

UNIVERSITY OF CALIFORNIA, SAN DIEGO

**Dynamic Scheduling of a Parallel Server System in Heavy Traffic with
Complete Resource Pooling: Asymptotic Optimality of a Threshold
Policy**

A dissertation submitted in partial satisfaction of the
requirements for the degree
Doctor of Philosophy

in

Mathematics

by

Steven L. Bell

Committee in charge:

Professor Ruth Williams, Chair
Professor Rene Cruz
Professor Bruce Driver
Professor Patrick Fitzsimmons
Professor Halbert White

2003

Copyright
Steven L. Bell, 2003
All rights reserved.

The dissertation of Steven L. Bell is approved, and
it is acceptable in quality and form for publication
on microfilm:

Chair

University of California, San Diego

2003

TABLE OF CONTENTS

Signature Page		iii
Table of Contents		iv
List of Figures		vi
Vita		vii
Abstract of the Dissertation		viii
1 Introduction		1
1.1 Notation and Terminology		5
2 The Parallel Server System		7
2.1 System Structure		7
2.2 Stochastic Primitives		9
2.3 Scheduling Control and Performance Measures		10
3 Sequence of Systems, Heavy Traffic, and the Cost Function		13
3.1 Sequence of Systems and Large Deviation Assumptions		14
3.2 The Fluid Model and Heavy Traffic		15
3.3 Diffusion Scaling and the Cost Function		18
4 Brownian Control Problem		21
4.1 Brownian Model		21
5 Solution of the Brownian Control Problem via Workload		24
5.1 Solution via Workload		25
5.2 Interpreting the Solution		28
6 Threshold Policy and Asymptotic Optimality Result		31
6.1 Threshold Policy		31
6.2 Asymptotic Optimality		36
7 Preliminaries and Description of the Proof		38
7.1 The Server-Buffer Tree \mathcal{G} : Layers and Buffer Renumbering		38
7.2 Threshold Sizes and Transient Nominal Activity Rates		41
7.3 State Space Collapse Result and Outline of its Proof		43
7.4 Residual Processes, Excursions, and Shifted Allocation Processes		45
7.5 Preliminaries on Stopped Arrival and Service Processes		47

7.6	Large Deviation Bounds for Renewal Processes	49
8	Proof of State Space Collapse	51
8.1	Main Technical Result	51
8.2	Auxiliary Constants for the Induction Proof	52
8.3	Induction Setup	57
8.4	Proof of Lemma 8.3.1: Estimates on Allocation Processes Imply Residual Processes Stay Near Zero	61
8.5	Proof of Lemma 8.3.2: Estimates on Allocations for Activities Immediately Below Each Buffer	78
8.6	Proof of Lemma 8.3.3: Estimate on Allocation for the Activity Immediately Above Each Buffer	80
8.7	Proof of Lemma 8.3.4: Transition Between Layers in the Server-Buffer Tree	82
8.8	Proofs of Theorems 7.3.1, 8.1.1, and 8.3.5	99
9	Weak Convergence under the Threshold Policy	104
9.1	Fluid Limits for Allocation Processes	105
9.2	Proof of Theorem 6.2.1: Convergence of Diffusion Scaled Performance Measures under the Threshold Policy	111
10	Asymptotic Optimality of the Threshold Policy	113
10.1	A Necessary Condition for Asymptotic Optimality	114
10.2	Proof of Asymptotic Optimality	116
A	Appendix: Proof of Multiparameter Stopping Time Property	124
	Bibliography	132

LIST OF FIGURES

2.1	Example of a parallel server system.	8
6.1	The server tree (on the left) and server-buffer tree (on the right) for the three-server network of Figure 2.1 with the data of Example 6.1.1.	36
7.1	Layers l and $l - 1$ of the server-buffer tree	40
9.1	Activity $b(k)$ above a server k that is in layer l	106

VITA

July 30, 1960	Born, Los Angeles, California
1987	B. A., University of California, Los Angeles
1994	M. A., San Diego State University
1994–2000	Teaching assistant, Department of Mathematics, University of California, San Diego
2003	Ph. D., University of California, San Diego

ABSTRACT OF THE DISSERTATION

Dynamic Scheduling of a Parallel Server System in Heavy Traffic with Complete Resource Pooling: Asymptotic Optimality of a Threshold Policy

by

Steven L. Bell

Doctor of Philosophy in Mathematics

University of California San Diego, 2003

Professor Ruth Williams, Chair

We consider a queueing system with a bank of non-identical servers working in parallel, exogenous arrivals classified into one of several infinite capacity buffers, and linear holding costs for each buffer. Jobs within a buffer are ordered on a first-in-first-out basis. Each incoming job requires a single service which may be provided by any of several different servers; the service time of a job may depend on both its buffer and the server providing the service. The system manager seeks to minimize holding costs by dynamically scheduling waiting jobs to available servers. In a recent work, under a complete resource pooling condition, Williams [39] proposed a continuous review threshold control policy for such a parallel server system by “interpreting” the analytic solution to the associated Brownian control problem (formal heavy traffic approximation). We show that this policy is asymptotically optimal in the heavy traffic limit and that the limiting cost is the same as the optimal cost in the Brownian control problem.

1

Introduction

We consider a dynamic scheduling problem for a parallel server queueing system. This system might be viewed as a model for a manufacturing or computer system, consisting of a bank of buffers for holding incoming jobs and a bank of flexible servers for processing these jobs (see e.g., [21]). Incoming jobs are classified into one of several different classes (or buffers). Jobs within a class are served on a first-in-first-out basis and may be served by any server from a given subset of the bank of servers (this subset may depend on the class). In addition, servers may have differing but overlapping capabilities and so may be able to service more than one class. Jobs of each class incur linear holding costs while present within the system. The system manager seeks to minimize holding costs by dynamically allocating waiting jobs to available servers.

The parallel server system is described in more detail in Chapter 2 below. With the exception of a few special cases, the dynamic scheduling problem for this system cannot be analyzed exactly and it is natural to consider more tractable approximations. One class of such approximations are the so-called Brownian control problems, first introduced by Harrison in [10] and further developed in [14, 17]. These are formal heavy traffic approximations to queueing control problems. Various authors (see for example [6, 18, 19, 22, 27, 28, 36]) have used analysis of these Brownian control problems, together with clever interpretation of their

optimal (analytic) solutions, to suggest “good” policies for the original queueing control problems.

For the parallel server system, Harrison and López [16] studied the associated Brownian control problem and identified a condition under which the solution of that problem exhibits *complete resource pooling*, i.e., in the Brownian model, the efforts of the individual servers can be efficiently combined to act as a single pooled resource or “superserver”. Under this condition, Harrison and López [16] conjecture that a “discrete review” scheduling policy (for the original parallel server system), obtained by using the BIGSTEP discretization procedure of Harrison [12], is asymptotically optimal in the heavy traffic limit.

Here, focusing on the parameter regime associated with the complete resource pooling condition of Harrison and López [16], we first review their formulation and solution of the Brownian control problem. Then we prove that a “continuous review” threshold policy proposed by Williams [39] is asymptotically optimal in the heavy traffic limit. Our treatment of the Brownian control problem and our description of the candidate threshold policy closely follows that presented in [39]. On the other hand, the proof of asymptotic optimality of this policy is new. In a related work [2], we have already proved that this policy is asymptotically optimal for a particular two-server, two-buffer system. Indeed, techniques developed in [2] have been useful for analysis of the more complex multiserver case treated here.

Since we began this work, three related works have appeared [1, 33, 29]. In [1], Ata and Kumar consider a dynamic scheduling problem for an open stochastic processing network that allows feedback routing. A parallel server system is a special case of such networks in which no routing occurs. Under heavy traffic and complete resource pooling conditions, Ata and Kumar prove asymptotic optimality of a discrete review policy for an open stochastic processing network with linear holding costs. Although this provides an asymptotically optimal policy for the parallel server problem considered here, we think it is still of interest to establish asymptotic optimality of a simple continuous review threshold policy, as we do here. The other related works are by Stolyar [33], who considers a generalized switch,

which operates in discrete time, and Mandelbaum and Stolyar [29], who consider a parallel server system. Although it does not allow routing, Stolyar's generalized switch is somewhat more general than a parallel server system operating in discrete time, in particular, it allows service rates that depend on the state of a random environment. Assuming heavy traffic and a resource pooling condition, which is slightly more general than a complete resource pooling condition, Stolyar [33] proves asymptotic optimality of a MaxWeight policy, for holding costs that are positive linear combinations of the individual queue lengths raised to the power $\beta + 1$ where $\beta > 0$. In particular, the holding costs are not linear. An advantage of the MaxWeight policy (which exploits the non-linear nature of the holding cost function) is that it does not require knowledge of the arrival rates for its execution, although checking the heavy traffic and resource pooling conditions does involve these rates. Following on from [33], in [29], Mandelbaum and Stolyar focus on a parallel server system (operating in continuous time). Assuming heavy traffic and complete resource pooling conditions they prove asymptotic optimality of a MaxWeight policy (called a generalized $c\mu$ -rule there), for holding costs that are sums of strictly increasing, strictly convex functions of the individual queue lengths. (They also prove a related result where queue lengths are replaced by sojourn times.) Again, the nonlinear nature of the holding cost function allows the authors to specify a policy that does not require knowledge of the arrival rates (nor of a solution of a certain dual linear program). Although Mandelbaum and Stolyar [29] conjecture a policy for linear holding costs (which like ours makes use of the solution of a dual linear program), they stop short of proving asymptotic optimality of that policy. Thus, our paper provides the only proof of asymptotic optimality of a continuous review policy for the parallel server system with linear holding costs. An additional difference between our work and that in [1, 29, 33] is that we impose finite exponential moment assumptions on our primitive stochastic processes, whereas only finite moment assumptions (of order $2 + \epsilon$) are needed for the results in [1, 29, 33]. We conjecture that our exponential moment assumptions could be relaxed at the expense of an increase in the size of our (logarithmic)

thresholds. We have not pursued this conjecture here, having chosen the tradeoff of smaller thresholds at the expense of higher moment assumptions.

This paper is organized as follows. In Chapter 2, we describe the model of a parallel server system considered here. In Chapter 3 we introduce a sequence of such systems, indexed by r (where r tends to infinity through a sequence of values in $[1, \infty)$), which is used in formulating the notion of heavy traffic asymptotic optimality. The cost function used in the r^{th} system is an average cumulative discounted linear holding cost, where the linear holding cost is per unit of normalized queue length (in diffusion scale). In Chapter 3, we also review the notion of heavy traffic defined in [14, 16] using a linear program, and recall its interpretation in terms of the behavior of an associated fluid model, as previously described in [39]. In Chapter 4, we describe the Brownian control problem associated with the sequence of parallel server systems. In Chapter 5, under the complete resource pooling condition of [16], we review the solution of the Brownian control problem obtained by Harrison and López [16], using a reduced form of the problem called the equivalent workload formulation [14, 17]. The complete resource pooling condition ensures that the Brownian workload process is one-dimensional. Moreover, from [16] we know that this condition is equivalent to uniqueness of a solution to the dual to the linear program described in Chapter 3. In Chapter 6, we describe the dynamic threshold policy proposed by Williams [39] for use in the original parallel server systems. We then state the main result (Theorem 6.2.2) which implies that this policy is asymptotically optimal in the heavy traffic limit and that the limiting cost is the same as the optimal cost in the Brownian control problem. An outline of our method of proof is given in Chapter 7. The details of the proof are contained in Chapters 8–10. Here a critical role is played by our analysis in Chapter 8 of what we call the residual processes, which measure the deviations of the queue lengths from the threshold levels, or from zero if a queue does not have a threshold on it, when the threshold policy is used. This allows us to establish a form of “state space collapse” (see Theorem 6.2.1) under this policy. The techniques used in proving state space collapse build on and extend those introduced

in [2]. In particular, a major new feature is the need to show that allocations of time to various activities stay close to their nominal allocations over sufficiently long time intervals (with high probability), which in turn is used to show that the residual processes stay close to zero (with high probability). Using a suitable numbering of the buffers, the proof of state space collapse proceeds by induction on the buffer number, highlighting the fact that the queue length for a particular buffer depends (via the threshold policy) on the queue lengths associated with lower numbered buffers.

1.1 Notation and Terminology

The set of non-negative integers will be denoted by \mathbb{N} and the value $+\infty$ will simply be denoted by ∞ . For any real number x , $\lfloor x \rfloor$ will denote the integer part of x , i.e., the greatest integer that is less than or equal to x , and $\lceil x \rceil$ will denote the smallest integer that is greater than or equal to x . The m -dimensional ($m \geq 1$) Euclidean space will be denoted by \mathbb{R}^m and \mathbb{R}_+ will denote $[0, \infty)$. Let $|\cdot|$ denote the norm on \mathbb{R}^m given by $|x| = \sum_{i=1}^m |x_i|$ for $x \in \mathbb{R}^m$. We define a sum over an empty index set to be zero. Vectors in \mathbb{R}^m should be treated as column vectors unless indicated otherwise, inequalities between vectors should be interpreted componentwise, the transpose of a vector a will be denoted by a' , the diagonal matrix with the entries of a vector a on its diagonal will be denoted by $\text{diag}(a)$, and the dot product of two vectors a and b in \mathbb{R}^m will be denoted by $a \cdot b$. For any set \mathcal{S} , let $|\mathcal{S}|$ denote the cardinality of \mathcal{S} .

For each positive integer m , let D^m be the space of “Skorokhod paths” in \mathbb{R}^m having time domain \mathbb{R}_+ . That is, D^m is the set of all functions $\omega : \mathbb{R}_+ \rightarrow \mathbb{R}^m$ that are right continuous on \mathbb{R}_+ and have finite left limits on $(0, \infty)$. The member of D^m that stays at the origin in \mathbb{R}^m for all time will be denoted by $\mathbf{0}$. For $\omega \in D^m$ and $t \geq 0$, let

$$\|\omega\|_t = \sup_{s \in [0, t]} |\omega(s)|. \quad (1.1)$$

Consider D^m to be endowed with the usual Skorokhod J_1 -topology (see [8]). Let \mathcal{M}^m denote the Borel σ -algebra on D^m associated with the J_1 -topology. All of the continuous-time processes in this paper will be assumed to have sample paths in D^m for some $m \geq 1$. (We shall frequently use the term process in place of stochastic process.)

Suppose $\{W^n\}_{n=1}^\infty$ is a sequence of processes with sample paths in D^m for some $m \geq 1$. Then we say that $\{W^n\}_{n=1}^\infty$ is tight if and only if the probability measures induced by the W^n 's on (D^m, \mathcal{M}^m) form a tight sequence, i.e., they form a weakly relatively compact sequence in the space of probability measures on (D^m, \mathcal{M}^m) . The notation " $W^n \Rightarrow W$ ", where W is a process with sample paths in D^m , will mean that the probability measures induced by the W^n 's on (D^m, \mathcal{M}^m) converge weakly to the probability measure on (D^m, \mathcal{M}^m) induced by W . If for each n , W^n and W are defined on the same probability space, we say that W^n converges to W uniformly on compact time intervals in probability (u.o.c. in prob.), if $\mathbf{P}(\|W^n - W\|_t \geq \varepsilon) \rightarrow 0$ as $n \rightarrow \infty$ for each $\varepsilon > 0$ and all $t \geq 0$. We note that if $\{W^n\}$ is a sequence of processes and W is a continuous deterministic process (all defined on the same probability space) then $W^n \Rightarrow W$ is equivalent to $W^n \rightarrow W$ u.o.c. in prob. This is implicitly used several times in the proofs below to combine statements involving convergence in distribution to deterministic processes.

2

The Parallel Server System

In this Chapter we review the description of a parallel server system given in [39].

2.1 System Structure

Our parallel server system consists of \mathbf{I} infinite capacity buffers for holding jobs awaiting service, indexed by $i \in \mathcal{I} \equiv \{1, \dots, \mathbf{I}\}$, and \mathbf{K} (non-identical) servers working in parallel, indexed by $k \in \mathcal{K} \equiv \{1, \dots, \mathbf{K}\}$ (see e.g., Figure 2.1, where a system considered by Harrison and López [16] is pictured). Each buffer has its own stream of jobs arriving from outside the system. Arrivals to buffer i are called class i jobs and jobs are ordered within each buffer according to their arrival times, with the earliest arrival being at the head of the line. Each entering job requires a single service before it exits the system. Several different servers may be capable of processing (or serving) a particular job class. Service of a given job class i by a given server k is called a processing activity. Although we shall not use this notation, such a processing activity can be associated with the ordered pair (i, k) and there is an upper bound of $\mathbf{I} \cdot \mathbf{K}$ on the number of processing activities. Due to system constraints, the actual number \mathbf{J} of processing activities available may be less than this upper bound. In any event, we assume there are $\mathbf{J} \leq \mathbf{I} \cdot \mathbf{K}$ possible




Figure 2.1: Example of a parallel server system.

processing activities labeled by $j \in \mathcal{J} \equiv \{1, \dots, \mathbf{J}\}$. The correspondences between activities and classes, and activities and servers, are described by two deterministic matrices \mathbf{C} , \mathbf{A} , where \mathbf{C} is an $\mathbf{I} \times \mathbf{J}$ matrix with

$$\mathbf{C}_{ij} = \begin{cases} 1 & \text{if activity } j \text{ processes class } i, \\ 0 & \text{otherwise,} \end{cases} \quad (2.1)$$

\mathbf{A} is a $\mathbf{K} \times \mathbf{J}$ matrix with

$$\mathbf{A}_{kj} = \begin{cases} 1 & \text{if server } k \text{ performs activity } j, \\ 0 & \text{otherwise.} \end{cases} \quad (2.2)$$

Note that each column of \mathbf{C} contains the number one exactly once and similarly for \mathbf{A} , since each activity j has exactly one class $i(j)$ and one server $k(j)$ associated with it. We also assume that each row of \mathbf{C} and each row of \mathbf{A} contains the number one at least once (i.e., each job class is capable of being processed by at least one activity and each server is capable of performing at least one activity).

For the example pictured in Figure 2.1, the matrices \mathbf{C} and \mathbf{A} are given by

$$\mathbf{C} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 1 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 \end{bmatrix} \quad \mathbf{A} = \begin{bmatrix} 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 1 \end{bmatrix}.$$

Once a job has commenced service at a server, it remains there until its service is complete, even if its service is interrupted for some time (e.g., by preemption by a job of another class). A server may not start on a new job of class i until it has finished serving any class i job that it is working on or that is in suspension at the server. In addition, a server cannot work unless it has a job to work on. When taking a new job from a buffer, a server always takes the job at the head-of-the-line. (For concreteness, we suppose that a deterministic tie-breaking rule is used when two (or more) servers want to simultaneously take jobs from the same buffer, e.g., there is an ordering of the servers and lower numbered servers take jobs before higher numbered ones.) This setup allows a job to be allocated to a server just before it begins service, rather than upon arrival to the system. We assume that the system is initially empty.

