# A LOWER BOUND ON THE QUEUEING DELAY IN RESOURCE CONSTRAINED LOAD BALANCING 

By David Gamarnik ${ }^{1}$, John N. Tsitsiklis ${ }^{2, *}$ and Martin Zubeldia ${ }^{2, * *}$<br>${ }^{1}$ Sloan School of Management, Massachusetts Institute of Technology, gamarnik@mit.edu<br>${ }^{2}$ Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, " ${ }^{*}$ int@mit.edu;<br>** mar.zubeldia@gmail.com


#### Abstract

We consider the following distributed service model: jobs with unit mean, general distribution, and independent processing times arrive as a renewal process of rate $\lambda n$, with $0<\lambda<1$, and are immediately dispatched to one of several queues associated with $n$ identical servers with unit processing rate. We assume that the dispatching decisions are made by a central dispatcher endowed with a finite memory, and with the ability to exchange messages with the servers.

We study the fundamental resource requirements (memory bits and message exchange rate), in order to drive the expected queueing delay in steadystate of a typical job to zero, as $n$ increases. We develop a novel approach to show that, within a certain broad class of "symmetric" policies, every dispatching policy with a message rate of the order of $n$, and with a memory of the order of $\log n$ bits, results in an expected queueing delay which is bounded away from zero, uniformly as $n \rightarrow \infty$. This complements existing results which show that, in the absence of such limitations on the memory or the message rate, there exist policies with vanishing queueing delay (at least with Poisson arrivals and exponential service times).


1. Introduction. Distributed processing systems are ubiquitous, from passport control at the airport and checkout lines at the supermarket, to call centers and server farms for cloud computing. Many of these systems involve a stream of incoming jobs dispatched to distributed queues, with each queue associated to a different server; see Figure 1 for a stylized model. Naturally, the performance of such systems depends critically on the policy used to dispatch jobs to queues.

In order to take full advantage of the multiple servers, the dispatcher can benefit from information about the current state of the queues (e.g., whether they are empty or not). For such information to be available when a job arrives to the system and a dispatching decision is to be made, it is necessary that the dispatcher periodically obtain information about the current queue states from the corresponding servers and/or have sufficient memory that allows it to extrapolate from the available information. In this paper, we explore the tradeoff between the performance of the system and the amount of resources used to gather and maintain relevant information.

There is a variety of ways in which the system described above can be operated; these correspond to different dispatching policies, and result in different performance and resource utilization. At one extreme, the dispatcher can route incoming jobs to a queue chosen uniformly at random. This policy requires no information on the state of the queues, but the jobs experience considerable delays. At the other extreme, the dispatcher can send incoming jobs to a queue with the smallest number of jobs or to a queue with the smallest workload. The jobs in these last two policies experience little or no delay, but substantial communication overhead between the dispatcher and the servers is required.

[^0]

Fig. 1. Parallel server queueing system with a central dispatcher.

Many intermediate policies have been proposed and analyzed in the past (e.g., the power-of- $d$-choices or join-an-idle-queue), and they use varying amounts of resources to achieve different levels of delay performance. A more detailed discussion of the relevant literature and the various schemes therein is deferred to Section 4.

Instead of focusing on yet another policy or decision making architecture, we step back and address a more fundamental question: what is the minimum amount of resources required to obtain the best possible performance, as the number of server increases? Regarding performance, we focus on the expected time that a job has to wait before starting service. Regarding resources, we focus on the average number of messages exchanged between the dispatcher and the servers per unit of time, and on the number of bits of "long term" memory that the dispatcher has at its disposal.
1.1. Relationship with the "balls into bins" model. Some performance-resources tradeoffs similar to the ones we study have been analyzed in the context of the balls into bins model [2], in which $n$ balls are to be placed sequentially into $n$ bins. In particular, the tradeoff between the number of messages exchanged and the maximum number of balls in any one bin was recently characterized in [10], showing that a constant maximum load can be obtained using $\log n$ rounds of communication for each ball. Furthermore, the tradeoff between memory size and maximum number of balls in any one bin was studied in [1]. There, the authors showed that if the dispatcher has $m$ bits of memory, and is only given $k$ random choices of bins for each ball as destinations, then a constant load can only be achieved if $k m \gg n$.

Note that the balls into bins model and the dynamic model that we consider are similar in that both involve sequential decisions about the destination of balls (or jobs). Furthermore, the performance metric of interest for the balls into bins model is the maximum number of balls placed in any one bin. Thus, there is the same incentive to send balls to bins with the least amount of balls in them, as in our dynamic model. As a consequence, there are many policies that are used and perform well in both settings. However, there are two fundamental differences that make the two models substantially different.

1. Balls do not leave the system. While this may not seem like a major difference at first glance, its impact is best highlighted by the following example: The Round-Robin policy is optimal for the balls into bins model (all bins end up with exactly one ball), but it can be far from optimal in the dynamic model.
2. Maximum load as a performance metric. This metric makes outliers extremely important in the balls into bins model, whereas in our dynamic model we are concerned with the
average load. Thus, a policy with a low average load and large outliers would be good for our dynamic model, but poor for the balls into bins model.

Because of these differences, results from one setting cannot be translated into the other.
1.2. Our contribution. We consider a broad family of decision making architectures and policies, which includes most of those considered in the earlier literature, and work toward characterizing the attainable delay performance for a given level of resources. We allow the dispatcher to have a limited memory, where it can store information on the state of the queues. We also allow the exchange of messages, either from the dispatcher to the servers (queries), or from the servers to the dispatcher (responses to queries or spontaneous status updates).

We show that if the average message rate is at most of the order of the arrival rate, and the memory size (in bits) is at most of the order of the logarithm of the number of servers, then every decision making architecture and policy, within a certain broad class of dispatching policies, results in a queueing delay that does not vanish as the system size increases. In particular, we show that the expected queueing delay in steady-state of a typical job is uniformly bounded away from zero as the number of servers goes to infinity. The main constraints that we impose on the policies that we consider are: (i) there is no queueing at the dispatcher, that is, each job is immediately dispatched to one of the parallel queues, and (ii) policies are symmetric, in a sense to be made precise later.

REMARK 1.1. For the case of Poisson arrivals and exponential service times, if either (i) the message rate is superlinear in the arrival rate [14], or if (ii) the memory size in bits is superlogarithmic in the number of servers and the message rate is greater than or equal to the arrival rate [7], then there exists a dispatching policy that results in a vanishing queueing delay as the system size increases. This gives sufficient conditions to achieve a vanishing queueing delay that are complementary to the necessary conditions obtained in this paper.

REMARK 1.2. The logarithmic threshold in the number of bits comes from the fact that the number of bits required to store the ID of a server is of the order of the logarithm of the number of servers. Thus, policies that have a logarithmic number of bits of memory can only store a constant number of server IDs in memory.
1.3. Outline of the paper. The remainder of the paper is organized as follows. In Section 2 we introduce some notation. The model and the main result are presented in Section 3. In Section 4 we discuss our result in the context of some concrete dispatching policies from the earlier literature. In Section 5 we provide the proof of our main result. Finally, in Section 6 we present our conclusions and suggestions for future work.
2. Notation. In this section we introduce some notation that will be used throughout the paper. For any positive functions $f$ and $g$, we write

$$
f(n) \in \omega(g(n)) \quad \text { if and only if } \quad \liminf _{n \rightarrow \infty} \frac{f(n)}{g(n)}=\infty
$$

We let $[\cdot]^{+} \triangleq \max \{\cdot, 0\}$. We let $\mathbb{Z}_{+}$and $\mathbb{R}_{+}$be the sets of nonnegative integers and real numbers, respectively. The indicator function is denoted by $\mathbb{1}$, so that $\mathbb{1}_{A}(x)$ is 1 if $x \in A$, and is 0 otherwise. Given a set $A$, its power set, the set of all subsets of $A$, is denoted by $\mathcal{P}(A)$. Random variables will always be denoted by upper case symbols. Nonrandom quantities will generally—but not always-be denoted by lower case symbols; exceptions will be pointed out as necessary.

We will use boldface fonts to denote vectors. If $\mathbf{v}$ is a vector, we denote its $i$ th component by $\mathbf{v}_{i}$. We will denote the (unordered) set of elements of a vector by using the superscript "set"; for example, if $\mathbf{v}=(2,1,3,1)$, then $\mathbf{v}^{\text {set }}=\{1,2,3\}$. Furthermore, we will use $|\mathbf{v}|$ to denote the dimension of a vector $\mathbf{v}$. If $\mathbf{v}=\left(\mathbf{v}_{1}, \ldots, \mathbf{v}_{m}\right)$ is a vector, and $\mathbf{u}$ is a vector with entries in $\{1, \ldots, m\}$, then $\mathbf{v}_{u}$ is a $|\mathbf{u}|$-dimensional vector whose $i$ th component is $\mathbf{v}_{\mathbf{u}_{i}}$; for example, if $\mathbf{u}=(3,1)$, then $\mathbf{v}_{\mathbf{u}}=\left(\mathbf{v}_{3}, \mathbf{v}_{1}\right)$.

For any positive integer $n$, we define the sets $\mathcal{N}_{n} \triangleq\{1, \ldots, n\}$, and

$$
\mathcal{S}_{n} \triangleq\left\{\mathbf{s} \in \bigcup_{i=0}^{n}\left(\mathcal{N}_{n}\right)^{i}: \text { there are no repeated elements in } \mathbf{s}\right\}
$$

where $\left(\mathcal{N}_{n}\right)^{0}=\{\varnothing\}$. We say that a permutation $\sigma: \mathcal{N}_{n} \rightarrow \mathcal{N}_{n}$ fixes a set $R \subset \mathcal{N}_{n}$ if $\sigma(i)=i$, for all $i \in R$. Furthermore, we say that a permutation $\sigma$ preserves the ordering of a subset $A \subset \mathcal{N}_{n}$ if $\sigma(i)<\sigma(j)$ whenever $i, j \in A$ and $i<j$. If $\mathbf{v}=\left(\mathbf{v}_{1}, \ldots, \mathbf{v}_{m}\right)$ is a vector in $\left(\mathcal{N}_{n}\right)^{m}$ and $\sigma$ is a permutation of $\mathcal{N}_{n}$, we denote by $\sigma(\mathbf{v})$ the vector $\left(\sigma\left(\mathbf{v}_{1}\right), \ldots, \sigma\left(\mathbf{v}_{m}\right)\right)$. Finally, for any function $X(\cdot)$ of time, and any $t \in \mathbb{R}$, we let $X\left(t^{-}\right)=\lim _{\tau \uparrow t} X(\tau)$, as long as the limit exists.
3. Model and main results. In this section we present our main result. We first present a unified framework that defines a broad set of dispatching policies, which includes most of the policies studied in previous literature. We then present our negative result on the expected queueing delay under resource constrained policies within this set of policies.
3.1. Modeling assumptions. We consider a system consisting of $n$ parallel servers, where each server has a processing rate equal to 1 . Furthermore, each server is associated with an infinite capacity FIFO queue. Jobs arrive to the system as a single renewal process of rate $\lambda n$, for some fixed $\lambda<1$. Job sizes are i.i.d., independent from the arrival process, and have an arbitrary distribution with mean 1 . We use the convention that a job that is being served remains in queue until its processing is completed. We assume that each server is workconserving: a server is idle if and only if the corresponding queue is empty.

A central controller (dispatcher) is responsible for routing each incoming job to a queue, immediately upon arrival. The dispatcher has limited information on the state of the queues; it can only rely on a limited amount of local memory and on messages that provide partial information about the state of the system. These messages (which are assumed to be instantaneous) can be sent from a server to the dispatcher at any time, or from the dispatcher to a server (in the form of queries) at the time of an arrival. Messages from a server can only contain information about the state of its own queue (number of remaining jobs and the remaining workload of each one). Within this context, a system designer has the freedom to choose a messaging policy, as well as the rules for updating the memory and for selecting the destination of an incoming job.

We are interested in the case where $n$ is very large, in the presence of constraints on the rate of message exchanges and on the memory size. The performance metric that we focus on is the expected queueing delay in steady-state of a typical job, that is, the expected time between its arrival and the time at which it starts receiving service. We will formalize this definition in Section 3.3.
3.2. Unified framework for dispatching policies. In this subsection we present a unified framework that describes memory-based dispatching policies. In order to do this, we introduce a sample path construction of the evolution of the system under an arbitrary policy.

Let $c_{n}$ be the number of memory bits available to the dispatcher. We define the corresponding set of memory states to be $\mathcal{M}_{n} \triangleq\left\{1, \ldots, 2^{c_{n}}\right\}$. Furthermore, we define the set of possible
states at a server as the set of nonnegative sequences $\mathcal{Q} \triangleq \mathbb{R}_{+}^{\mathbb{Z}_{+}}$, where a sequence specifies the remaining workload of each job in that queue, including the one that is being served. (In particular, an idle server is represented by the zero sequence.) As long as a queue has a finite number of jobs, the queue state is a sequence that has only a finite number of nonzero entries. The reason that we include the workload of the jobs in the state is that we wish to allow for a broad class of policies, that can take into account the remaining workload in the queues. In particular, we allow for information-rich messages that describe the full workload sequence at the server that sends the message. We are interested in the process

$$
\mathbf{Q}(\cdot)=\left(\mathbf{Q}_{1}(\cdot), \ldots, \mathbf{Q}_{n}(\cdot)\right)=\left(\left(\mathbf{Q}_{1, j}(\cdot)\right)_{j=1}^{\infty}, \ldots,\left(\mathbf{Q}_{n, j}(\cdot)\right)_{j=1}^{\infty}\right) \in \mathcal{Q}^{n}
$$

which describes the evolution of the workload of each job in each queue. Here $\mathbf{Q}_{i, j}(t)$ is the remaining workload of the $j$ th job in the $i$ th queue, at time $t$, which for $j \geq 2$ is simply the job's service time. We are also interested in the process $M(\cdot)$ that takes values in the set $\mathcal{M}_{n}$ and describes the evolution of the memory state, and in the process $Z(\cdot)$ that describes the remaining time until the next arrival of a job.
3.2.1. Fundamental processes and initial conditions. The processes of interest will be driven by certain common fundamental processes.

1. Arrival process: A delayed renewal counting process $A_{n}(\cdot)$ with rate $\lambda n$, and event times $\left\{T_{k}\right\}_{k=1}^{\infty}$, defined on a probability space $\left(\Omega_{A}, \mathcal{A}_{A}, \mathbb{P}_{A}\right)$.
2. Spontaneous messages process: A Poisson counting process $R_{n}(\cdot)$ with rate $\mu n$, and event times $\left\{T_{k}^{s}\right\}_{k=1}^{\infty}$, defined on a probability space $\left(\Omega_{R}, \mathcal{A}_{R}, \mathbb{P}_{R}\right)$.
3. Job sizes: A sequence of i.i.d. random variables $\left\{W_{k}\right\}_{k=1}^{\infty}$ with mean one, defined on a probability space $\left(\Omega_{W}, \mathcal{A}_{W}, \mathbb{P}_{W}\right)$.
4. Randomization variables: Four independent and individually i.i.d. sequences of random variables $\left\{U_{k}\right\}_{k=1}^{\infty},\left\{V_{k}\right\}_{k=1}^{\infty},\left\{X_{k}\right\}_{k=1}^{\infty}$, and $\left\{Y_{k}\right\}_{k=1}^{\infty}$, uniform on [0, 1], defined on a common probability space $\left(\Omega_{U}, \mathcal{A}_{U}, \mathbb{P}_{U}\right)$.
5. Initial conditions: Random variables $\mathbf{Q}(0), M(0)$, and $Z(0)$, defined on a common probability space $\left(\Omega_{0}, \mathcal{A}_{0}, \mathbb{P}_{0}\right)$.

The whole system will be defined on the associated product probability space

$$
\begin{aligned}
& \left(\Omega_{A} \times \Omega_{R} \times \Omega_{W} \times \Omega_{U} \times \Omega_{0}, \mathcal{A}_{A} \times \mathcal{A}_{R} \times \mathcal{A}_{W} \times \mathcal{A}_{U} \times \mathcal{A}_{0}\right. \\
& \left.\quad \mathbb{P}_{A} \times \mathbb{P}_{R} \times \mathbb{P}_{W} \times \mathbb{P}_{U} \times \mathbb{P}_{0}\right)
\end{aligned}
$$

to be denoted by $(\Omega, \mathcal{A}, \mathbb{P})$. All of the randomness in the system originates from these fundamental processes, and everything else is a deterministic function of them.
3.2.2. A construction of sample paths. We consider some fixed $n$, and provide a construction of a Markov process $(\mathbf{Q}(\cdot), M(\cdot), Z(\cdot))$, that takes values in the set $\mathcal{Q}^{n} \times \mathcal{M}_{n} \times \mathbb{R}_{+}$. The memory process $M(\cdot)$ is piecewise constant, and can only jump at the time of an event. All processes to be considered will have the càdlàg property (right-continuous with left limits) either by assumption (e.g., the underlying fundamental processes) or by construction.

There are three types of events: job arrivals, spontaneous messages, and service completions. We now describe the sources of these events, and what happens when they occur.

Job arrivals: At the time of the $k$ th event of the arrival process $A_{n}(\cdot)$, which occurs at time $T_{k}$ and involves a job with size $W_{k}$, the following transitions happen sequentially but instantaneously.

