficult to show that the *upper bound* on the service times at the Residual queues (Proposition 2) depend only on the batch size $\frac{\rho}{r}b(n)$, and are independent of the specific values of the $\lambda_{n,i}$'s. Therefore, the only factor that may significantly change the delay distribution at the i th Residual queue is its associated Residual batch size, $R_i(k)$ (Eq. (7)). Since the arrival rates are no longer uniform, this distribution will in principle be different from the uniform setting. However, we will show that the difference is small and does not change the scaling of delay.

Denote by p_i the probability that an incoming job to the system happens to arrive at queue i . We have that

$$p_i = \frac{\lambda_{n,i}}{\sum_{i \in I} \lambda_{n,i}} \leq \frac{u(n)}{\tilde{\rho} n} = \frac{1}{\frac{\tilde{\rho}}{u(n)} n}, \quad (28)$$

where $u(n)$ is the upper bound on $\lambda_{n,i}$ defined in Condition 1, and the inequality follows from the assumptions that $\lambda_{n,i} \leq u(n)$ and that $\sum_{i \in I} \lambda_{n,i} \geq \tilde{\rho} n$ (Eq. (26)). Using Eq. (28), and following the same steps as in the proof of Proposition 2, one can show that

$$\max_{i \in I} \mathbb{E}(\tilde{S}_i^R(k)) \lesssim \frac{b^2(n)u^2(n)}{n^2},$$

and

$$\max_{i \in I} \text{Var}(\tilde{S}_i^R(k)) \lesssim \frac{b^2(n)u^2(n)}{n^2}.$$

We repeat the arguments in Proposition 3 by replacing, e.g., the bounds on $\text{Var}(\tilde{S}_i^R(k))$ with $\max_{i \in I} \text{Var}(\tilde{S}_i^R(k)) \lesssim \frac{b^2(n)u^2(n)}{n^2}$. Letting $b(n) = K \frac{n \ln n}{d(n)}$, we have that, whenever $u(n) \ll \sqrt{\frac{d(n)}{\ln n}}$,

$$\begin{aligned} \max_{i \in I} \mathbb{E}(W_i^R) &\lesssim \frac{n}{b(n)} \left(\frac{b(n)}{n^2} + \frac{b^2(n)u^2(n)}{n^2} \right) \lesssim \frac{b(n)u^2(n)}{n} \\ &\lesssim \frac{\ln n}{d(n)} u^2(n). \end{aligned} \quad (29)$$

3. Service times for the Matching queue. Surprisingly, it turns out that the bounds on the service times statistics for the Matching queue in the uniform arrival-rate case (Lemma 8) will carry over to the general setting, unchanged. Recall from Eq. (12) that the service time of a Matching batch is composed of two elements: the length of a COLLECT phase, S_{col} , and, with probability $q(g_n)$, the length of a clearing phase S_{cle} . Since the size of the Matching batch is at most $\frac{\rho}{r}b(n)$ by definition, the characterizations of S_{col} and S_{cle} given by Lemmas 3 and 4 continue to hold, with λ being replaced by $\lambda = \frac{\rho}{r}$. It remains to verify that the bound on $q(g_n)$ stays unchanged. While the $\lambda_{n,i}$'s are now different and the subgraph induced by Δ and Γ is no longer uniformly distributed, the upper bound on $q(g_n)$ given in Lemma 7 still holds, because Lemma 5 applies to arbitrary distributions of subgraphs. Therefore, the scaling in Proposition 1

$$\mathbb{E}(W^M) \lesssim \frac{1}{y(n)} = \frac{\ln n}{d(n)}, \quad (30)$$

continues to hold under non-uniform arrival rates.

Combining Eqs. (25), (29) and (30), we have, analogous to Eq. (22), that for all $n \geq 1$,

$$\begin{aligned} \mathbb{E}_\pi(W|g_n, \lambda_n) &\leq \mathbb{E}(A(1)) + \max\left\{\mathbb{E}(W^M), \mathbb{E}(W_1^R)\right\} \\ &\leq \frac{K_1 \ln n}{d(n)} + \frac{K_2 u^2(n) \ln n}{d(n)} \leq \frac{K u^2(n) \ln n}{d(n)}, \end{aligned} \quad (31)$$

if $g_n \in \mathcal{H}_n$, where K_1, K_2 and K are positive constants independent of n, g_n and λ_n .

Finally, we justify the assumption made in Eq. (26), by showing that the scheduling policy could be easily modified such that Eq. (26) always holds: we simply insert to each queue an independent Poisson stream of “dummy packets” at rate δ , where δ is chosen to be any constant in $(0, \frac{1-\rho}{r})$. The dummy packets are merely place-holders: when a server begins serving a dummy packet, it simply goes on a vacation of duration $\text{Expo}(1)$. This trivially guarantees the validity of Eq. (26), with $\tilde{\lambda} = \delta$, and since $\sum_{i \in I} \lambda_{n,i} \leq \rho n$, doing so is equivalent to having a new system with ρ replaced by $\rho' = \rho + \delta r < 1$. Hence, all results continue to hold.¹² This concludes the proof of Theorem 1.

7. CONCLUSIONS AND FUTURE WORK

The main message of this paper is that the benefits of diminishing delay and large capacity region can be jointly achieved in a system where the level of processing flexibility of each server is small compared to the system size. The main result, Theorem 1, proves this using a randomly constructed interconnection topology and a virtual-queue-based scheduling policy. As a by-product, it also provides an explicit upper bound (Eq. (3)) that captures a

¹²This maneuver of inserting dummy packets is of course a mere technical device for the proof. In reality, delay and capacity would both become much less of a concern when the traffic intensity, ρ , is too small, at which point less sophisticated scheduling policies may suffice.

trade-off between the delay ($\mathbb{E}(W)$), level of robustness ($u(n)$), and degree of processing flexibility ($d(n)$).