2.2 Stochastic Primitives

All random variables and stochastic processes used in our model description are assumed to be defined on a given complete probability space $(\Omega, \mathcal{F}, \mathbf{P})$. The expectation operator under \mathbf{P} will be denoted by \mathbf{E} . For $i \in \mathcal{I}$, we take as given a sequence of strictly positive i.i.d. random variables $\{u_i(\ell), \ell = 1, 2, \dots\}$ with mean $\lambda_i^{-1} \in (0, \infty)$ and squared coefficient of variation (variance divided by the square of the mean) $a_i^2 \in [0, \infty)$. We interpret $u_i(\ell)$ as the interarrival time between the $(\ell - 1)^{\text{st}}$ and the ℓ^{th} arrival to class i where, by convention, the “0th arrival” is assumed to occur at time zero. Setting $\xi_i(0) = 0$ and

$$\xi_i(n) = \sum_{\ell=1}^n u_i(\ell), \quad n = 1, 2, \dots, \quad (2.4)$$

we define

$$A_i(t) = \sup\{n \geq 0 : \xi_i(n) \leq t\} \quad \text{for all } t \geq 0. \quad (2.5)$$

Then $A_i(t)$ is the number of arrivals to class i that have occurred in $[0, t]$, and λ_i is the long run arrival rate to class i . For $j \in \mathcal{J}$, we take as given a sequence of strictly

positive i.i.d. random variables $\{v_j(\ell), \ell = 1, 2, \dots\}$ with mean $\mu_j^{-1} \in (0, \infty)$ and squared coefficient of variation $b_j^2 \in [0, \infty)$. We interpret $v_j(\ell)$ as the amount of service time required by the ℓ^{th} job processed by activity j , and μ_j as the long run rate at which activity j could process its associated class of jobs $i(j)$ if the associated server $k(j)$ worked continuously and exclusively on this class. For $j \in \mathcal{J}$, let $\eta_j(0) = 0$,

$$\eta_j(n) = \sum_{\ell=1}^n v_j(\ell), \quad n = 1, 2, \dots, \quad (2.6)$$

and

$$S_j(t) = \sup\{n \geq 0 : \eta_j(n) \leq t\} \quad \text{for all } t \geq 0. \quad (2.7)$$

Then $S_j(t)$ is the number of jobs that activity j could process in $[0, t]$ if the associated server worked continuously and exclusively on the associated class of jobs during this time interval. The interarrival time sequences $\{u_i(\ell), \ell = 1, 2, \dots\}$, $i \in \mathcal{I}$, and service time sequences $\{v_j(\ell), \ell = 1, 2, \dots\}$, $j \in \mathcal{J}$, are all assumed to be mutually independent.

2.3 Scheduling Control and Performance Measures

Scheduling control is exerted by specifying a \mathbf{J} -dimensional allocation process $T = \{T(t), t \geq 0\}$ where

$$T(t) = (T_1(t), \dots, T_{\mathbf{J}}(t))' \quad \text{for } t \geq 0, \quad (2.8)$$

and $T_j(t)$ is the cumulative amount of service time devoted to activity $j \in \mathcal{J}$ by the associated server $k(j)$ in the time interval $[0, t]$. Now T must satisfy certain properties that go along with its interpretation. Indeed, one could give a discrete-event type description of the properties that T must have, including any system specific constraints such as no preemption of service. However, for our analysis, we shall only need the properties of T described in (2.12)–(2.17) below.

Let

$$I(t) = \mathbf{1}t - \mathbf{A}T(t), \quad t \geq 0, \quad (2.9)$$

where $\mathbf{1}$ is the \mathbf{K} -dimensional vector of all ones. Then for each $k \in \mathcal{K}$, $I_k(t)$ is interpreted as the cumulative amount of time that server k has been idle up to time t . A natural constraint on T is that each component of the cumulative idletime process I must be continuous and non-decreasing. This immediately implies the property that each component of T is Lipschitz continuous with Lipschitz constant less than or equal to one. For each $j \in \mathcal{J}$, $S_j(T_j(t))$ is interpreted as the number of complete jobs processed by activity j in $[0, t]$. For $i \in \mathcal{I}$, let

$$Q_i(t) = A_i(t) - \sum_{j=1}^{\mathbf{J}} \mathbf{C}_{ij} S_j(T_j(t)), \quad t \geq 0, \quad (2.10)$$

which we write in vector form (with a slight abuse of notation for $S(T(t))$) as

$$Q(t) = A(t) - \mathbf{C}S(T(t)), \quad t \geq 0. \quad (2.11)$$

Then $Q_i(t)$ is interpreted as the number of class i jobs that are either in queue or “in progress” (i.e., being served or in suspension) at time t . We regard Q and I as performance measures for our system.

We shall use the following minimal set of properties of any scheduling control T with associated queue length process Q and idletime process I . For all $i \in \mathcal{I}$, $j \in \mathcal{J}$, $k \in \mathcal{K}$,

$$T_j(t) \in \mathcal{F} \text{ for each } t \geq 0, \quad (2.12)$$

$$T_j(\cdot) \text{ is Lipschitz continuous with Lipschitz constant less than or equal to one, non-decreasing, and } T_j(0) = 0, \quad (2.13)$$

$$I_k(\cdot) \text{ is continuous, non-decreasing, and } I_k(0) = 0, \quad (2.14)$$

$$Q_i(t) \geq 0 \text{ for all } t \geq 0. \quad (2.15)$$

For later reference, we collect here the queueing system equations satisfied by Q

and I :

$$Q(t) = A(t) - \mathbf{C}S(T(t)), \quad t \geq 0, \quad (2.16)$$

$$I(t) = \mathbf{1}t - \mathbf{A}T(t), \quad t \geq 0, \quad (2.17)$$

where Q , T and I satisfy properties (2.12)–(2.15). We emphasize that these are descriptive equations satisfied by the queueing system, given \mathbf{C} , \mathbf{A} , A , S and a control T , which suffice for the purposes of our analysis. In particular, we do not intend them to be a complete discrete-event type description of the dynamics.

Remark. The reader might expect that T should satisfy some additional non-anticipating property. Although this is a reasonable assumption to make, and indeed the policy we propose in Chapter 6 satisfies such a condition, we have not restricted T a priori in this way. Indeed, we shall see that, for the parallel server system under the complete resource pooling condition, our policy is asymptotically optimal even when anticipating policies are allowed. This is related to the fact that the Brownian control problem has a so-called “pathwise solution”, cf. [16].

The cost function we shall use involves linear holding costs associated with the expense of holding jobs of each class in the system until they have completed service. We defer the precise description of this cost function to the next chapter, since it is formulated in terms of normalized queue lengths, where the normalization is in diffusion scale. Indeed, in the next chapter, we describe the sequence of parallel server systems to be used in formulating the notion of heavy traffic asymptotic optimality.

3

Sequence of Systems, Heavy Traffic, and the Cost Function

For the parallel server system described in the last chapter, the problem of finding a control policy that minimizes a cost associated with holding jobs in the system is notoriously difficult. One possible means for discriminating between policies is to look for policies that outperform others in some asymptotic regime. Here we regard the parallel server system as a member of a sequence of systems indexed by r that is approaching heavy traffic (this notion is defined below). In this asymptotic regime, the queue length process is normalized with diffusive scaling – this corresponds to viewing the system over long intervals of time of order r^2 (where r will tend to infinity in the asymptotic limit) and regarding a single job as only having a small contribution to the overall cost of storage, where this is quantified to be of order $1/r$. The setup in this chapter is a generalization of that used in [2] to multiserver systems.

3.1 Sequence of Systems and Large Deviation Assumptions

Consider a sequence of parallel server systems indexed by r , where r tends to infinity through a sequence of values in $[1, \infty)$. These systems all have the same basic structure as that described in Chapter 2, except that the arrival and service rates, scheduling control policy, and form of the cost function (which is defined below in Section 3.3) are allowed to depend on r . Accordingly, we shall indicate the dependence of relevant parameters and processes on r by appending a superscript to them. We assume that the interarrival and service times are given for each $r \geq 1$, $i \in \mathcal{I}$, $j \in \mathcal{J}$, by

$$u_i^r(\ell) = \frac{1}{\lambda_i^r} \check{u}_i(\ell), \quad v_j^r(\ell) = \frac{1}{\mu_j^r} \check{v}_j(\ell), \quad \text{for } \ell = 1, 2, \dots, \quad (3.1)$$

where the $\check{u}_i(\ell)$, $\check{v}_j(\ell)$, do not depend on r , have mean one and squared coefficients of variation a_i^2 , b_j^2 , respectively. The sequences $\{\check{u}_i(\ell), \ell = 1, 2, \dots\}$, $\{\check{v}_j(\ell), \ell = 1, 2, \dots\}$, $i \in \mathcal{I}$, $j \in \mathcal{J}$ are all mutually independent sequences of i.i.d. random variables. (The above structure is a convenient means of allowing the sequence of systems to approach heavy traffic by simply changing arrival and service rates while keeping the underlying sources of variability $\check{u}_i(\ell)$, $\check{v}_j(\ell)$ fixed. For a first reading, the reader may like to simply choose $\lambda^r = \lambda$ and $\mu^r = \mu$ for all r . Indeed, that simplification is made in the paper [39]. This type of setup has been used previously by others in treating heavy traffic limits, see e.g., Peterson [30].)

We make the following assumption on the first order parameters associated with our sequence of systems.

Assumption 3.1.1 *There are vectors $\lambda \in \mathbb{R}_+^{\mathbf{I}}$, $\mu \in \mathbb{R}_+^{\mathbf{J}}$, such that*

- (i) $\lambda_i > 0$ for all $i \in \mathcal{I}$, $\mu_j > 0$ for all $j \in \mathcal{J}$,
- (ii) $\lambda^r \rightarrow \lambda$, $\mu^r \rightarrow \mu$, as $r \rightarrow \infty$.

In addition, we make the following exponential moment assumptions to ensure that certain *large deviation estimates* hold for the renewal processes $A_i^r, i \in \mathcal{I}$, and $S_j^r, j \in \mathcal{J}$ (cf. Lemma 7.6.1 below and Appendix A in [2]).

Assumption 3.1.2 For $i \in \mathcal{I}, j \in \mathcal{J}$, and all $\ell \geq 1$, let

$$u_i(\ell) = \frac{1}{\lambda_i} \check{u}_i(\ell), \quad v_j(\ell) = \frac{1}{\mu_j} \check{v}_j(\ell). \quad (3.2)$$

Assume that there is a non-empty open neighborhood \mathcal{O}_0 of $0 \in \mathbb{R}$ such that for all $l \in \mathcal{O}_0$,

$$\Lambda_i^a(l) \equiv \log \mathbf{E}[e^{lu_i(1)}] < \infty, \text{ for all } i \in \mathcal{I}, \text{ and} \quad (3.3)$$

$$\Lambda_j^s(l) \equiv \log \mathbf{E}[e^{lv_j(1)}] < \infty, \text{ for all } j \in \mathcal{J}. \quad (3.4)$$

Note that (3.3) and (3.4) hold with ℓ in place of 1 for all $\ell = 1, 2, \dots$, since $\{u_i(\ell), \ell = 1, 2, \dots\}, i \in \mathcal{I}$, and $\{v_j(\ell), \ell = 1, 2, \dots\}, j \in \mathcal{J}$, are each sequences of i.i.d. random variables.

Remark. This finiteness of exponential moments assumption allows us to prove asymptotic optimality of a threshold policy with thresholds of order $\log r$. We conjecture that this condition could be relaxed to a sufficiently high finite moment assumption, and our method of proof would still work, provided larger thresholds are used to allow for sufficient accomodation of larger deviations of the primitive renewal processes from their rate processes. In this case, Lemma 7.6.1 would need to be modified to use estimates based on sufficiently high finite moments, rather than exponential moments (cf. [4]). Here we have chosen the tradeoff of smaller thresholds and exponential moment assumptions, rather than larger thresholds and certain finite moment assumptions.

3.2 The Fluid Model and Heavy Traffic

There is no conventional notion of heavy traffic for our model, since the nominal (or average) load on a server depends on the scheduling policy. Harrison [14] (see

also Laws [27] and Harrison-Van Mieghem [17]) has proposed a notion of heavy traffic for stochastic networks with scheduling control. Following [39], we attempt to motivate that definition here (in the context of our parallel server system) via desired behavior of an associated fluid model. This fluid model also plays a role in establishing asymptotic optimality of control policies.

Here we regard a flow as a continuous deterministic function. The *fluid model* corresponding to our sequence of parallel server systems is a formal limit under law of large numbers scaling in which the primitive stochastic processes A^r and S^r are replaced by linear flows which move at rates equal to the limiting average rates λ and μ , respectively. In addition, the processes Q^r , I^r and T^r are replaced by flows \bar{Q} , \bar{I} and \bar{T} which must satisfy the *fluid model equations* (cf. (2.16)–(2.17)):

$$\bar{Q}(t) = \lambda t - \mathbf{R}\bar{T}(t), \quad (3.5)$$

$$\bar{I}(t) = \mathbf{1}t - \mathbf{A}\bar{T}(t), \quad (3.6)$$

where $\mathbf{R} = \mathbf{C} \text{diag}(\mu)$ and for all i, j, k ,

$$\bar{T}_j(\cdot) \text{ is Lipschitz continuous with Lipschitz constant less than or } (3.7)$$

$$\text{equal to one, non-decreasing, and } \bar{T}_j(0) = 0,$$

$$\bar{I}_k(\cdot) \text{ is continuous, non-decreasing, and } \bar{I}_k(0) = 0, \quad (3.8)$$

$$\bar{Q}_i(t) \geq 0 \text{ for all } t \geq 0. \quad (3.9)$$

Equations (3.5)–(3.6) and (3.7)–(3.9) define the fluid model (with zero initial condition). A flow, \bar{T} , satisfying (3.5)–(3.6) and (3.7)–(3.9) will be called a *fluid control*. One might interpret \bar{T} as representing average (at a law of large numbers scale) or “nominal” allocations of service to the processing activities. For a given fluid control \bar{T} , we say the fluid system is *balanced* if the associated fluid “queue length” \bar{Q} does not change with time (cf. Harrison [11]). Here, since the system starts empty, that means $\bar{Q} \equiv 0$. In addition, we say the fluid system *incurs no idleness* (or all fluid servers are fully occupied) if $\bar{I} \equiv 0$, i.e., $\mathbf{A}\bar{T}(t) = \mathbf{1}t$ for all $t \geq 0$.

Definition 3.2.1 *The fluid model is in heavy traffic if the following two conditions hold:*

(i) *there is a unique fluid control \bar{T}^* under which the fluid system is balanced, and*

(ii) *under \bar{T}^* , the fluid system incurs no idleness.*

Remark. The question of how to perform a heavy traffic analysis of a model when the above conditions are not satisfied is an interesting open problem. (For a start, see the related work by Laws [26] and recent work of Harrison [15].)

Since any fluid control is differentiable at almost every time (by (3.7)), we can convert the above notion of heavy traffic into one involving the rates $x^*(t) = \dot{\bar{T}}^*(t)$, where “ $\dot{\cdot}$ ” denotes time derivative. This leads to the notion of heavy traffic as formulated by Harrison [14] in terms of the following linear program.

Linear Program

$$\text{minimize } \rho \text{ subject to } \mathbf{R}x = \lambda, \quad \mathbf{A}x \leq \rho \mathbf{1} \text{ and } x \geq 0. \quad (3.10)$$

Remark. Harrison [14] calls this the *static allocation problem*.

The following is used by Harrison to define heavy traffic.

Assumption 3.2.2 *There is a unique optimal solution (ρ^*, x^*, \cdot) of the linear program (3.10). Moreover, that solution is such that $\rho^* = 1$ and $\mathbf{A}x^* = \mathbf{1}$.*

Remark. It will turn out that x_j^* determines the average fraction of time that server $k(j)$ should devote to activity j . For this reason, x^* is called the *nominal allocation vector*.

The following lemma is proved in [39].

Lemma 3.2.3 *(Williams [39]) The fluid model is in heavy traffic if and only if Assumption 3.2.2 holds.*

Example 3.2.4 *Harrison and López [16] considered the system pictured in Figure 2.1 with*

$$\lambda = \left(\frac{1}{2}, 1, 2, 3\right)', \quad \mu = \left(\frac{5}{4}, \frac{5}{2}, 1, 2, 5, \frac{5}{2}, 4, \frac{4}{5}\right)'. \quad (3.11)$$

With these data, Assumption 3.2.2 holds, and the unique optimal solution of (3.10) has

$$x^* = \left(\frac{2}{5}, \frac{2}{5}, \frac{1}{5}, \frac{9}{10}, \frac{1}{10}, 1, 0, 0\right)'. \quad (3.12)$$

We impose the following heavy traffic assumption on our sequence of parallel server systems, henceforth.

Assumption 3.2.5 (*Heavy Traffic Assumption*) *For the sequence of parallel server systems defined in Section 3.1, assume that Assumption 3.2.2 holds, and that there is a vector $\theta \in \mathbb{R}^{\mathbf{I}}$ such that*

$$r(\lambda^r - \mathbf{R}^r x^*) \rightarrow \theta, \quad \text{as } r \rightarrow \infty, \quad (3.13)$$

where $\mathbf{R}^r = \mathbf{C} \text{diag}(\mu^r)$.

For the formulation of the Brownian control problem, it will be helpful to distinguish *basic activities* j which have a strictly positive nominal fluid allocation level x_j^* from *non-basic activities* j for which $x_j^* = 0$. By relabeling the activities if necessary, we may and do assume henceforth that $x_j^* > 0$ for $j = 1, \dots, \mathbf{B}$ and $x_j^* = 0$ for $j = \mathbf{B} + 1, \dots, \mathbf{J}$. Thus there are \mathbf{B} basic activities and $\mathbf{J} - \mathbf{B}$ non-basic activities.

3.3 Diffusion Scaling and the Cost Function

For a fixed r and control policy T^r , the associated queue length process $Q^r = (Q_1^r, \dots, Q_{\mathbf{I}}^r)'$ and idletime process $I^r = (I_1^r, \dots, I_{\mathbf{K}}^r)'$ are given by (2.16)–(2.17) where the superscript r needs to be appended to A , S , Q , I and T there. The diffusion scaled queue length and idletime processes are defined by

$$\hat{Q}^r(t) = r^{-1}Q^r(r^2t), \quad \hat{I}^r(t) = r^{-1}I^r(r^2t), \quad \text{for all } t \geq 0. \quad (3.14)$$

We consider an average cumulative discounted holding cost for the diffusion scaled queue length process and control T^r :

$$\hat{J}^r(T^r) = \mathbf{E} \left(\int_0^\infty e^{-\gamma t} h \cdot \hat{Q}^r(t) dt \right), \quad (3.15)$$

where $\gamma > 0$ is a fixed constant (discount factor) and $h = (h_1, \dots, h_{\mathbf{I}})'$, $h_i > 0$ for all $i \in \mathcal{I}$, is a constant vector of holding costs per unit time per unit of diffusion scaled queue length. Recall that “ \cdot ” denotes the dot product between two vectors.