1. First, the dispatcher chooses a vector of distinct servers $\mathbf{S}_{k}$, from which it solicits information about their state, according to

$$
\mathbf{S}_{k}=f_{1}\left(M\left(T_{k}^{-}\right), W_{k}, U_{k}\right)
$$

where $f_{1}: \mathcal{M}_{n} \times \mathbb{R}_{+} \times[0,1] \rightarrow \mathcal{S}_{n}$ is a measurable function defined by the policy. Note that the set of servers that are sampled only depends on the current memory state and on the size of the incoming job, but it is chosen in a randomized way, thanks to the independent random variable $U_{k}$. Thus, we allow for randomized policies; for example, the dispatcher might choose to sample a fixed number of servers uniformly at random.
2. Then, messages are sent to the servers in the vector $\mathbf{S}_{k}$, and the servers respond with messages containing their queue states; thus, the information received by the dispatcher is the vector $\mathbf{Q}_{\mathbf{s}_{k}}$. This results in $2\left|\mathbf{S}_{k}\right|$ messages exchanged. Using this information, the destination of the incoming job is chosen to be

$$
D_{k}=f_{2}\left(M\left(T_{k}^{-}\right), W_{k}, \mathbf{S}_{k}, \mathbf{Q}_{\mathbf{S}_{k}}\left(T_{k}^{-}\right), V_{k}\right),
$$

where $f_{2}: \mathcal{M}_{n} \times \mathbb{R}_{+} \times \mathcal{S}_{n} \times\left(\bigcup_{i=0}^{n} \mathcal{Q}^{i}\right) \times[0,1] \rightarrow \mathcal{N}_{n}$ is a measurable function defined by the policy. Note that the destination of a job can only depend on the current memory state, the job size, as well as the vector of queried servers and the state of their queues, but it is chosen in a randomized way, thanks to the independent random variable $V_{k}$. Once again, we allow for randomized policies that, for example, dispatch jobs uniformly at random.
3. Finally, the memory state is updated according to

$$
M\left(T_{k}\right)=f_{3}\left(M\left(T_{k}^{-}\right), W_{k}, \mathbf{S}_{k}, \mathbf{Q}_{\mathbf{S}_{k}}\left(T_{k}^{-}\right), D_{k}\right),
$$

where $f_{3}: \mathcal{M}_{n} \times \mathbb{R}_{+} \times \mathcal{S}_{n} \times\left(\bigcup_{i=0}^{n} \mathcal{Q}^{i}\right) \times \mathcal{N}_{n} \rightarrow \mathcal{M}_{n}$ is a measurable function defined by the policy. Note that the new memory state is obtained using the same information as for selecting the destination, plus the destination of the job, but without randomization.

Spontaneous messages: At the time of the $k$ th event of the spontaneous message process $R_{n}(\cdot)$, which occurs at time $T_{k}^{s}$, the $i$ th server sends a spontaneous message to the dispatcher if and only if

$$
g_{1}\left(\mathbf{Q}\left(T_{k}^{s}\right), X_{k}\right)=i,
$$

where $g_{1}: \mathcal{Q}^{n} \times[0,1] \rightarrow\{0\} \cup \mathcal{N}_{n}$ is a measurable function defined by the policy. In that case, the memory is updated to the new memory state

$$
M\left(T_{k}^{S}\right)=g_{2}\left(M\left(T_{k}^{s-}\right), i, \mathbf{Q}_{i}\left(T_{k}^{S}\right)\right)
$$

where $g_{2}: \mathcal{M}_{n} \times \mathcal{N}_{n} \times \mathcal{Q} \rightarrow \mathcal{M}_{n}$ is a measurable function defined by the policy, and which prescribes the server who sends a message. On the other hand, no message is sent when $g_{1}\left(\mathbf{Q}\left(T_{k}^{s}\right), X_{k}\right)=0$. Note that the dependence of $g_{1}$ on $\mathbf{Q}$ allows the message rate at each server to depend on the server's current workload. For example, we could let idle servers send repeated spontaneous messages (as a Poisson process) to inform the dispatcher of their idleness.

Service completions: As time progresses, the remaining workload of each job that is at the head of line in a queue decreases at a constant, unit rate. When a job's workload reaches zero, the job leaves the system and every other job advances one slot. Let $\left\{T_{k}^{d}(i)\right\}_{k=1}^{\infty}$ be the sequence of departure times at the $i$ th server. At those times, the $i$ th server sends a message to the dispatcher if and only if

$$
h_{1}\left(\mathbf{Q}_{i}\left(T_{k}^{d}(i)\right), Y_{k}\right)=1
$$

where $h_{1}: \mathcal{Q} \times[0,1] \rightarrow\{0,1\}$ is a measurable function defined by the policy. In that case, the memory is updated to the new memory state

$$
M\left(T_{k}^{d}(i)\right)=h_{2}\left(M\left(T_{k}^{d}(i)^{-}\right), i, \mathbf{Q}_{i}\left(T_{k}^{d}(i)\right)\right)
$$

where $h_{2}: \mathcal{M}_{n} \times \mathcal{N}_{n} \times \mathcal{Q} \rightarrow \mathcal{M}_{n}$ is a measurable function defined by the policy. On the other hand, no message is sent when $h_{1}\left(\mathbf{Q}_{i}\left(T_{k}^{d}(i)\right), Y_{k}\right)=0$.

REMARK 3.1. We have chosen to describe the collection of queried servers by a vector, implying an ordering of the servers in that collection. We could have described this collection as an (unordered) set. These two options are essentially equivalent but it turns out that the ordering provided by the vector description allows for a simpler presentation of the proof.

REMARK 3.2. For any given $n$, a policy is completely determined by the spontaneous message rate $\mu$, and the functions $f_{1}, f_{2}, f_{3}, g_{1}, g_{2}, h_{1}$ and $h_{2}$. Furthermore, many policies in the literature that are described without explicit mention of memory or messages can be cast within our framework, as we will see in Section 4.

REMARK 3.3. The memory update functions $f_{3}, g_{2}$ and $h_{2}$ do not involve randomization, even though our main result could be extended in that direction. We made this choice because none of the policies introduced in earlier literature require such randomization, and because it simplifies notation and the proofs.

REMARK 3.4. We only consider the memory used to store information in between arrivals or messages. Thus, when counting the memory resources used by a policy, we do not take into account information that is used in zero time (e.g., the responses from the queries at the time of an arrival) or the memory required to evaluate the various functions that describe the policy. If that additional memory were to be accounted for, then any memory constraints would be more severe, and therefore our negative result would still hold.

The dispatching policies that we have introduced obey certain constraints:
(i) The dispatcher can only send messages to the servers at the time of an arrival, and in a single round of communication. This eliminates the possibility of policies that sequentially poll the servers uniformly at random until they find an idle one. Indeed, it can be shown that such sequential polling policies may lead to asymptotically vanishing delays, without contradicting our lower bounds. On the other hand, in practice, queries involve some processing and travel time $\epsilon$. Were we to consider a more realistic model with $\epsilon>0$, sequential polling would also incur positive delay.
(ii) We assume that the dispatcher must immediately send an incoming job to a server upon arrival. This prevents the dispatcher from maintaining a centralized queue and operating the system as a $\mathrm{G} / \mathrm{G} / n$ queue.

We now introduce a symmetry assumption on the policies. In essence it states that at the time of a job arrival, and given the current memory state, if certain sampling and dispatching decisions and a certain memory update are possible, then a permuted version of these decisions and updates is also possible (and equally likely), starting with a suitably permuted memory state.

AsSumption 3.1 (Symmetric policies). We assume that the dispatching policy is symmetric, in the following sense. For any given permutation of the servers $\sigma$, there exists a corresponding (not necessarily unique) permutation $\sigma_{M}$ of the memory states $\mathcal{M}_{n}$ that satisfies all of the following properties.

1. For every $m \in \mathcal{M}_{n}$ and $w \in \mathbb{R}_{+}$, and if $U$ is a uniform random variable on $[0,1]$, then

$$
\sigma\left(f_{1}(m, w, U)\right) \stackrel{d}{=} f_{1}\left(\sigma_{M}(m), w, U\right)
$$

where $\stackrel{d}{=}$ stands for equality in distribution.
2. For every $m \in \mathcal{M}_{n}, w \in \mathbb{R}_{+}, \mathbf{s} \in \mathcal{S}_{n}$, and $\mathbf{q} \in \mathcal{Q}^{|\mathbf{s}|}$, and if $V$ is a uniform random variable on $[0,1]$, then ${ }^{1}$

$$
\sigma\left(f_{2}(m, w, \mathbf{s}, \mathbf{q}, V)\right) \stackrel{d}{=} f_{2}\left(\sigma_{M}(m), w, \sigma(\mathbf{s}), \mathbf{q}, V\right)
$$

3. For every $m \in \mathcal{M}_{n}, w \in \mathbb{R}_{+}, \mathbf{s} \in \mathcal{S}_{n}$, and $\mathbf{q} \in \mathcal{Q}^{|s|}$, and $d \in \mathcal{N}_{n}$, we have

$$
\sigma_{M}\left(f_{3}(m, w, \mathbf{s}, \mathbf{q}, d)\right)=f_{3}\left(\sigma_{M}(m), w, \sigma(\mathbf{s}), \mathbf{q}, \sigma(d)\right)
$$

As a concrete illustration, our symmetry assumption implies the following. If a certain memory state mandates that the vector $(2,4,5)$ of servers must be sampled (with probability 1 ), independently from the incoming job size, then there exists some other memory state which mandates that the vector $(1,5,7)$ will be sampled, independently from the incoming job size, and the same holds for every 3-element vector with distinct entries. Since there are $n(n-1)(n-2)$ different vectors, there must be at least so many different memory states. This suggests that if we have too few memory states, the number of "distinguished" servers, that is, servers that are treated in a special manner is severely limited. This is a key element of the proof of the delay lower bound that we present in the next subsection.

REMARK 3.5. One may contemplate a different (stronger) definition of symmetry. For example, in the first part, we could have required that $\sigma\left(f_{1}(m, w, u)\right)=f_{1}\left(\sigma_{M}(m), w, u\right)$, for all $u \in[0,1]$. While this would lead to a simpler proof, this stronger definition would be too restrictive, as explained in Appendix A.

REMARK 3.6. Note that a symmetry assumption is imposed on the memory update function $f_{3}$ at the time that a job is dispatched. However, we do not need to impose a similar assumption on the memory update functions $g_{2}$ and $h_{2}$ at the times that the dispatcher receives a message. Similarly, there is no symmetry assumption on the functions $g_{1}$ and $h_{1}$ that govern the generation of server messages. In particular, we allow each server to send spontaneous messages at its own identity-dependent, and hence asymmetric, rate.
3.3. Delay lower bound for resource constrained policies. Before stating the main result of this paper, we introduce formal definitions for the average message rate between the dispatcher and the servers, and for our performance metric for the delay. Furthermore, we introduce an assumption on the arrival process.

First, given a policy of the form specified in the previous subsection, we define the average message rate between the dispatcher and the servers as

$$
\begin{align*}
& \limsup _{t \rightarrow \infty} \frac{1}{t}\left[\sum_{k=1}^{A_{n}(t)} 2\left|\mathbf{S}_{k}\right|+\sum_{k=1}^{R_{n}(t)} \mathbb{1}_{\mathcal{N}_{n}}\left(g_{1}\left(\mathbf{Q}\left(T_{k}^{S}\right), X_{k}\right)\right)\right. \\
& \left.\quad+\sum_{i=1}^{n} \sum_{k: T_{k}^{d}(i)<t} \mathbb{1}_{\{1\}}\left(h_{1}\left(\mathbf{Q}_{i}\left(T_{k}^{d}(i)\right), Y_{k}\right)\right)\right] \tag{3.1}
\end{align*}
$$

[^1]Second, we provide a formal definition of our performance metric for the delay. We assume that the process $(\mathbf{Q}(\cdot), M(\cdot), Z(\cdot))$ is stationary, with invariant probability measure $\pi$. Since the destinations of jobs (and their queueing delays) are deterministic functions of the state and i.i.d. randomization variables, the point process of arrivals with the queueing delays as marks, is also stationary. Using this, we define the expected queueing delay in steady-state $\pi$ of a typical job, denoted by $\mathbb{E}_{\pi}^{0}\left[L_{0}\right]$, as follows. If $L_{k}$ is the queueing delay of the $k$ th job under the stationary process $(\mathbf{Q}(\cdot), M(\cdot), Z(\cdot))$, then

$$
\begin{equation*}
\mathbb{E}_{\pi}^{0}\left[L_{0}\right] \triangleq \mathbb{E}_{\pi}\left[\frac{1}{\lambda n t} \sum_{k=1}^{A_{n}(t)} L_{k}\right] \tag{3.2}
\end{equation*}
$$

where the right-hand side is independent from $t$ due to the stationarity of the processes involved (see [3]). Furthermore, if the stationary process $(\mathbf{Q}(\cdot), M(\cdot), Z(\cdot))$ is ergodic (in the sense that every invariant set has measure either 0 or 1 under $\pi$ ), we have

$$
\mathbb{E}_{\pi}^{0}\left[L_{0}\right]=\lim _{t \rightarrow \infty} \frac{1}{A_{n}(t)} \sum_{k=1}^{A_{n}(t)} L_{k} \quad \text { a.s. }
$$

Finally, we introduce an assumption on the arrival process.

ASSUMPTION 3.2. Let $I_{n}$ be distributed as the typical inter-arrival times of the delayed renewal process $A_{n}(\cdot)$. We assume that there exists a constant $\bar{\epsilon}>0$, independent from $n$, such that the following holds. For every $\epsilon \in(0, \bar{\epsilon}]$, there exists a positive constant $\delta_{\epsilon}$ such that

$$
\delta_{\epsilon} \leq \mathbb{P}\left(I_{n} \leq \frac{\epsilon}{n}\right) \leq 1-\delta_{\epsilon},
$$

for all $n$.
This assumption implies that arbitrarily small inter-arrival times of order $\Theta(1 / n)$ occur with a probability that is bounded away from 0 , and from 1 , for all $n$. In particular, this excludes deterministic inter-arrival times, and inter-arrival times that can take values of order $o(1 / n)$ with probability of order $1-o(1)$. On the other hand, if $A(\cdot)$ is a delayed renewal process, where the typical inter-arrival times are continuous random variables with positive density around 0 , then the process $A_{n}(\cdot)$, defined as $A_{n}(t) \triangleq A(n t)$ for all $t \geq 0$, satisfies Assumption 3.2.

We are now ready to state the main result. It asserts that within the class of symmetric policies that we consider, and under some upper bounds on the memory size (logarithmic) and the message rate (linear), the expected queueing delay in steady-state of a typical job is bounded below by a positive constant.

THEOREM 3.1 (Positive delay for resource constrained policies). For any constants $\lambda \in$ $(0,1), c, \alpha>0$, and for every arrival process that satisfies Assumption 3.2, there exists $a$ constant $\zeta(\lambda, c, \alpha)>0$ with the following property. For any fixed $n$, consider a symmetric memory-based dispatching policy, that is, that satisfies Assumption 3.1, with at most $c \log _{2} n$ bits of memory, with an average message rate (cf. equation (3.1)) upper bounded by $\alpha n$ in expectation, and under which the process $(\mathbf{Q}(\cdot), M(\cdot), Z(\cdot))$ admits at least one invariant probability measure $\pi_{n}$. Then, for all $n$ large enough, we have

$$
\mathbb{E}_{\pi_{n}}^{0}\left[L_{0}\right] \geq \zeta(\lambda, c, \alpha)
$$

where $\mathbb{E}_{\pi_{n}}^{0}\left[L_{0}\right]$ is the expected queueing delay in steady-state $\pi_{n}$ of a typical job.

The proof is given in Section 5.
REMARK 3.7. Theorem 3.1 only deals with memory-based dispatching policies under which $(\mathbf{Q}(\cdot), M(\cdot), Z(\cdot))$ admits at least one invariant probability measure. This is not a substantial restriction since there are many policies with this property, for example, all policies reviewed in Section 4. On the other hand, policies without this property are not of interest.
4. Dispatching policies in the literature. In this section we put our results in perspective by showing that various dispatching policies considered earlier in the literature are special cases of the class of symmetric dispatching policies described above. Most policies have only been studied for the case of Poisson arrivals and exponential service times, so this review is restricted to that case unless stated otherwise.

### 4.1. Open-loop policies.

4.1.1. Random routing. The simplest policy is to dispatch each arriving job to a random queue, with each queue being equally likely to be selected. In this case, the system behaves as $n$ independent parallel $\mathrm{M} / \mathrm{M} / 1$ queues. This policy needs no messages or memory, and has a positive queueing delay independent of $n$.
4.1.2. Round Robin $(R R)$. When the dispatcher has no access to the workload of incoming jobs and no messages are allowed, it is optimal to dispatch arriving jobs to the queues in a round-robin fashion [15]. This policy does not require messages but needs $\left\lceil\log _{2} n\right\rceil$ bits of memory to store the ID of the next queue to receive a job. In the limit, each queue behaves like a D/M/1 queue (see [15]). While random routing is a symmetric policy, Round Robin is not. To see this, note that a memory state $i$ must be followed by state $i+1$, and such a transition is not permutation-invariant; in particular, the memory update function $f_{3}$ does not satisfy the symmetry assumption. Round Robin can be made symmetric by using an additional $n\left\lceil\log _{2} n\right\rceil$ bits of memory to specify the order with which the different servers are selected. But in any case, this policy also has a positive queueing delay, that does not vanish as $n$ increases.