The scaling regime considered in this paper assumes that the traffic intensity is fixed as n increases, which fails to capture system performance in the heavy-traffic regime ($\rho \approx 1$). It would be interesting to consider a scaling regime in which ρ and n scale simultaneously (e.g., the celebrated Halfin-Whitt regime [25]), but it is unclear at this stage what exact formulations and analytical techniques are appropriate. At a higher level, it would be interesting to understand how long-term input characteristics in a queueing network (e.g., how arrival rates are drawn or evolve over time) impact its temporal performance (e.g., delay), and what role the network's intrinsic structure (e.g., flexibility) has to play here.

Theorem 1 also leaves open several promising directions for future improvements and extensions, which we discuss briefly. It is currently necessary to have that $d(n) \gg \ln n$ in order to achieve a diminishing delay. The $\ln n$ factor is essentially tight if a random-graph-based connectivity structure is to be used, because $d(n) = \Omega(\ln n)$ is necessary for an Erdős-Rényi type random graph to be connected (i.e., not have isolated queues). We suspect that the $\ln n$ factor can be reduced using a different graph generation procedure.

Fixing $u(n)$ to a constant, we observe that when $d(n) = n$ (fully connected graph), the system behaves approximately as a $M/M/n$ queue with total arrival rate $n\rho$; if ρ is held fixed, then the expected delay is known to vanish exponentially fast in n . Thus, there is a gap between this fact and the polynomial scaling of $\mathcal{O}(\ln n/d(n))$ given in Theorem 1. We believe that this gap is due to an intrinsic inefficiency of the batching procedure used in our scheduling policy. The batching procedure provides analytical convenience by allowing us to leverage the regularity in the arrival traffic in obtaining delay upper bounds. However, the use of batching also mandates that all jobs wait for a prescribed period of time before receiving any service. This can be highly inefficient in the scaling regime that we consider, where the traffic intensity is fixed at ρ , because a $1 - \rho$ fraction of all servers (approximately) are actually idle at any moment in time. We therefore conjecture that a scheduling policy that tries to direct a job to an available server immediately upon its arrival (such as a greedy policy) can achieve a significantly better delay scaling, but the induced queueing dynamics, and the appropriate analytical techniques, may be very different from those presented in this paper.

Finally, the parameter $u(n)$ captures the level of robustness against uncertainties in the arrival rates. In terms of inducing a diminishing waiting time as $n \rightarrow \infty$, there is still a gap between the current requirement that $u(n) \ll \sqrt{d(n)/\ln n}$ and the upper bound that $u(n) = \mathcal{O}(d(n))$ (since each queue is connected to only $d(n)$ servers on average). The $u^2(n)$ factor in the scaling of $\mathbb{E}(W)$ is due to the impact on delay by the arrival rate fluctuations in the Residual queues (cf. Eq. (29)). It is unclear what the optimal dependence of $\mathbb{E}(W)$ on $u(n)$ is. It is also possible that the condition $u(n) \ll d(n)$ alone is sufficient for achieving a diminishing delay; we conjecture that this is the case.