To write equations for \hat{Q}^r, \hat{I}^r , we introduce centered and diffusion scaled versions \hat{A}^r, \hat{S}^r of the primitive processes A^r, S^r :

$$\hat{A}^r(t) = r^{-1} (A^r(r^2t) - \lambda^r r^2t), \quad \hat{S}^r(t) = r^{-1} (S^r(r^2t) - \mu^r r^2t), \quad (3.16)$$

a *deviation process* \hat{Y}^r (which measures normalized deviations of server time allocations from the nominal allocations given by x^*):

$$\hat{Y}^r(t) = r^{-1} (x^* r^2t - T^r(r^2t)), \quad (3.17)$$

and a fluid scaled allocation process \bar{T}^r :

$$\bar{T}^r(t) = r^{-2} T^r(r^2t). \quad (3.18)$$

On substituting the above into (2.16)–(2.17), we obtain

$$\hat{Q}^r(t) = \hat{A}^r(t) - \mathbf{C} \hat{S}^r(\bar{T}^r(t)) + r(\lambda^r - \mathbf{R}^r x^*)t + \mathbf{R}^r \hat{Y}^r(t), \quad (3.19)$$

$$\hat{I}^r(t) = \mathbf{A} \hat{Y}^r(t), \quad (3.20)$$

where by (2.13)–(2.15) we have

$$\hat{I}_k^r(\cdot) \text{ is continuous, non-decreasing and } \hat{I}_k^r(0) = 0, \text{ for all } k \in \mathcal{K}, \quad (3.21)$$

$$\hat{Q}_i^r(t) \geq 0 \text{ for all } t \geq 0 \text{ and } i \in \mathcal{I}. \quad (3.22)$$

On combining Assumption 3.1.1 with the finite variance and mutual independence of the stochastic primitive sequences of i.i.d. random variables $\{\tilde{u}_i(\ell)\}_{\ell=1}^\infty$,

$i \in \mathcal{I}$, $\{\check{v}_j(\ell)\}_{\ell=1}^{\infty}$, $j \in \mathcal{J}$, we may deduce from renewal process functional central limit theorems (cf. [20]) that

$$(\hat{A}^r, \hat{S}^r) \Rightarrow (\tilde{A}, \tilde{S}), \quad \text{as } r \rightarrow \infty, \quad (3.23)$$

where \tilde{A} , \tilde{S} are independent, \tilde{A} is an \mathbf{I} -dimensional driftless Brownian motion that starts from the origin and has a diagonal covariance matrix whose i^{th} diagonal entry is $\lambda_i a_i^2$, and \tilde{S} is a \mathbf{J} -dimensional driftless Brownian motion that starts from the origin and has diagonal covariance matrix whose j^{th} diagonal entry is $\mu_j b_j^2$.

4

Brownian Control Problem

Given the heavy traffic assumption of the previous chapter and the fluid model interpretation of this condition, to keep queue lengths from growing on average, it seems desirable to choose a control for the sequence of parallel server systems that asymptotically on average allocates service to the processing activities in accordance with the proportions given by x^* . At first one might be tempted to set $T^r(t) = x^*t$. However, this is not a valid control, since in general (2.15) will be violated due to the random fluctuations in the arrival process A^r and service process S^r . To see how to achieve the proportions x^* on average, and to do so in an optimal manner, we turn to a finer approximate model than the fluid model. This is called a Brownian model, with an associated Brownian control problem that may be regarded as a formal diffusion approximation to the queueing control problem. The relationship between the Brownian model and the fluid model is analogous to the relationship between the central limit theorem and the law of large numbers.

4.1 Brownian Model

To formally describe the Brownian model, we consider a sequence of parallel server systems as described in Section 3.1. Following the method proposed by

Harrison [10, 13, 16], one arrives at the following formal Brownian control problem approximation (under diffusive scaling) to the control problem for a parallel server system. The Brownian control problem described in this chapter is a slight variant of that used by Harrison and López [16]. In particular, we allow \tilde{Y} to anticipate the future of \tilde{X} . One can obtain this formulation by formally passing to the limit as $r \rightarrow \infty$ in the control problem for the r^{th} parallel server system. An important assumption in this formal procedure is that in fluid or law of large numbers scale, the allocation processes achieve the average levels for a balanced system in the heavy traffic limit, i.e., formally we have that as $r \rightarrow \infty$, $\bar{T}^r \Rightarrow \bar{T}^*$, where

$$\bar{T}^*(t) \equiv x^*t, \quad t \geq 0. \quad (4.1)$$

The Brownian motion \tilde{X} appearing in the Brownian control problem defined below is the formal limit in distribution of \hat{X}^r , where for $t \geq 0$,

$$\hat{X}^r(t) \equiv \hat{A}^r(t) - \mathbf{C}\hat{S}^r(\bar{T}^r(t)) + r(\lambda^r - \mathbf{R}^r x^*)t. \quad (4.2)$$

The functional central limit theorems for the independent renewal processes A^r, S^r (cf. (3.23)), a time change theorem (together with the assumption that $\bar{T}^r \Rightarrow \bar{T}^*$), and (3.13), are used to derive the covariance matrix and drift for the Brownian motion \tilde{X} . The control process \tilde{Y} in the Brownian control problem arises as a formal limit of the deviation processes \hat{Y}^r (cf. (3.17)), where convergence of $\hat{Y}^r(0)$ to $\tilde{Y}(0)$ may be not required. Recall from the end of Section 3.2 that the first \mathbf{B} components of \hat{Y}^r correspond to basic activities and the last $\mathbf{J} - \mathbf{B}$ components correspond to non-basic activities. The non-increasing assumption in property (4.7) below corresponds to the fact that $\hat{Y}_j^r(t) = -r^{-1}T_j^r(r^2t)$ is non-increasing whenever j is a non-basic activity. The initial conditions on \tilde{I} and \tilde{Y}_j , $j = \mathbf{B} + 1, \dots, \mathbf{J}$, in (4.6)–(4.7) are relaxed from those in the prelimit to allow for the possibility of an initial jump in the queue length in the Brownian control problem. (In fact, for the optimal solution of the Brownian control problem, under the complete resource pooling condition to be assumed later, such a jump will

not occur and then the Brownian control problem is equivalent to one in which $\tilde{I}(0) = 0, \tilde{Y}(0) = 0$.)

Definition 4.1.1 (Brownian control problem)

$$\text{minimize } \mathbf{E} \left(\int_0^\infty e^{-\gamma t} h \cdot \tilde{Q}(t) dt \right) \quad (4.3)$$

using a \mathbf{J} -dimensional control process $\tilde{Y} = (\tilde{Y}_1, \dots, \tilde{Y}_{\mathbf{J}})'$ such that

$$\tilde{Q}(t) = \tilde{X}(t) + \mathbf{R}\tilde{Y}(t) \quad \text{for all } t \geq 0, \quad (4.4)$$

$$\tilde{I}(t) = \mathbf{A}\tilde{Y}(t) \quad \text{for all } t \geq 0, \quad (4.5)$$

$$\tilde{I}_k \text{ is non-decreasing and } \tilde{I}_k(0) \geq 0, \quad \text{for all } k \in \mathcal{K}, \quad (4.6)$$

$$\tilde{Y}_j \text{ is non-increasing and } \tilde{Y}_j(0) \leq 0, \quad \text{for } j = \mathbf{B} + 1, \dots, \mathbf{J}, \quad (4.7)$$

$$\tilde{Q}_i(t) \geq 0 \text{ for all } t \geq 0, \quad i \in \mathcal{I}, \quad (4.8)$$

where \tilde{X} is an \mathbf{I} -dimensional Brownian motion that starts from the origin, has drift θ (cf. (3.13)) and a diagonal covariance matrix whose i^{th} diagonal entry is equal to $\lambda_i a_i^2 + \sum_{j=1}^{\mathbf{J}} \mathbf{C}_{ij} \mu_j b_j^2 x_j^*$.

5

Solution of the Brownian Control Problem via Workload

In this chapter, we summarize the results from Harrison [14], Harrison and Van Mieghem [17], Harrison and López [16], and Williams [39] (see also Bramson and Williams [5] for the validity of results in [14] for parallel server systems), concerning the reduction of the dimension of the Brownian control problem to an “equivalent workload formulation”, and a resulting solution of the Brownian control problem. For this, let

$$\mathcal{N} = \{\delta \in \mathbb{R}^{\mathbf{I}} : \delta = \mathbf{R}x, \mathbf{A}x = 0 \text{ and } x_N = 0\}, \quad (5.1)$$

where x_N is the $(\mathbf{J}-\mathbf{B})$ -dimensional vector consisting of the last $(\mathbf{J}-\mathbf{B})$ components of the \mathbf{J} -dimensional vector x . One may intuitively think of the elements of \mathcal{N} as reversible displacements in the sense that, in the Brownian control problem, when the “queue length” $\tilde{Q}(t)$ is strictly positive, one can adjust this queue length using a small displacement $\delta \in \mathcal{N}$ without incurring any additional “idleness” and without using any non-basic activities. Since \mathcal{N} is a vector space, such changes are “reversible”. (For further discussion of the interpretation of \mathcal{N} , we refer the reader to Harrison [14].) In a sense, the idea of the equivalent workload formulation is to focus on controlling non-reversible displacements of “queue length”, i.e., those in the orthogonal complement \mathcal{N}^\perp of \mathcal{N} .

5.1 Solution via Workload

The following result claimed in Harrison [14] and validated, in particular, for parallel server systems in [5], is a key to this. The statement involves the dual to the linear program (3.10) introduced in Section 3.2.

Dual Program

$$\text{maximize } y \cdot \lambda \quad \text{subject to } y' \mathbf{R} \leq z' \mathbf{A}, \quad z \cdot \mathbf{1} \leq 1 \quad \text{and} \quad z \geq 0. \quad (5.2)$$

Theorem 5.1.1 (Harrison [14], Bramson-Williams [5]) *In the dual program (5.2), let $(y^1, z^1), \dots, (y^{\mathbf{L}}, z^{\mathbf{L}})$ be the extreme points of the feasible set of solutions that satisfy $y^\ell \cdot \lambda = 1$, $\ell = 1, \dots, \mathbf{L}$, (i.e., the maximum in the dual program (5.2) is achieved at these extreme points). Then,*

$$\mathcal{N}^\perp = \text{span}\{y^1, \dots, y^{\mathbf{L}}\}. \quad (5.3)$$

We focus here on the case in which there is a unique optimal solution of the dual program (5.2) and hence \mathcal{N}^\perp is one-dimensional. Theorem 5.1.3 below, due to Harrison and López [16], provides a convenient characterization of this case. For the statement of their result, we need the following notion of communicating servers.

Definition 5.1.2 *Consider the graph \mathcal{G} in which servers and buffers form the nodes and (undirected) edges between nodes are given by basic activities. We say that “all servers communicate via basic activities” if, for each pair of servers, there is a path in \mathcal{G} joining all of the servers.*

Theorem 5.1.3 (Harrison-López [16]) *The following are equivalent:*

- (i) *the dual program (5.2) has a unique optimal solution (y^*, z^*) ,*
- (ii) *the number of basic activities \mathbf{B} is equal to $\mathbf{I} + \mathbf{K} - 1$,*
- (iii) *all servers communicate via basic activities.*

Example 5.1.4 *For Example 3.2.4, there is a unique solution of the dual program given by*

$$y^* = \left(\frac{1}{5}, \frac{1}{10}, \frac{1}{4}, \frac{1}{10} \right)', \quad z^* = \left(\frac{1}{4}, \frac{1}{2}, \frac{1}{4} \right)'. \quad (5.4)$$

Since $\lambda = \mathbf{R}x^*$ and $\lambda_i > 0$ for each i , each buffer i is connected to some server by a basic activity (i.e., there is a basic activity j such that $i(j) = i$). It follows that condition (iii) is equivalent to the condition that the entire graph \mathcal{G} is connected. The following corollary to Theorem 5.1.3 is proved in [39].

Corollary 5.1.5 *(Williams [39]) The graph \mathcal{G} is a tree if and only if the equivalent conditions of Theorem 5.1.3 hold.*

Henceforth we make the following assumption.

Assumption 5.1.6 *(Complete Resource Pooling Condition) The equivalent conditions (i)–(iii) of Theorem 5.1.3 hold.*

Let (y^*, z^*) be the unique optimal solution of (5.2). By complementary slackness, $((y^*)'\mathbf{R})_j = ((z^*)'\mathbf{A})_j$ for $j = 1, \dots, \mathbf{B}$. Let u^* be the $(\mathbf{J} - \mathbf{B})$ -dimensional vector of dual “slack variables” defined by $((y^*)'\mathbf{R} - (z^*)'\mathbf{A})_j + u_{j-\mathbf{B}}^* = 0$ for $j = \mathbf{B} + 1, \dots, \mathbf{J}$. Then the following lemma is proved in [16] (Corollary to Proposition 3) using the relation between the primal and dual linear programs.

Lemma 5.1.7 (Harrison-López [16]) *We have $y^* > 0$, $z^* > 0$, $u^* > 0$,*

$$(y^*)'\mathbf{R} = (z^*)'\mathbf{A} - [0' \ (u^*)'], \quad \text{and} \quad z^* \cdot \mathbf{1} = 1, \quad (5.5)$$

where $\mathbf{1}$ is the \mathbf{K} -dimensional vector of ones, $0'$ is a \mathbf{B} -dimensional row vector of zeros.

Now, we review a solution of the Brownian control problem obtained by Harrison and López [16]. For \tilde{Q} satisfying (4.4)–(4.8), define $\tilde{W} = y^* \cdot \tilde{Q}$, which Harrison

[14] calls the (Brownian) workload. Let \tilde{Y}_N be the $(\mathbf{J} - \mathbf{B})$ -dimensional process whose components are \tilde{Y}_j , $j = \mathbf{B} + 1, \dots, \mathbf{J}$. By Lemma 5.1.7 and (4.4)–(4.8),

$$\tilde{W}(t) = y^* \cdot \tilde{X}(t) + \tilde{V}(t) \quad \text{for all } t \geq 0, \quad (5.6)$$

where

$$\tilde{V} \equiv z^* \cdot \tilde{I} - u^* \cdot \tilde{Y}_N \text{ is non-decreasing and } \tilde{V}(0) \geq 0, \quad (5.7)$$

$$\tilde{W}(t) \geq 0 \text{ for all } t \geq 0. \quad (5.8)$$

Now, for each $t \geq 0$, since the holding cost vector $h > 0$ and $y^* > 0$, we have

$$h \cdot \tilde{Q}(t) = \sum_{i=1}^{\mathbf{I}} \left(\frac{h_i}{y_i^*} \right) y_i^* \tilde{Q}_i(t) \geq h^* \tilde{W}(t) \quad (5.9)$$

where

$$h^* \equiv \min_{i=1}^{\mathbf{I}} \left(\frac{h_i}{y_i^*} \right). \quad (5.10)$$

It is well-known and straightforward to see that any solution pair (\tilde{W}, \tilde{V}) of (5.6)–(5.8) must satisfy for all $t \geq 0$,

$$\tilde{V}(t) \geq \tilde{V}^*(t) \equiv \sup_{0 \leq s \leq t} \left(-y^* \cdot \tilde{X}(s) \right), \quad (5.11)$$

and hence $\tilde{W}(t) \geq \tilde{W}^*(t)$ where

$$\tilde{W}^*(t) = y^* \cdot \tilde{X}(t) + \tilde{V}^*(t). \quad (5.12)$$

The process \tilde{W}^* is a one-dimensional *reflected Brownian motion* driven by the one-dimensional Brownian motion $y^* \cdot \tilde{X}$, and \tilde{V}^* is its local time at zero (see e.g., [7], Chapter 8). In particular, \tilde{V}^* can have a point of increase at time t only if $\tilde{W}^*(t) = 0$.

Now, let i^* be a class index such that $h^* = h_{i^*}/y_{i^*}^*$, i.e., the minimum in (5.10) is achieved at $i = i^*$ (i.e., $h_{i^*}/y_{i^*}^* \leq h_i/y_i^*$ for all $i = 1, 2, \dots, \mathbf{I}$), and let k^* be a server that can serve class i^* via a basic activity. Note that neither i^* nor k^* need be unique in general. Then the following choices \tilde{Q}^* and \tilde{I}^* for \tilde{Q} and \tilde{I} ensure

that for each $t \geq 0$, properties (4.6)–(4.8) hold and the inequality in (5.9) is an equality with $\tilde{W}(t) = \tilde{W}^*(t)$ there:

$$\tilde{Q}_{i^*}^*(t) = \tilde{W}^*(t)/y_{i^*}^*, \quad \tilde{Q}_i^*(t) = 0 \text{ for all } i \neq i^*, \quad (5.13)$$

$$\tilde{I}_{k^*}^*(t) = \tilde{V}^*(t)/z_{k^*}^*, \quad \tilde{I}_k^*(t) = 0 \text{ for } k \neq k^*, \quad \tilde{Y}_N^* = \mathbf{0}. \quad (5.14)$$

A control process \tilde{Y}^* such that (4.4)–(4.8) hold with $\tilde{Q}^*, \tilde{Y}^*, \tilde{I}^*$ in place of $\tilde{Q}, \tilde{Y}, \tilde{I}$ there is given in [39]. It can be readily verified that this is an optimal solution for the Brownian control problem (cf. [16]) and the associated minimum cost is

$$J^* \equiv \mathbf{E} \left(\int_0^\infty e^{-\gamma t} h \cdot \tilde{Q}^*(t) dt \right) = h^* \mathbf{E} \left(\int_0^\infty e^{-\gamma t} \tilde{W}^*(t) dt \right). \quad (5.15)$$

The quantity J^* is finite and can be computed as in Section 5.3 of [9].

5.2 Interpreting the Solution

Now, even though the Brownian control problem can be analyzed exactly (as above), the solution obtained does not automatically translate to a policy for the sequence of parallel server systems. However, some desirable features are suggested by the form of (5.13)–(5.14), namely,

- (a) try to keep the bulk of the work in the class i^* with the lowest (or equal lowest) ratio of holding cost to workload contribution, i.e., the class i with the lowest value of h_i/y_i^* ,
- (b) try to ensure that the bulk of the idleness is incurred only when there is almost no work in the entire system, and
- (c) try to ensure that the bulk of the idletime is incurred by server k^* alone.

Harrison [12] has proposed a general scheme (called BIGSTEP) for obtaining candidate policies for a queueing control problem from a solution of the associated Brownian control problem. The policies obtained in this manner are so-called

discrete-review policies which allow review of the system status and changes in the control policy only at a fixed discrete set of times. For a two-server example (with Poisson arrivals, deterministic service times and particular values for λ, μ), a family of discrete-review policies was constructed and shown to be asymptotically optimal in Harrison [13]. Based on their solution of the Brownian control problem and the general scheme laid out by Harrison [12], Harrison and López [16] proposed the use of a family of discrete review policies for the multiserver problem considered here, but they did not prove asymptotic optimality of these policies. Recently, Ata and Kumar [1] have proved asymptotic optimality of a family of discrete review policies for open stochastic processing networks that include parallel server systems, under heavy traffic and complete resource pooling conditions.

Another approach to translation of solutions of Brownian control problems into viable policies has been proposed by Kushner et al. (see e.g., [23, 24, 25]). However this also involves discretization by way of numerical approximation. We note that, of the works by Kushner et al. mentioned above, the paper by Kushner and Chen [23] is the closest to the current one in that it considers a parallel server model. However, it is in a very different parameter regime, namely one that corresponds to heavy traffic but with *no resource pooling*.