### 4.2. Policies based on queue lengths.

4.2.1. Join a shortest queue (SQ). If we wish to minimize the queueing delay and have access to the queue lengths but not to the job sizes, an optimal policy is to have each incoming job join a shortest queue, breaking any ties uniformly at random [18]. When $n$ goes to infinity, the queueing delay vanishes, but this policy requires a message rate of $2 \lambda n^{2}$ ( $n$ queries and $n$ responses for each arrival), and no memory. This policy is symmetric and achieves vanishing delay, but uses a superlinear number of messages.
4.2.2. Join a shortest of $d$ random queues $(S Q(d))$. In order to sharply decrease the number of messages sent, Mitzenmacher [12] and Vvedenskaya et al. [17] introduced the power-of- $d$-choices policy. When there is an arrival, $d$ servers are chosen uniformly at random, and the job is sent to a shortest queue among those $d$ servers. This policy fits our framework, and in particular is symmetric; it uses $2 \lambda d n$ messages per unit of time, and zero memory. This policy was also analyzed in the case of heavy-tailed service times by Bramson et al. [5], yielding similar results. In any case, this policy has positive delay, which is consistent with Theorem 3.1.
4.2.3. Join a shortest of $d_{n}$ random queues $\left(S Q\left(d_{n}\right)\right)$. More recently, Mukherjee et al. [14] analyzed a variation of the $\mathrm{SQ}(d)$ policy, which lets $d$ be a function of the system size $n$. This policy is symmetric, uses $2 \lambda d_{n} n$ messages per unit of time and zero memory, and has zero delay as long as $d_{n} \rightarrow \infty$, which is consistent with Theorem 3.1.
4.2.4. Join a shortest of $d$ queues, with memory $(S Q(d, b))$. Another improvement over the power-of- $d$-choices, proposed by Mitzenmacher et al. in [13], is obtained by using extra memory to store the IDs of the $b$ (with $b \leq d$ ) least loaded queues known at the time of the previous arrival. When a new job arrives, $d$ queues are sampled uniformly at random and the job is sent to a least loaded queue among the $d$ sampled and the $b$ stored queues. This policy is symmetric, needs $2 \lambda d n$ messages per unit of time and $\Omega\left(b \log _{2} n\right)$ bits of memory, and has positive delay, consistent with Theorem 3.1.
4.2.5. $S Q(d)$ for divisible jobs. Recently, Ying et al. [19] considered the case of jobs of size $m_{n}$ (with $m_{n} \in \omega(1)$ and $m_{n} / n \rightarrow 0$ ) arriving as a Poisson process of rate $n \lambda / m_{n}$, where each job can be divided into $m_{n}$ tasks with mean size 1 . Then, the dispatcher samples $d m_{n}$ queues and does a water-filling of those queues with the $m_{n}$ tasks. In this case, the number of messages sent per unit of time is $2 \lambda d n$ and no memory is used. Even though this was not mentioned in [19], this policy can be shown to drive the queueing delay to 0 if $d \geq 1 /(1-\lambda)$. However, this model does not fall into our framework because it involves divisible jobs.

### 4.3. Policies based on remaining workload.

4.3.1. Join a least loaded queue $(L L)$. An appealing policy is the one that sends incoming jobs to a queue with the least remaining workload, in which case the whole system behaves as an $\mathrm{M} / \mathrm{M} / n$ queue. This policy is symmetric and achieves a vanishing delay as $n \rightarrow \infty$, but it has the same quadratic messaging requirements as SQ.
4.3.2. Join a least loaded of $d$ queues $(L L(d))$. A counterpart of $\operatorname{SQ}(d)$ is $\operatorname{LL}(d)$, in which the dispatcher upon arrival chooses $d$ queues uniformly at random and sends the job to one of those queues with the least remaining workload, breaking any ties uniformly at random. This setting was studied in [9], and it does not result in asymptotically vanishing delay, consistent with Theorem 3.1.
4.4. Policies based on job size. The previous policies dispatched the incoming jobs based on information about the state of the queues, obtained by dynamically exchanging messages with the servers. Such information could include the remaining workload at the different queues. On the other hand, if the dispatcher only knows the size of an incoming job (which might be difficult in practice [6]), it could use a static and memoryless policy that selects a target server based solely on the job size. Harchol-Balter et al. [8] showed that delay is minimized over all such static policies by a nonsymmetric policy that partitions the set of possible job sizes into consecutive intervals and assigns each interval to a different server. This is especially effective when the jobs have highly variable sizes (e.g., heavy-tailed), yet the resulting delay can be no better than that of an M/D/1 queue, and does not vanish as $n \rightarrow \infty$. This scheme does not require any message exchanges, and could be made symmetric by using the memory to store a list of the $n$ intervals of job sizes corresponding to each of the $n$ servers.

### 4.5. Pull-based load balancing.

4.5.1. Join-idle-queue (JIQ). In order to reduce the message rate, Badonnel and Burgess [4], Lu et al. [11] and Stolyar [16] propose a scheme where messages are sent from a server to the dispatcher whenever the server becomes idle, so that the dispatcher can keep track of the set of idle servers in real time. Then, an arriving job is to be sent to an empty queue (if there is one) or to a queue chosen uniformly at random (if all queues are nonempty). This policy requires at most $\lambda n$ messages per unit of time and exactly $n$ bits of memory (one bit for each queue, indicating whether it is empty or not). Stolyar [16] has shown that when $n$ goes to infinity, the average delay vanishes. This policy is symmetric. It has a vanishing delay and a linear message rate, but uses superlogarithmic memory, consistent with Theorem 3.1.
4.5.2. Resource constrained pull-based $(R C P B)$. In order to reduce the message rate and the memory usage, Gamarnik et al. [7] propose a family of dispatching policies, similar to Join-Idle-Queue, where the dispatcher keeps a small list of up to $c_{n}$ idle servers, and where messages are sent from each idle server to the dispatcher as a Poisson process of rate $v_{n}$. Then, an arriving job is sent to an empty queue (if the dispatcher knows the ID of an idle server) or to a queue chosen uniformly at random (if the dispatcher's list is empty). This policy requires $(1-\lambda) v_{n} n$ messages per unit of time and $c_{n} \log _{2} n$ bits of memory (the size of the list, $c_{n}$, times the number of bits required to store one ID, $\log _{2} n$ bits). Gamarnik et al. [7] showed that, if we either have a high message rate regime (RCPB-HMess) with $c_{n} \geq 1$ and $v_{n} \rightarrow \infty$, or a high memory regime (RCPB-HMem) with $c_{n} \rightarrow \infty$ and $v_{n} \geq \lambda /(1-\lambda)$, the expected queueing delay vanishes as $n \rightarrow \infty$. This policy is symmetric, and it only has a vanishing delay when either the message rate is superlinear, or the memory is superlogarithmic. This is consistent with Theorem 3.1.
4.6. Memory, messages and queueing delay. We now summarize the resource requirements (memory and message rate) and the asymptotic delay of the policies reviewed in this section that fall within our framework.

Note that any one of the listed policies that achieves vanishing queueing delay (see Table 1) falls into one (or both) of the following two categories:
(a) Those requiring $\omega(n)$ message rate, namely, $\mathrm{SQ}, \mathrm{SQ}\left(d_{n}\right)$ and LL.
(b) Those requiring $\omega\left(\log _{2} n\right)$ bits of memory (JIQ, and RCPB-HMem).

Our main result effectively establishes this fundamental limitation of symmetric policies.

Table 1
Memory usage, average message rate, and limiting delay for various policies in the literature

| Policy | Memory (bits) | Message rate | Limiting delay |
| :--- | :---: | :---: | :---: |
| Random | 0 | 0 | $>0$ |
| RR [15] | $\log _{2} n$ | 0 | $>0$ |
| SQ [18] | 0 | $2 \lambda n^{2}$ | 0 |
| SQ $(d)[12]$ | 0 | $2 d \lambda n$ | $>0$ |
| $\operatorname{SQ}\left(d_{n}\right)[14]$ | 0 | $\omega(n)$ | 0 |
| SQ $(d, b)[13]$ | $b \log _{2} n$ | $2 d \lambda n$ | $>0$ |
| LL | 0 | $2 \lambda n^{2}$ | 0 |
| LL $(d)[9]$ | 0 | $2 d \lambda n$ | $>0$ |
| JIQ [16] | $n$ | $\lambda n$ | 0 |
| RCPB-HMess [7] | $\log _{2} n$ | $\omega(n)$ | 0 |
| RCPB-HMem [7] | $\omega\left(\log _{2} n\right)$ | $\lambda n$ | 0 |

5. Proof of main result. Let us fix some $n$. In the sequel, we will assume that $n$ is large enough whenever needed for certain inequalities to hold. We fix a memory-based policy that satisfies Assumption 3.1 (symmetry), with at most $n^{c}$ memory states, and which results in the process $(\mathbf{Q}(\cdot), M(\cdot), Z(\cdot))$ having at least one invariant probability measure. Let us fix such an invariant probability measure $\pi_{n}$. We consider the process in steady-state; that is, we assume that $(\mathbf{Q}(0), M(0), Z(0))$ is distributed according to $\pi_{n}$. Accordingly, probabilities $\mathbb{P}(\cdot)$ and expectations $\mathbb{E}[\cdot]$ encountered in the sequel will always refer to the process in steady-state.

The high-level outline of the proof is as follows. In Section 5.1 we show that under our symmetry assumption, the dispatcher can give special treatment to at most $c$ servers, which we call distinguished servers. The treatment of all other servers, is symmetric, in some appropriate sense.

In Section 5.2 we consider a sequence of bad events under which, over a certain time interval, there are $c+1$ consecutive arrivals, no service completions or messages from the servers, and all sampled servers are "busy" with a substantial workload. Then, in Section 5.3, we show that this sequence of bad events has nonnegligible probability.

In Section 5.4, we develop some further consequences of the symmetry assumption, which we use to constrain the information available to the dispatcher at the time of the $(c+1) \mathrm{st}$ arrival. Loosely speaking, the idea is that during the interval of interest, the server only has information on $c$ distinguished servers together with (useless) information on some busy servers. This in turn implies (Section 5.5) that at least one of the first $c+1$ arrivals must be dispatched to a server on which no useful information is available, and which therefore has a nonnegligible probability of inducing a nonnegligible delay, thus completing the proof.
5.1. Local limitations of finite memory. We consider the (typical) case where a relatively small number of servers ( $\sqrt{n}$ or less) are sampled. We will use the symmetry assumption to show that except for a small set of distinguished servers, of size at most $c$, all other servers must be treated as indistinguishable.

Proposition 5.1. Let $U$ be a uniform random variable over $[0,1]$. For all $n$ large enough, for every memory state $m \in \mathcal{M}_{n}$ and every possible job size $w \in \mathbb{R}_{+}$, the following holds. Consider any vector of servers $\mathbf{s} \in \mathcal{S}_{n}$ (and its associated set of servers $\mathbf{s}^{\text {set }}$ ) with $|\mathbf{s}| \leq \sqrt{n}$, and any integer $\ell$ with $|\mathbf{s}|+1 \leq \ell \leq n$. Consider the event $B(m, w ; \mathbf{s}, \ell)$ that exactly $\ell$ servers are sampled and that the first $|\mathbf{s}|$ of them are the same as the vector $\mathbf{s}$, that is,

$$
B(m, w ; \mathbf{s}, \ell)=\left\{\left|f_{1}(m, w, U)\right|=\ell\right\} \cap \bigcap_{i=1}^{|\mathbf{s}|}\left\{f_{1}(m, w, U)_{i}=\mathbf{s}_{i}\right\}
$$

and assume that the conditional probability measure

$$
\mathbb{P}(\cdot \mid B(m, w ; \mathbf{s}, \ell))
$$

is well defined. Then, there exists a unique set $R(m, w, \mathbf{s}, \ell) \subset \mathcal{N}_{n} \backslash \mathbf{s e t}^{\text {set }}$ of minimal cardinality such that

$$
\begin{equation*}
\mathbb{P}\left(f_{1}(m, w, U)_{|\mathbf{s}|+1}=j \mid B(m, w ; \mathbf{s}, \ell)\right) \tag{5.1}
\end{equation*}
$$

is the same for all $j \notin R(m, w, \mathbf{s}, \ell) \cup \mathbf{s}^{\text {set }}$. Furthermore, $|R(m, w, \mathbf{s}, \ell)| \leq c$.
REmark 5.1. With some notational abuse, the measure $\mathbb{P}$ in Proposition 5.1 need not correspond to the measure $\mathbb{P}$ that describes the process. We are simply considering probabilities associated with a deterministic function of the uniform random variable $U$.

Proof of Proposition 5.1. Throughout the proof, we fix a particular memory state $m$, job size $w$, vector of servers $\mathbf{s}$ with $|\mathbf{s}| \leq \sqrt{n}$, and an integer $\ell$ in the range $|\mathbf{s}|+1 \leq \ell \leq n$. To simplify notation, we will suppress the dependence on $w$.

Consider the random vector $\mathbf{S}(m) \triangleq f_{1}(m, U)$. Let $\mathbf{v}$ be the vector whose components are indexed by $j$ ranging in the set $\left(\mathbf{s}^{\text {set }}\right)^{c}=\mathcal{N}_{n} \backslash \mathbf{s}^{\text {set }}$, and defined for any such $j$, by

$$
\mathbf{v}_{j}=\mathbb{P}\left(\mathbf{S}(m)_{|\mathbf{s}|+1}=j \mid B(m ; \mathbf{s}, \ell)\right)
$$

We need to show that for $j$ outside a "small" set, all of the components of $\mathbf{v}$ are equal. Let $z_{1}, \ldots, z_{d}$ be the distinct values of $\mathbf{v}_{j}$, as $j$ ranges over $\left(\mathbf{s}^{\text {set }}\right)^{c}$, and let $A_{\alpha}=\left\{j \in\left(\mathbf{s}^{\text {set }}\right)^{c} \mid\right.$ $\left.\mathbf{v}_{j}=z_{\alpha}\right\}$. The sequence of sets $\left(A_{1}, \ldots, A_{d}\right)$ provides a partition of $\left(\mathbf{s}^{\text {set }}\right)^{c}$ into equivalence classes, with $\mathbf{v}_{j}=\mathbf{v}_{j^{\prime}}=z_{\alpha}$, for all $j, j^{\prime}$ in the $\alpha$ th equivalence class $A_{\alpha}$. Let $k_{1}, \ldots, k_{d}$ be the cardinalities of the equivalence classes $A_{1}, \ldots, A_{d}$. Without loss of generality, assume that $k_{d}$ is a largest such cardinality. We define

$$
R=\left\{j \in\left(\mathbf{s}^{\text {set }}\right)^{c} \mid \mathbf{v}_{j} \neq \mathbf{v}_{d}\right\}=A_{1} \cup \cdots \cup A_{d-1}
$$

so that $R^{c} \cap\left(\mathbf{s}^{\text {set }}\right)^{c}=A_{d}$. For every $j, j^{\prime} \in A_{d}$, we have $\mathbf{v}_{j}=\mathbf{v}_{j^{\prime}}=\mathbf{v}_{d}$, and therefore the condition (5.1) is satisfied by $R$. Note that by choosing $\mathbf{v}_{d}$ to be the most common value, we are making the cardinality of the set $R^{c} \cap\left(\mathbf{s}^{\text {set }}\right)^{c}=\left\{j \notin \mathbf{s}^{\text {set }} \mid \mathbf{v}_{j}=\mathbf{v}_{d}\right\}$ as large as possible, from which it follows that the set $R \cap\left(\mathbf{s}^{\text {set }}\right)^{c}$ is as small as possible, and therefore $R$, as defined, is indeed a minimal cardinality subset of $\left(\mathbf{s}^{\text {set }}\right)^{c}$ that satisfies (5.1).