8. REFERENCES

- [1] S. Kandula, S. Sengupta, A. Greenberg, P. Patel, R. Chaiken, "Nature of datacenter traffic: measurements and analysis," *IMC*, 2009.
- [2] G. Soundararajan, C. Amza, and A. Goel, "Database replication policies for dynamic content applications," *Proc. of EuroSys*, 2006.
- [3] W. Jordan and S. C. Graves, "Principles on the benefits of manufacturing process flexibility," *Management Science*, 41(4):577–594, 1995.
- [4] S. Gurumurthi and S. Benjaafar, "Modeling and analysis of flexible queueing systems," *Management Science*, 49:289–328, 2003.
- [5] S. M. Iravani, M. P. Van Oyen, and K. T. Sims, "Structural flexibility: A new perspective on the design of manufacturing and service operations," *Management Science*, 51(2):151–166, 2005.
- [6] M. Chou, G. A. Chua, C-P. Teo, and H. Zheng, "Design for process flexibility: efficiency of the long chain and sparse structure," *Operations Research*, 58(1):43–58, 2010.
- [7] D. Simchi-Levi and Y. Wei, "Understanding the performance of the long chain and sparse designs in process flexibility," submitted, 2011.
- [8] M. Chou, C-P. Teo, and H. Zheng, "Process flexibility revisited: the graph expander and its applications," *Operations Research*, 59:1090–1105, 2011.
- [9] M. Leconte, M. Lelarge and L. Massoulié, "Bipartite Graph Structures for Efficient Balancing of Heterogeneous Loads," *ACM Sigmetrics*, London, 2012.
- [10] R. Talreja, W. Whitt, "Fluid models for overloaded multiclass many-service queueing systems with FCFS routing," *Management Sci.*, 54(8):1513–1527, 2008.
- [11] J. Visschers, I. Adan, and G. Weiss, "A product form solution to a system with multi-type jobs and multi-type servers," *Queueing Systems*, 70:269–298, 2012.
- [12] A. Mandelbaum and M. I. Reiman, "On pooling in queueing networks," *Management Science*, 44(7):971–981, 1998.
- [13] J. M. Harrison and M. J. Lopez, "Heavy traffic resource pooling in parallel-server systems," *Queueing Systems*, 33:39–368, 1999.
- [14] S. L. Bell and R. J. Williams, "Dynamic scheduling of a system with two parallel servers in heavy traffic with resource pooling: asymptotic optimality of a threshold policy," *Ann. Appl. Probab.*, 11(3): 608–649, 2001.
- [15] R. Wallace and W. Whitt, "A staffing algorithm for call centers with skill-based routing," *Manufacturing and Service Operations Management*, 7:276–294, 2005.
- [16] A. Bassamboo, R. S. Randhawa, and J. A. Van Mieghem, "A little flexibility is all you need: on the asymptotic value of flexible capacity in parallel queueing systems," submitted, 2011.
- [17] J. N. Tsitsiklis and K. Xu, "On the power of (even a little) resource pooling," *Stochastic Systems*, 2: 1–66 (electronic), 2012.
- [18] M. Neely, E. Modiano and Y. Cheng, "Logarithmic delay for $n \times n$ packet switches under the crossbar constraint," *IEEE/ACM Trans. Netw.*, 15(3):657–668, 2007.
- [19] J. N. Tsitsiklis and D. Shah, "Bin packing with queues," *J. Appl. Prob.*, 45(4):922–939, 2008.
- [20] N. McKeown, A. Mekkittikul, V. Anantharam and J. Walrand, "Achieving 100% throughput in an input-queued switch," *IEEE Trans. on Comm.*, 47(8):1260–1267, 1999.
- [21] S. Kunniyur and R. Srikant, "Analysis and design of an adaptive virtual queue," *ACM SIGCOMM*, 2001.
- [22] P. Erdős and A. Rényi, "On random matrices," *Magyar Tud. Akad. Mat. Kutató Int. Közl.*, 8:455–461, 1964.
- [23] S. R. Bodas, "High-performance scheduling algorithms for wireless networks," Ph.D. dissertation, University of Texas at Austin, Dec. 2010.
- [24] V. F. Kolchin, B. A. Sevast'yanov, and V. P. Chistyakov, *Random Allocation*, John Wiley & Sons, 1978.
- [25] S. Halfin and W. Whitt, "Heavy-traffic limits for queues with many exponential servers," *Operations Research*, 29: 567–588, 1981.
- [26] J. N. Tsitsiklis and K. Xu, "Queueing system topologies with limited flexibility," *Sigmetrics*, Pittsburgh, PA, June, 2013.

APPENDIX

A. ADDITIONAL PROOFS

A.1 Proof of Lemma 1

Proof. (Lemma 1) Let G_n be an $(rn, n, p(n))$ random bipartite graph. We first assume that $r \leq 1$. Since $d(n) \gg \ln n$, it is well known that (cf. [22]; see also Lemma 6),

$$\mathbb{P}(G_n \text{ admits a full matching}) \rightarrow 1,$$

as $n \rightarrow \infty$. We let \hat{G}_n be the set of all graphs that admit a full matching. For each $g \in \hat{G}_n$, choose one full matching M , and let $L(i) = j$ such that $(i, j) \in M$. L is an efficient assignment function by the definition of a full matching, since each j is matched with at most one i thorough L .

We now consider the case of $r > 1$. For each n , partition I into r disjoint subsets, I_1, \dots, I_r , where each subset has size at most n . Using again Lemma 6, we have that

$$\mathbb{P}(G_n|_{I_k \cup J} \text{ admits a full matching for all } k = 1, \dots, r) \rightarrow 1, \quad (32)$$

as $n \rightarrow \infty$, where $G_n|_{I_k \cup J}$ is the graph G_n restricted to the vertex set $I_k \cup J$. Similar to the previous case, we let \hat{G}_n be the set of graphs that satisfy the event in Eq. (32). For each n and k , let M_k be a full matching in $G_n|_{I_k \cup J}$, and for all $i \in I_k$, let $L(i) = j$ where $(i, j) \in M_k$. We then have $|L^{-1}(j)| \leq r \leq r + 1$ and L is an efficient assignment function. \square

A.2 Proof of Lemma 4

The following technical lemma is useful (see [24], Chapter 1, §2).

Lemma 9. *Consider throwing balls uniformly into n bins. Let N_m be the number of balls required so that at least m bins are non-empty. If $m = o(n)$, then $\mathbb{E}(N_m) \sim m$, and $\text{Var}(N_m) = o(m)$.*

Proof. (Lemma 4) The length of S_{col} is equal to the time it takes for $\lambda b(n)$ many servers to enter the state STANDBY. Denote by $T_\lambda(j)$ the time of the j th jump in a rate- λ Poisson process. We have that

$$\begin{aligned} \mathbb{E}(S_{col}) &= \mathbb{E}(T_{(\rho+\epsilon)n}(N_{\lambda b(n)})) \\ &= \sum_{j=1}^{\infty} \mathbb{E}(T_{(\rho+\epsilon)n}(j)) \mathbb{P}(N_{\lambda b(n)} = j) \\ &\stackrel{(a)}{=} \frac{1}{(\rho+\epsilon)n} \sum_{j=1}^{\infty} j \mathbb{P}(N_{\lambda b(n)} = j) \\ &= \frac{1}{(\rho+\epsilon)n} \mathbb{E}(N_{\lambda b(n)}), \end{aligned} \quad (33)$$

where the equality (a) follows from the fact that $T_{(\rho+\epsilon)n}(j) \stackrel{d}{=} \sum_{i=1}^j E_i$, where the E_i 's are i.i.d., with distribution $E_i \stackrel{d}{=} \text{Expo}((\rho+\epsilon)n)$. By Lemma 9, and since $b(n) \ll n$, we have that

$$\mathbb{E}(S_{col}) \sim \frac{1}{(\rho+\epsilon)n} \lambda b(n) = \frac{\lambda}{(\rho+\epsilon)} \cdot \frac{b(n)}{n}. \quad (34)$$