Assuming the complete resource pooling condition, in the next chapter we describe a simple “continuous review” policy for the sequence of parallel server systems, which allows changes in the control to be made at random times and in particular at times when the system status changes. Our policy is a dynamic priority policy in which priorities for certain “transition” activities depend on the number of jobs in the associated class relative to certain threshold or “safety-stock” levels. Changes in the priorities only occur as a threshold is crossed. We prove in Chapter 10 that the proposed policy is asymptotically optimal. In [2], we have already proved that this is so for the special case of a two-server two-buffer system. An important feature of that proof was that the limiting (under diffusion scale) queue length and idleness processes were effectively one-dimensional, i.e., a form of state-space collapse occurred in the diffusion limit. A similar phenomenon occurs

here for our multiserver system (cf. Theorem 6.2.1).

6

Threshold Policy and Asymptotic Optimality Result

In this chapter, we describe a dynamic threshold policy which Williams [39] proposed as a candidate for an asymptotically optimal control policy for parallel server systems under the heavy traffic (Assumption 3.2.5) and complete resource pooling (Assumption 5.1.6) conditions. (A notion of asymptotic optimality is described formally at the end of this chapter.) This threshold policy takes advantage of the tree structure of the server-buffer graph \mathcal{G} (cf. Corollary 5.1.5). We note that there can be many asymptotically optimal policies. The one described here is simply proposed as one that is intuitively appealing and is easy to describe. Indeed, independently, Squillante et al. [32] proposed a tree-based threshold priority policy for a parallel server system. However, their policy appears to be different from that described here.

6.1 Threshold Policy

In reviewing the description of the threshold policy, it may help readers to keep the following intuitive characteristics in mind. The policy only involves the use of basic activities (and so in the following, the word “activity” will be synonymous

with “basic activity”). The aim of the policy is to achieve the desirable properties (a)–(c) listed in the previous chapter. A key to the description of the policy is a hierarchical structure of the server-buffer tree \mathcal{G} and an associated protocol for the dynamic allocation of class priorities at each server. This protocol is described in an iterative manner, working from the bottom of the tree up towards the root. A server tree \mathcal{S} , which results from suppressing the buffers in \mathcal{G} , is helpful in describing this iterative procedure. See Figure 6.1 for an example of a \mathcal{G} and associated \mathcal{S} .

In this paper, each tree with a distinguished root will be assumed to grow downwards from its root and the root will be at the highest level. The server k^* at the root is chosen to be one that can service a “cheapest” class i^* , i.e., a class that achieves the minimum in (5.10), via a basic activity. One may now think of placing the root at the highest level and letting the servers and buffers in \mathcal{G} (or just the servers in the server tree) cascade below it in a hierarchy of alternating levels of servers and buffers where lower levels are farther from the root. The idea behind the threshold policy is to keep servers below the root level busy the bulk of the time (indeed, they should only be rarely idle and their idletimes should vanish on diffusion scale as the heavy traffic limit is approached), while simultaneously preventing the queue lengths of all classes except i^* from growing appreciably on diffusion scale. Transition buffers which link one level of servers to those at the next highest level are used to achieve these two competing goals. We call the classes associated with these buffers *transition classes*. In brief, when the queue length for a transition class gets to or below its threshold, any service of that class by the server immediately above it in the tree is suspended and this causes temporary overloading (on average) of the servers below, which prevents these servers from incurring much idleness. When the queue length for the transition class builds up to a level above its threshold, then assistance from the higher server is again permitted and the servers below are temporarily underloaded (on average) and so queue lengths for classes serviced by these servers are prevented from growing too large. The effect of this policy is to allow the movement (via the transition

buffers) of excess work from lower level to higher level buffers (and eventually, by an upwards cascade, to the buffer for class i^*), while simultaneously keeping all servers busy the bulk of the time, unless the entire system is nearly empty and even then to ensure that the vast majority of the idleness is incurred by the server k^* at the root of the tree.

Some of the features of the structure of the server-buffer tree that have been touched on in the above discussion warrant further comment, as these features will play an integral role in the arguments to follow. In particular, the tree structure dictates that each buffer $i \in \mathcal{I}$ can have only one activity $a(i)$ servicing the buffer from above (here, $a(\cdot)$ is mnemonic for “activity above”). Furthermore, the server (server k , say) conducting activity $a(i)$ can have at most one activity $b(k)$ servicing a buffer immediately above the server in the tree. In other words, each buffer and server in the server-buffer tree \mathcal{G} can have at most one activity immediately above it. If buffer i is a non-transition class, then there are no servers servicing buffer i from below, i.e., buffer i has only one associated basic activity, namely $a(i)$.

Recall that we are assuming throughout that the Assumption 3.2.5 of heavy traffic and Assumption 5.1.6 of complete resource pooling both hold. In the following, when describing the server-buffer tree \mathcal{G} , the terms class and buffer will be used interchangeably and only basic activities will be utilized. In reviewing the description of the threshold policy, the reader may find it helpful to refer to Example 6.1.1 which follows the description.

Threshold Policy. To describe the threshold policy, we first focus on the server tree \mathcal{S} and imagine it arranged in levels with the root at the highest level. Consider a server at the lowest level. As a server within the server-buffer tree \mathcal{G} , this server is to service its classes according to a priority scheme that gives lowest priority to the class that is immediately above the server in \mathcal{G} . This class is also served by a server in the next level up in \mathcal{S} and so is a transition class. There will always be such a class unless the server is at the root of the tree. The priority ranking of the other classes that a server at the lowest level serves is not so important. These are all terminal classes in that there are no servers below them. Here, for concreteness,

we rank the classes so that for a given server, the lower numbered classes receive priority over the higher numbered ones. For future use, we place a threshold on the transition classes associated with each of the servers in the lowest level of \mathcal{S} .

Now go to the next level up the server tree \mathcal{S} . This level may have “terminal” server nodes and server nodes that lead to server nodes lower down the server tree. As servers in the server-buffer tree \mathcal{G} , each server at this level performs its activities in the following prioritized manner. Activities leading (via transition buffers) to server nodes lower down the tree are given highest priority (if there is more than one such activity, rank the activities so that activities serving lower numbered classes are served first). However, if the number of jobs in a transition class associated with such an activity is at or below the threshold for that class, service of that activity is suspended. The next priority is given to activities that service (terminal) classes that are only served by that server (again ranked according to a scheme that gives lower numbered buffers higher priority), and lowest priority is given to the activity serving the transition class leading to the next highest level of servers. This transition class should again have a threshold placed on it such that service of that class by the server at the next highest level is suspended when the number of jobs in the class is equal to or below the threshold level. If two or more servers simultaneously begin to serve a particular transition buffer, a tie breaking rule is used to decide which server takes a job first. For concreteness we suppose that the lowest numbered server shall select a job before the next higher numbered server, and so on. Note that two servers in the same level cannot both serve the same buffer below them since \mathcal{G} is a tree.

This process is repeated until the root of the server tree is reached. At the root, the same procedure is applied as for lower levels in the tree, except that there are no activities above the root server in the server-buffer tree and an overriding rule is that lowest priority is given to the “cheapest” class i^* .

Threshold Sizes. For each r , the size of the thresholds in the r^{th} parallel server system is of order $\log r$. However, while each threshold is of order $\log r$, for our proofs, the threshold sizes need to increase moderately as one moves up the tree

to compensate for an associated accumulation of stochastic variability. For an intuitive understanding of this, consider a transition buffer. The amount of time allocated by a server k to an activity that serves the buffer from below depends on the other (higher priority) activities performed by server k (servicing buffers from the level immediately below that server). The times that these higher priority activities are in use depend on their associated buffer levels which may in turn depend on activities farther down the tree, and so on. The hierarchical dependence just described may cause the allocation processes associated with activities located farther up the tree to deviate more from their long term averages than their counterparts below. Larger threshold values for the transition buffers above will allow related activities more time to approach their long term averages before the associated queue lengths approach zero or twice the threshold size. For similar reasons, the threshold size for a buffer belonging to a group of transition buffers served by one server from above, should be larger the lower the priority of the buffer. To facilitate the description of this increase in threshold sizes, in section 7.1 we show how the buffers can be renumbered, so that higher priority buffers for a given server have lower numbers and buffers higher up the tree are assigned higher numbers. This renumbering does not change the threshold policy, it simply allows us to streamline its description. In particular, under this scheme, the size of the threshold for each transition buffer increases with its numbering. The details of this scheme are provided in the next chapter.

Remark. Note that our threshold policy prescribes that a threshold be placed on class i^* if that class is a transition class. We conjecture that the policy which removes the threshold in this case is also asymptotically optimal, but we keep it here as a means to simplify our proof. In particular, servers below i^* may experience significant idletime if the threshold is removed. Our threshold policy also involves preemption of service. There is a corresponding policy without preemption that we conjecture has the same behavior in the heavy traffic limit, since in that regime a maximum of \mathbf{J} jobs (in suspension or not) should not impact the asymptotic behavior of the system.




Figure 6.1: The server tree (on the left) and server-buffer tree (on the right) for the three-server network of Figure 2.1 with the data of Example 6.1.1.

Example 6.1.1 Consider the system pictured in Figure 2.1 with the data described in Example 3.2.4, and suppose that $h = (1, \frac{3}{5}, 2, 1)'$. Then $h^* = 5$, $i^* = 1$ and $k^* = 1$, cf. [16, 39]. The associated server tree and server-buffer tree are pictured in Figure 6.1.

6.2 Asymptotic Optimality

In Chapters 8 and 9, we prove the following limit theorem for a certain sequence of threshold policies, $\{T^{r,*}\}$. For each r , $T^{r,*}$ is a threshold policy as described above, where the thresholds on the transition buffers are of order $\log r$. The specific threshold sizes, which, as mentioned above, increase as one moves across and up the tree, are specified precisely in the next chapter (cf. (7.4)).

Theorem 6.2.1 Consider the sequence of parallel server systems indexed by r , where the r^{th} system operates under the threshold policy $T^{r,*}$ described above. Then the associated normalized queue length and idletime processes satisfy

$$(\hat{Q}^r, \hat{I}^r) \Rightarrow (\tilde{Q}^*, \tilde{I}^*), \quad \text{as } r \rightarrow \infty, \quad (6.1)$$

where $\tilde{Q}_i^* = \mathbf{0}$ for all $i \in \mathcal{I} \setminus \{i^*\}$, $\tilde{I}_k^* = \mathbf{0}$ for all $k \in \mathcal{K} \setminus \{k^*\}$, $\tilde{Q}_{i^*}^*$ is a one-dimensional reflected Brownian motion that starts from zero and has drift $(y^* \cdot \theta) / y_{i^*}^*$ and variance $\sum_{i=1}^{\mathbf{I}} (y_i^*)^2 (\lambda_i a_i^2 + \sum_{j=1}^{\mathbf{J}} \mathbf{C}_{ij} \mu_j b_j^2 x_j^*) / (y_{i^*}^*)^2$, and $\tilde{I}_{k^*}^*$ is a specific multiple of the local time at the origin of $\tilde{Q}_{i^*}^*$ (In fact, $\tilde{Q}_{i^*}^*$, $\tilde{I}_{k^*}^*$ are equivalent in law to the processes given by (5.11)–(5.14)).

Recall the definitions of J^* and \hat{J}^r from (5.15) and (3.15), respectively. The following theorem is the main result of this paper. It is proved in Chapter 10 using Theorem 6.2.1. It shows that J^* is the best that one can achieve asymptotically and that this asymptotically minimal cost is achieved by the sequence of dynamic threshold policies $\{T^{r,*}\}$. Thus we conclude that our sequence of threshold policies $\{T^{r,*}\}$ is asymptotically optimal.

Theorem 6.2.2 *Suppose that $\{T^r\}$ is any sequence of scheduling control policies (one for each member of the sequence of parallel server systems). Then*

$$\liminf_{r \rightarrow \infty} \hat{J}^r(T^r) \geq J^* = \lim_{r \rightarrow \infty} \hat{J}^r(T^{r,*}), \quad (6.2)$$

and $J^* < \infty$.

Remark. The notion of asymptotic optimality used here is also used for example in Puhalskii-Reiman [31].

7

Preliminaries and Description of the Proof

In this chapter, we will precisely specify the threshold sizes used for $\{T^{r,*}\}$, give some preliminary definitions and results, and outline the proofs of Theorems 6.2.1 and 6.2.2.

7.1 The Server-Buffer Tree \mathcal{G} : Layers and Buffer Renumbering

Recall from Chapter 5 that the graph \mathcal{G} , consisting of server and buffer nodes linked by basic activity edges, is a tree (cf. Corollary 5.1.5). In the sequel, to facilitate our proof, we will decompose the server-buffer tree, \mathcal{G} , somewhat differently from how we described it in Chapter 6. We say that the server-buffer tree consists of one or more *layers*, where a layer consists of a server level along with the buffer level immediately below it. (At the lowest layer, the buffer level may be empty.) Activities that serve the buffers (from above and below) in a particular layer are also considered part of that same layer. We enumerate the layers from 1 to l^* , l^* being the top layer of the tree consisting of a single server, server k^* , the buffers

served by this server (including buffer i^*), and the corresponding activities. (The activities which serve these buffers from above are all performed by server k^* .) Recall from Chapter 2 that \mathcal{I} , \mathcal{K} , and \mathcal{J} index job classes, servers, and activities, respectively.

Convention. *Since our threshold policy only uses basic activities, to simplify notation, in this chapter and the next (Chapters 7 and 8) only, the index set \mathcal{J} will just include the basic activities $1, 2, \dots, \mathbf{B}$. Then \mathbf{J} will be the same as \mathbf{B} .*

Figure 7.1 depicts two layers l ($l > 1$) and $l - 1$ in the server-buffer tree. We refer to the server level in layer l as \mathcal{K}^l and the buffer level as \mathcal{I}^l . Consider a server k in layer l . The collection of buffers in layer l served by server k is denoted by $\underline{\mathcal{I}}_k$ (in Figure 7.1, $i \in \underline{\mathcal{I}}_k$).

Arrangement Convention. *If $k \neq k^*$, we assume, without loss of generality, that the buffers in $\underline{\mathcal{I}}_k$ are arranged in a manner which ensures that the transition buffers are positioned to the left of the non-transition buffers, and within the groups of transition and non-transition buffers, the lower numbered buffers are to the left of the higher numbered ones. If $k = k^*$, the arrangement in the previous sentence holds with the exception that buffer i^* is placed at the far right in level \mathcal{I}^l . (The simplest way to think of doing this rearrangement is to work from the top of the tree downwards.)*

The activity that serves buffer $i \in \underline{\mathcal{I}}_k$ using server k is labeled $a(i)$, and the activities that serve buffer i using servers from layer $l - 1$ are denoted by $\underline{\mathcal{J}}_i$, while the servers themselves are denoted by $\underline{\mathcal{K}}_i$ (the underscore indicates that the servers are “below” i while $a(\cdot)$ is a mnemonic for “above”). We will denote the collection of activities which service buffer i (activity $a(i)$ together with $\underline{\mathcal{J}}_i$) by \mathcal{J}_i . The servers in $\underline{\mathcal{K}}_i$ are in layer $l - 1$ and so repeating the above scheme in this layer, we have that the buffers served by server $k' \in \underline{\mathcal{K}}_i$ from above are indexed by $\underline{\mathcal{I}}_{k'}$, the transition buffers in $\underline{\mathcal{I}}_{k'}$ are to the left of the non-transition buffers, within the groups of transition and non-transition buffers, the lower numbered buffers are to the left of the higher numbered ones, the activity of serving buffer $i' \in \underline{\mathcal{I}}_{k'}$ by server

k' is called $a(i')$, and so on. Note also that for each buffer $i \in \mathcal{I}^1$, $\underline{\mathcal{J}}_i = \emptyset$, $\underline{\mathcal{K}}_i = \emptyset$, i.e., layer 1 contains no activities that serve buffers in this layer from below. In fact, we could have $\mathcal{I}^1 = \emptyset$.

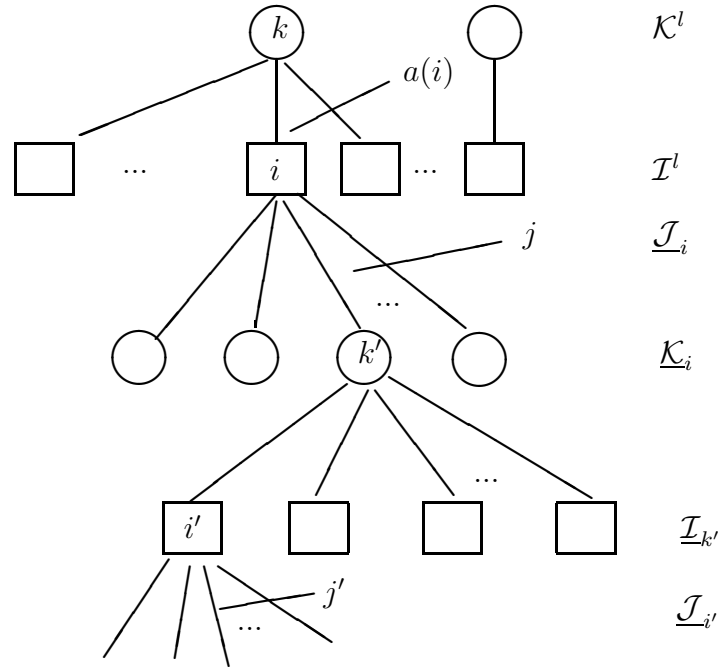


Figure 7.1: Layers l and $l - 1$ of the server-buffer tree

Renumbering Convention. *To simplify our proofs, we renumber the buffers in the server-buffer tree: starting from layer 1, if $\mathcal{I}^1 \neq \emptyset$, we enumerate the buffers in \mathcal{I}^1 from left to right, i.e., the buffer farthest to the left is buffer 1 and the one farthest to the right is buffer $|\mathcal{I}^1|$. Continuing with layer 2, the buffer farthest to the left is labeled $|\mathcal{I}^1| + 1$, and the one farthest to the right is labeled $|\mathcal{I}^1| + |\mathcal{I}^2|$, and so on ending with layer l^* . If $\mathcal{I}^1 = \emptyset$, we use the same methodology except that buffer 1 will then be the buffer farthest to the left in layer 2. Note that by our arrangement convention described above, for any server $k \neq k^*$, with this renumbering, the transition buffers in $\underline{\mathcal{I}}_k$ are numbered lower than the non-transition buffers in $\underline{\mathcal{I}}_k$ and hence have higher priority (for server k) and within the groups of transition*

and non-transition buffers in $\underline{\mathcal{I}}_k$, the higher priority buffers have lower numbers. If $k = k^*$, the previous statement holds with the exception that i^* is the highest numbered buffer and has the lowest priority among all buffers in $\underline{\mathcal{I}}_{k^*}$. It also follows from our numbering scheme that $i^* = \mathbf{I}$.