We now establish the desired upper bound on the cardinality of $R$. Let $\Sigma_{\mathbf{s}^{\text {set }}}$ be the set of permutations that fix the set $\mathbf{s}^{\text {set }}$. Consider an arbitrary permutation $\sigma \in \Sigma_{\mathbf{s}}$ set and let $\sigma_{M}$ be a corresponding permutation of the memory states, as defined by Assumption 3.1. We let $\mathbf{v}_{\sigma^{-1}}$ be the vector with components $\left(\mathbf{v}_{\sigma^{-1}}\right)_{j}=\mathbf{v}_{\sigma^{-1}(j)}$, for $j \notin \mathbf{s}^{\text {set }}$. Note that as we vary $\sigma$ over the set $\Sigma_{\mathbf{s}^{\text {set }}}, \mathbf{v}_{\sigma^{-1}}$ ranges over all possible permutations of the vector $\mathbf{v}$. We also have, for $j \notin \mathbf{s}^{\text {set }}$,

$$
\begin{aligned}
& \left(\mathbf{v}_{\sigma^{-1}}\right)_{j} \\
& \quad=\mathbf{v}_{\sigma^{-1}(j)} \\
& \quad=\mathbb{P}\left(\mathbf{S}(m)_{|\mathbf{s}|+1}=\sigma^{-1}(j) \mid B(m ; \mathbf{s}, \ell)\right) \\
& \quad=\mathbb{P}\left(\mathbf{S}(m)_{|\mathbf{s}|+1}=\sigma^{-1}(j) \mid\{|\mathbf{S}(m)|=\ell\} \cap \bigcap_{i=1}^{|\mathbf{s}|}\left\{\mathbf{S}(m)_{i}=\mathbf{s}_{i}\right\}\right) \\
& \quad=\mathbb{P}\left(\sigma\left(\mathbf{S}(m)_{|\mathbf{s}|+1}\right)=j \mid\{|\mathbf{S}(m)|=\ell\} \cap \bigcap_{i=1}^{|\mathbf{s}|}\left\{\mathbf{S}(m)_{i}=\mathbf{s}_{i}\right\}\right) \\
& \quad=\mathbb{P}\left(\sigma\left(\mathbf{S}(m)_{|\mathbf{s}|+1}\right)=j \mid\{|\sigma(\mathbf{S}(m))|=\ell\} \cap \bigcap_{i=1}^{|\mathbf{s}|}\left\{\sigma\left(\mathbf{S}(m)_{i}\right)=\sigma\left(\mathbf{s}_{i}\right)\right\}\right) \\
& \quad=\mathbb{P}\left(\sigma\left(\mathbf{S}(m)_{|\mathbf{s}|+1}\right)=j \mid\{|\sigma(\mathbf{S}(m))|=\ell\} \cap \bigcap_{i=1}^{|\mathbf{s}|}\left\{\sigma\left(\mathbf{S}(m)_{i}\right)=\mathbf{s}_{i}\right\}\right) \\
& \quad=\mathbb{P}\left(\mathbf{S}\left(\sigma_{M}(m)\right)_{|\mathbf{s}|+1}=j \mid\left\{\left|\mathbf{S}\left(\sigma_{M}(m)\right)\right|=\ell\right\} \cap \bigcap_{i=1}^{|\mathbf{s}|}\left\{\mathbf{S}\left(\sigma_{M}(m)\right)_{i}=\mathbf{s}_{i}\right\}\right) .
\end{aligned}
$$

Note that in the above expressions, the only random variables are $\mathbf{S}(m)$ and $\mathbf{S}\left(\sigma_{M}(m)\right)$, while $\mathbf{s}$ is a fixed vector. The next to last equality above holds because $\sigma$ fixes the elements in the vector $\mathbf{s}$; the last equality follows because the random variables $\sigma(\mathbf{S}(m))$ and $\mathbf{S}\left(\sigma_{M}(m)\right)$ are
identically distributed, according to Part 1 of the symmetry Assumption 3.1. The equality that was established above implies that $\sigma_{M}(m)$ completely determines the vector $\mathbf{v}_{\sigma^{-1}}$. As $\sigma \in \Sigma_{\mathbf{s}^{\text {set }}}$ changes, $\sigma_{M}(m)$ can take at most $n^{c}$ distinct values, due to the assumed bound on the memory size, and this leads to a bound on the number of possible permutations of the vector $\mathbf{v}$ :

$$
\left|\left\{\mathbf{v}_{\sigma^{-1}}: \sigma \in \Sigma_{\mathbf{s}^{\mathrm{set}}}\right\}\right| \leq n^{c}
$$

We now argue that since $\mathbf{v}$ has relatively few distinct permutations, most of its entries $\mathbf{v}_{i}$ must be equal. Recall the partition of the set $\left(\mathbf{s}^{\text {set }}\right)^{c}$ of indices into equivalence classes, of sizes $k_{1}, \ldots, k_{d}$, with $k_{d}$ being the largest cardinality. Note that there is a one-to-one correspondence between distinct permutations $\mathbf{v}_{\sigma^{-1}}$ of the vector $\mathbf{v}$ and distinct partitions of ( $\left.\mathbf{s}^{\text {set }}\right)^{c}$ into a sequence of subsets of cardinalities $k_{1}, \ldots, k_{d}$, with the value $z_{\alpha}$ being taken on the $\alpha$ th subset. It follows that the number of different partitions of $S^{c}$ into sets with the given cardinalities, which is given by the multinomial coefficient, satisfies

$$
\binom{n-|\mathbf{s}|}{k_{1}!k_{2}!\cdots k_{d}!}=\left|\left\{\mathbf{v}_{\sigma^{-1}}: \sigma \in \Sigma_{\mathbf{s}^{\mathrm{set}}}\right\}\right| \leq n^{c}
$$

The number of choices of a $k_{d}$-element subset is no larger than the number of partitions. Therefore,

$$
\binom{n-|\mathbf{s}|}{k_{d}} \leq n^{c} .
$$

An elementary calculation (cf. Lemma B.1) implies that when $n$ is large enough, we must have either (i) $k_{d} \geq n-|\mathbf{s}|-c$ or (ii) $k_{d} \leq c$. We argue that the second possibility cannot occur. Indeed, if $k_{d} \leq c$, and since $k_{d}$ is the largest cardinality, it follows that $k_{\alpha} \leq c$ for every $\alpha$. Since $k_{1}+\cdots+k_{d}=n-|\mathbf{s}|$, we obtain that the number of classes, $d$, is at least $\lceil(n-|\mathbf{s}|) / c\rceil$. When dealing with $d$ different classes, the number of possible partitions is at least $d!$; this can be seen by focusing on the least-indexed entry in each of the $d$ classes and noting that these $d$ entries may appear in an arbitrary order. Since $|\mathbf{s}| \leq \sqrt{n}$, we have $n-|\mathbf{s}| \geq n / 2$, and putting everything together, we obtain

$$
\lceil(n / 2 c)\rceil!\leq\left\lceil\frac{n-|\mathbf{s}|}{c}\right\rceil!\leq n^{c} .
$$

This is clearly impossible when $n$ is large enough, and case (ii) can therefore be eliminated. We conclude that $\left|A_{d}\right|=k_{d} \geq n-|\mathbf{s}|-c$. Since $\left|A_{1} \cup \cdots \cup A_{d}\right|=\left|\left(\mathbf{s}^{\text {set }}\right)^{c}\right|=n-|\mathbf{s}|$, it follows that $|R|=\left|A_{1} \cup \cdots \cup A_{d-1}\right| \leq c$, which is the desired cardinality bound on $R$.

It should be apparent that any minimal cardinality set $R$ that satisfies (5.1) must be constructed exactly as our set $A_{d}$. Thus, nonuniqueness of the set $R$ with the desired properties will arise if and only if there is another subset $A_{\alpha}$, with $\alpha \neq d$, with the same maximal cardinality $k_{d}$. On the other hand, since $|\mathbf{s}| \leq \sqrt{n}$, we have $k_{d} \geq n-|\mathbf{s}|-c>n / 2$, when $n$ is large enough. But having two disjoint subsets, $A_{d}$ and $A_{\alpha}$, each of cardinality larger than $n / 2$ is impossible, which proves uniqueness.

Using a similar argument, we can also show that the distribution of the destination of the incoming job is uniform (or zero) outside the set of sampled servers and a set of at most $c$ distinguished servers.

PROPOSITION 5.2. Let $V$ be a uniform random variable over $[0,1]$. For all $n$ large enough, for every memory state $m \in \mathcal{M}_{n}$, every vector of indices $\mathbf{s} \in \mathcal{S}_{n}$ with $|\mathbf{s}| \leq \sqrt{n}$, every
queue vector state $\mathbf{q} \in \mathcal{Q}^{|\mathbf{s}|}$, and every job size $w \in \mathbb{R}_{+}$, the following holds. There exists $a$ unique set $R^{\prime}(m, w, \mathbf{s}, \mathbf{q}) \subset \mathcal{N}_{n} \backslash \mathbf{s}^{\text {set }}$ of minimal cardinality such that

$$
\mathbb{P}\left(f_{2}(m, w, \mathbf{s}, \mathbf{q}, V)=j\right)=\mathbb{P}\left(f_{2}(m, w, \mathbf{s}, \mathbf{q}, V)=k\right),
$$

for all $j, k \notin R^{\prime}(m, w, \mathbf{s}, \mathbf{q}) \cup \mathbf{s}^{\text {set }}$. Furthermore, $\left|R^{\prime}(m, w, \mathbf{s}, \mathbf{q})\right| \leq c$.
Proof. The proof is analogous to the proof of the previous proposition. We start by defining a vector $\mathbf{v}$, whose components are again indexed by $j$ ranging in the set $\mathcal{N}_{n} \backslash \mathbf{s}^{\text {set }}$, by

$$
\mathbf{v}_{j}=\mathbb{P}\left(f_{2}(m, w, \mathbf{s}, \mathbf{q}, V)=j\right)
$$

Other than this new definition of the vector $\mathbf{v}$, the rest of the proof follows verbatim the one for Proposition 5.1.
5.2. A sequence of "bad" events. In this subsection we introduce a sequence of "bad" events that we will be focusing on in order to establish a positive lower bound on the delay.

Recall that $T_{1}^{s}$ is the time of the first event of the underlying Poisson process of rate $\mu n$ that generates the spontaneous messages from the servers. Recall also that we denote by $\mathbf{Q}_{i, 1}(t)$ the remaining workload of the job being serviced in server $i$, at time $t$, with $\mathbf{Q}_{i, 1}(t)=0$ if no job is present at server $i$. Let

$$
B \triangleq\left\{i: \sum_{j=1}^{\infty} \mathbf{Q}_{i, j}(0) \geq 2 \gamma\right\}
$$

which is the set of servers with at least $2 \gamma$ remaining workload in their queues, and let $N_{b}=$ $|B|$, where $\gamma \leq 1$ is a small positive constant, independent of $n$, to be specified later.

Consider the following events:
(i) the first $c+1$ jobs after time 0 are all of size at least $2 \gamma$,

$$
\mathcal{A}_{w} \triangleq\left\{W_{1}, \ldots, W_{c+1} \geq 2 \gamma\right\}
$$

(ii) the first potential spontaneous message occurs after time $\gamma / n$, and the $(c+1)$ st arrival occurs before time $\gamma / n$,

$$
\mathcal{A}_{a} \triangleq\left\{T_{1}^{s}>\frac{\gamma}{n}\right\} \cap\left\{T_{c+1}<\frac{\gamma}{n}\right\}
$$

(iii) there are no service completions before time $\gamma / n$,

$$
\mathcal{A}_{s} \triangleq\left\{\mathbf{Q}_{i, 1}(0) \notin\left(0, \frac{\gamma}{n}\right), \forall i\right\}
$$

(iv) there are at least $\gamma n$ servers that each have at least $2 \gamma$ remaining workload at time zero,

$$
\mathcal{A}_{b} \triangleq\left\{N_{b} \geq \gamma n\right\} .
$$

For an interpretation, the event

$$
\mathcal{H}_{0}^{+} \triangleq \mathcal{A}_{w} \cap \mathcal{A}_{a} \cap \mathcal{A}_{s} \cap \mathcal{A}_{b},
$$

corresponds to an unfavorable situation for the dispatcher. This is because, at time zero, the dispatcher's memory contains possibly useful information on at most $c$ distinguished servers (Propositions 5.1 and 5.2), and has to accommodate $c+1$ arriving jobs by time $\gamma / n$. On the other hand, a nontrivial fraction of the servers are busy and will remain so until time $\gamma / n$ (event $\mathcal{A}_{b}$ ), and it is possible that sampling will not reveal any idle servers (as long as
the number of sampled servers is not too large). Thus, at least one of the jobs may end up at a busy server, resulting in positive expected delay. In what follows, we go through the just outlined sequence of unfavorable events, and then, in Section 5.3, we lower bound its probability.

Starting with $\mathcal{H}_{0}^{+}$, we define a nested sequence of events, after first introducing some more notation. For $k=1, \ldots, c+1$, let $\mathbf{S}_{k}$ be the random (hence denoted by an upper case symbol) vector of servers that are sampled upon the arrival of the $k$ th job; its components are denoted by $\left(\mathbf{S}_{k}\right)_{i}$. For $i=0,1, \ldots,\left|\mathbf{S}_{k}\right|$, we let

$$
R_{k, i} \triangleq R\left(M\left(T_{k}^{-}\right), W_{k},\left(\left(\mathbf{S}_{k}\right)_{1}, \ldots,\left(\mathbf{S}_{k}\right)_{i-1}\right),\left|\mathbf{S}_{k}\right|\right)
$$

be the (random) subset of servers defined in Proposition 5.1 (whenever $M\left(T_{k}^{-}\right), W_{k}$, $\left(\left(\mathbf{S}_{k}\right)_{1}, \ldots,\left(\mathbf{S}_{k}\right)_{i-1}\right)$, and $\left|\mathbf{S}_{k}\right|$ are such that the proposition applies), with the convention that $\left(\left(\mathbf{S}_{k}\right)_{1}, \ldots,\left(\mathbf{S}_{k}\right)_{i-1}\right)=\varnothing$ when $i=1$. Otherwise, we let $R_{k, i} \triangleq \varnothing$. Furthermore, we define

$$
R_{k} \triangleq \bigcup_{i=1}^{\left|\mathbf{S}_{k}\right|} R_{k, i}
$$

Moreover, let $D_{k}$ be the destination of the $k$ th job, and let

$$
R_{k}^{\prime} \triangleq R^{\prime}\left(M\left(T_{k}^{-}\right), W_{k}, \mathbf{S}_{k}, \mathbf{Q}_{\mathbf{s}_{k}}\left(T_{k}^{-}\right)\right)
$$

be the (random) subset of servers defined in Proposition 5.2 (whenever $M\left(T_{k}^{-}\right), W_{k}, \mathbf{S}_{k}$ and $\mathbf{Q}_{\mathbf{S}_{k}}\left(T_{k}^{-}\right)$are such that the proposition applies). Otherwise, we let $R_{k}^{\prime} \triangleq \varnothing$. Finally, given a collection of constants $\xi_{1}, \ldots, \xi_{c+1}$, independent of $n$ and to be determined later, we define a nested sequence of events recursively, by

$$
\begin{align*}
& \mathcal{H}_{k}^{-} \triangleq \mathcal{H}_{k-1}^{+} \cap\left\{\left|\mathbf{S}_{k}\right| \leq \xi_{k}\right\} \\
& \mathcal{H}_{k} \triangleq \mathcal{H}_{k}^{-} \cap\left\{\left(\mathbf{S}_{k}\right)_{i} \in R_{k, i} \cup B, i=1, \ldots,\left|\mathbf{S}_{k}\right|\right\}  \tag{5.2}\\
& \mathcal{H}_{k}^{+} \triangleq \mathcal{H}_{k} \cap\left\{D_{k} \in \mathbf{S}_{k}^{\text {set }} \cup R_{k}^{\prime} \cup B\right\}
\end{align*}
$$

for $k=1, \ldots, c+1$.
5.3. Lower bound on the probability of "bad" events. In this subsection, we establish a positive lower bound, valid for all $n$ large enough, for the probability of the event $\mathcal{H}_{c+1}^{+}$. In order to do this, we will obtain such uniform lower bounds for the probability of $\mathcal{H}_{k}^{+}$, for $k \geq 0$, by induction. We start with the base case.

LEmma 5.3. There exists a constant $\alpha_{0}^{+}>0$, independent of $n$, such that

$$
\mathbb{P}\left(\mathcal{H}_{0}^{+}\right) \geq \alpha_{0}^{+}
$$

Proof. Note that the event $\mathcal{A}_{a}$ only depends on the processes of arrivals and spontaneous messages after time zero, $\mathcal{A}_{w}$ only depends on the i.i.d. workloads $W_{1}, \ldots, W_{c+1}$, and $\mathcal{A}_{s} \cap \mathcal{A}_{b}$ only depends on the initial queue length vector $\mathbf{Q}(0)$. It follows that

$$
\mathbb{P}\left(\mathcal{H}_{0}^{+}\right)=\mathbb{P}\left(\mathcal{A}_{a}\right) \mathbb{P}\left(\mathcal{A}_{w}\right) \mathbb{P}\left(\mathcal{A}_{s} \cap \mathcal{A}_{b}\right)
$$

We will now lower bound each of these probabilities.
Note that $\mathbb{P}\left(\mathcal{A}_{a}\right)$ is the intersection of two independent events. The first is the event that the first arrival in a Poisson process with rate $\mu n$ happens after time $\gamma / n$, or equivalently, it is the event that the first arrival of a Poisson process of rate $\mu$ happens after time $\gamma$, which
has positive probability that does not depend on $n$. The second is the event that $c+1$ arrivals of the delayed renewal process $A_{n}(\cdot)$ occur before time $\gamma / n$, that is, the event that $T_{c+1}<\gamma / n$. Since the process $(\mathbf{Q}(\cdot), M(\cdot), Z(\cdot))$ is stationary, the first arrival time $\left(T_{1}\right)$ is distributed according to the residual time of typical inter-arrival times. In particular, if $F$ is the cumulative distribution function of typical inter-arrival times of $A_{n}(t)$ (which have mean $1 / \lambda n)$, the well-known formula for the distribution of residual times gives

$$
\begin{aligned}
\mathbb{P}\left(T_{1}<\frac{\gamma}{n(c+1)}\right) & =\lambda n \int_{0}^{\frac{\gamma}{n(c+1)}}(1-F(u)) d u \\
& =\lambda \int_{0}^{\frac{\gamma}{c+1}}\left(1-F\left(\frac{v}{n}\right)\right) d v .
\end{aligned}
$$

Recall that Assumption 3.2 states that $1-F(v / n) \geq \delta_{v}>0$, for all $v>0$ sufficiently small, and for all $n$. As a result, we have

$$
\begin{align*}
\lambda \int_{0}^{\frac{\gamma}{c+1}}\left(1-F\left(\frac{v}{n}\right)\right) d v & \geq \lambda \int_{0}^{\frac{\gamma}{c+1}}\left(1-F\left(\frac{\gamma}{n(c+1)}\right)\right) d v  \tag{5.3}\\
& =\frac{\lambda \gamma}{c+1} \delta_{\frac{\gamma}{c+1}}
\end{align*}
$$

for all $\gamma$ sufficiently small. On the other hand, for $k=2, \ldots, c+1$, Assumption 3.2 also implies that

$$
\mathbb{P}\left(T_{k}-T_{k-1} \leq \frac{\gamma}{n(c+1)}\right) \geq \delta_{\frac{\gamma}{(c+1)}} .
$$