To compute the second moment, by Lemma 14 in Appendix B, we have that

$$\mathbb{E}(T_{(\rho+\epsilon)n}^2(j)) = \frac{j(j+1)}{(\rho+\epsilon)^2 n^2}. \quad (35)$$

Using a similar argument to that in the calculation of $\mathbb{E}(S_{col})$, we have

$$\begin{aligned} &\mathbb{E}(S_{col}^2) \\ &= \mathbb{E}(T_{(\rho+\epsilon)n}^2(N_{\lambda b(n)})) \\ &= \sum_{j=1}^{\infty} \mathbb{E}(T_{(\rho+\epsilon)n}(j)^2) \mathbb{P}(N_{\lambda b(n)} = j) \\ &\stackrel{(a)}{=} \frac{1}{(\rho+\epsilon)^2 n^2} \sum_{j=1}^{\infty} j(j+1) \mathbb{P}(N_{\lambda b(n)} = j) \\ &= \frac{1}{(\rho+\epsilon)^2 n^2} (\mathbb{E}(N_{\lambda b(n)}) + \mathbb{E}(N_{\lambda b(n)}^2)) \\ &= \frac{1}{(\rho+\epsilon)^2 n^2} [\mathbb{E}(N_{\lambda b(n)}) + (\mathbb{E}(N_{\lambda b(n)})^2 + \text{Var}(N_{\lambda b(n)}))] \\ &\lesssim \frac{b^2(n)}{n^2}, \end{aligned} \quad (36)$$

where the equality (a) follows from Eq. (35), and the last step follows from Lemma 9. \square

A.3 Proof of Lemma 7

Proof. (Lemma 7) Let \tilde{G}_n be as in Lemma 5. The size of the random subset $\Gamma \subset I$ induced by a Matching batch is at most $\lambda b(n)$, and the size of the random subset $\Delta \subset J$ induced by the STANDBY servers is exactly $\lambda b(n)$. Hence, let $m(n) = \lambda b(n)$, and there exists a distribution $\mathbf{P}_{\mathcal{M}_{m(n),n}}$ over $\mathcal{M}_{m(n),n}$, with $\mathcal{M}_{m(n),n}$ defined as in Lemma 5, such that $\Gamma \cup \Delta$ is covered by the random subset generated by $\mathbf{P}_{\mathcal{M}_{m(n),n}}$ almost surely. Note that Γ and Δ are generated independently of g_n . We now apply Lemma 5 to obtain that for all $g_n \in \tilde{G}_n$

$$\begin{aligned} q(g_n) &\leq l(g_n) \\ &\lesssim b^2(n) (1-p(n))^{\lambda b(n)} = b^2(n) \exp[-\lambda b(n) \cdot \ln(1-p(n))] \\ &\stackrel{(a)}{=} b^2(n) \exp\left(-\lambda b(n)p(n) - \lambda b(n) \frac{p(n)^2}{2} + \mathcal{O}(b(n)p(n)^3)\right) \\ &\stackrel{(b)}{\lesssim} \frac{n^2}{y(n)^2} \exp\left(\ln \frac{1}{n^{\lambda K/2}}\right) = \frac{1}{y(n)^2} n^{-(\lambda K-4)/2}, \end{aligned} \quad (37)$$

where step (a) is based on the Taylor expansion: $\ln(1-x) = -x - \frac{x^2}{2} + \mathcal{O}(x^3)$, and (b) is based on the substitution $b(n) = K \frac{n}{y(n)}$ and $p(n) = \frac{d(n)}{n} = \frac{y(n) \ln n}{n}$. The constant $\frac{1}{2}$ in (b) can be replaced by any other positive value less than 1. \square

A.4 Proof of Lemma 8

Proof. (Lemma 8) Combining Eqs. (12) and Lemmas 3 (for S_{cle}), 4 (for S_{col}), and 7 (for $q(g_n)$), we have

$$\begin{aligned} \mathbb{E}(S^M(k)) &\stackrel{(a)}{=} \mathbb{E}(S_{col}) + q(g_n) \mathbb{E}(S_{cle}) \\ &\leq \frac{\lambda}{\rho+\epsilon} \cdot \frac{K}{y(n)} + C \left(\frac{1}{y(n)^2} n^{-\frac{1}{2}(\lambda K-4)} \right) \frac{n}{y(n)} \\ &= \frac{\lambda}{\rho+\epsilon} \cdot \frac{K}{y(n)} + \frac{C}{y(n)^3} n^{-\frac{1}{2}(\lambda K-6)}, \\ &\sim \frac{\lambda}{\rho+\epsilon} \cdot \frac{K}{y(n)}, \end{aligned} \quad (38)$$

whenever $K > \frac{6}{\lambda}$, where C is some positive constant. Note that in (a) we have used the fact that $X_{q(g_n)}$ and S_{cle} are independent.