7.2 Threshold Sizes and Transient Nominal Activity Rates

It will be convenient to define the following *transient nominal activity rates* which apply to activities processed by a single server from above and that serve buffers at or to the left of a given buffer when that buffer is either a transition buffer above its threshold or a non-empty non-transition buffer. For $k \in \mathcal{K}$ such that $\underline{\mathcal{I}}_k \neq \emptyset$ and $i \in \underline{\mathcal{I}}_k$, let

$$x_i^+ = \sum_{\substack{i' \in \underline{\mathcal{I}}_k \\ i' \leq i}} x_{a(i')}^*, \quad (7.1)$$

and for each $i' \in \underline{\mathcal{I}}_k$, $i' \leq i$, let

$$\hat{x}_{i,i'} = \frac{x_{a(i')}^*}{x_i^+}. \quad (7.2)$$

(Note that this defines x_i^+ and $\hat{x}_{i,i}$ for all $i \in \mathcal{I}$ since each buffer is below exactly one server.) Since $x_{a(i')}^*$ determines the overall average fraction of time that server k should devote to activity $a(i')$, and since activities $a(i')$, $i' > i$, $i' \in \underline{\mathcal{I}}_k$, will be turned off during a period in which buffer i either exceeds its threshold if it is a transition buffer or is non-empty if it is a non-transition buffer, $\hat{x}_{i,i'}$, $i' \leq i$, $i' \in \underline{\mathcal{I}}_k$, might be interpreted as the average fraction of time that server k should devote to activity $a(i')$ during such a period of time. Note that

$$\sum_{\substack{i' \in \underline{\mathcal{I}}_k \\ i' \leq i}} \hat{x}_{i,i'} = 1, \quad (7.3)$$

and if $i \neq i^*$, $i' \in \underline{\mathcal{I}}_k$, $i' \leq i$, then $\hat{x}_{i,i'} > x_{a(i')}^*$ since then $x_i^+ < 1$ (if $k \in \mathcal{K}^l$ and $l \neq l^*$, server k serves a buffer from layer $l + 1$, by an activity $b(k)$ that is above k , and so by the heavy traffic condition, $x_{b(k)}^* + \sum_{i' \in \underline{\mathcal{I}}_k} x_{a(i')}^* = 1$, and if $k \in \mathcal{K}^{l^*}$, then $k = k^*$ and $\sum_{i' \in \underline{\mathcal{I}}_{k^*}} x_{a(i')}^* = 1$, where $i' \leq i^*$ for all i' in $\underline{\mathcal{I}}_{k^*}$). If buffer i is a transition buffer, during a period in which it is above its threshold, or if it is a non-transition buffer, during a period in which it has jobs waiting or in service, the activities served by server k that have lower priority than $a(i)$ will not be in use. Consequently, the higher priority activities will (on average) be on for a higher proportion of time during this period than given by the long run averages $x_{a(i')}^*$, $i' \leq i$, $i' \in \underline{\mathcal{I}}_k$. The constants $\hat{x}_{i,i'}$ are used to account for these higher anticipated average allocations.

We now define the size of the thresholds to be used with our threshold policy. For each $r \geq 1$, let $L_0^r = \lceil c \log r \rceil$ for a sufficiently large constant c . The minimum size of c is determined by the proofs of Lemmas 8.3.1–8.3.4 and Theorem 6.2.2 (see the remark below). For $1 \leq i \leq \mathbf{I}$, let

$$L_i^r = \left\lceil \frac{L_{i-1}^r}{\epsilon_{i-1}^3} \right\rceil, \quad (7.4)$$

where $\{\epsilon_i\}_{i=0}^{\mathbf{I}-1}$ is defined as follows. We also define a constant $\epsilon_{\mathbf{I}}$ in this process. First we choose $\hat{\epsilon} > 0$ such that

$$\hat{\epsilon} < \min \left\{ \frac{d_{\min}^{\delta} \delta_{\min} \mu_{\min} \lambda_{\min} x_{\min}^*}{2048(\mathbf{J} + 2) \mu_{\max} \lambda_{\max} \delta_{\max} \mu_{\text{sum}}}, \frac{\prod_{i=1}^{\mathbf{I}} \gamma_i}{\mathbf{I}^{\mathbf{I}}} \right\}, \quad (7.5)$$

where $\lambda_{\min} = \min\{1, \lambda_i : i \in \mathcal{I}\}$, $\mu_{\min} = \min\{1, \mu_j : j \in \mathcal{J}\}$, $\lambda_{\max} = \max\{1, \lambda_i : i \in \mathcal{I}\}$, $\mu_{\max} = \max\{1, \mu_j : j \in \mathcal{J}\}$, $d_{\min}^{\delta} = \min\{1, \lambda_i - \sum_{j \in \underline{\mathcal{I}}_i} x_j^* \mu_j : i \in \mathcal{I}\}$, $\delta_{\min} = \min\{1, \sum_{j \in \underline{\mathcal{I}}_i} x_j^* \mu_j + \hat{x}_{i,i} \mu_{a(i)} - \lambda_i : i \in \mathcal{I} \setminus \{i^*\}\}$, $\delta_{\max} = \max\{1, \sum_{j \in \underline{\mathcal{I}}_i} x_j^* \mu_j + \hat{x}_{i,i} \mu_{a(i)} - \lambda_i : i \in \mathcal{I}\}$, $x_{\min}^* = \min\{x_j^* : j \in \mathcal{J}\}$, $\mu_{\text{sum}} = \max\{1, \sum_{j \in \mathcal{J}} \mu_j\}$, and

$$\gamma_i = \frac{\mu_{a(i)}}{64 \sum_{j \in \mathcal{J}_i} \mu_j}, \quad 1 \leq i \leq \mathbf{I}. \quad (7.6)$$

Then we define by backward induction on i ,

$$\epsilon_{\mathbf{I}} = \frac{\hat{\epsilon}}{\mathbf{I}}, \quad \epsilon_i = \frac{\gamma_{i+1} \epsilon_{i+1}}{\mathbf{I}}, \quad 0 \leq i < \mathbf{I}. \quad (7.7)$$

(Note that $\gamma_i < 1$ for all $i \in \mathcal{I}$, and $\epsilon_i < \epsilon_{i+1} \leq \hat{\epsilon} < 1$, $i = 0, 1, 2, \dots, \mathbf{I} - 1$.) Finally, in the r^{th} system, if i is a transition buffer, we let L_i^r be the threshold for buffer i . If i is a non-transition buffer, L_i^r is not used to define a threshold, it is simply defined to make the iterative definition of L_i^r simple as i increases.

Remark. For our method of proof to work, the constant c must be sufficiently large. In the proofs of Lemmas 8.3.1–8.3.4 and of uniform integrability in the proof of Theorem 6.2.2, a means for determining a value c^* is described such that our method works provided $c > c^*$. This value is determined from several applications of large deviation estimates for the renewal processes associated with the interarrival and service time sequences (cf. Assumption 3.1.2). As in [2], we have not attempted to give a concise formula for c^* nor to optimize its value, since the relevant fact is that sufficiently large thresholds of order $\log r$ work and this order is the smallest for which our proof works. (For an analysis of the effects of different threshold sizes for some dynamic scheduling problems, see [34, 35].)

7.3 State Space Collapse Result and Outline of its Proof

A key element in the proof of Theorem 6.2.1 is to first show the following “state space collapse” result.

Theorem 7.3.1 *Consider the sequence of parallel server systems indexed by r , where the r^{th} system operates under the threshold policy, $T^{r,*}$, described in Chapter 6 and Section 7.2. Then*

$$\left(\hat{Q}_i^r, \hat{I}_k^r : i \in \mathcal{I} \setminus \{i^*\}, k \in \mathcal{K} \setminus \{k^*\} \right) \Rightarrow \mathbf{0} \quad \text{as } r \rightarrow \infty, \quad (7.8)$$

where $\mathbf{0}$ is the function in $D^{\mathbf{I}+\mathbf{K}-2}$ that remains at the origin of $\mathbb{R}^{\mathbf{I}+\mathbf{K}-2}$ for all time.

The idea behind the proof of this theorem is that, for sufficiently large r , under the threshold control $T^{r,*}$, and for a transition class $i \in \mathcal{I} \setminus \{i^*\}$, once the queue

length process Q_i^r , has first reached its threshold level L_i^r (cf. (7.4)), the normalized class i queue length process \hat{Q}_i^r “keeps close” to the normalized threshold level $\hat{L}_i^r \equiv L_i^r/r$, since when it is above this level, it is driven down towards the level at an “average” rate of $(\hat{x}_{i,i}\mu_{a(i)}^r + \sum_{j \in \underline{\mathcal{J}}_i} x_j^* \mu_j^r - \lambda_i^r)r > 0$, where $x_{a(i)}^* < \hat{x}_{i,i} \leq 1$ (cf. (7.2)), and when it is below the level, it is driven up towards the level at an average rate of $(\lambda_i^r - \sum_{j \in \underline{\mathcal{J}}_i} x_j^* \mu_j^r)r > 0$. For a non-transition class $i \neq i^*$, $\underline{\mathcal{J}}_i = \emptyset$ and \hat{Q}_i^r is driven down towards zero at an average rate of $(\hat{x}_{i,i}\mu_{a(i)}^r - \lambda_i^r)r > 0$. The claimed positivity of the quantities above holds for large r since $\mathbf{R}x^* = \lambda$ (cf. Assumptions 3.1.1 and 3.2.5) and $\hat{x}_{i,i} > x_{a(i)}^*$ for $i \neq i^*$.

Indeed, large deviation estimates for the primitive renewal processes and estimates for the cumulative allocation processes (cf. Lemma 8.3.1) are used below to show that the probability that, on any given compact time interval,

- (a) for a transition class $i \neq i^*$, \hat{Q}_i^r deviates by at least $\hat{L}_i^r - |\underline{\mathcal{J}}_i|r^{-1}$ from the threshold level \hat{L}_i^r , or
- (b) for a non-transition class $i \neq i^*$, \hat{Q}_i^r reaches the level \hat{L}_i^r , or higher,

goes to zero as $r \rightarrow \infty$. From (a) we then have, on any compact time interval, that, with high probability for large r , once Q_i^r reaches L_i^r , for a transition buffer $i \in \mathcal{I} \setminus \{i^*\}$, the idleness processes \hat{I}_k^r , for $k \in \underline{\mathcal{K}}_i$, cannot increase since the queue length process \hat{Q}_i^r stays away from zero. These results, as well as an estimate on the idletime of server $k \in \underline{\mathcal{K}}_i$ prior to Q_i^r reaching L_i^r , are stated formally in Theorem 8.1.1. The proof of this theorem proceeds by an induction on the buffers in the server-buffer tree, starting from $i = 1$ and iterating to buffer $i = \mathbf{I} - 1$. The induction setup is described in Section 8.3 and the proofs are in Sections 8.4–8.8. Theorem 7.3.1 follows from Theorem 8.1.1 using the fact that the threshold L_i^r being of order $\log r$ implies that $\hat{L}_i^r - |\underline{\mathcal{J}}_i|r^{-1}$ goes to zero as r goes to infinity.

Once Theorem 7.3.1 is established, one can show, using the model equations for queue length and idletime (cf. (2.16)–(2.17)), that the fluid scaled allocations

$\bar{T}^{r,*}(t) \equiv r^{-2}T^{r,*}(r^2t)$ associated with $T^{r,*}$ satisfy

$$\bar{T}^{r,*} \Rightarrow \bar{T}^*, \quad \text{as } r \rightarrow \infty, \quad (7.9)$$

where \bar{T}^* was defined in (4.1). One can then combine the above to prove Theorem 6.2.1. These results are proved in Chapter 9.

For the proof of Theorem 6.2.2, we first show (cf. Lemma 10.1.1) that for any subsequence that achieves the “liminf” on the left side of (6.2) as a limit and for which the “liminf” is finite, the fluid level asymptotic behavior described in (7.9) must hold along the subsequence. This, together with a pathwise lower bound for $h^r \cdot \hat{Q}^r$, where h^r is a perturbation of h given in (10.12), allows us to establish the inequality in (6.2). The equality in (6.2) follows from Theorem 6.2.1 after showing that a certain uniform integrability condition holds.

7.4 Residual Processes, Excursions, and Shifted Allocation Processes

Key to our proof of Theorem 7.3.1 is the behavior of what we call the *residual processes* defined for $i \in \mathcal{I}$, $r \geq 1$, $s \geq 0$, by

$$R_i^r(s) = \begin{cases} Q_i^r(s) - L_i^r, & \text{if } i \text{ is a transition class,} \\ Q_i^r(s), & \text{otherwise.} \end{cases} \quad (7.10)$$

For the case when $i \neq i^*$ is a transition buffer, the idea of our proof is to move the center of one’s attention to the threshold and to show that Q_i^r reaches the threshold level L_i^r relatively quickly and then “chatters” back and forth across this threshold, not frequently deviating “far” from it, so that Q_i^r rarely again goes as low as the level $|\underline{\mathcal{J}}_i|$, or as high as the level $2L_i^r - |\underline{\mathcal{J}}_i|$. When translated into the behavior of R_i^r , this means that we show that R_i^r reaches the level zero relatively quickly and then it chatters back and forth across the zero level, rarely deviating by as much as $\pm(L_i^r - |\underline{\mathcal{J}}_i|)$ from this level (cf. Theorem 8.1.1). If $i \neq i^*$ is a

non-transition buffer, we show that Q_i^r (or equivalently R_i^r) rarely goes above the level L_i^r . We have defined R_i^r for $i = i^*$ for convenience only.

For describing the excursions of R_i^r above zero, we introduce the following notation.

Definition 7.4.1 For $i \in \mathcal{I}$, let $\tau_{i,0}^r = \inf\{s \geq 0 : R_i^r(s) \geq 0\}$, $\tau_{i,1}^r = \inf\{s \geq \tau_{i,0}^r : R_i^r(s) \geq 1\}$, $\tau_{i,2}^r = \inf\{s \geq \tau_{i,1}^r : R_i^r(s) \leq 0\}$ and define recursively $\tau_{i,2n-1}^r = \inf\{s \geq \tau_{i,2n-2}^r : R_i^r(s) \geq 1\}$, $\tau_{i,2n}^r = \inf\{s \geq \tau_{i,2n-1}^r : R_i^r(s) \leq 0\}$, for $n = 2, 3, \dots$. We say that $[\tau_{i,2n-1}^r, \tau_{i,2n}^r]$ is the n^{th} “up” excursion interval for $R_i^r(\cdot)$. On $\{\tau_{i,2n-1}^r < \infty\}$, we let $\beta_{i,n}^r = \tau_{i,2n}^r - \tau_{i,2n-1}^r$, the length of this interval, and on $\{\tau_{i,2n-1}^r = \infty\}$, we let $\beta_{i,n}^r = 0$.

For describing the “down” excursion intervals of R_i^r when $i \in \mathcal{I}$ is a transition class, we define

$${}^dR_i^r \equiv -R_i^r = L_i^r - Q_i^r. \quad (7.11)$$

If i is not a transition class, we set ${}^dR_i^r \equiv \mathbf{0}$, and hence the following definition is only non-trivial for a transition class i .

Definition 7.4.2 For $i \in \mathcal{I}$, let ${}^d\tau_{i,1}^r = \inf\{s \geq \tau_{i,0}^r : {}^dR_i^r(s) \geq 1\}$, ${}^d\tau_{i,2}^r = \inf\{s \geq {}^d\tau_{i,1}^r : {}^dR_i^r(s) \leq 0\}$ and define recursively ${}^d\tau_{i,2n-1}^r = \inf\{s \geq {}^d\tau_{i,2n-2}^r : {}^dR_i^r(s) \geq 1\}$, ${}^d\tau_{i,2n}^r = \inf\{s \geq {}^d\tau_{i,2n-1}^r : {}^dR_i^r(s) \leq 0\}$, for $n = 2, 3, \dots$. We say that $[{}^d\tau_{i,2n-1}^r, {}^d\tau_{i,2n}^r]$ is the n^{th} “down” excursion interval for $R_i^r(\cdot)$ and on $\{{}^d\tau_{i,2n-1}^r < \infty\}$ we let ${}^d\beta_{i,n}^r = {}^d\tau_{i,2n}^r - {}^d\tau_{i,2n-1}^r$, and on $\{{}^d\tau_{i,2n-1}^r = \infty\}$ we define ${}^d\beta_{i,n}^r = 0$.

Estimates of the on-time of the activities in \mathcal{J}_i , that process class i jobs in the n^{th} (up/down) excursion interval for R_i^r , are needed to obtain estimates of the value of R_i^r during such an interval (cf. (8.40) and (8.63)). The estimates for the on-time of the activities in $\underline{\mathcal{J}}_i$, in turn depend on estimates for the on-time of activities farther down the tree, whereas the estimate for the on-time of activity $a(i)$ depends on estimates for the on-time of the (higher priority) activities that are served from above by the same server that serves buffer i from above. To

keep track of all relevant on-times, for each $i \in \mathcal{I}$, $j \in \mathcal{J}$, define shifted allocation processes for $n \geq 1$ and $s \geq 0$, by

$$T_{i,j}^{r,n}(s) = T_j^r(\tau_{i,2n-1}^r + s) - T_j^r(\tau_{i,2n-1}^r), \quad \text{on } \{\tau_{i,2n-1}^r < \infty\}, \quad (7.12)$$

$$dT_{i,j}^{r,n}(s) = T_j^r(d\tau_{i,2n-1}^r + s) - T_j^r(d\tau_{i,2n-1}^r), \quad \text{on } \{d\tau_{i,2n-1}^r < \infty\}, \quad (7.13)$$

and on $\{\tau_{i,2n-1}^r = \infty\}$ let $T_{i,j}^{r,n} \equiv \mathbf{0}$, on $\{d\tau_{i,2n-1}^r = \infty\}$ let $dT_{i,j}^{r,n} \equiv \mathbf{0}$. We have that on $\{\tau_{i,2n-1}^r < \infty\}$, $T_{i,j}^{r,n}$ measures the on-time of activity $j \in \mathcal{J}$ following an up-crossing to or above level one by R_i^r , and, for a transition class i , on $\{d\tau_{i,2n-1}^r < \infty\}$, $dT_{i,j}^{r,n}$ measures the on-time of activity j following a down-crossing to or below level minus one by R_i^r . Note that $dT_{i,j}^{r,n} \equiv \mathbf{0}$ if i is not a transition class.

7.5 Preliminaries on Stopped Arrival and Service Processes

For the proof of (8.1) in Theorem 8.1.1, we need to establish some preliminary results concerning the properties of the arrival and service processes stopped at certain hitting times, so that we can apply the results of Appendix A in [2] to shifted versions of these processes.

Let $\mathbb{N}_\infty = \mathbb{N} \cup \{\infty\}$. Consider $\mathbb{N}_\infty^{\mathbf{I}} \times \mathbb{N}_\infty^{\mathbf{J}}$ to be partially ordered by componentwise inequality, i.e., $(n, m) \leq (p, q)$ if and only if $n_i \leq p_i$, and $m_j \leq q_j$ for all $i \in \mathcal{I}$, $j \in \mathcal{J}$. Recall from Section 2.2 the definition, for the r^{th} system, of the cumulative interarrival time process for class $i \in \mathcal{I}$, ξ_i^r , and the cumulative service time process for activity $j \in \mathcal{J}$, η_j^r . For each $(p, q) \in \mathbb{N}_\infty^{\mathbf{I}} \times \mathbb{N}_\infty^{\mathbf{J}}$ let

$$\mathcal{F}_{pq}^r = \sigma\{\xi_i^r(\cdot \wedge (p_i + 1)), \eta_j^r(\cdot \wedge (q_j + 1)) : i \in \mathcal{I}, j \in \mathcal{J}\} \vee \mathcal{P}_0,$$

where \mathcal{P}_0 denotes the collection of \mathbf{P} -null sets in the complete probability space $(\Omega, \mathcal{F}, \mathbf{P})$. Then $\{\mathcal{F}_{pq}^r : (p, q) \in \mathbb{N}_\infty^{\mathbf{I}} \times \mathbb{N}_\infty^{\mathbf{J}}\}$ is a multiparameter filtration (cf. Ethier and Kurtz [8], p. 85).