Combining this with equation (5.3), and using the fact that the first arrival time and the subsequent inter-arrival times are independent, we obtain

$$
\begin{aligned}
\mathbb{P}\left(T_{c+1}<\frac{\gamma}{n}\right) & \geq \mathbb{P}\left(\left\{T_{1}<\frac{\gamma}{n(c+1)}\right\} \cap \bigcap_{k=2}^{c+1}\left\{T_{k}-T_{k-1} \leq \frac{\gamma}{n(c+1)}\right\}\right) \\
& \geq \mathbb{P}\left(T_{1}<\frac{\gamma}{n(c+1)}\right) \prod_{k=2}^{c+1} \mathbb{P}\left(T_{k}-T_{k-1} \leq \frac{\gamma}{n(c+1)}\right) \\
& \geq \frac{\lambda \gamma}{c+1}\left(\delta_{c+1}\right)^{c+1},
\end{aligned}
$$

which is a positive constant independent from $n$.
We also have

$$
\begin{aligned}
\mathbb{P}\left(\mathcal{A}_{w}\right) & =\prod_{i=1}^{c+1} \mathbb{P}\left(W_{i} \geq 2 \gamma\right) \\
& =\mathbb{P}\left(W_{i} \geq 2 \gamma\right)^{c+1}
\end{aligned}
$$

which is independent of $n$, and positive for $\gamma$ small enough.
We now consider the event $\mathcal{A}_{s}$. If $\mathcal{A}_{s}^{c}$ holds, then there exists a server $i$ such that $0<$ $\mathbf{Q}_{i, 1}(0) \leq \gamma / n$, and thus we have a job departure during $\left(0, \frac{\gamma}{n}\right]$. Let $X$ be the number of service completions during ( $0, \frac{\gamma}{n}$ ]. The occurrence of $\mathcal{A}_{s}^{c}$ implies $X \geq 1$. Furthermore, the expected number of service completions in steady-state during any fixed interval must be equal to the expected number of arrivals, so that

$$
\begin{equation*}
\mathbb{P}\left(\mathcal{A}_{s}^{c}\right) \leq \mathbb{E}[X]=(n \lambda) \frac{\gamma}{n}=\lambda \gamma . \tag{5.4}
\end{equation*}
$$

We now consider the event $\mathcal{A}_{b}$. Recall that

$$
N_{b}=\left|\left\{i: \sum_{j=1}^{\infty} \mathbf{Q}_{i, j}(0) \geq 2 \gamma\right\}\right|
$$

Let

$$
N_{I}=\left|\left\{i: \sum_{j=1}^{\infty} \mathbf{Q}_{i, j}(0)=0\right\}\right|
$$

and

$$
N_{d}=\left|\left\{i: 0<\sum_{j=1}^{\infty} \mathbf{Q}_{i, j}(0)<2 \gamma\right\}\right|
$$

Then, $n=N_{b}+N_{I}+N_{d}$. Furthermore, all servers with $0<\sum_{j=1}^{\infty} \mathbf{Q}_{i, j}(0)<2 \gamma$ will have a departure in $(0,2 \gamma)$. Let $Y$ be the number of departures (service completions) during ( $0,2 \gamma$ ). Then, $Y \geq N_{d}$. We use once more that the expected number of service completions in steadystate during any fixed interval must be equal to the expected number of arrivals, to obtain

$$
n \lambda 2 \gamma=\mathbb{E}[Y] \geq \mathbb{E}\left[N_{d}\right]
$$

Furthermore, by applying Little's law to the number of busy servers, in steady-state, we obtain

$$
\mathbb{E}\left[N_{I}\right]=(1-\lambda) n .
$$

Hence

$$
\mathbb{E}\left[N_{b}\right]=n-\mathbb{E}\left[N_{I}\right]-\mathbb{E}\left[N_{d}\right] \geq n(\lambda-2 \lambda \gamma)
$$

On the other hand, we have

$$
\begin{aligned}
\mathbb{E}\left[N_{b}\right] & \leq \mathbb{P}\left(N_{b} \leq \gamma n\right) \gamma n+\mathbb{P}\left(N_{b}>\gamma n\right) n \\
& \leq \gamma n+\mathbb{P}\left(N_{b} \geq \gamma n\right) n \\
& =\gamma n+\mathbb{P}\left(\mathcal{A}_{b}\right) n .
\end{aligned}
$$

Combining these last two inequalities, we obtain

$$
\begin{equation*}
\mathbb{P}\left(\mathcal{A}_{b}\right) \geq \lambda-2 \lambda \gamma-\gamma \tag{5.5}
\end{equation*}
$$

Finally, using equations (5.4) and (5.5), we have

$$
\begin{aligned}
\mathbb{P}\left(\mathcal{A}_{s} \cap \mathcal{A}_{b}\right) & =\mathbb{P}\left(\mathcal{A}_{b}\right)-\mathbb{P}\left(\mathcal{A}_{b} \cap \mathcal{A}_{s}^{c}\right) \\
& \geq \mathbb{P}\left(\mathcal{A}_{b}\right)-\mathbb{P}\left(\mathcal{A}_{s}^{c}\right) \\
& \geq \lambda-2 \lambda \gamma-\gamma-\gamma \lambda,
\end{aligned}
$$

which is a positive constant if $\gamma$ is chosen small enough.
We now carry out the inductive step, from $k-1$ to $k$, in a sequence of three lemmas. We make the induction hypothesis that there exists a positive constant $\alpha_{k-1}^{+}$such that $\mathbb{P}\left(\mathcal{H}_{k-1}^{+}\right) \geq$ $\alpha_{k-1}^{+}$, and we sequentially prove that there exist positive constants $\alpha_{k}^{-}, \alpha_{k}$, and $\alpha_{k}^{+}$such that $\mathbb{P}\left(\mathcal{H}_{k}^{-}\right) \geq \alpha_{k}^{-}$(Lemma 5.4), $\mathbb{P}\left(\mathcal{H}_{k}\right) \geq \alpha_{k}$ (Proposition 5.5), and $\mathbb{P}\left(\mathcal{H}_{k}^{+}\right) \geq \alpha_{k}^{+}$(Lemma 5.7).

LEMmA 5.4. Suppose that $\mathbb{P}\left(\mathcal{H}_{k-1}^{+}\right) \geq \alpha_{k-1}^{+}>0$ and that the constant $\xi_{k}$ is chosen to be large enough. Then, there exists a constant $\alpha_{k}^{-}>0$, such that for all n large enough, we have $\mathbb{P}\left(\mathcal{H}_{k}^{-}\right) \geq \alpha_{k}^{-}$.

Proof. First, recall our assumption that the average message rate (cf. equation (3.1)) is upper bounded by $\alpha n$ in expectation. Therefore,

$$
\mathbb{E}\left[\limsup _{t \rightarrow \infty} \frac{1}{t} \sum_{j=1}^{A_{n}(t)} 2\left|\mathbf{S}_{j}\right|\right] \leq \alpha n
$$

where $A_{n}(t)$ is the number of arrivals until time $t$. By Fatou's lemma, we also have

$$
\limsup _{t \rightarrow \infty} \mathbb{E}\left[\frac{1}{t} \sum_{j=1}^{A_{n}(t)} 2\left|\mathbf{S}_{j}\right|\right] \leq \alpha n
$$

Recall that the process $(\mathbf{Q}(\cdot), M(\cdot), Z(\cdot))$ is stationary. Then, since the sampled vectors are a deterministic function of the state, and i.i.d. randomization variables, the point process of arrivals with the sampled vectors as marks, is also stationary. As a result, the expression

$$
\mathbb{E}\left[\frac{1}{t} \sum_{j=1}^{A_{n}(t)} 2\left|\mathbf{S}_{j}\right|\right]
$$

is independent from $t$ (see equation (1.2.9) of [3]). In particular, for $t=\gamma / n$, we have that

$$
\begin{equation*}
\mathbb{E}\left[\frac{1}{\gamma} \sum_{j=1}^{A_{n}\left(\frac{\gamma}{n}\right)} 2\left|\mathbf{S}_{j}\right|\right] \leq \alpha \tag{5.6}
\end{equation*}
$$

Moreover, since $k \leq c+1$, we have

$$
\begin{aligned}
\mathbb{E}\left[\sum_{j=1}^{A_{n}\left(\frac{\gamma}{n}\right)}\left|\mathbf{S}_{j}\right|\right] & \geq \mathbb{E}\left[\left.\sum_{j=1}^{A_{n}\left(\frac{\gamma}{n}\right)}\left|\mathbf{S}_{j}\right| \right\rvert\, A_{n}\left(\frac{\gamma}{n}\right) \geq c+1\right] \mathbb{P}\left(A_{n}\left(\frac{\gamma}{n}\right) \geq c+1\right) \\
& \geq \mathbb{E}\left[\left|\mathbf{S}_{k}\right| \left\lvert\, A_{n}\left(\frac{\gamma}{n}\right) \geq c+1\right.\right] \mathbb{P}\left(A_{n}\left(\frac{\gamma}{n}\right) \geq c+1\right)
\end{aligned}
$$

Combining this with equation (5.6), we obtain

$$
\mathbb{E}\left[\left.\frac{2}{\gamma}\left|\mathbf{S}_{k}\right| \right\rvert\, A_{n}\left(\frac{\gamma}{n}\right) \geq c+1\right] \mathbb{P}\left(A_{n}\left(\frac{\gamma}{n}\right) \geq c+1\right) \leq \alpha
$$

This yields the upper bound

$$
\begin{equation*}
\mathbb{E}\left[\left|\mathbf{S}_{k}\right| \left\lvert\, A_{n}\left(\frac{\gamma}{n}\right) \geq c+1\right.\right] \leq \frac{\alpha \gamma}{2 \mathbb{P}\left(A_{n}\left(\frac{\gamma}{n}\right) \geq c+1\right)} \tag{5.7}
\end{equation*}
$$

On the other hand, using the fact that $\mathcal{H}_{k-1}^{+} \subset\left\{A_{n}(\gamma / n) \geq c+1\right\}$, we have

$$
\begin{align*}
& \mathbb{P}\left(\mathcal{H}_{k}^{-}\right) \\
& \quad=\mathbb{P}\left(\mathcal{H}_{k-1}^{+} \cap\left\{\left|\mathbf{S}_{k}\right| \leq \xi_{k}\right\}\right) \\
& \quad=\mathbb{P}\left(\mathcal{H}_{k-1}^{+} \cap\left\{A_{n}\left(\frac{\gamma}{n}\right) \geq c+1\right\} \cap\left\{\left|\mathbf{S}_{k}\right| \leq \xi_{k}\right\}\right)  \tag{5.8}\\
& \quad=\mathbb{P}\left(\mathcal{H}_{k-1}^{+} \cap\left\{\left|\mathbf{S}_{k}\right| \leq \xi_{k}\right\} \left\lvert\, A_{n}\left(\frac{\gamma}{n}\right) \geq c+1\right.\right) \mathbb{P}\left(A_{n}\left(\frac{\gamma}{n}\right) \geq c+1\right) \\
& \quad \geq \mathbb{P}\left(\mathcal{H}_{k-1}^{+}\right)-\mathbb{P}\left(\left|\mathbf{S}_{k}\right|>\xi_{k} \left\lvert\, A_{n}\left(\frac{\gamma}{n}\right) \geq c+1\right.\right) \mathbb{P}\left(A_{n}\left(\frac{\gamma}{n}\right) \geq c+1\right) .
\end{align*}
$$

Furthermore, for any constant $\xi_{k}>0$, Markov's inequality implies

$$
\mathbb{P}\left(\left|\mathbf{S}_{k}\right|>\xi_{k} \left\lvert\, A_{n}\left(\frac{\gamma}{n}\right) \geq c+1\right.\right) \leq \frac{\mathbb{E}\left[\left|\mathbf{S}_{k}\right| \left\lvert\, A_{n}\left(\frac{\gamma}{n}\right) \geq c+1\right.\right]}{\xi_{k}}
$$

which combined with equation (5.8) yields

$$
\mathbb{P}\left(\mathcal{H}_{k}^{-}\right) \geq \mathbb{P}\left(\mathcal{H}_{k-1}^{+}\right)-\frac{\mathbb{E}\left[\left|\mathbf{S}_{k}\right| \left\lvert\, A_{n}\left(\frac{\gamma}{n}\right) \geq c+1\right.\right]}{\xi_{k}} \mathbb{P}\left(A_{n}\left(\frac{\gamma}{n}\right) \geq c+1\right)
$$

Applying the inequality (5.7) to the equation above, we obtain

$$
\mathbb{P}\left(\mathcal{H}_{k}^{-}\right) \geq \mathbb{P}\left(\mathcal{H}_{k-1}^{+}\right)-\frac{\alpha \gamma}{2 \xi_{k}}
$$

Finally, combining this with the fact that $\mathbb{P}\left(\mathcal{H}_{k-1}^{+}\right) \geq \alpha_{k-1}^{+}>0$, we have that

$$
\mathbb{P}\left(\mathcal{H}_{k}^{-}\right) \geq \alpha_{k-1}^{+}-\frac{\alpha \gamma}{2 \xi_{k}} \triangleq \alpha_{k}^{-}
$$

which is positive for all $\xi_{k}$ large enough.
Proposition 5.5. Suppose that $\mathbb{P}\left(\mathcal{H}_{k}^{-}\right) \geq \alpha_{k}^{-}$, and that the constant $\xi_{k}$ is chosen large enough. Then, there exists a constant $\alpha_{k}>0$, such that for all $n$ large enough, we have $\mathbb{P}\left(\mathcal{H}_{k}\right) \geq \alpha_{k}$.

Proof. Recall the definitions

$$
\mathcal{H}_{k}=\mathcal{H}_{k}^{-} \cap\left\{\left(\mathbf{S}_{k}\right)_{i} \in R_{k, i} \cup B, i=1, \ldots,\left|\mathbf{S}_{k}\right|\right\}
$$

and

$$
\mathcal{H}_{k}^{-}=\mathcal{H}_{k-1}^{+} \cap\left\{\left|\mathbf{S}_{k}\right| \leq \xi_{k}\right\} .
$$

For $i=1, \ldots,\left|\mathbf{S}_{k}\right|$, let us denote

$$
H_{k, i} \triangleq\left\{\left(\mathbf{S}_{k}\right)_{i} \in R_{k, i} \cup B\right\} .
$$

Then,

$$
\begin{align*}
& \mathbb{P}\left(\mathcal{H}_{k}\right) \\
& \quad=\mathbb{P}\left(\mathcal{H}_{k}^{-} \cap\left\{\left(\mathbf{S}_{k}\right)_{i} \in R_{k, i} \cup B, i=1, \ldots,\left|\mathbf{S}_{k}\right|\right\}\right) \\
& \quad=\sum_{\ell} \mathbb{P}\left(\mathcal{H}_{k-1}^{+} \cap\left\{\left|\mathbf{S}_{k}\right|=\ell\right\} \cap \bigcap_{i=1}^{\ell} H_{k, i}\right)  \tag{5.9}\\
& \quad=\sum_{\ell} \mathbb{P}\left(\bigcap_{i=1}^{\ell} H_{k, i} \mid \mathcal{H}_{k-1}^{+} \cap\left\{\left|\mathbf{S}_{k}\right|=\ell\right\}\right) \mathbb{P}\left(\mathcal{H}_{k-1}^{+} \cap\left\{\left|\mathbf{S}_{k}\right|=\ell\right\}\right) \\
& \quad=\sum_{\ell} \mathbb{P}\left(\mathcal{H}_{k-1}^{+} \cap\left\{\left|\mathbf{S}_{k}\right|=\ell\right\}\right) \prod_{i=1}^{\ell} \mathbb{P}\left(H_{k, i} \mid \mathcal{H}_{k-1}^{+} \cap\left\{\left|\mathbf{S}_{k}\right|=\ell\right\} \cap \bigcap_{j=1}^{i-1} H_{k, j}\right),
\end{align*}
$$

where the sum is over all integers $\ell$ such that the conditional probabilities above are well defined. Intuitively, in the last step, we are treating the selection of the random vector $\mathbf{S}_{k}$ as a sequential selection of its components, which leads us to consider the product of suitable conditional probabilities. The next lemma provides a lower bound for the factors in this product.

LEMMA 5.6. For all $n$ large enough, we have

$$
\mathbb{P}\left(H_{k, i} \mid \mathcal{H}_{k-1}^{+} \cap\left\{\left|\mathbf{S}_{k}\right|=\ell\right\} \cap \bigcap_{j=1}^{i-1} H_{k, j}\right) \geq \frac{\gamma}{2}
$$

for all $\ell \leq \xi_{k}$ and $i \leq \ell$ such that the conditional probability above is well defined.

The idea of the proof of this lemma is that when a next component, $\left(\mathbf{S}_{k}\right)_{i}$ is chosen, it is either a "distinguished" server, in the set $R_{k, i}$, or else it is a server chosen uniformly outside the set $R_{k, i}$ (cf. Proposition 5.1), in which case it has a substantial probability of being a busy server, in the set $B$. Although the intuition is clear, the formal argument is rather tedious and is deferred to Appendix C.