Similarly,

$$\begin{aligned}
& \mathbb{E} \left(S^M(k)^2 \right) \\
&= (1 - q(g_n)) \mathbb{E} (S_{col}^2) + q(g_n) \mathbb{E} (S_{col}^2 + S_{cle}^2 + 2S_{col}S_{cle}) \\
&= \mathbb{E} (S_{col}^2) + 2q(g_n) \mathbb{E} (S_{col}) \mathbb{E} (S_{cle}) + q(g_n) \mathbb{E} (S_{cle}^2) \\
&\lesssim \frac{1}{y(n)^2} + \frac{1}{y(n)^2} n^{-(\lambda K - 4)/2} \frac{1}{y(n)} \cdot \frac{n}{y(n)} \\
&\quad + \frac{1}{y(n)^2} n^{-(\lambda K - 4)/2} \frac{n^2}{y(n)^2}, \\
&\lesssim \frac{1}{y(n)^2} + \frac{1}{y(n)^4} n^{-(\lambda K - 8)/2}, \\
&\lesssim \frac{1}{y(n)^2}, \tag{39}
\end{aligned}$$

whenever $K > \frac{8}{\lambda}$. Note that $\text{Var} (S^M(k)) \leq \mathbb{E} (S^M(k)^2) \lesssim \frac{1}{y(n)^2}$. This completes the proof. \square

A.5 Proof of Proposition 2

Proof. (Proposition 2) We will explicitly construct the sequence $\{\tilde{S}_i^R(k)\}$ and show that it satisfies all three claims. Throughout the proof, we will omit the superscript R and the subscript i in our notation (e.g., write $S(k)$ in place of $S_i^R(k)$).

Before describing the coupling procedure in detail, we explain in words what is to be accomplished. Recall from Section 4.2.1 that the type of an event point for the server can be either “matching” or “residual”. The times between adjacent event points (i.e., event periods) are an i.i.d. exponential random variables, with *one exception*: if the event point is of type “matching” and the Matching queue is in state COLLECT. In this case, the server has to remain in STANDBY state until the *end of the COLLECT period*, and this induces dependency among the event periods.

The main idea of the coupling procedure is to simply “prolong” the lengths of *all* event periods, so that the resulting lengths are i.i.d. (Step 1). As a technical maneuver, we will also append an additional event period, if the first event point in a Residual batch is not of a Matching type (Step 2). This guarantees that the *number of event periods*, and hence the total service time for a Residual batch, is i.i.d. among different batches.

Let L be an efficient assignment function for g_n (Section 4.2.1), and $j = L(i)$ the index of the server assigned to queue i . Fix any $k \geq 1$. The service period (Definition 5) of the k th batch, $[E_k, D_k]$, is composed of a sequence of event periods for server j (Section 4.2.1). Let $\{e_m\}_{1 \leq m \leq M_k}$ be the times of all event points that fall within $[E_i, D_k]$, and let e_0 be the event point that immediately precedes E_k . We have

$$e_0 \leq E_k \leq e_1 \leq e_2 \leq \dots \leq e_{M_k} = D_k, \tag{40}$$

We will define $\{\tilde{S}(k)\}_{k \geq 1}$, for every sample path of the system, as follows

A coupled construction of $\tilde{S}(k)$

For all $k \geq 1$:

- **Step 1:** For all $m = 0, 1, \dots, M_k - 1$:

(a) If e_m is of type “matching,” and the Matching queue is in state COLLECT at time e_m , we let

$$p_m = (e_{m+1} - e_m) + h - (l - e_m), \tag{41}$$

where l is the first time after e_m when the Matching leaves state COLLECT, and h is the total amount of time it has spent in COLLECT before exiting. Note that by definition, $h \geq (l - e_m)$ almost surely, and hence $p_m \geq e_{m+1} - e_m$ almost surely.

- (b) Otherwise, let

$$p_m = (e_{m+1} - e_m) + \tilde{h}, \tag{42}$$

where \tilde{h} is drawn according to the distribution of the length of a COLLECT period in the Matching queue (See Lemma 4),

$$\tilde{h} \stackrel{d}{=} \sum_{i=1}^{N_{\lambda b(n)}} \text{Expo} \left(\frac{1}{n(\rho + \epsilon)} \right), \tag{43}$$

where N_m was defined in Lemma 9.

- **Step 2:** Draw \tilde{p} from the distributions

$$\tilde{p} \stackrel{d}{=} \begin{cases} 0, & \text{if } e_0 \text{ is of type “matching”;} \\ \text{Expo}(1) + \sum_{i=1}^{N_{\lambda b(n)}} \text{Expo} \left(\frac{1}{n(\rho + \epsilon)} \right) & \text{otherwise,} \end{cases} \tag{44}$$

- **Step 3:** Set

$$\tilde{S}(k) = \tilde{p} \cdot \mathbb{I}(R(k) > 0) + \sum_{m=0}^{M_k-1} p_m, \tag{45}$$

where $R(k)$, defined in Eq. (7), is the size of the k th Residual batch. The indicator function ensures that the service time is zero if the batch is empty.

By construction, $\tilde{p} \geq 0$, and $p_m \geq e_{m+1} - e_m$ for all $0 \leq m \leq M_k - 1$, almost surely. Therefore, the first claim of Proposition 2, that $\tilde{S}(k) \geq S(k)$ for all k almost surely, holds. We now show that the elements in $\{\tilde{S}(k)\}$ are i.i.d., and independent from the arrival times $\{T(k)\}$. The following lemma is useful.

Lemma 10. *For any k , the random variables $p_0, p_1, \dots, p_{M_k-1}, \tilde{p}$ are i.i.d., with distribution*

$$p \stackrel{d}{=} \text{Expo}(1) + \sum_{i=1}^{N_{\lambda b(n)}} \text{Expo}(n(\rho + \epsilon)). \tag{46}$$

Furthermore,

$$\mathbb{E}(p) \rightarrow 1, \tag{47}$$

as $n \rightarrow \infty$.