Definition 7.5.1 A (multiparameter) stopping time relative to $\{\mathcal{F}_{pq}^r : (p, q) \in \mathbb{N}_\infty^{\mathbf{I}} \times \mathbb{N}_\infty^{\mathbf{J}}\}$ is a random variable \mathcal{T} taking values in $\mathbb{N}_\infty^{\mathbf{I}} \times \mathbb{N}_\infty^{\mathbf{J}}$ such that

$$\{\mathcal{T} = (p, q)\} \in \mathcal{F}_{pq}^r \text{ for all } (p, q) \in \mathbb{N}_\infty^{\mathbf{I}} \times \mathbb{N}_\infty^{\mathbf{J}}. \quad (7.14)$$

The σ -algebra associated with such a stopping time \mathcal{T} is

$$\mathcal{F}_{\mathcal{T}}^r = \{B \in \mathcal{F} : B \cap \{\mathcal{T} = (p, q)\} \in \mathcal{F}_{pq}^r \text{ for all } (p, q) \in \mathbb{N}_\infty^{\mathbf{I}} \times \mathbb{N}_\infty^{\mathbf{J}}\}. \quad (7.15)$$

Lemma 7.5.2 Suppose $r \geq 1$ is such that $L_0^r \geq \mathbf{J} + 1$. Then, for each $\iota \in \mathcal{I}$, $n \geq 1$,

$$\begin{aligned} \mathcal{T}_{n,\iota}^r &\equiv (A_i^r(\tau_{\iota,2n-1}^r), S_j^r(T_j^r(\tau_{\iota,2n-1}^r)) : i \in \mathcal{I}, j \in \mathcal{J}), \\ {}^d\mathcal{T}_{n,\iota}^r &\equiv (A_i^r({}^d\tau_{\iota,2n-1}^r), S_j^r(T_j^r({}^d\tau_{\iota,2n-1}^r)) : i \in \mathcal{I}, j \in \mathcal{J}) \end{aligned}$$

are (multiparameter) stopping times relative to the filtration $\{\mathcal{F}_{pq}^r : (p, q) \in \mathbb{N}_\infty^{\mathbf{I}} \times \mathbb{N}_\infty^{\mathbf{J}}\}$, where we adopt the convention that each of $A_i^r(\cdot)$, $S_j^r(T_j^r(\cdot))$, $i \in \mathcal{I}$, $j \in \mathcal{J}$, takes the value ∞ when evaluated at the time ∞ .

Remark. Here we need r large enough to simplify the proof of Lemma 7.5.2. Since, in the sequel, we let r approach infinity, this result will suffice for our purposes.

Proof. This lemma can be proved in a similar manner to Lemma 8.3 in [38]. For completeness, we include a proof in Appendix A. \square

Lemma 7.5.3 Let \mathcal{T} be a (multiparameter) stopping time relative to the filtration $\{\mathcal{F}_{pq}^r : (p, q) \in \mathbb{N}_\infty^{\mathbf{I}} \times \mathbb{N}_\infty^{\mathbf{J}}\}$. Let $\mathcal{T}^{\mathcal{I}}$ denote the first \mathbf{I} components of \mathcal{T} and $\mathcal{T}^{\mathcal{J}}$ denote the other \mathbf{J} ($= \mathbf{B}$) components of \mathcal{T} so that $\mathcal{T} = (\mathcal{T}^{\mathcal{I}}, \mathcal{T}^{\mathcal{J}})$. In the following, for notational convenience, we make the convention that each of $u_i^r(\cdot)$, $v_j^r(\cdot)$, $i \in \mathcal{I}$, $j \in \mathcal{J}$, takes the value ∞ when its argument takes the value ∞ . Then,

$$(u_i^r(\mathcal{T}_i^{\mathcal{I}} + 1), v_j^r(\mathcal{T}_j^{\mathcal{J}} + 1) : i \in \mathcal{I}, j \in \mathcal{J}) \in \mathcal{F}_{\mathcal{T}}^r, \quad (7.16)$$

and, on $\{\mathcal{T} \in \mathbb{N}^{\mathbf{I}} \times \mathbb{N}^{\mathbf{J}}\}$, the conditional distribution of $\{(u_i^r(\mathcal{T}_i^{\mathcal{I}} + n), v_j^r(\mathcal{T}_j^{\mathcal{J}} + n) : i \in \mathcal{I}, j \in \mathcal{J}), n = 2, 3, \dots\}$ given $\mathcal{F}_{\mathcal{T}}^r$ is the same as the (unconditioned)

distribution of the original family of i.i.d. random variables $\{(u_i^r(n), v_j^r(n) : i \in \mathcal{I}, j \in \mathcal{J}), n = 1, 2, \dots\}$.

Proof. For a proof with $|\mathcal{I}| = 1$, $|\mathcal{J}| = 2$, see the proof of Lemma 7.6 in [2]. The general proof is similar. \square

7.6 Large Deviation Bounds for Renewal Processes

The following lemma, which will be used extensively in proving state space collapse in Chapter 8, summarizes the results of the discussion in Appendix A in [2].

Lemma 7.6.1 *Let $\{\zeta(i)\}_{i=1}^\infty$ be a sequence of strictly positive independent random variables, where $\{\zeta(i)\}_{i=2}^\infty$ are identically distributed with finite mean $1/\nu$, for some $\nu \in (0, \infty)$, and $\zeta(1)$ may have a different distribution from $\zeta(i)$ for $i > 1$. Assume that there is a nonempty open neighborhood \mathcal{O} of $0 \in \mathbb{R}$ such that for $i = 2, 3, \dots$,*

$$\Lambda(l) \equiv \log \mathbf{E}(e^{l\zeta(i)}) < \infty \quad \text{for all } l \in \mathcal{O}. \quad (7.17)$$

Let the values of $r \geq 1$ range through a sequence that increases to infinity. For each r , let $\nu^r > 0$ and suppose that $\lim_{r \rightarrow \infty} \nu^r = \nu$. For each r and $i = 2, 3, \dots$, let

$$\zeta^r(i) = \frac{\nu}{\nu^r} \zeta(i). \quad (7.18)$$

Given $0 < \epsilon < \nu/2$, let $r_\epsilon \geq 1$ be such that for $r \geq r_\epsilon$,

$$|\nu^r - \nu| < \epsilon, \quad (7.19)$$

$$\frac{\nu^r}{\nu} \left(\frac{1}{\nu^r + \frac{\epsilon}{2}} \right) \leq \frac{1}{\nu} \left(\frac{1}{1 + \frac{\epsilon}{3\nu}} \right) < \frac{1}{\nu}, \quad (7.20)$$

$$\frac{1}{\nu} \left(1 + \frac{\epsilon}{2(\nu^r - \epsilon)} \right) \geq \frac{1}{\nu} \left(1 + \frac{\epsilon}{2\nu} \right) > \frac{1}{\nu}. \quad (7.21)$$

For each $r \geq 1$, $s \geq 0$, let

$$N^r(s) = \sup \left\{ n \geq 0 : \sum_{i=1}^n \zeta^r(i) \leq s \right\}. \quad (7.22)$$

Then for $r \geq r_\epsilon$, $s > 2/\epsilon$,

$$\begin{aligned} \mathbf{P}(N^r(s) > (\nu^r + \epsilon)s) &\leq \exp \left(-((\nu^r + \epsilon)s - 1) \Lambda^* \left(\frac{1}{\nu} \left(\frac{1}{1 + \frac{\epsilon}{3\nu}} \right) \right) \right) \\ &\leq \exp \left(-(\nu s - 1) \Lambda^* \left(\frac{1}{\nu} \left(\frac{1}{1 + \frac{\epsilon}{3\nu}} \right) \right) \right), \end{aligned} \quad (7.23)$$

and for $r \geq r_\epsilon$, $s \geq 0$,

$$\begin{aligned} \mathbf{P}(N^r(s) < (\nu^r - \epsilon)s) &\leq \exp \left(-(\nu^r - \epsilon)s \Lambda^* \left(\frac{1}{\nu} \left(1 + \frac{\epsilon}{2\nu} \right) \right) \right) \\ &\quad + \mathbf{P} \left(\zeta^r(1) > \frac{\epsilon}{2\nu^r} s \right) \\ &\leq \exp \left(-(\nu - 2\epsilon)s \Lambda^* \left(\frac{1}{\nu} \left(1 + \frac{\epsilon}{2\nu} \right) \right) \right) \\ &\quad + \mathbf{P} \left(\zeta^r(1) > \frac{\epsilon}{2\nu^r} s \right), \end{aligned} \quad (7.24)$$

where

$$\Lambda^*(x) \equiv \sup_{l \in \mathbb{R}} (lx - \Lambda(l)), \quad (7.25)$$

and where the values of the quantities involving Λ^* in the above are strictly positive. (The function Λ^* is called the Legendre-Fenchel transform of Λ .)

In the sequel, we will use the following result from Appendix A of [2] for estimating the probability involving $\zeta^r(1)$ in (7.24). Assuming $\zeta(1)$ has the same distribution as $\{\zeta(i)\}_{i=2}^\infty$, let $\zeta^r(1) = (\nu/\nu^r)\zeta(1)$, and $0 < l_0 \in \mathcal{O}$. Then, for each $r \geq 1$ and $s \geq 0$, for any $n \geq 1$,

$$\mathbf{P} \left(\max_{i=1}^n \zeta^r(i) > \frac{\epsilon}{2\nu^r} s \right) \leq n \exp \left(-\frac{l_0 \epsilon s}{2\nu} \right) \exp(\Lambda(l_0)). \quad (7.26)$$

8

Proof of State Space Collapse

Throughout this chapter, it is assumed that in the r^{th} parallel server system we use the allocation process $T^{r,*}$ associated with the threshold policy described in Chapter 6 and Section 7.2. To simplify notation, here we shall simply write T^r in place of $T^{r,*}$, since no other policy is considered in this chapter. The associated queue length and idletime processes will be denoted by Q^r, I^r , respectively.

8.1 Main Technical Result

Recall the definition of the residual processes (cf. (7.10)), and the role of c in the definition of L_0^r (cf. Section 7.2). For the following theorem, which is the main technical result of this chapter and from which Theorem 7.3.1 will follow, there is $c^* > 0$ such that the results hold provided the fixed constant c is greater than c^* (cf. the proof of Theorem 8.1.1).

Theorem 8.1.1 *For each $i \in \mathcal{I} \setminus \{i^*\}$, $k \in \mathcal{K}_i$, $t \geq 0$, and $\epsilon > 0$,*

$$\mathbf{P} \left(\sup_{\tau_{i,0}^r \leq s \leq r^2 t} |R_i^r(s)| \geq L_i^r - |\underline{\mathcal{J}}_i| \right) \rightarrow 0 \quad \text{as } r \rightarrow \infty, \quad (8.1)$$

$$\mathbf{P} (I_k^r(\tau_{i,0}^r) \geq r\epsilon) \rightarrow 0 \quad \text{as } r \rightarrow \infty, \quad (8.2)$$

$$\mathbf{P} (I_k^r(r^2 t) - I_k^r(\tau_{i,0}^r) > 0, \tau_{i,0}^r < r^2 t) \rightarrow 0 \quad \text{as } r \rightarrow \infty, \quad (8.3)$$

and, in addition, if i^* is a transition class, then for each $k \in \underline{\mathcal{K}}_{i^*}$, $t \geq 0$, and $\epsilon > 0$,

$$\mathbf{P}\left(\inf_{\tau_{i^*,0}^r \leq s \leq r^2 t} Q_{i^*}^r(s) \leq |\underline{\mathcal{J}}_{i^*}| \right) \rightarrow 0 \quad \text{as } r \rightarrow \infty, \quad (8.4)$$

$$\mathbf{P}\left(I_k^r(\tau_{i^*,0}^r) \geq r\epsilon\right) \rightarrow 0 \quad \text{as } r \rightarrow \infty, \quad (8.5)$$

$$\mathbf{P}\left(I_k^r(r^2 t) - I_k^r(\tau_{i^*,0}^r) > 0, \tau_{i^*,0}^r < r^2 t\right) \rightarrow 0 \quad \text{as } r \rightarrow \infty. \quad (8.6)$$

Here, $\tau_{i,0}^r = \inf\{s \geq 0 : Q_i^r(s) \geq L_i^r\}$ if class i is a transition class, $\tau_{i,0}^r \equiv 0$ if class i is a non-transition class, and $|\underline{\mathcal{J}}_i|$ is the number of basic activities that serve class i from below in the server-buffer tree (cf. Figure 7.1).

Here we have used the convention in (8.2)–(8.3) and (8.5)–(8.6) that $I_k^r(\tau_{i,0}^r) = \lim_{t \rightarrow \infty} I_k^r(t)$, on $\{\tau_{i,0}^r = \infty\}$, and in (8.1) (respectively, (8.4)) that the supremum (respectively, infimum) over an empty set is defined to equal $-\infty$ (respectively, ∞).

Remark. Although the results in Theorem 8.1.1 suffice for the proof of Theorem 7.3.1 and subsequently for Theorem 6.2.1, a refinement of Theorem 8.1.1, with estimates of the left members in (8.1)–(8.3), (cf. Theorem 8.3.5) is needed to establish certain uniform integrability used in the proof of Theorem 6.2.2.

8.2 Auxiliary Constants for the Induction Proof

In the proofs of Theorems 8.1.1 and 8.3.5, we will use the following constants which grow logarithmically with r . For $i \in \mathcal{I}$, $r \geq 1$, L_i^r , and ϵ_i as defined in (7.4) and (7.7), respectively, let

$$s_i^r = \frac{L_i^r - (|\underline{\mathcal{J}}_i| + 2)}{(\lambda_i^r + \epsilon_i)}, \quad (8.7)$$

$$t_i^r = \frac{8L_i^r}{\lambda_i - \sum_{j \in \underline{\mathcal{J}}_i} x_j^* \mu_j}. \quad (8.8)$$

Note that if i is a non-transition class then the denominator in (8.8) is equal to $\lambda_i > 0$ (by our convention that a sum over an empty set is zero), and if i is a transition

class then the denominator is also positive since $\lambda_i = x_{a(i)}^* \mu_{a(i)} + \sum_{j \in \underline{\mathcal{J}}_i} x_j^* \mu_j$ (cf. (3.10)). Let

$$s_0^r = t_0^r = d_{s_0}^r = L_0^r, \quad (8.9)$$

and, for each $i \in \mathcal{I}$, $r \geq 1$, let

$$d_{s_i}^r = \frac{L_i^r - (|\underline{\mathcal{J}}_i| + 2)}{\sum_{j \in \underline{\mathcal{J}}_i} (\mu_j^r + \epsilon_i)(x_j^* + \epsilon_i)}, \quad \text{if } i \text{ is a transition class,} \quad (8.10)$$

otherwise let $d_{s_i}^r = L_i^r$. Finally, for each $r \geq 1$, let

$$M^r = \max \left\{ s_i^r, t_i^r, d_{s_i}^r : i \in \mathcal{I} \right\}. \quad (8.11)$$

Several lemmas are used in establishing Theorems 8.1.1 and 8.3.5 (including Lemmas 8.3.1–8.3.4 of Section 8.3). For the proofs of these lemmas we require that $r \geq 1$ is large enough so that various relations involving the auxiliary constants defined above and the parameters for the r^{th} parallel server system hold. It is important to note that, for this, the size of r should not depend on the variable t referred to in those theorems and lemmas. We first require that r is large enough that

$$s_i^r \geq s_{i-1}^r, \quad t_i^r \geq t_{i-1}^r, \quad d_{s_i}^r \geq d_{s_{i-1}}^r, \quad \text{for all } i \in \mathcal{I}. \quad (8.12)$$

For each $i \in \mathcal{I}$, let

$$\delta_i = \sum_{j \in \underline{\mathcal{J}}_i} x_j^* \mu_j + \hat{x}_{i,i} \mu_{a(i)} - \lambda_i. \quad (8.13)$$

Then, $\delta_i > 0$ for $i \in \mathcal{I} \setminus \{i^*\}$ since $\lambda_i = \sum_{j \in \underline{\mathcal{J}}_i} x_j^* \mu_j + x_{a(i)}^* \mu_{a(i)} < \sum_{j \in \underline{\mathcal{J}}_i} x_j^* \mu_j + \hat{x}_{i,i} \mu_{a(i)}$, using the fact that $x_{a(i)}^* < \hat{x}_{i,i}$ for $i \neq i^*$ (cf. (7.2)). For the proof of part of Lemma 8.3.1 (associated with (I.1) and up excursions) we require that, for all $i \in \mathcal{I} \setminus \{i^*\}$,

$$|\lambda_i^r - \lambda_i| < \min \left\{ \frac{\epsilon_i}{4}, \frac{(\mu_{a(i)} - \epsilon_i) \epsilon_i}{32 |\underline{\mathcal{J}}_i|} \right\}, \quad (8.14)$$

$$|\mu_j^r - \mu_j| < \min \left\{ \frac{\epsilon_i}{4}, \frac{(\mu_{a(i)} - \epsilon_i) \epsilon_i}{32 |\underline{\mathcal{J}}_i|} \right\}, \quad \text{for all } j \in \underline{\mathcal{J}}_i, \quad (8.15)$$

$$\mu_j^r - \epsilon_i > \frac{4\lambda_i + \delta_i}{4\lambda_i + 2\delta_i} \mu_j \quad \text{for all } j \in \underline{\mathcal{J}}_i, \quad (8.16)$$

$$\mu_{a(i)}^r - \epsilon_i > \frac{8\lambda_i + 3\delta_i}{8\lambda_i + 4\delta_i} \mu_{a(i)}, \quad (8.17)$$

$$\lambda_i^r + \epsilon_i < \left(\frac{2\lambda_i}{2\lambda_i + \delta_i} \right) \left(\sum_{j \in \underline{\mathcal{J}}_i} x_j^* \mu_j + \hat{x}_{i,i} \mu_{a(i)} \right) = \lambda_i + \frac{\lambda_i \delta_i}{2\lambda_i + \delta_i}, \quad (8.18)$$

$$1 < \frac{\delta_i}{4\lambda_i + 2\delta_i} \hat{x}_{i,i} \mu_{a(i)} s_i^r. \quad (8.19)$$

(Note that (8.16)–(8.18) do not hold for $i = i^*$, for large $r \geq 1$, since, in this case, $\delta_{i^*} = 0$, $\epsilon_{i^*} > 0$, and $\lambda_{i^*}^r \rightarrow \lambda_{i^*}$, $\mu_j^r \rightarrow \mu_j$, $j \in \mathcal{J}_{i^*}$, as $r \rightarrow \infty$.)