Applying Lemma 5.6 to equation (5.9), and using the fact that $\mathbb{P}\left(\mathcal{H}_{k}^{-}\right) \geq \alpha_{k}^{-}>0$, we obtain

$$
\begin{aligned}
\mathbb{P}\left(\mathcal{H}_{k}\right) & \geq \sum_{\ell} \mathbb{P}\left(\mathcal{H}_{k-1}^{+} \cap\left\{\left|\mathbf{S}_{k}\right|=\ell\right\}\right)\left(\frac{\gamma}{2}\right)^{\ell} \\
& \geq \mathbb{P}\left(\mathcal{H}_{k-1}^{+} \cap\left\{\left|\mathbf{S}_{k}\right| \leq \xi_{k}\right\}\right)\left(\frac{\gamma}{2}\right)^{\xi_{k}} \\
& =\mathbb{P}\left(\mathcal{H}_{k}^{-}\right)\left(\frac{\gamma}{2}\right)^{\xi_{k}} \\
& \geq \alpha_{k}^{-}\left(\frac{\gamma}{2}\right)^{\xi_{k}} \triangleq \alpha_{k}>0,
\end{aligned}
$$

for all $n$ large enough.
LEMMA 5.7. Suppose that $\mathbb{P}\left(\mathcal{H}_{k}\right) \geq \alpha_{k}$. Then, there exist a constant $\alpha_{k}^{+}>0$, such that for all $n$ large enough, we have $\mathbb{P}\left(\mathcal{H}_{k}^{+}\right) \geq \alpha_{k}^{+}$.

The proof is similar to the proof of Proposition 5.5 but with $\xi_{k}=1$, and it is omitted. Intuitively, choosing the destination of a job has the same statistical properties as choosing one more server to sample, which brings us back to the setting of Proposition 5.5.

This concludes the induction step. It follows that there exists a constant $\alpha_{c+1}^{+}>0$, which is independent of $n$, and such that $\mathbb{P}\left(\mathcal{H}_{c+1}^{+}\right) \geq \alpha_{c+1}^{+}$.
5.4. Upper bound on the number of useful distinguished servers. Let us provide some intuition on what comes next. The dispatcher initially may treat in a nontypical manner the servers in an initial set of at most $c$ distinguished servers. As servers get sampled, the dispatcher acquires and possibly stores information about other servers. Ultimately, at the time of the $(c+1)$ st arrival, the dispatcher may have acquired information and therefore treat in a special manner (i.e., asymmetrically) the servers in the set

$$
\bar{R} \triangleq \bigcup_{k=1}^{c+1}\left(R_{k} \cup R_{k}^{\prime}\right)
$$

Recall that, for $k=1, \ldots, c+1$, we have

$$
R_{k}=\bigcup_{i=1}^{\left|\mathbf{S}_{k}\right|} R_{k, i}
$$

where each of the sets in the union has cardinality at most $c$, by Proposition 5.1. Furthermore, for $k=1, \ldots, c+1$, the cardinality of $R_{k}^{\prime}$ is also at most $c$, by Proposition 5.2. It follows that

$$
\begin{equation*}
|\bar{R}| \leq c \sum_{k=1}^{c+1}\left(1+\left|\mathbf{S}_{k}\right|\right) \tag{5.10}
\end{equation*}
$$

If we are to rely solely on this upper bound, the size of $\bar{R}$ can be larger than $c+1$, and it is possible in principle that the knowledge of so many "distinguished" servers (in the set $\bar{R}$ ) is enough for the dispatcher to identify $c+1$ idle servers to which to route the first $c+1$ jobs. On the other hand, under the event $\mathcal{H}_{c+1}^{+}$, all new information comes from servers that are "busy" (in the set $B$ ), and hence cannot be useful for the dispatching decisions. The next proposition states that for every sample path $\omega \in \mathcal{H}_{c+1}^{+}$, the set of idle (and therefore, potentially useful) servers on which information is available, namely, the set $\bar{R} \backslash B$, has cardinality of at most $c$.

Proposition 5.8. The event $\mathcal{H}_{c+1}^{+}$implies the event $|\bar{R} \backslash B| \leq c$.
Proof. Let us fix a realization $\omega \in \mathcal{H}_{c+1}^{+}$. We will upper bound the number of distinct images of the set $\bar{R} \backslash B$ under permutations of the set $\mathcal{N}_{n}$ of servers, which will lead to an upper bound on the cardinality of the set itself. In order to simplify notation, we will suppress the notational dependence on $\omega$ of all random variables for the rest of this proof.

We introduce a subset of the set of all possible permutations of $\mathcal{N}_{n}$, with this subset being rich enough to lead to the desired bound. Toward this goal, we define the set

$$
F \triangleq \bigcup_{k=1}^{c+1}\left(\bigcup_{i=1}^{\left|\mathbf{S}_{k}\right|}\left[\left\{\left(\mathbf{S}_{k}\right)_{i}\right\} \backslash R_{k, i}\right] \cup\left[\left\{D_{k}\right\} \backslash\left(R_{k}^{\prime} \cup \mathbf{S}_{k}^{\mathrm{set}}\right)\right]\right)
$$

This is the set of servers that were sampled, or that were chosen as the destination for a job, which were not in the distinguished sets $R_{k, i}$, or $R_{k}^{\prime} \cup \mathbf{S}_{k}^{\text {set }}$, respectively.

Using our assumption $\omega \in \mathcal{H}_{c+1}^{+}$and the definition of $\mathcal{H}_{c+1}^{+}$, we have

$$
\bigcup_{k=1}^{c+1} \bigcup_{i=1}^{\left|\mathbf{S}_{k}\right|}\left\{\left(\mathbf{S}_{k}\right)_{i}\right\} \backslash R_{k, i} \subset B \quad \text { and } \quad \bigcup_{k=1}^{c+1}\left\{D_{k}\right\} \backslash\left(R_{k}^{\prime} \cup \mathbf{S}_{k}^{\text {set }}\right) \subset B .
$$

As a result, we have $F \subset B$, and thus

$$
(\bar{R} \backslash B) \cap F=\varnothing
$$

Let $\Sigma$ be the set of permutations $\sigma$ of the server set $\mathcal{N}_{n}$ that:
(i) preserve the ordering of $\bar{R} \backslash B$ in the sense defined in Section 2,
(ii) fix the set $(\bar{R} \cap B) \cup F$ and
(iii) satisfy $\sigma(\bar{R} \backslash B) \cap(\bar{R} \backslash B)=\varnothing$.

Consider two permutations $\sigma, \tau \in \Sigma$ such that $\sigma(\bar{R} \backslash B)=\tau(\bar{R} \backslash B)$. Then, the fact that $\sigma$ and $\tau$ both preserve the order of $\bar{R} \backslash B$ implies that $\sigma(i)=\tau(i)$, for all $i \in \bar{R} \backslash B$.

Lemma 5.9. Let $\sigma, \tau \in \Sigma$, and let $\sigma_{M}$ and $\tau_{M}$, respectively, be associated permutations of the memory states as specified in Assumption 3.1 (Symmetry). Let $m(0)$ be the initial memory state, at time 0. If $\sigma_{M}(m(0))=\tau_{M}(m(0))$, then $\sigma(\bar{R})=\tau(\bar{R})$.

Loosely speaking, Lemma 5.9 asserts that for the given sample path, permutations $\sigma, \tau$ in $\Sigma$ that lead to different sets $\bar{R}$ of distinguished servers must also lead (through $\sigma_{M}$ and $\tau_{M}$ )
to different initial memory states. The proof is an elementary consequence of our symmetry assumption on the underlying dynamics. However, it is tedious and is deferred to Appendix D.

By Lemma 5.9, and for $\sigma \in \Sigma$, distinct images $\sigma(\bar{R})$ must correspond to distinct memory states $\sigma_{M}(m(0))$. Since the number of different memory states is upper bounded by $n^{c}$, this implies that

$$
|\{\sigma(\bar{R}): \sigma \in \Sigma\}| \leq n^{c}
$$

Furthermore, since every $\sigma \in \Sigma$ fixes the set $\bar{R} \cap B$, we have

$$
\begin{equation*}
|\{\sigma(\bar{R} \backslash B): \sigma \in \Sigma\}|=|\{\sigma(\bar{R}): \sigma \in \Sigma\}| \leq n^{c} \tag{5.11}
\end{equation*}
$$

Recall now that the only restrictions on the image $\sigma(\bar{R} \backslash B)$ under permutations in $\sigma \in \Sigma$ is that the set $(\bar{R} \cap B) \cup F$ is fixed, and that $\sigma(\bar{R} \backslash B) \cap(\bar{R} \backslash B)=\varnothing$. This implies that $\sigma(\bar{R} \backslash B)$ can be any set of the same cardinality within $(\bar{R} \cup F)^{c}$. It follows that

$$
\begin{equation*}
|\{\sigma(\bar{R} \backslash B): \sigma \in \Sigma\}| \geq\binom{ n-|\bar{R} \cup F|}{|\bar{R} \backslash B|} . \tag{5.12}
\end{equation*}
$$

Recall also that under the event $\mathcal{H}_{c+1}^{+}$we must have $\left|\mathbf{S}_{k}\right| \leq \xi_{k}$, for $k=1, \ldots, c+1$. Thus $|F| \leq \xi_{1}+\cdots+\xi_{c+1}+c+1 \triangleq f$, and using equation (5.10), $|\bar{R}| \leq c\left(\xi_{1}+\cdots+\xi_{c+1}\right)+c+$ $1 \triangleq \theta$. Combining these two upper bounds, we obtain

$$
\binom{n-|\bar{R} \cup F|}{|\bar{R} \backslash B|} \geq\binom{ n-(f+\theta)}{|\bar{R} \backslash B|} .
$$

Combining this with equations (5.11) and (5.12), we obtain the inequality

$$
n^{c} \geq\binom{ n-(f+\theta)}{|\bar{R} \backslash B|} .
$$

Finally, using the bound $|\bar{R}| \leq \theta$, and applying Lemma B.1, we conclude that in order for this equation to hold for all $n$ large enough, we must have $|\bar{R} \backslash B| \leq c$.
5.5. Completing the proof. We are now ready to complete the proof, by arguing that at least one of the first $c+1$ arrivals must be sent to a server that is either known to be busy or to a server on which no information is available, and therefore has positive probability of being busy.

Recall that for any fixed sample path in $\mathcal{H}_{c+1}^{+}$, we have (cf. equation (5.2))

$$
\left\{D_{1}, \ldots, D_{c+1}\right\} \subset B \cup \bigcup_{k=1}^{c+1}\left(\mathbf{S}_{k}^{\text {set }} \cup R_{k}^{\prime}\right)
$$

Furthermore the event $\mathcal{H}_{c+1}^{+}$implies that $\left(\mathbf{S}_{k}\right)_{i} \in R_{k, i} \cup B$, for $i=1, \ldots,\left|\mathbf{S}_{k}\right|$ and $k=$ $1, \ldots, c+1$. Therefore,

$$
\mathbf{S}_{k}^{\mathrm{set}} \subset \bigcup_{i=1}^{\left|\mathbf{S}_{k}\right|} R_{k, i} \cup B=R_{k} \cup B
$$

for $k=1, \ldots, c+1$. It follows that

$$
\begin{aligned}
\left\{D_{1}, \ldots, D_{c+1}\right\} & \subset B \cup \bigcup_{k=1}^{c+1}\left(\mathbf{S}_{k}^{\text {set }} \cup R_{k}^{\prime}\right) \\
& \subset B \cup \bigcup_{k=1}^{c+1}\left(R_{k} \cup R_{k}^{\prime}\right) \\
& =B \cup \bar{R} .
\end{aligned}
$$

Moreover, Proposition 5.8 states that $|\bar{R} \backslash B| \leq c$. Thus, either (a) there exists $k$ such that $D_{k} \in B$, or (b) $D_{i} \in \bar{R} \backslash B$ for $i=1, \ldots, c+1$, and hence there exists a pair $k, l$, with $k<l$, such that $D_{k}=D_{l}$. We will now show that in both cases, the queueing delay is at least $\gamma$.

Let $L_{k}$ be the queueing delay of the $k$ th arrival. Recall that for $i \in B$, we have $\mathbf{Q}_{i, 1}(0)>$ $2 \gamma$. Then, for case (a), with $D_{k}=i \in B$ we have

$$
L_{k}=\left(\mathbf{Q}_{i, 1}(0)-T_{k}\right)^{+} \geq 2 \gamma-\frac{\gamma}{n} \geq \gamma>0 .
$$

On the other hand, for case (b), we have

$$
L_{l} \geq\left[W_{k}-\left(T_{l}-T_{k}\right)\right]^{+} \geq 2 \gamma-\left(\frac{\gamma}{n}-0\right) \geq \gamma>0
$$

In both cases, we have

$$
L_{1}+\cdots+L_{c+1} \geq \gamma
$$

Since this is true for every sample path in $\mathcal{H}_{c+1}^{+}$, we obtain

$$
\begin{equation*}
\mathbb{E}\left[L_{1}+\cdots+L_{c+1} \mid \mathcal{H}_{c+1}^{+}\right] \geq \gamma \tag{5.13}
\end{equation*}
$$

Finally, recall that the process $(\mathbf{Q}(\cdot), M(\cdot), Z(\cdot))$ is stationary, with invariant probability measure $\pi_{n}$. Then, setting $t=\gamma / n$ in equation (3.2), we obtain

$$
\begin{aligned}
\mathbb{E}_{\pi_{n}}^{0}\left[L_{0}\right] & =\frac{1}{\lambda \gamma} \mathbb{E}\left[\sum_{j=1}^{A_{n}\left(\frac{\gamma}{n}\right)} L_{j}\right] \\
& \geq \frac{1}{\lambda \gamma} \mathbb{E}\left[\left.\sum_{j=1}^{A_{n}\left(\frac{\gamma}{n}\right)} L_{j} \right\rvert\, \mathcal{H}_{c+1}^{+}\right] \mathbb{P}\left(\mathcal{H}_{c+1}^{+}\right) \\
& \geq \frac{1}{\lambda \gamma} \mathbb{E}\left[L_{1}+\cdots+L_{c+1} \mid \mathcal{H}_{c+1}^{+}\right] \mathbb{P}\left(\mathcal{H}_{c+1}^{+}\right)
\end{aligned}
$$

where the last inequality comes from the fact that $\mathcal{H}_{c+1}^{+} \subset\left\{A_{n}(\gamma / n) \geq c+1\right\}$. Combining this with equation (5.13) and the fact that $\mathbb{P}\left(\mathcal{H}_{c+1}^{+}\right) \geq \alpha_{c+1}^{+}>0$, we obtain

$$
\mathbb{E}_{\pi_{n}}^{0}\left[L_{0}\right] \geq \frac{\alpha_{c+1}^{+}}{\lambda}>0
$$

As the constant in the lower bound does not depend on $n$, this completes the proof of the theorem.
6. Conclusions and future work. We showed that when we have a limited amount of memory and a modest budget of messages per unit of time, and under a symmetry assumption, all dispatching policies result in queueing delay that is uniformly bounded away from zero. In particular, this implies that the queueing delay does not vanish as the system size increases.

Our result complements the results in [7], in which the authors showed that if we have a little more of either resource, that is, if the number of memory bits or the message rate grows faster with $n$, then there exists a symmetric policy that drives the queueing delay to zero as $n \rightarrow \infty$. Consequently, we now have necessary and sufficient conditions on the amount of resources available to a central dispatcher, in order to achieve a vanishing queueing delay as the system size increases.

There are several interesting directions for future research. For example:
(i) All the policies in the literature that achieve a vanishing queueing delay need a message rate at least equal to the arrival rate $\lambda n$. We conjecture that this is not a necessary condition for a policy to have a vanishing queueing delay, as long as it has access to the incoming job sizes.
(ii) In light of the symmetry assumption in Theorem 3.1, an immediate open question is whether the result still holds without this assumption. Our proof relies heavily on symmetry and is hard to generalize. However, perhaps (nonsymmetric) policies that use the memory to store the beginning and the end of streaks of idle servers could achieve a vanishing queueing delay in the low memory and low message rate regime where symmetric policies cannot do it.
(iii) We have focused on a system with homogeneous servers. For the case of nonhomogeneous servers, even stability can become an issue, and there are interesting tradeoffs between the resources used and the stability region. In this setting, we expect a result similar to our lower bound for queueing delay, stating that a resource constrained policy cannot be stable for every stabilizable system.

## APPENDIX A: COMPARISON WITH A MORE RESTRICTIVE SYMMETRY ASSUMPTION

In this appendix we explain why the stronger symmetry assumption

$$
\begin{equation*}
\sigma\left(f_{1}(m, w, u)\right)=f_{1}\left(\sigma_{M}(m), w, u\right) \quad \forall u \in[0,1] \tag{A.1}
\end{equation*}
$$

would be unduly restrictive.
Consider a policy that samples a fixed number $d$ of servers, uniformly at random (regardless of the memory state and of the incoming job size), and that satisfies this stronger symmetry assumption. Then, $f_{1}(m, w, u)$ is a vector of dimension $d$, for all $m \in \mathcal{M}_{n}, w \in \mathbb{R}_{+}$, and $u \in[0,1]$. Let $\sigma, \tau$ be a pair of permutations such that $\sigma\left(f_{1}(m, w, u)\right) \neq \tau\left(f_{1}(m, w, u)\right)$. The stronger symmetry assumption in equation (A.1) implies that there exists a pair of associated permutations $\sigma_{M}, \tau_{M}$ of the memory states such that

$$
f_{1}\left(\sigma_{M}(m), w, u\right)=\sigma\left(f_{1}(m, w, u)\right) \neq \tau\left(f_{1}(m, w, u)\right)=f_{1}\left(\tau_{M}(m), w, u\right)
$$

It follows that $\sigma_{M}(m) \neq \tau_{M}(m)$, and thus there must be at least as many memory states as the number of different vectors of dimension $d$ with different entries. There are $\binom{n}{d} d$ ! such vectors, and therefore a large memory would be required to implement such a uniform sampling policy if equation (A.1) were to be enforced.