Proof. At server j , no two event points of type “match” can occur during the *same* COLLECT period of the Matching queue, since after the first such event point, server j will remain in state STANDBY until the end of the COLLECT period. Therefore, Step 1 of the coupling does not introduce any dependency among the p_m 's, and the definitions of \tilde{p} and all the p_m 's involve only independent random variables. The specific distribution of p is immediate from the construction, and the fact that all jobs sizes and vacation lengths are distributed according to Expo(1).

To show Eq. (47), by the same argument as that in the proof of Lemma 4, we have that

$$\begin{aligned}
\mathbb{E} \left(\sum_{i=1}^{N_{\lambda b(n)}} \text{Expo}(n(\rho + \epsilon)) \right) &= \mathbb{E}(N_{\lambda b(n)}) \cdot \mathbb{E}(\text{Expo}(n(\rho + \epsilon))) \\
&\lesssim \frac{b(n)}{n}, \tag{48}
\end{aligned}$$

and hence $\mathbb{E}(p) \rightarrow \mathbb{E}(\text{Expo}(1)) = 1$ as $n \rightarrow \infty$. \square

We say that an event point at server j is of type “residual- i ,” if

1. The event point is of type “residual”.
2. A job from queue i is selected to be processed.

By the definition of a Residual batch, if $R(k) > 0$, the departure time of the k th Residual batch, D_k , coincides with the $R_i(k)$ th event point of type “residual- i ” after E_k . Recall that an event period is of type “residual” with probability $1 - (\rho + \epsilon)$, and when this occurs, queue i will be selected with probability $\frac{1}{|L^{-1}(j)|}$. Denote by \tilde{M}_k the number of non-zero terms in $\tilde{S}(k)$ (Eq. (45)). The preceding arguments show that that \tilde{M}_k 's are independent, and

$$\tilde{M}_k \stackrel{d}{=} \mathbb{I}(R(k) > 0) + \sum_{u=1}^{R(k)} G_{k,u}, \quad \forall k \geq 0. \quad (49)$$

where the $G_{k,u}$'s are i.i.d., with distribution

$$G_{k,u} \stackrel{d}{=} \text{Geo}\left(\frac{1 - (\rho + \epsilon)}{|L^{-1}(j)|}\right). \quad (50)$$

Finally, by Lemma 10, we have that

$$\tilde{S}(k) = p \cdot \mathbb{I}(R(k) > 0) + \sum_{u=1}^{R(k)} \sum_{v=1}^{G_{k,u}} p_{k,u,v}, \quad (51)$$

where p and the $p_{k,u,v}$'s are i.i.d., with distribution given in Eq. (46). Since the size of Residual batches ($R(k)$) are i.i.d., and independent from the arrival times, we have established the second claim of Proposition 2.

To prove the last claim of the Proposition 2 that concerns the first two moments of $\tilde{S}(k)$, we begin with the following fact on the balls-into-bins problem.

Lemma 11. *Denote by $H_i(m, n)$ be the number of balls in bin i after throwing m balls uniformly into n bins. Define*

$$R_i(m, n) = (H_i(m, n) - 1)^+. \quad (52)$$

If $m = o(n)$, then

$$\mathbb{E}(R_i(m, n)) \lesssim \frac{m^2}{n^2}, \quad (53)$$

$$\mathbb{E}(R_i(m, n)^2) \lesssim \frac{m^2}{n^2}, \quad (54)$$

as $n \rightarrow \infty$.

Proof. By the definition of R_i , we have the following two cases

1. Conditioning on $H_i(m, n) \geq 1$, then

$$R_i(m, n) \stackrel{d}{=} H_i(m - 1, n) \stackrel{d}{=} \text{Bino}\left(m - 1, \frac{1}{n}\right), \quad (55)$$

where $\text{Bino}(n, p)$ is the binomial distribution with parameter n and p .

2. Conditioning on $H_i(m, n) = 0$, then $R_i(m, n) = 0$.

Therefore, we have that

$$\mathbb{E}(R_i(m, n)) = \mathbb{P}(H_i(m, n) \geq 1) \cdot \left(\frac{m-1}{n}\right), \quad (56)$$

and

$$\begin{aligned} & \mathbb{E}(R_i(m, n)^2) \\ &= \mathbb{P}(H_i(m, n) \geq 1) \cdot \left[\left(\frac{m-1}{n}\right)^2 + \frac{m-1}{n} \left(1 - \frac{1}{n}\right) \right]. \end{aligned} \quad (57)$$

Since $H_i(m, n) \stackrel{d}{=} \text{Bino}(m, \frac{1}{n})$, we have

$$\begin{aligned} \mathbb{P}(H_i(m, n) \geq 1) &= 1 - \left(1 - \frac{1}{n}\right)^m \\ &= 1 - \exp\left(m \ln\left(1 - \frac{1}{n}\right)\right) \\ &\stackrel{(a)}{\lesssim} 1 - \exp\left(-\frac{m}{n}\right) \\ &\stackrel{(b)}{\lesssim} 1 - \left(1 - \frac{m}{n}\right) \\ &\lesssim \frac{m}{n}. \end{aligned} \quad (58)$$

where steps (a) and (b) follow from the Taylor expansions $\ln(1 - x) \gtrsim -x$ and $\exp(x) \gtrsim 1 - x$ as $x \downarrow 0$, respectively. The claim is proved by substituting Eq. (58) into Eqs. (56) and (57). \square