For each $i \in \mathcal{I}$, define

$$d\delta_i = \lambda_i - \sum_{j \in \underline{\mathcal{J}}_i} x_j^* \mu_j > 0, \quad (8.20)$$

since $\lambda_i = \sum_{j \in \underline{\mathcal{J}}_i} x_j^* \mu_j + x_{a(i)}^* \mu_{a(i)} > \sum_{j \in \underline{\mathcal{J}}_i} x_j^* \mu_j \geq 0$, using the fact that $x_{a(i)}^* \mu_{a(i)} > 0$. For the proof of part of Lemma 8.3.1 (associated with (I.1) and down excursions) we require that, if i is a transition class,

$$\mu_j^r + \epsilon_i < \frac{8\lambda_i - 3d\delta_i}{8\lambda_i - 4d\delta_i} \mu_j, \quad \text{for all } j \in \underline{\mathcal{J}}_i, \quad (8.21)$$

$$\lambda_i^r - \epsilon_i > \frac{2\lambda_i}{2\lambda_i - d\delta_i} \sum_{j \in \underline{\mathcal{J}}_i} x_j^* \mu_j, \quad (8.22)$$

$$1 < \left(\frac{\lambda_i d\delta_i}{(4\lambda_i - d\delta_i)(2\lambda_i - d\delta_i)} \sum_{j \in \underline{\mathcal{J}}_i} x_j^* \mu_j \right) d_s^r. \quad (8.23)$$

(Note that (8.23) does not hold for a non-transition class i since the right hand side there is zero for all $r \geq 1$ by our convention that a sum over an empty index set is zero.)

For each $i \in \mathcal{I}$, let

$$\tilde{\epsilon}_i = \frac{\epsilon_i \sum_{j \in \underline{\mathcal{J}}_i} \mu_j}{1 - 2|\underline{\mathcal{J}}_i| \epsilon_i}, \quad \text{if } \underline{\mathcal{J}}_i \neq \emptyset, \quad (8.24)$$

and set $\tilde{\epsilon}_i = 1$ if $\underline{\mathcal{J}}_i = \emptyset$. To use Lemma 7.6.1, we need $\tilde{\epsilon}_i < \min\{\lambda_i/2, \mu_j/2 : j \in \underline{\mathcal{J}}_i\}$ for a transition class i ($\tilde{\epsilon}_i$ will only be used in this case). To see

that this condition is satisfied, we note that by (7.5) and (7.7), $\epsilon_i < 1/(4\mathbf{J}) < 1/(4|\underline{\mathcal{J}}_i|)$ so that $\tilde{\epsilon}_i < 2\epsilon_i \sum_{j \in \underline{\mathcal{J}}_i} \mu_j < \min\{\lambda_i/2, \mu_j/2\}$, for all $j \in \mathcal{J}$, since $\epsilon_i < \min\{\lambda_{\min}, \mu_{\min}\}/4\mu_{\text{sum}}$. Then, for the proof of part of Lemma 8.3.1 (associated with (I.2)), we require that

$$\lambda_i^r - \sum_{j \in \underline{\mathcal{J}}_i} x_j^* \mu_j^r \geq \frac{1}{2} \left(\lambda_i - \sum_{j \in \underline{\mathcal{J}}_i} x_j^* \mu_j \right), \quad (8.25)$$

$$|\lambda_i - \lambda_i^r| < \tilde{\epsilon}_i, \quad |\mu_j - \mu_j^r| < \tilde{\epsilon}_i, \quad \text{for all } j \in \underline{\mathcal{J}}_i. \quad (8.26)$$

For the proof of Lemma 8.3.4, for all $\iota \geq i$, $i \in \mathcal{I}$, $\iota \in \mathcal{I}$, we require that

$$s_\iota^r > \left[\frac{\epsilon_i}{4} \left(\mu_{a(i)} - \frac{\mu_{a(i)} \epsilon_i}{16} \right) \right]^{-1}, \quad \frac{\mu_{a(i)}^r}{\mu_{a(i)}} > \frac{1}{2}, \quad \frac{\lambda_\iota^r}{\lambda_\iota} < \frac{3}{2}, \quad (8.27)$$

$$\left| \lambda_i^r - x_{a(i)}^* \mu_{a(i)}^r - \sum_{j \in \underline{\mathcal{J}}_i} x_j^* \mu_j^r \right| < \frac{\mu_{a(i)} \epsilon_i}{32}, \quad (8.28)$$

$$L_i^r > \frac{\epsilon_i (|\underline{\mathcal{J}}_\iota| + 2) \max\{\mu_{a(i)}, \lambda_i - \sum_{j \in \underline{\mathcal{J}}_i} x_j^* \mu_j\}}{\min\{\lambda_\iota, \sum_{j \in \underline{\mathcal{J}}_i} (\mu_j + \epsilon_\iota)\}}, \quad (8.29)$$

and, in addition, if ι is a transition class,

$$\frac{\sum_{j \in \underline{\mathcal{J}}_\iota} (\mu_j + \epsilon_\iota) (x_j^* + \epsilon_\iota)}{\sum_{j \in \underline{\mathcal{J}}_\iota} (\mu_j^r + \epsilon_\iota) (x_j^* + \epsilon_\iota)} > \frac{1}{2}, \quad d_{S_\iota}^r \geq \frac{4}{\mu_{a(i)}^r \epsilon_i}. \quad (8.30)$$

For each $i \in \mathcal{I}$, let

$$\check{\epsilon}_i = \frac{(\mu_{a(i)} - \epsilon_i) \epsilon_i}{16\mathbf{J}} < \min \left\{ \frac{\lambda_i}{2}, \frac{\mu_j}{2} : j \in \mathcal{J}_i \right\}. \quad (8.31)$$

The inequality here holds by (7.5). In addition, if i is a transition class, let

$$\epsilon_{1,i} = \left(\sum_{j \in \underline{\mathcal{J}}_i} (\mu_j + \epsilon_i) \right)^{-1} \frac{\epsilon_i}{16} \left(\mu_{a(i)} - \frac{\mu_{a(i)} \epsilon_i}{16} \right), \quad (8.32)$$

and set $\epsilon_{1,i} = 1$ if i is not a transition class. In either case, we also define

$$\epsilon_{2,i} = \frac{\mu_{a(i)} \epsilon_i}{16} (x_{a(i)}^* + \epsilon_i)^{-1} < \frac{\mu_{a(i)}}{2}. \quad (8.33)$$

If $i \in \mathcal{I}$ is a transition class, for each $j \in \underline{\mathcal{J}}_i$ let

$$\epsilon_{3,j} = \frac{\mu_{a(i)}\epsilon_i}{16|\underline{\mathcal{J}}_i|}(x_j^* + \epsilon_{1,i})^{-1} < \frac{\mu_j}{2}. \quad (8.34)$$

The inequalities in (8.33) and (8.34) hold since $\epsilon_i < \min\{x_{\min}^*, \mu_{\min}x_{\min}^*/\mu_{\max}\}$, by (7.5). To apply the large deviation bounds of Lemma 7.6.1 in the proofs of Lemmas 8.3.1 and 8.3.4, we require that

$$L_0^r > \max \left\{ \frac{2}{\epsilon_i}, \frac{2}{\check{\epsilon}_i}, \frac{2}{(x_{a(i)}^* - \epsilon_i)\epsilon_{2,i}}, \frac{4}{(x_j^* + \epsilon_i)\epsilon_i}, \right. \\ \left. \frac{2}{(x_j^* + \epsilon_i)\tilde{\epsilon}_i}, \frac{2}{(x_j^* + \epsilon_{1,i})\epsilon_{3,j}} : i \in \mathcal{I}, j \in \underline{\mathcal{J}}_i \right\}. \quad (8.35)$$

Assumption 8.2.1 *We henceforth assume that $r^* \geq 1$ is defined such that for all $r \geq r^*$, the following hold:*

(i) *conditions (8.12), (8.14)–(8.19), (8.21)–(8.23), (8.25)–(8.30), and (8.35) hold for all $i \in \mathcal{I}$, $\iota \in \mathcal{I}$, with the proviso that (8.14)–(8.19) do not need to hold if $i = i^*$, (8.21)–(8.23) do not need to hold if i is a non-transition class, and (8.30) does not need to hold if ι is a non-transition class,*

(ii) *for each $i \in \mathcal{I}$, $j \in \underline{\mathcal{J}}_i$, (7.19)–(7.21) hold with*

(a) λ_i^r in place of ν^r , λ_i in place of ν , and any of ϵ_i , $\tilde{\epsilon}_i$, $\check{\epsilon}_i$, or $\check{\epsilon}_i/2$ in place of ϵ there,

(b) μ_j^r in place of ν^r , μ_j in place of ν , and any of ϵ_i , $\epsilon_i/2$, $\tilde{\epsilon}_i$, $\check{\epsilon}_i$, or $\epsilon_{3,j}$, in place of ϵ there, and

(c) $\mu_{a(i)}^r$ in place of ν^r , $\mu_{a(i)}$ in place of ν , and $\epsilon_{2,i}$ in place of ϵ there.

Remark. The condition before (7.19)–(7.21) that $0 < \epsilon < \nu/2$ is automatically satisfied for the choices of ϵ , ν in Assumption 8.2.1.

8.3 Induction Setup

In the sequel we will use induction on i to show that the following (I)–(II) hold for each $i \in \mathcal{I} \setminus \{i^*\}$, and (III) holds for $i = i^*$. Recall the definitions of r^* , s_i^r , t_i^r , d_s^r , M^r from Section 8.2. Note in particular that r^* is independent of t (see below).

(I) For all $r \geq r^*$, $t > 0$ satisfying $r^2 t \geq M^r$, and each $k \in \underline{\mathcal{K}}_i$,

$$(I.1) \quad \mathbf{P} \left(\sup_{\tau_{i,0}^r \leq s \leq r^2 t} |R_i^r(s)| \geq L_i^r - |\underline{\mathcal{J}}_i| \right) \\ \leq p_{1,i}(r^2 t) \left(C_{1,i}^{(1)} \exp(-C_{1,i}^{(2)} L_0^r) + C_{1,i}^{(3)} \exp(-C_{1,i}^{(4)} r^2 t) \right),$$

$$(I.2) \quad \mathbf{P} \left(I_k^r(\tau_{i,0}^r) \geq t_i^r \right) \\ \leq p_{2,i}(r^2 t) \left(C_{2,i}^{(1)} \exp(-C_{2,i}^{(2)} L_0^r) + C_{2,i}^{(3)} \exp(-C_{2,i}^{(4)} r^2 t) \right),$$

$$(I.3) \quad \mathbf{P} \left(I_k^r(r^2 t) - I_k^r(\tau_{i,0}^r) > 0, \tau_{i,0}^r < r^2 t \right) \\ \leq p_{3,i}(r^2 t) \left(C_{3,i}^{(1)} \exp(-C_{3,i}^{(2)} L_0^r) + C_{3,i}^{(3)} \exp(-C_{3,i}^{(4)} r^2 t) \right),$$

where $p_{1,i}$, $p_{2,i}$, $p_{3,i}$ are polynomials with non-negative coefficients, and $C_{l,i}^{(m)}$, $l = 1, 2, 3$, $m = 1, 2, 3, 4$ are positive constants; the polynomials and constants do not depend on t or r . The polynomials $p_{1,i}$ and $p_{3,i}$ have degree at most $i + 1$ and $p_{2,i}$ has degree at most i .

(II) For all $r \geq r^*$, $t > 0$ satisfying $r^2 t \geq M^r$, and each $\iota \in \mathcal{I}$ with $\iota > i$,

$$(II.1) \quad \sup_{n \geq 1} \mathbf{P} \left(T_{\iota, a(i)}^{r,n}(s_\iota^r) \geq (x_{a(i)}^* + \epsilon_i) s_\iota^r, s_\iota^r \leq \beta_{\iota,n}^r, \tau_{\iota, 2n-1}^r \leq r^2 t \right) \\ \leq p_{4,i}(r^2 t) \left(C_{4,i}^{(1)} \exp(-C_{4,i}^{(2)} L_0^r) + C_{4,i}^{(3)} \exp(-C_{4,i}^{(4)} r^2 t) \right),$$

(II.2) if ι is a transition class,

$$\sup_{n \geq 1} \mathbf{P} \left({}^d T_{\iota, a(i)}^{r,n}(d_s^r) \leq (x_{a(i)}^* - \epsilon_i) d_s^r, d_{\tau_{\iota, 2n-1}^r}^r \leq r^2 t \right) \\ \leq p_{5,i}(r^2 t) \left(C_{5,i}^{(1)} \exp(-C_{5,i}^{(2)} L_0^r) + C_{5,i}^{(3)} \exp(-C_{5,i}^{(4)} r^2 t) \right),$$

$$(II.3) \quad \mathbf{P} \left(T_{a(i)}^r(t_\iota^r) \leq (x_{a(i)}^* - \epsilon_i) t_\iota^r \right) \\ \leq p_{6,i}(r^2 t) \left(C_{6,i}^{(1)} \exp(-C_{6,i}^{(2)} L_0^r) + C_{6,i}^{(3)} \exp(-C_{6,i}^{(4)} r^2 t) \right),$$

where $p_{4,i}, p_{5,i}, p_{6,i}$ are polynomials (of degree at most $i+1$) with non-negative coefficients, and $C_{l,i}^{(m)}$, $l = 4, 5, 6$, $m = 1, 2, 3, 4$ are positive constants; the polynomials and constants do not depend on t or r .

Remark. The variable t appearing on the right side of (II.3) stems from an estimate obtained in (8.139) in the proof of Lemma 8.3.4, involving the number of class i jobs in the system at time $t_i^r \leq r^2 t$ for sufficiently large $r \geq 1$.

In addition to (I) and (II), if i^* is a transition class, we will show the following.

(III) For all $r \geq r^*$, $t > 0$ satisfying $r^2 t \geq M^r$, and each $k \in \underline{\mathcal{K}}_{i^*}$,

$$(III.1) \quad \mathbf{P} \left(\inf_{\tau_{i^*,0}^r \leq s \leq r^2 t} R_{i^*}^r(s) \leq -L_{i^*}^r + |\underline{\mathcal{J}}_{i^*}| \right) \\ \leq p_{1,i^*}(r^2 t) \left(C_{1,i^*}^{(1)} \exp(-C_{1,i^*}^{(2)} L_0^r) + C_{1,i^*}^{(3)} \exp(-C_{1,i^*}^{(4)} r^2 t) \right),$$

$$(III.2) \quad \mathbf{P} \left(I_k^r(\tau_{i^*,0}^r) \geq t_{i^*}^r \right) \\ \leq p_{2,i^*}(r^2 t) \left(C_{2,i^*}^{(1)} \exp(-C_{2,i^*}^{(2)} L_0^r) + C_{2,i^*}^{(3)} \exp(-C_{2,i^*}^{(4)} r^2 t) \right),$$

$$(III.3) \quad \mathbf{P} \left(I_k^r(r^2 t) - I_k^r(\tau_{i^*,0}^r) > 0, \tau_{i^*,0}^r < r^2 t \right) \\ \leq p_{3,i^*}(r^2 t) \left(C_{3,i^*}^{(1)} \exp(-C_{3,i^*}^{(2)} L_0^r) + C_{3,i^*}^{(3)} \exp(-C_{3,i^*}^{(4)} r^2 t) \right),$$

where $p_{1,i^*}, p_{2,i^*}, p_{3,i^*}$ are polynomials with non-negative coefficients, and $C_{l,i^*}^{(m)}$, $l = 1, 2, 3$, $m = 1, 2, 3, 4$ are positive constants; the polynomials and constants do not depend on t or r . The polynomials p_{1,i^*} and p_{3,i^*} have degree at most $\mathbf{I} + 1$ and p_{2,i^*} has degree at most \mathbf{I} .

Property (I) is used to obtain the conclusions (8.1)–(8.3) in Theorem 8.1.1. Property (II) describes the properties associated with buffer i that are carried forward to prove (I) for buffers $\iota > i$ and to prove Property (III), which in turn is used to prove (8.4)–(8.6). The induction proof depends on the following lemmas that are proved in Sections 8.4–8.7 below.

Lemma 8.3.1 *Fix $i \in \mathcal{I} \setminus \{i^*\}$. Suppose that for all $r \geq r^*$, $t > 0$ satisfying*

$r^2t \geq M^r$, and each $j \in \underline{\mathcal{I}}_i$, we have

$$\begin{aligned}
(i) \quad & \sup_{n \geq 1} \mathbf{P} \left(T_{i,j}^{r,n}(s_i^r) \leq (x_j^* - \epsilon_i) s_i^r, s_i^r \leq \beta_{i,n}^r, \tau_{i,2n-1}^r \leq r^2t \right) \\
& \leq p_{7,i}(r^2t) \left(C_{7,i}^{(1)} \exp(-C_{7,i}^{(2)} L_0^r) + C_{7,i}^{(3)} \exp(-C_{7,i}^{(4)} r^2t) \right), \\
(ii) \quad & \text{if } i \text{ is a transition buffer,} \\
& \sup_{n \geq 1} \mathbf{P} \left(dT_{i,j}^{r,n}(d s_i^r) \geq (x_j^* + \epsilon_i) d s_i^r, d\tau_{i,2n-1}^r \leq r^2t \right) \\
& \leq p_{8,i}(r^2t) \left(C_{8,i}^{(1)} \exp(-C_{8,i}^{(2)} L_0^r) + C_{8,i}^{(3)} \exp(-C_{8,i}^{(4)} r^2t) \right), \\
(iii) \quad & \mathbf{P} \left(T_j^r(t_i^r) \geq (x_j^* + \epsilon_i) t_i^r \right) \\
& \leq p_{9,i}(r^2t) \left(C_{9,i}^{(1)} \exp(-C_{9,i}^{(2)} L_0^r) + C_{9,i}^{(3)} \exp(-C_{9,i}^{(4)} r^2t) \right), \\
(iv) \quad & \sup_{n \geq 1} \mathbf{P} \left(T_{i,a(i)}^{r,n}(s_i^r) \leq (\hat{x}_{i,i} - \epsilon_i) s_i^r, s_i^r \leq \beta_{i,n}^r, \tau_{i,2n-1}^r \leq r^2t \right) \\
& \leq p_{10,i}(r^2t) \left(C_{10,i}^{(1)} \exp(-C_{10,i}^{(2)} L_0^r) + C_{10,i}^{(3)} \exp(-C_{10,i}^{(4)} r^2t) \right),
\end{aligned}$$

where $p_{7,i}$, $p_{8,i}$, $p_{9,i}$, $p_{10,i}$ are polynomials (of degree at most i) with non-negative coefficients, and $C_{l,i}^{(m)}$, $l = 7, 8, 9, 10$, $m = 1, 2, 3, 4$ are positive constants; the polynomials and constants are independent of t , n , and r . Then, (I) holds for i .