On the other hand, the symmetry assumption that we have adopted in this paper only requires equality in distribution, and uniform sampling can be achieved with only one memory state (i.e., with no bits of memory). Indeed, since the sampling of servers is done uniformly at random, we have

$$
f_{1}(m, w, U) \stackrel{d}{=} \sigma\left(f_{1}(m, w, U)\right)
$$

for all permutations $\sigma$.
This example shows that the symmetry assumption that we have adopted can be substantially weaker (and thus easier to satisfy), and allows small-memory implementation of simple natural policies.

## APPENDIX B: A COMBINATORIAL INEQUALITY

We record here an elementary fact.
Lemma B.1. Let us fix positive integer constants $a$ and $c$. Suppose that $b$ satisfies

$$
\begin{equation*}
\binom{n-a}{b} \leq n^{c} \tag{B.1}
\end{equation*}
$$

As long as $n$ is large enough, we must have $b \leq c$ or $b \geq n-a-c$.
Proof. Suppose that $b=c+1$. The quantity $\binom{n-a}{c+1}$ is a polynomial in $n$ of degree $c+1$ and therefore, when $n$ is large, (B.1) cannot hold. In the range $c+1 \leq b \leq(n-a) / 2$, the quantity $\binom{n-a}{b}$ increases with $b$, and hence (B.1) cannot hold either. Using the symmetry of the binomial coefficient, a similar argument is used to exclude the possibility that $(n-a) / 2 \leq$ $b \leq n-a-c-1$.

## APPENDIX C: PROOF OF LEMMA 5.6

In order to simplify notation, we introduce the following. For any $m \in \mathcal{M}_{n}, w \in \mathbb{R}_{+}$, and $b \in \mathcal{P}\left(\mathcal{N}_{n}\right)$, we define the event

$$
\mathcal{A}_{m, w, b} \triangleq\left\{M\left(T_{k}^{-}\right)=m, B=b, W_{k}=w\right\}
$$

and we let $\mathbb{P}_{m, w, b}$ be the conditional probability measure

$$
\mathbb{P}_{m, w, b}(\cdot) \triangleq \mathbb{P}\left(\cdot \mid \mathcal{A}_{m, w, b}\right)
$$

Let us fix some $\ell \leq \xi_{k}$ and some $i \leq \ell$. We have

$$
\begin{aligned}
& \mathbb{P}\left(H_{k, i} \mid \mathcal{H}_{k-1}^{+} \cap\left\{\left|\mathbf{S}_{k}\right|=\ell\right\} \cap \bigcap_{j=1}^{i-1} H_{k, j}\right) \\
& =\int_{m, w, b} \mathbb{P}_{m, w, b}\left(H_{k, i} \mid \mathcal{H}_{k-1}^{+} \cap\left\{\left|\mathbf{S}_{k}\right|=\ell\right\} \cap \bigcap_{j=1}^{i-1} H_{k, j}\right) \\
& \quad \cdot d \mathbb{P}\left(\mathcal{A}_{m, w, b} \mid \mathcal{H}_{k-1}^{+} \cap\left\{\left|\mathbf{S}_{k}\right|=\ell\right\} \cap \bigcap_{j=1}^{i-1} H_{k, j}\right) .
\end{aligned}
$$

Moreover,

$$
\begin{aligned}
& \mathbb{P}_{m, w, b}\left(H_{k, i} \mid \mathcal{H}_{k-1}^{+} \cap\left\{\left|\mathbf{S}_{k}\right|=\ell\right\} \cap \bigcap_{j=1}^{i-1} H_{k, j}\right) \\
& =\sum_{\mathbf{s}} \mathbb{P}_{m, w, b}\left(H_{k, i} \mid \mathcal{H}_{k-1}^{+} \cap\left\{\left|\mathbf{S}_{k}\right|=\ell\right\} \cap \bigcap_{j=1}^{i-1}\left\{\left(\mathbf{S}_{k}\right)_{j}=\mathbf{s}_{j}\right\}\right) \\
& \quad \cdot \mathbb{P}_{m, w, b}\left(\bigcap_{j=1}^{i-1}\left\{\left(\mathbf{S}_{k}\right)_{j}=\mathbf{s}_{j}\right\} \mid \mathcal{H}_{k-1}^{+} \cap\left\{\left|\mathbf{S}_{k}\right|=\ell\right\} \cap \bigcap_{j=1}^{i-1} H_{k, j}\right),
\end{aligned}
$$

where the sum is over all $(i-1)$-dimensional vectors $\mathbf{s}$ whose components are distinct indices of servers, and such that the conditional probabilities above are well defined.

It is not hard to see that the desired result follows immediately once we establish the following claim.

Claim C.1. For all $n$ large enough, we have

$$
\mathbb{P}_{m, w, b}\left(H_{k, i} \mid \mathcal{H}_{k-1}^{+} \cap\left\{\left|\mathbf{S}_{k}\right|=\ell\right\} \cap \bigcap_{j=1}^{i-1}\left\{\left(\mathbf{S}_{k}\right)_{j}=\mathbf{s}_{j}\right\}\right) \geq \frac{\gamma}{2}
$$

for all $(m, w, b, \mathbf{s})$ such that the conditional probability above is well defined.
Proof. Let us fix some $(m, w, b, \mathbf{s})$. Since $\mathcal{H}_{k-1}^{+}$implies $|B| \geq \gamma n$, we have

$$
\begin{equation*}
|b| \geq \gamma n . \tag{C.1}
\end{equation*}
$$

On the other hand, recall that

$$
H_{k, i}=\left\{\left(\mathbf{S}_{k}\right)_{i} \in R_{k, i} \cup B\right\},
$$

where

$$
\mathbf{S}_{k}=f_{1}\left(M\left(T_{k}^{-}\right), W_{k}, U_{k}\right)
$$

and $R_{k, i}$ is equal to the set

$$
R\left(M\left(T_{k}^{-}\right), W_{k},\left(\left(\mathbf{S}_{k}\right)_{1}, \ldots,\left(\mathbf{S}_{k}\right)_{i-1}\right),\left|\mathbf{S}_{k}\right|\right)
$$

defined in Proposition 5.1, whenever the proposition applies. Otherwise, we have $R_{k, j}=\varnothing$. In any case, $R_{k, j}$ is a deterministic function of the same random variables. Then, conditioned on $M\left(T_{k}^{-}\right)=m, W_{k}=w, B=b,\left(\left(\mathbf{S}_{k}\right)_{1}, \ldots,\left(\mathbf{S}_{k}\right)_{j-1}\right)=\mathbf{s}$, and $\left|\mathbf{S}_{k}\right|=\ell$, we have

$$
H_{k, i}=\left\{\left(f_{1}\left(m, w, U_{k}\right)\right)_{i} \in r_{k, i} \cup b\right\}
$$

where $r_{k, i}$ denotes the corresponding realization of the random set $R_{k, i}$. Note that the only randomness left in this event comes from $U_{k}$, which is a randomization random variable that is chosen independent from all the events prior to time $T_{k}^{-}$. It follows that $H_{k, i}$ is conditionally independent from $\mathcal{H}_{k-1}^{+}$, and thus

$$
\begin{aligned}
& \mathbb{P}_{m, w, b}\left(H_{k, i} \mid \mathcal{H}_{k-1}^{+} \cap\left\{\left|\mathbf{S}_{k}\right|=\ell\right\} \cap \bigcap_{j=1}^{i-1}\left\{\left(\mathbf{S}_{k}\right)_{j}=\mathbf{s}_{j}\right\}\right) \\
& \quad=\mathbb{P}_{m, w, b}\left(H_{k, i} \mid\left\{\left|\mathbf{S}_{k}\right|=\ell\right\} \cap \bigcap_{j=1}^{i-1}\left\{\left(\mathbf{S}_{k}\right)_{j}=\mathbf{s}_{j}\right\}\right) .
\end{aligned}
$$

We now define the event $G_{k, \mathbf{s}, i, \ell}$ to be

$$
G_{k, \mathbf{s}, i, \ell} \triangleq\left\{\left|\mathbf{S}_{k}\right|=\ell\right\} \cap \bigcap_{j=1}^{i-1}\left\{\left(\mathbf{S}_{k}\right)_{j}=\mathbf{s}_{j}\right\} .
$$

We are interested in bounding $\mathbb{P}_{m, w, b}\left(H_{k, i} \mid G_{k, \mathbf{s}, i, \ell}\right)$, which we decompose into two terms:

$$
\begin{align*}
\mathbb{P}_{m, w, b}\left(H_{k, i} \mid G_{k, \mathbf{s}, i, \ell}\right)= & \mathbb{P}_{m, w, b}\left(\left(\mathbf{S}_{k}\right)_{i} \in r_{k, i} \cup b \mid G_{k, \mathbf{s}, i, \ell}\right) \\
= & \mathbb{P}_{m, w, b}\left(\left(\mathbf{S}_{k}\right)_{i} \in r_{k, i} \mid G_{k, \mathbf{s}, i, \ell}\right)  \tag{C.2}\\
& +\mathbb{P}_{m, w, b}\left(\left(\mathbf{S}_{k}\right)_{i} \in b \backslash r_{k, i} \mid G_{k, \mathbf{s}, i, \ell}\right) .
\end{align*}
$$

Since the conditional probability measure $\mathbb{P}_{m, w, b}\left(\cdot \mid G_{k, \mathbf{s}, i, \ell}\right)$ is well defined, and since $\ell \leq \xi_{k}$ and $\xi_{l} \leq \sqrt{n}$ for all $n$ large enough, Proposition 5.1 applies and yields

$$
\mathbb{P}_{m, w, b}\left(\left(\mathbf{S}_{k}\right)_{i}=s \mid G_{k, \mathbf{s}, i, \ell}\right)=\mathbb{P}_{m, w, b}\left(\left(\mathbf{S}_{k}\right)_{i}=s^{\prime} \mid G_{k, \mathbf{s}, i, \ell}\right),
$$

for all $s, s^{\prime} \notin r_{k, i} \cup\left\{\mathbf{s}_{1}, \ldots, \mathbf{s}_{i-1}\right\}$. As a result,

$$
\begin{aligned}
& \mathbb{P}_{m, w, b}\left(\left(\mathbf{S}_{k}\right)_{i} \in b \backslash r_{k, i} \mid G_{k, \mathbf{s}, i, \ell}\right) \\
& \quad \geq \mathbb{P}_{m, w, b}\left(\left(\mathbf{S}_{k}\right)_{i} \in b \backslash\left(r_{k, i} \cup\left\{\mathbf{s}_{1}, \ldots, \mathbf{s}_{i-1}\right\}\right) \mid G_{k, \mathbf{s}, i, \ell}\right) \\
& \quad=\frac{\left|b \backslash\left(r_{k, i} \cup\left\{\mathbf{s}_{1}, \ldots, \mathbf{s}_{i-1}\right\}\right)\right|}{n-\left|r_{k, i} \cup\left\{\mathbf{s}_{1}, \ldots, \mathbf{s}_{i-1}\right\}\right|} \mathbb{P}_{m, w, b}\left(\left(\mathbf{S}_{k}\right)_{i} \notin r_{k, i} \cup\left\{\mathbf{s}_{1}, \ldots, \mathbf{s}_{i-1}\right\} \mid G_{k, \mathbf{s}, i, \ell}\right) .
\end{aligned}
$$

Moreover, using the facts that $|b| \geq \gamma n$ (equation (C.1)), $\left|r_{k, i}\right| \leq c$ (Proposition 5.1), and $i \leq \ell$, we obtain

$$
\begin{aligned}
& \frac{\left|b \backslash\left(r_{k, i} \cup\left\{\mathbf{s}_{1}, \ldots, \mathbf{s}_{i-1}\right\}\right)\right|}{n} \cdot\left|r_{k, i} \cup\left\{\mathbf{s}_{1}, \ldots, \mathbf{s}_{i-1}\right\}\right| \\
& \quad \geq \frac{\gamma n-c-\ell}{n} \cdot \mathbb{P}_{m, w, b}\left(\left(\mathbf{S}_{k}\right)_{i} \notin r_{k, i} \cup\left\{\mathbf{s}_{1}, \ldots, \mathbf{s}_{i-1}\right\}\left|G_{k, \mathbf{s}, i, \ell} \notin r_{k, i} \cup\left\{\mathbf{s}_{1}, \ldots, \mathbf{s}_{i-1}\right\}\right| G_{k, \mathbf{s}, i, \ell}\right) \\
& \quad \geq \frac{\gamma}{2} \cdot \mathbb{P}_{m, w, b}\left(\left(\mathbf{S}_{k}\right)_{i} \notin r_{k, i} \cup\left\{\mathbf{s}_{1}, \ldots, \mathbf{s}_{i-1}\right\} \mid G_{k, \mathbf{s}, i, \ell}\right)
\end{aligned}
$$

when $n$ is large enough. Finally, since the elements of the vector $\mathbf{S}_{k}$ are distinct,

$$
\mathbb{P}_{m, w, b}\left(\left(\mathbf{S}_{k}\right)_{i} \notin r_{k, i} \cup\left\{\mathbf{s}_{1}, \ldots, \mathbf{s}_{i-1}\right\} \mid G_{k, \mathbf{s}, i, \ell}\right)=\mathbb{P}_{m, w, b}\left(\left(\mathbf{S}_{k}\right)_{i} \notin r_{k, i} \mid G_{k, \mathbf{s}, i, \ell}\right)
$$

and therefore

$$
\mathbb{P}_{m, w, b}\left(\left(\mathbf{S}_{k}\right)_{i} \in b \backslash r_{k, i} \mid G_{k, \mathbf{s}, i, \ell}\right) \geq \frac{\gamma}{2} \mathbb{P}_{m, w, b}\left(\left(\mathbf{S}_{k}\right)_{i} \notin r_{k, i} \mid G_{k, \mathbf{s}, i, \ell}\right)
$$

We now substitute into equation (C.2), and obtain

$$
\begin{aligned}
& \mathbb{P}_{m, w, b}\left(H_{k, i} \mid G_{k, \mathbf{s}, i, \ell}\right) \\
& \quad \geq \mathbb{P}_{m, w, b}\left(\left(\mathbf{S}_{k}\right)_{i} \in r_{k, i} \mid G_{k, \mathbf{s}, i, \ell}\right)+\frac{\gamma}{2} \mathbb{P}_{m, w, b}\left(\left(\mathbf{S}_{k}\right)_{i} \notin r_{k, i} \mid G_{k, \mathbf{s}, i, \ell}\right) \\
& \quad \geq \frac{\gamma}{2}
\end{aligned}
$$

for all $n$ large enough.

## APPENDIX D: PROOF OF LEMMA 5.9

We first prove a claim about the set-valued functions $R$ and $R^{\prime}$ introduced in Propositions 5.1 and 5.2 , respectively.

CLAim D.1. For every $m \in \mathcal{M}_{n}, w \in \mathbb{R}_{+}, \mathbf{s} \in \mathcal{S}_{n}$ with $|\mathbf{s}| \leq \sqrt{n}, \mathbf{q} \in \mathcal{Q}^{|\mathbf{s}|}$, and for $\ell=|\mathbf{s}|+1, \ldots, n$, and for every permutation $\sigma$, we have $R\left(\sigma_{M}(m), w, \sigma(\mathbf{s}), \ell\right)=$ $\sigma(R(m, w, \mathbf{s}, \ell))$ and $R^{\prime}\left(\sigma_{M}(m), w, \sigma(\mathbf{s}), \mathbf{q}\right)=\sigma\left(R^{\prime}(m, w, \mathbf{s}, \mathbf{q})\right)$.

Proof. In order to simplify notation, we suppress the dependence on $w$ of the functions $R, R^{\prime}$, and $f_{1}$ throughout the proof of the lemma.