Recall that $R(k)$ is the number of jobs in the k th Residual batch, defined in Eq. (7). We have that

$$R(k) \stackrel{d}{=} R_1(\lambda b(n), \lambda n), \quad (59)$$

where $R(\cdot, \cdot)$ on the right-hand side is defined in Lemma 11. By Lemma 11, we have

$$\mathbb{E}(R(k)) \lesssim \frac{b^2(n)}{n^2}, \quad (60)$$

$$\mathbb{E}(R(k)^2) \lesssim \frac{b^2(n)}{n^2}. \quad (61)$$

By Eq. (51), we have

$$\begin{aligned} \mathbb{E}(\tilde{S}(k)) &= \mathbb{E}\left(p \cdot \mathbb{I}(R(k) > 0) + \sum_{u=1}^{R(k)} \sum_{v=1}^{G_{k,u}} p_{k,u,v}\right) \\ &\stackrel{(a)}{=} \mathbb{E}(p) \mathbb{P}(R(k) > 0) + \mathbb{E}(p) \mathbb{E}(R(k)) \mathbb{E}(G_{k,n}) \\ &\stackrel{(b)}{\leq} \mathbb{E}(p) \mathbb{E}(R(k)) + \mathbb{E}(p) \mathbb{E}(R(k)) \mathbb{E}(G_{k,n}) \\ &= [\mathbb{E}(p) (\mathbb{E}(G_{k,n}) + 1)] \mathbb{E}(R(k)) \\ &\stackrel{(c)}{\lesssim} \mathbb{E}(R(k)), \\ &\lesssim \frac{b^2(n)}{n^2}, \end{aligned} \quad (62)$$

where the step (a) is based on Lemma 13 in Appendix B, step (b) from the fact that $\mathbb{P}(X > 0) \leq \mathbb{E}(X)$ for any random variable X defined on \mathbb{Z}_+ , step (c) from Eqs. (47) and (50), and step (d) from

Eq. (60). Similarly,

$$\begin{aligned}
\mathbb{E} \left(\tilde{S}(k)^2 \right) &= \mathbb{E} \left(\left(p \cdot \mathbb{I}(R(k) > 0) + \sum_{u=1}^{R(k)} \sum_{v=1}^{G_{k,n}} p_{k,u,v} \right)^2 \right) \\
&= \mathbb{E}(p^2) \mathbb{P}(R(k) > 0) \\
&\quad + 2\mathbb{E}(p) \mathbb{P}(R(k) > 0) \mathbb{E} \left(\sum_{u=1}^{R(k)} \sum_{v=1}^{G_{k,n}} p_{k,u,v} \right) \\
&\quad + \mathbb{E} \left(\left(\sum_{u=1}^{R(k)} \sum_{v=1}^{G_{k,n}} p_{k,u,v} \right)^2 \right) \\
&\stackrel{(a)}{\lesssim} \mathbb{E}(R(k)) + \mathbb{E}(R(k))^2 + (\mathbb{E}(R(k)) + \mathbb{E}(R(k)^2)) \\
&\lesssim \frac{b^2(n)}{n^2}, \tag{63}
\end{aligned}$$

where step (a) follows from Lemma 13 from Appendix B, and the fact that $\mathbb{P}(R(k) > 0) \leq \mathbb{E}(R(k))$. With Eqs. (62) and (63), we have completed the proof of Proposition 2. \square

A.6 Proof of Proposition 3

Proof. (Proposition 3) We begin by stating a technical lemma which formalizes the obvious fact that, on a per-sample-path basis, delays are monotone with respect to the service times. The proof is elementary and omitted.

Lemma 12. *Consider an initially empty first-come-first-serve (FIFO) queue with a single server. Let $\{T(k)\}_{k \geq 1}$ be a sequence of arrival times, and let $\{S(k)\}_{k \geq 1}$ and $\{\tilde{S}(k)\}_{k \geq 1}$ be sequences of service times, all defined on a common probability space. Denote by $\{W(k)\}_{k \geq 1}$ and $\{\tilde{W}(k)\}_{k \geq 1}$ the resulted waiting times induced by $(\{T(k)\}, \{S(k)\})$ and $(\{T(k)\}, \{\tilde{S}(k)\})$, respectively. If $\mathbb{P}(\tilde{S}(k) \geq S(k), \forall k \in N) = 1$, then*

$$\mathbb{P}(\tilde{W}(k) \geq W(k), \forall k \in N) = 1, \tag{64}$$

which also implies that

$$\mathbb{E}(\tilde{W}(k)) \geq \mathbb{E}(W(k)), \quad \forall k \geq 1. \tag{65}$$

Fix $i \in I$. By Lemma 12 and Proposition 2, $\mathbb{E}(W_i^R)$ is upper-bounded by the expected waiting time in steady-state of an $GI/GI/1$ queue with arrival times $\{A(k)\}$ and service times $\{\tilde{S}_i^R(k)\}$. Using the Kingman's formula, Lemma 2, and Proposition 2, we have $\mathbb{E}(W_i^R) \leq \tilde{\lambda} \frac{\sigma_a^2 + \sigma_s^2}{2(1-\tilde{\rho})}$, where

$$\tilde{\lambda} = \frac{1}{\mathbb{E}(A(k))} \sim \frac{n}{b(n)}, \quad \tilde{\rho} = \frac{\mathbb{E}(\tilde{S}_1^R(k))}{\mathbb{E}(A(k))} \lesssim \frac{b^2(n)}{b(n)n} \lesssim \frac{b(n)}{n}, \tag{66}$$

$$\sigma_a^2 = \text{Var}(A(k)) \lesssim \frac{b(n)}{n^2}, \quad \sigma_s^2 = \text{Var}(\tilde{S}_1^R(k)) \lesssim \frac{b^2(n)}{n^2}, \tag{67}$$

which yield

$$\mathbb{E}(W_i^R) \lesssim \frac{n}{b(n)} \left(\frac{b(n)}{n^2} + \frac{b^2(n)}{n^2} \right) \lesssim \frac{b(n)}{n}. \tag{68}$$

In particular, if $b(n) = K \frac{n}{y(n)}$, then $\mathbb{E}(W_i^R) \lesssim \frac{1}{y(n)}$, $\forall i \in I$. It is also clear that the scaling in the above equation holds uniformly over all $i \in I$. This completes the proof. \square

A.7 Proof of Eq. (4)

We show in this section that the constant in Eq. (3) can be further expanded to incorporate the traffic intensity, ρ . In particular, there exists a constant K' , independent of ρ , such that

$$K \leq \frac{K'}{1-\rho}. \tag{69}$$

To see why this is true, we go back to Propositions 1 and 3, which characterize the expected delays in the Matching queue and Residual queues, respectively. It is easy to verify, from the proof of Proposition 3, that the bounds on the delays in the Residual queues can be made to be independent of ρ (This is intuitively consistent with the fact that the Residual queues are used to only absorb the ‘‘discrepancies’’). By Proposition 1, the expected delay in the Matching queue does depend on ρ , but only through the value of the effective traffic intensity, $\tilde{\rho}$. In particular, by Eq. (20), we have

$$\tilde{\rho} \leq \frac{\rho}{\rho + \epsilon} \leq \frac{1}{1 + \epsilon}, \quad \forall \rho \in (0, 1). \tag{70}$$

Setting $\epsilon = \frac{1}{2}(1 - \rho)$, we have

$$\frac{1}{1 - \tilde{\rho}} \leq \frac{1}{1 - \frac{1}{1 + \frac{1}{2}(1-\rho)}} = \frac{2}{1 - \rho}. \tag{71}$$

Eq. (71) combined with the Kingman's formula proves our claim, for the case with uniform traffic rates ($\lambda_{i,n} = \rho/r$ for all i). The validity of Eq. (4) in the general case (λ satisfying Condition 1) can be established with an identical argument.

A.8 Proof of Eq. (6)

Define

$$f_{n,K'}(g_n, \lambda_n) = \mathbb{I} \left(\mathbb{E}_{\pi_n}(W|g_n, \lambda_n) \leq K' \frac{u^2(n) \ln n}{d(n)} \right), \tag{72}$$

where π_n be the scheduling policy given in Theorem 1. Let G_n be an $(rn, n, \frac{d(n)}{n})$ random bipartite graph. By the first two claims of Theorem 1, there exists $K > 0$ such that whenever

$$\mathbb{P}_{n, \frac{d(n)}{n}}(f_{n,K}(G_n, \lambda_n) = 1) \geq 1 - \delta_n, \tag{73}$$

for all n and $\lambda_n \in \tilde{\Lambda}_n$, where $\delta_n \downarrow 0$ and do not depend on λ_n . Now let λ_n be distributed as μ_n , and G_n as $\mathbb{P}_{n, \frac{d(n)}{n}}$. We have, by Eq. (73), that

$$\begin{aligned}
&\mathbb{P}(f_{n,K}(G_n, \lambda_n) = 1) \\
&= \int_{\lambda_n} \mathbb{P}_{n, \frac{d(n)}{n}}(f_{n,K}(G_n, \lambda_n) = 1) d\mu_n \\
&\geq \int_{\lambda_n} (1 - \delta_n) d\mu_n \\
&= 1 - \delta_n. \tag{74}
\end{aligned}$$

This proves Eq. (6).

B. TECHNICAL LEMMAS

Lemma 13. *Let X, Y and Z be random variables taking values in \mathbb{Z}_+ . Let random variables $X, \{Y_i\}_{i \geq 1}$ and $\{Z_{i,j}\}_{i,j \geq 1}$ be drawn*

independently according to the distributions of X, Y and Z , respectively. Define

$$W = \sum_{i=1}^X \sum_{j=1}^{Y_i} Z_{i,j}.$$

We have

$$\mathbb{E}(W) = \mu_X \mu_Y \mu_Z,$$

$$\mathbb{E}(W^2) \leq \mu_X \mu_Y s_Z + \mu_X s_Y \mu_Z^2 + s_X \mu_Y^2 \mu_Z^2,$$

where $\mu_X \triangleq \mathbb{E}(X)$, $s_X \triangleq \mathbb{E}(X^2)$, and similarly for Y and Z .

Lemma 14. Let X_i be a sequence of i.i.d. random variables, with $\mathbb{E}(X) = \mu$ and $\mathbb{E}(X^2) = s$, then

$$\mathbb{E}\left(\left(\sum_{i=1}^n X_i\right)^2\right) = n s + n(n-1) \mu^2. \quad (75)$$

In the special case where $X_i \stackrel{d}{=} \text{Expo}(\lambda)$, we have

$$\mathbb{E}\left(\left(\sum_{i=1}^n X_i\right)^2\right) = \frac{2}{\lambda^2} n + \frac{1}{\lambda^2} n(n-1) = \frac{n(n+1)}{\lambda^2}. \quad (76)$$