Let $U$ be a uniform random variable over $[0,1]$. For every $m \in \mathcal{M}_{n}$, we define the random vector $\mathbf{S}(m)=f_{1}(m, U)$. Recall that $R(m, \mathbf{s}, \ell) \subset \mathcal{N}_{n} \backslash \mathbf{s}^{\text {set }}$ is the unique set of minimal cardinality such that

$$
\begin{aligned}
& \mathbb{P}\left(\mathbf{S}(m)_{|\mathbf{s}|+1}=j \mid\{|\mathbf{S}(m)|=\ell\} \cap \bigcap_{i=1}^{|\mathbf{s}|}\left\{\mathbf{S}(m)_{i}=\mathbf{s}_{i}\right\}\right) \\
& \quad=\mathbb{P}\left(\mathbf{S}(m)_{|\mathbf{s}|+1}=j^{\prime} \mid\{|\mathbf{S}(m)|=\ell\} \cap \bigcap_{i=1}^{|\mathbf{s}|}\left\{\mathbf{S}(m)_{i}=\mathbf{s}_{i}\right\}\right),
\end{aligned}
$$

for all $j, j^{\prime} \notin R(m, \mathbf{s}, \ell) \cup \mathbf{s}^{\text {set }}$. It is not hard to see, for example, by replacing $j, j^{\prime}$ in the above equality by $\sigma^{-1}(j), \sigma^{-1}\left(j^{\prime}\right) \notin R(m, \mathbf{s}, \ell) \cup \mathbf{s}^{\text {set }}$, that $\sigma(R(m, \mathbf{s}, \ell)) \subset \mathcal{N}_{n} \backslash \sigma\left(\mathbf{s}^{\text {set }}\right)$ is the unique set of minimal cardinality such that

$$
\begin{aligned}
& \mathbb{P}\left(\sigma\left(\mathbf{S}(m)_{|\mathbf{s}|+1}\right)=j \mid\{|\sigma(\mathbf{S}(m))|=\ell\} \cap \bigcap_{i=1}^{|\mathbf{s}|}\left\{\sigma\left(\mathbf{S}(m)_{i}\right)=\sigma\left(\mathbf{s}_{i}\right)\right\}\right) \\
& \quad=\mathbb{P}\left(\sigma\left(\mathbf{S}(m)_{|\mathbf{s}|+1}\right)=j^{\prime} \mid\{|\sigma(\mathbf{S}(m))|=\ell\} \cap \bigcap_{i=1}^{|\mathbf{s}|}\left\{\sigma\left(\mathbf{S}(m)_{i}\right)=\sigma\left(\mathbf{s}_{i}\right)\right\}\right),
\end{aligned}
$$

for all $j, j^{\prime} \notin \sigma(R(m, \mathbf{s}, \ell)) \cup \sigma\left(\mathbf{s}^{\text {set }}\right)$. On the other hand, the symmetry assumption states that

$$
\sigma(\mathbf{S}(m)) \stackrel{d}{=} \mathbf{S}\left(\sigma_{M}(m)\right)
$$

Combining the last two equalities we get that $\sigma(R(m, \mathbf{s}, \ell)) \subset \mathcal{N}_{n} \backslash \sigma\left(\mathbf{s}^{\text {set }}\right)$ is the unique set of minimal cardinality such that

$$
\begin{aligned}
& \mathbb{P}\left(\mathbf{S}\left(\sigma_{M}(m)\right)_{|\mathbf{s}|+1}=j \mid\left\{\left|\mathbf{S}\left(\sigma_{M}(m)\right)\right|=\ell\right\} \cap \bigcap_{i=1}^{|\mathbf{s}|}\left\{\mathbf{S}\left(\sigma_{M}(m)\right)_{i}=\sigma\left(\mathbf{s}_{i}\right)\right\}\right) \\
& \quad=\mathbb{P}\left(\mathbf{S}\left(\sigma_{M}(m)\right)_{|\mathbf{s}|+1}=j^{\prime} \mid\left\{\left|\mathbf{S}\left(\sigma_{M}(m)\right)\right|=\ell\right\} \cap \bigcap_{i=1}^{|\mathbf{s}|}\left\{\mathbf{S}\left(\sigma_{M}(m)\right)_{i}=\sigma\left(\mathbf{s}_{i}\right)\right\}\right),
\end{aligned}
$$

for all $i, j \notin \sigma(R(m, \mathbf{s}, \ell)) \cup \sigma\left(\mathbf{s}^{\text {set }}\right)$. However, this is exactly the definition of $R\left(\sigma_{M}(m)\right.$, $\sigma(\mathbf{s}), \ell)$ (uniqueness is crucial at this point), so we have

$$
\sigma(R(m, \mathbf{s}, \ell))=R\left(\sigma_{M}(m), \sigma(\mathbf{s}), \ell\right)
$$

The proof of $R^{\prime}\left(\sigma_{M}(m), \sigma(\mathbf{s}), \mathbf{q}\right)=\sigma\left(R^{\prime}(m, \mathbf{s}, \mathbf{q})\right)$ is analogous (this time making use of the symmetry of the mapping $f_{2}$ ) and is omitted.

We continue with the proof of Lemma 5.9. Under the event $\mathcal{H}_{c+1}^{+}$, we have $\left(\mathbf{S}_{1}\right)_{i} \in R_{1, i} \cup$ $B$, for $i=1, \ldots,\left|\mathbf{S}_{1}\right|$. Applying Claim D. 1 and the fact $m\left(t_{1}^{-}\right)=m(0)$, which implies that $\sigma_{M}\left(m\left(t_{1}^{-}\right)\right)=\tau_{M}\left(m\left(t_{1}^{-}\right)\right)$, we obtain
(D.1)

$$
\begin{aligned}
\sigma\left(R_{1,1}\right) & =\sigma\left(R\left(m\left(t_{1}^{-}\right), w_{1}, \varnothing,\left|\mathbf{S}_{1}\right|\right)\right) \\
& =R\left(\sigma_{M}\left(m\left(t_{1}^{-}\right)\right), w_{1}, \varnothing,\left|\mathbf{S}_{1}\right|\right) \\
& =R\left(\tau_{M}\left(m\left(t_{1}^{-}\right)\right), w_{1}, \varnothing,\left|\mathbf{S}_{1}\right|\right) \\
& =\tau\left(R\left(m\left(t_{1}^{-}\right), w_{1}, \varnothing,\left|\mathbf{S}_{1}\right|\right)\right) \\
& =\tau\left(R_{1,1}\right) .
\end{aligned}
$$

Now recall that $\sigma$ and $\tau$ preserve the order of $\bar{R} \backslash B$ and fix $\bar{R} \cap B$, so in particular they preserve the order of $R_{1,1} \backslash B \subset \bar{R} \backslash B$ and fix $R_{1,1} \cap B \subset \bar{R} \cap B$. Combining this with equation (D.1), we must have $\sigma(i)=\tau(i)$, for all $i \in R_{1,1}$. If $\left(\mathbf{S}_{1}\right)_{1} \in R_{1,1}$, this implies that

$$
\begin{equation*}
\tau\left(\left(\mathbf{S}_{1}\right)_{1}\right)=\sigma\left(\left(\mathbf{S}_{1}\right)_{1}\right) . \tag{D.2}
\end{equation*}
$$

On the other hand, if $\left(\mathbf{S}_{1}\right)_{1}$ does not belong to $R_{1,1}$, then, from the definition of $F$, we must have $\left(\mathbf{S}_{1}\right)_{1} \in F$. Since $\sigma$ and $\tau$ fix the set $F$, we conclude that equation (D.2) must hold in all cases.

Proceeding inductively, and using the same argument, we obtain

$$
\sigma\left(R_{1, i}\right)=\tau\left(R_{1, i}\right),
$$

for $i=1, \ldots,\left|\mathbf{S}_{1}\right|$, and $\sigma(i)=\tau(i)$, for all $i \in \mathbf{S}_{1}^{\text {set }}$. It follows that $\sigma\left(\mathbf{S}_{1}\right)=\tau\left(\mathbf{S}_{1}\right)$. Combining this with the fact that $\sigma_{M}\left(m\left(t_{1}^{-}\right)\right)=\tau_{M}\left(m\left(t_{1}^{-}\right)\right)$, and applying Claim D. 1 twice, we obtain

$$
\begin{align*}
\sigma\left(R_{1}^{\prime}\right) & =\sigma\left(R^{\prime}\left(m\left(t_{1}^{-}\right), w_{1}, \mathbf{S}_{1}, \mathbf{q}_{\mathbf{S}_{1}}\left(t_{1}^{-}\right)\right)\right) \\
& =R^{\prime}\left(\sigma_{M}\left(m\left(t_{1}^{-}\right)\right), w_{1}, \sigma\left(\mathbf{S}_{1}\right), \mathbf{q}_{\mathbf{S}_{1}}\left(t_{1}^{-}\right)\right) \\
& =R^{\prime}\left(\tau_{M}\left(m\left(t_{1}^{-}\right)\right), w_{1}, \tau\left(\mathbf{S}_{1}\right), \mathbf{q}_{\mathbf{S}_{1}}\left(t_{1}^{-}\right)\right)  \tag{D.3}\\
& =\tau\left(R^{\prime}\left(m\left(t_{1}^{-}\right), w_{1}, \mathbf{S}_{1}, \mathbf{q}_{\mathbf{S}_{1}}\left(t_{1}^{-}\right)\right)\right) \\
& =\tau\left(R_{1}^{\prime}\right) .
\end{align*}
$$

Now recall that $\sigma$ and $\tau$ preserve the order of $\bar{R} \backslash B$ and fix $\bar{R} \cap B$, so in particular they preserve the order of $R_{1}^{\prime} \backslash B \subset \bar{R} \backslash B$ and fix $R_{1}^{\prime} \cap B \subset \bar{R} \cap B$. Combining this with equation (D.3), we must have $\sigma(i)=\tau(i)$, for all $i \in R_{1}^{\prime}$. Furthermore, recall that we also have that $\sigma(i)=\tau(i)$, for all $i \in \mathbf{S}_{1}^{\text {set }}$. If $D_{1} \in R_{1}^{\prime} \cup \mathbf{S}_{1}^{\text {set }}$, this implies that

$$
\begin{equation*}
\sigma\left(D_{1}\right)=\tau\left(D_{1}\right) . \tag{D.4}
\end{equation*}
$$

On the other hand, if $D_{1}$ does not belong to $R_{1}^{\prime} \cup \mathbf{S}_{1}^{\text {set }}$, then, from the definition of $F$, we must have $D_{1} \in F$. Since $\sigma$ and $\tau$ fix the set $F$, we conclude that equation (D.4) must hold in all cases.

We now consider a memory update. Using the symmetry assumption, we have

$$
\begin{aligned}
\sigma_{M}\left(m\left(t_{1}\right)\right) & =\sigma_{M}\left(f_{3}\left(m\left(t_{1}^{-}\right), w_{1}, \mathbf{S}_{1}, \mathbf{q}_{\mathbf{S}_{1}}\left(t_{1}^{-}\right), D_{1}\right)\right) \\
& =f_{3}\left(\sigma_{M}\left(m\left(t_{1}^{-}\right)\right), w_{1}, \sigma\left(\mathbf{S}_{1}\right), \mathbf{q}_{\mathbf{S}_{1}}\left(t_{1}^{-}\right), \sigma\left(D_{1}\right)\right)
\end{aligned}
$$

Then, since $\sigma_{M}\left(m\left(t_{1}^{-}\right)\right)=\tau_{M}\left(m\left(t_{1}^{-}\right)\right), \sigma\left(\mathbf{S}_{1}\right)=\tau\left(\mathbf{S}_{1}\right)$, and $\tau\left(D_{1}\right)=\sigma\left(D_{1}\right)$, we have

$$
\begin{aligned}
& f_{3}\left(\sigma_{M}\left(m\left(t_{1}^{-}\right)\right), w_{1}, \sigma\left(\mathbf{S}_{1}\right), \mathbf{q}_{\mathbf{S}_{1}}\left(t_{1}^{-}\right), \sigma\left(D_{1}\right)\right) \\
& \quad=f_{3}\left(\tau_{M}\left(m\left(t_{1}^{-}\right)\right), w_{1}, \tau\left(\mathbf{S}_{1}\right), \mathbf{q}_{\mathbf{S}_{1}}\left(t_{1}^{-}\right), \tau\left(D_{1}\right)\right)
\end{aligned}
$$

Using the symmetry assumption once again, we obtain

$$
\begin{aligned}
f_{3} & \left(\tau_{M}\left(m\left(t_{1}^{-}\right)\right), w_{1}, \tau\left(\mathbf{S}_{1}\right), \mathbf{q}_{\mathbf{S}_{1}}\left(t_{1}^{-}\right), \tau\left(D_{1}\right)\right) \\
& =\tau_{M}\left(f_{3}\left(m\left(t_{1}^{-}\right), w_{1}, \mathbf{S}_{1}, \mathbf{q}_{\mathbf{S}_{1}}\left(t_{1}^{-}\right), D_{1}\right)\right) \\
& =\tau_{M}\left(m\left(t_{1}\right)\right) .
\end{aligned}
$$

We conclude that

$$
\sigma_{M}\left(m\left(t_{1}\right)\right)=\tau_{M}\left(m\left(t_{1}\right)\right)
$$

Finally, since the memory states at time $t_{1}$ are still equal, we can proceed inductively by applying the same argument to obtain that, for $k=2, \ldots, c+1$, we have $\sigma\left(R_{k, i}\right)=\tau\left(R_{k, i}\right)$ for $i=1, \ldots,\left|\mathbf{S}_{k}\right|$, and $\sigma\left(R_{k}^{\prime}\right)=\tau\left(R_{k}^{\prime}\right)$. It follows that $\sigma(\bar{R})=\tau(\bar{R})$.

Acknowledgments. This work was supported in part by an MIT Jacobs Presidential Fellowship, NSF Grant CMMI-1234062 and ONR Grant N00014-17-1-2790.

## REFERENCES

[1] Alon, N., Lubetzky, E. and Gurel-Gurevich, O. (2009). Choice-memory tradeoff in allocations. In 2009 50th Annual IEEE Symposium on Foundations of Computer Science—FOCS 2009 230-238. IEEE Computer Soc., Los Alamitos, CA. MR2648405 https://doi.org/10.1109/FOCS.2009.49
[2] AZar, Y., Broder, A. Z., Karlin, A. R. and Upfal, E. (1999). Balanced allocations. SIAM J. Comput. 29 180-200. MR1710347 https://doi.org/10.1137/S0097539795288490
[3] Baccelli, F. and Brémaud, P. (2003). Elements of Queueing Theory: Palm Martingale Calculus and Stochastic Recurrences, 2nd ed. Applications of Mathematics (New York) 26. Springer, Berlin. MR1957884 https://doi.org/10.1007/978-3-662-11657-9
[4] Badonnel, R. and Burgess, M. (2008). Dynamic pull-based load balancing for autonomic servers. In Network Operations and Management Symposium (NOMS).
[5] Bramson, M., Lu, Y. and Prabhakar, B. (2013). Decay of tails at equilibrium for FIFO join the shortest queue networks. Ann. Appl. Probab. 23 1841-1878. MR3114919 https://doi.org/10.1214/12-AAP888
[6] Feitelson, D. and Jette, M. A. (1997). Improved utilization and responsiveness with gang scheduling. In IPPS '97 Proceedings of the Job Scheduling Strategies for Parallel Processing 238-261.
[7] Gamarnik, D., Tsitsiklis, J. N. and Zubeldia, M. (2018). Delay, memory, and messaging tradeoffs in distributed service systems. Stoch. Syst. 8 45-74. MR3775991 https://doi.org/10.1287/stsy.2017.0008
[8] Harchol-Balter, M., Crovella, M. E. and Murta, C. D. (1999). On choosing a task assignment policy for a distributed server system. J. Parallel Distrib. Comput. 59 204-228.
[9] Hellemans, T. and Van Houdt, B. (2018). On the power-of-d-choices with least loaded server selection. Proc. ACM Meas. Anal. Comput. Syst. 2 Article No. 27.
[10] LENZEN, C. and WATtENHOFER, R. (2016). Tight bounds for parallel randomized load balancing. Distrib. Comput. 29 127-142. MR3488178 https://doi.org/10.1007/s00446-014-0225-4
[11] Lu, Y., Xie, Q., Kliot, G., Geller, A., Larus, J. R. and Greenberg, A. (2011). Join-idle-queue: A novel load balancing algorithm for dynamically scalable web services. Perform. Eval. 68 1056-1071.
[12] Mitzenmacher, M. D. (1996). The Power of Two Choices in Randomized Load Balancing. ProQuest LLC, Ann Arbor, MI. Thesis (Ph.D.)—Univ. California, Berkeley. MR2695522
[13] Mitzenmacher, M., Prabhakar, B. and Shah, D. (2002). Load balancing with memory. In Proceedings of the 43 rd Annual IEEE Symposium on Foundations of Computer Science.
[14] Mukherjee, D., Borst, S., Van Leeuwaarden, J. and Whiting, P. (2016). Universality of power-of-d load balancing schemes. In Workshop on Mathematical Performance Modeling and Analysis (MAMA).
[15] Stamoulis, G. D. and Tsitsiklis, J. N. (1991). Optimal distributed policies for choosing among multiple servers. In Proceedings of the 30th Conference on Decision and Control 815-820.
[16] Stolyar, A. L. (2015). Pull-based load distribution in large-scale heterogeneous service systems. Queueing Syst. 80 341-361. MR3367704 https://doi.org/10.1007/s11134-015-9448-8
[17] Vvedenskaya, N. D., Dobrushin, R. L. and Karpelevich, F. I. (1996). Queueing system with selection of the shortest of two queues: An asymptotic approach. Probl. Inf. Transm. 32 15-27. MR1384927
[18] Winston, W. (1977). Optimality of the shortest line discipline. J. Appl. Probab. 14 181-189. MR0428516 https://doi.org/10.1017/s0021900200104772
[19] Ying, L., Srikant, R. and Kang, X. (2015). The power of slightly more than one sample in randomized load balancing. In Proceedings of the IEEE Conference on Computer Communications.


[^0]:    Received July 2018; revised May 2019.
    MSC2010 subject classifications. 60G99.
    Key words and phrases. Load balancing, queueing delay, resource constraints.

[^1]:    ${ }^{1}$ Note that the argument on the right-hand side of the relation below involves $\mathbf{q}$ rather than a permuted version of $\mathbf{q}$, even though the vector $\mathbf{s}$ gets permuted. We are essentially comparing a situation where the dispatcher queries a vector $\mathbf{s}$ and receives certain numerical values $\mathbf{q}$ with the situation where the dispatcher queries a vector $\sigma(\mathbf{s})$ and receives the same numerical values $\mathbf{q}$.