Lemma 8.3.2 *Let $i \in \mathcal{I} \setminus \{i^*\}$. Suppose that (I) and (II) hold with i' in place of i , for all $i' < i$. Then (i)–(iii) of Lemma 8.3.1 hold for all $r \geq r^*$, $t > 0$ satisfying $r^2t \geq M^r$, and each $j \in \underline{\mathcal{I}}_i$.*

Lemma 8.3.3 *Let $i \in \mathcal{I} \setminus \{i^*\}$. Suppose that (I) and (II) hold with i' in place of i , for all $i' < i$. Then (iv) of Lemma 8.3.1 holds for all $r \geq r^*$, and $t > 0$ satisfying $r^2t \geq M^r$.*

Lemma 8.3.4 *Let $i \in \mathcal{I} \setminus \{i^*\}$. Suppose that (I) and (II) hold with i' in place of i , for all $i' < i$. Then (II) holds for i .*

Lemmas 8.3.1–8.3.4 combined with a formal induction yield the following.

Theorem 8.3.5 (I) and (II) hold for each $i \in \mathcal{I} \setminus \{i^*\}$. In addition, (III) holds if i^* is a transition class.

With Lemmas 8.3.1–8.3.4 in place, the steps in the induction argument used in proving Theorem 8.3.5 are as follows (assuming the ordering of the buffers is as described in Section 7.1).

1. Fix $i \in \mathcal{I} \setminus \{i^*\}$, and assume that (I) and (II) hold for all $i' < i$ (for $i = 1$, this is a vacuous assumption).
2. Use Lemmas 8.3.2 and 8.3.3 to show that (i)–(iv) in Lemma 8.3.1 hold for i .
3. Apply Lemma 8.3.1 to conclude that (I) holds for i .
4. Apply Lemma 8.3.4 to conclude that (II) holds for i .

It then follows that (I) and (II) hold for all $i \in \mathcal{I} \setminus \{i^*\}$. Then, if i^* is a transition buffer, we combine the proof of Lemma 8.3.2 and parts of the proof of Lemma 8.3.1 (adapted for $i = i^*$) with the fact that (II) holds for $i = i^*$ and $i < i^*$, to prove that (III) holds.

The formal proof of Theorem 8.3.5 is given in Section 8.8. As a guide to the reader, before beginning the proofs of the lemmas and Theorems 8.3.5 and 8.1.1, we briefly describe some of the ideas involved in proving Theorems 8.3.5 and 8.1.1.

Lemma 8.3.1 is the main lemma that drives the induction. The result of this lemma yields that all queue length processes and idletime processes (except $Q_{i^*}^r$ and $I_{k^*}^r$) vanish (on diffusion scale) as r goes to infinity (for c sufficiently large). To obtain (I) for buffer $i \in \mathcal{I} \setminus \{i^*\}$, our proof requires that, in addition to large deviation estimates on the primitive renewal processes, there are estimates on the allocation processes (in each excursion interval for the residual process associated with buffer i) corresponding to the activities which process class i jobs, i.e., (i)–(iv) in Lemma 8.3.1 hold. The estimates (i)–(iii) are derived from the induction assumptions for (II) (with i replaced by $i' < i$ in (II)), using the fact that, for each $j \in \underline{\mathcal{J}}_i$, the utilization of activity j is constrained by the (higher priority) activities for server $k(j)$ which serve buffers in the layer below that server. These estimates are given by the proof of Lemma 8.3.2 which is contained in Section

8.5. Similarly, the on-time of the activity, $a(i)$, that serves buffer i from above can be estimated (in an (up) excursion for R_i^r) by having estimates, derived from the induction assumption for (II), on the (higher priority) activities associated with server $k(a(i))$ (recall that buffers which have higher priority for server $k(a(i))$ are all numbered lower than i). These estimates are proved in Lemma 8.3.3 in Section 8.6. Finally, the proof of Lemma 8.3.4 in Section 8.7 (which uses the fact that (I) holds for buffer i together with the induction assumption for (II), with i replaced by $i' < i$ in (II), completes the induction step by showing how to transition between layers.

For (III.1)–(III.3), assuming that i^* is a transition buffer, we first use the proof of Lemma 8.3.2 (cf. Section 8.5) to show that (ii) and (iii) in Lemma 8.3.1 hold for i^* , $r \geq r^*$, $t > 0$ satisfying $r^2 t \geq M^r$, and each $j \in \underline{\mathcal{J}}_{i^*}$, given that (I) and (II) hold for all $i < i^*$ (cf. (8.90)–(8.91) and (8.92)–(8.93), respectively). Then, we use the proof of Lemma 8.3.1 in Section 8.4 (with i^* in place of i , there) to show that (III.1)–(III.3) hold. (For (III.1) we use the second part of the proof of (I.1) (for the down excursions of $R_{i^*}^r$, i.e., (8.59)–(8.77)), for (III.2) we use the proof of (I.2) (cf. (8.78)–(8.86)), and for (III.3) we use the proof of (I.3) (cf. (8.87)).

Properties (I) and (III) will be used to prove (8.1)–(8.3) and (8.4)–(8.6), respectively, of Theorem 8.1.1, for a sufficiently large constant c appearing in the definition of L_0^r .

8.4 Proof of Lemma 8.3.1: Estimates on Allocation Processes Imply Residual Processes Stay Near Zero

Proof of Lemma 8.3.1. Fix $i \in \mathcal{I} \setminus \{i^*\}$. Suppose that for all $r \geq r^*$, $t > 0$ satisfying $r^2 t \geq M^r$, and each $j \in \underline{\mathcal{J}}_i$, (i)–(iv) hold. The proof that follows is an extension of the proof of Theorem 7.2 in [2].

For the special case when class i is not a transition class, we have $\underline{\mathcal{K}}_i = \emptyset$, $\underline{\mathcal{J}}_i = \emptyset$, $\tau_{i,0}^r \equiv 0$, and ${}^dR_i^r \equiv \mathbf{0}$. This implies that (I.2) holds trivially and the part of the proof of (I.1) below, from (8.59) through (8.77), is not needed. Consequently, (8.59)-(8.77) and (8.78)-(8.86) below will only be of importance when i is a transition class.

Proof of (I.1). We first consider the ‘‘up’’ excursions of R_i^r . For $n \geq 1$, $r \geq 1$, $j \in \mathcal{J}_i$, on $\{\tau_{i,2n-1}^r < \infty\}$, we define shifted renewal processes, for $s \geq 0$, as follows.

$$A_{i,i}^{r,n}(s) = A_i^r(\tau_{i,2n-1}^r + s) - A_i^r(\tau_{i,2n-1}^r), \quad (8.36)$$

$$S_{i,j}^{r,n}(s) = S_j^r(T_j^r(\tau_{i,2n-1}^r) + s) - S_j^r(T_j^r(\tau_{i,2n-1}^r)), \quad (8.37)$$

$$\check{A}_{i,i}^{r,n}(s) = \sup\{m \geq 0 : \xi_i^r(A_i^r(\tau_{i,2n-1}^r) + m) - \xi_i^r(A_i^r(\tau_{i,2n-1}^r)) \leq s\}, \quad (8.38)$$

$$\check{S}_{i,j}^{r,n}(s) = \sup\{m \geq 0 : \eta_j^r(S_j^r(T_j^r(\tau_{i,2n-1}^r)) + m) - \eta_j^r(S_j^r(T_j^r(\tau_{i,2n-1}^r))) \leq s\}, \quad (8.39)$$

and, for concreteness, on $\{\tau_{i,2n-1}^r = \infty\}$ we define $A_{i,i}^{r,n}$, $S_{i,j}^{r,n}$, $\check{A}_{i,i}^{r,n}$, $\check{S}_{i,j}^{r,n}$ to be identically zero. Recall the definition of $T_{i,j}^{r,n}$ from (7.12).

Consider the n^{th} ‘‘up’’ excursion interval for R_i^r . We have that on $\{\tau_{i,2n-1}^r < \infty\}$, for $0 \leq s \leq \beta_{i,n}^r$,

$$R_i^r(\tau_{i,2n-1}^r + s) = 1 + A_{i,i}^{r,n}(s) - \sum_{j \in \mathcal{J}_i} S_{i,j}^{r,n}(T_{i,j}^{r,n}(s)), \quad (8.40)$$

and, taking account of the fact that a new arrival to class i occurs at $\tau_{i,2n-1}^r < \infty$ and a job may have been partially served by activity $j \in \mathcal{J}_i$ at $\tau_{i,2n-1}^r < \infty$, we also have that, for $s \geq 0$,

$$A_{i,i}^{r,n}(s) = \check{A}_{i,i}^{r,n}(s), \quad \text{and} \quad S_{i,j}^{r,n}(s) \geq \check{S}_{i,j}^{r,n}(s), \quad j \in \mathcal{J}_i. \quad (8.41)$$

By (7.5), we have that

$$\epsilon_i < \min \left\{ \frac{\mu_j}{2}, \frac{\lambda_i}{2}, \left(1 - \frac{8\lambda_i + 2\delta_i}{8\lambda_i + 3\delta_i}\right) \hat{x}_{i,i}, \left(1 - \frac{4\lambda_i}{4\lambda_i + \delta_i}\right) x_j^* : j \in \mathcal{J}_i \right\}, \quad (8.42)$$

where δ_i and $\hat{x}_{i,i}$ are given in (8.13) and (7.2), respectively. For the third term in (8.42), we have that

$$\begin{aligned}
\left(1 - \frac{8\lambda_i + 2\delta_i}{8\lambda_i + 3\delta_i}\right) \hat{x}_{i,i} &= \frac{\delta_i \hat{x}_{i,i}}{8\lambda_i + 3\delta_i} \\
&\geq \frac{\delta_i \hat{x}_{i,i}}{8(\lambda_i + \delta_i)} \\
&= \frac{\delta_i \hat{x}_{i,i}}{8(\sum_{j \in \mathcal{J}_i} x_j^* \mu_j + \hat{x}_{i,i} \mu_{a(i)})} \\
&\geq \frac{\delta_i \hat{x}_{i,i}}{8 \sum_{j \in \mathcal{J}_i} \mu_j} \\
&\geq \frac{\delta_i x_{a(i)}^*}{8 \sum_{j \in \mathcal{J}} \mu_j} \\
&\geq \frac{\delta_{\min} x_{\min}^*}{8\mu_{\text{sum}}} \\
&> \hat{\epsilon} > \epsilon_i, \text{ by (7.5),} \tag{8.43}
\end{aligned}$$

where the second inequality follows since $x_j^* \leq 1$, for all $j \in \mathcal{J}$, and $\hat{x}_{i,i} \leq 1$, for all $i \in \mathcal{I}$. The inequalities involving the remaining terms in (8.42) can be justified in a similar manner.

In the following, it is assumed that $r \geq r^*$ and $t > 0$ satisfies $r^2 t \geq M^r$. Then $r^2 t > 2/\epsilon_i$ by (8.9)–(8.12), and (8.35). Define $n_i^r = \lfloor (\lambda_i^r + \epsilon_i) r^2 t \rfloor + 1$. Then since each “up” excursion of R_i^r is initiated by an arrival to class i , using the large deviations bounds for renewal processes given in Lemma 7.6.1 we have the following estimate, for all $r \geq r^*$, of the probability that there are at least $n_i^r - 1$ complete “up” excursions of R_i^r in $[0, r^2 t]$ ($\tau_{i, 2n_i^r - 1}^r$ is the beginning of the $(n_i^r)^{\text{th}}$ (up) excursion interval):

$$\begin{aligned}
\mathbf{P}(\tau_{i, 2n_i^r - 1}^r \leq r^2 t) &\leq \mathbf{P}(A_i^r(r^2 t) \geq n_i^r) \\
&\leq \mathbf{P}(A_i^r(r^2 t) > (\lambda_i^r + \epsilon_i) r^2 t) \\
&\leq K_1 \exp(-K_1' r^2 t), \tag{8.44}
\end{aligned}$$

where $K_1 = \exp(\Lambda_i^{a,*}((\lambda_i(1 + \epsilon_i/3\lambda_i))^{-1})) > 0$, $K_1' = \lambda_i \Lambda_i^{a,*}((\lambda_i(1 + \epsilon_i/3\lambda_i))^{-1}) > 0$, depend on the Legendre-Fenchel transform $\Lambda_i^{a,*}$ of the logarithmic moment gener-

ating function Λ_i^a of $u_i(1)$ (cf. (3.3)–(3.4)), λ_i , and ϵ_i , but are independent of t and r .

Now,

$$\begin{aligned}
& \mathbf{P}\left(R_i^r(s) \geq L_i^r - |\underline{\mathcal{J}}_i| \text{ some } s \in [0, r^2t]\right) \\
& \leq \mathbf{P}\left(\tau_{i,2n_i^r-1}^r \leq r^2t\right) \\
& \quad + \mathbf{P}\left(\tau_{i,2n_i^r-1}^r > r^2t, R_i^r(s) \geq L_i^r - |\underline{\mathcal{J}}_i| \text{ some } s \in [0, r^2t]\right) \\
& \leq \mathbf{P}\left(\tau_{i,2n_i^r-1}^r \leq r^2t\right) \\
& \quad + \sum_{n=1}^{n_i^r-1} \mathbf{P}\left(R_i^r(s) \geq L_i^r - |\underline{\mathcal{J}}_i| \text{ some } s \in (\tau_{i,2n-1}^r, \tau_{i,2n}^r), \right. \\
& \qquad \qquad \qquad \left. \tau_{i,2n-1}^r \leq r^2t\right). \tag{8.45}
\end{aligned}$$

For each positive integer n , let

$$\begin{aligned}
\Upsilon_i^{r,n} & = \left\{ A_{i,i}^{r,n}(s_i^r) \leq (\lambda_i^r + \epsilon_i)s_i^r; \right. \\
& \quad S_{i,j}^{r,n}((x_j^* - \epsilon_i)s_i^r) \geq (\mu_j^r - \epsilon_i)(x_j^* - \epsilon_i)s_i^r, j \in \underline{\mathcal{J}}_i; \\
& \quad T_{i,j}^{r,n}(s_i^r) > (x_j^* - \epsilon_i)s_i^r, j \in \underline{\mathcal{J}}_i; \\
& \quad S_{i,a(i)}^{r,n}((\hat{x}_{i,i} - \epsilon_i)s_i^r) \geq (\mu_{a(i)}^r - \epsilon_i)(\hat{x}_{i,i} - \epsilon_i)s_i^r; \\
& \quad \left. T_{i,a(i)}^{r,n}(s_i^r) > (\hat{x}_{i,i} - \epsilon_i)s_i^r; \tau_{i,2n-1}^r \leq r^2t \right\}. \tag{8.46}
\end{aligned}$$

The set $\Upsilon_i^{r,n}$ is a “good” set, in the sense that, on it, various shifted stochastic processes can be bounded on one side at time s_i^r by certain linear functions.

Let

$$\begin{aligned}
\rho_i^{r,n} & = \xi_i^r(A_i^r(\tau_{i,2n-1}^r) + L_i^r - |\underline{\mathcal{J}}_i| - 1) \\
& \quad - \xi_i^r(A_i^r(\tau_{i,2n-1}^r)) \text{ on } \{\tau_{i,2n-1}^r < \infty\}, \tag{8.47}
\end{aligned}$$

and let $\rho_i^{r,n} \equiv 0$ on $\{\tau_{i,2n-1}^r = \infty\}$.

Now,

$$\begin{aligned}
& \left\{ R_i^r(s) \geq L_i^r - |\underline{\mathcal{J}}_i| \text{ some } s \in (\tau_{i,2n-1}^r, \tau_{i,2n}^r), \tau_{i,2n-1}^r < \infty \right\} \\
& = \left\{ R_i^r(s) \geq L_i^r - |\underline{\mathcal{J}}_i| \text{ some } s \in (\tau_{i,2n-1}^r, \tau_{i,2n}^r), \right. \\
& \qquad \qquad \qquad \left. \rho_i^{r,n} \leq \beta_{i,n}^r, \tau_{i,2n-1}^r < \infty \right\}, \tag{8.48}
\end{aligned}$$

since, on $\{\tau_{i,2n-1}^r < \infty\}$, $\rho_i^{r,n}$ is the minimum possible amount of time required for R_i^r to reach the level $L_i^r - |\underline{\mathcal{J}}_i|$ in the n^{th} (up) excursion. Then,

$$\begin{aligned} & \mathbf{P}\left(R_i^r(s) \geq L_i^r - |\underline{\mathcal{J}}_i| \text{ some } s \in (\tau_{i,2n-1}^r, \tau_{i,2n}^r), \right. \\ & \quad \left. \tau_{i,2n-1}^r \leq r^2 t\right) \\ & \leq \mathbf{P}\left(\tau_{i,2n-1}^r \leq r^2 t, (\Upsilon_i^{r,n})^c, \rho_i^{r,n} \leq \beta_{i,n}^r\right) \\ & \quad + \mathbf{P}\left(R_i^r(s) \geq L_i^r - |\underline{\mathcal{J}}_i| \text{ some } s \in (\tau_{i,2n-1}^r, \tau_{i,2n}^r), \Upsilon_i^{r,n}\right). \end{aligned} \quad (8.49)$$

Now, on $\Upsilon_i^{r,n}$, we have

$$\begin{aligned} & 1 + A_{i,i}^{r,n}(s_i^r) - \sum_{j \in \mathcal{J}_i} S_{i,j}^{r,n}(T_{i,j}^{r,n}(s_i^r)) \\ & \leq 1 + A_{i,i}^{r,n}(s_i^r) - \sum_{j \in \underline{\mathcal{J}}_i} S_{i,j}^{r,n}((x_j^* - \epsilon_i)s_i^r) - S_{i,a(i)}^{r,n}((\hat{x}_{i,i} - \epsilon_i)s_i^r) \\ & \leq 1 + (\lambda_i + \epsilon_i)s_i^r - \sum_{j \in \underline{\mathcal{J}}_i} (x_j^* - \epsilon_i)(\mu_j^r - \epsilon_i)s_i^r - (\hat{x}_{i,i} - \epsilon_i)(\mu_{a(i)}^r - \epsilon_i)s_i^r \\ & \leq 1 + \left(\frac{2\lambda_i}{2\lambda_i + \delta_i} \left(\sum_{j \in \underline{\mathcal{J}}_i} x_j^* \mu_j + \hat{x}_{i,i} \mu_{a(i)} \right) \right. \\ & \quad \left. - \sum_{j \in \underline{\mathcal{J}}_i} \left(\frac{4\lambda_i}{4\lambda_i + \delta_i} x_j^* \frac{4\lambda_i + \delta_i}{4\lambda_i + 2\delta_i} \mu_j \right) - \frac{8\lambda_i + 2\delta_i}{8\lambda_i + 3\delta_i} \hat{x}_{i,i} \frac{8\lambda_i + 3\delta_i}{8\lambda_i + 4\delta_i} \mu_{a(i)} \right) s_i^r \\ & = 1 + \left(\frac{2\lambda_i}{2\lambda_i + \delta_i} \left(\sum_{j \in \underline{\mathcal{J}}_i} x_j^* \mu_j + \hat{x}_{i,i} \mu_{a(i)} \right) - \frac{2\lambda_i}{2\lambda_i + \delta_i} \left(\sum_{j \in \underline{\mathcal{J}}_i} x_j^* \mu_j + \hat{x}_{i,i} \mu_{a(i)} \right) \right. \\ & \quad \left. - \frac{\delta_i}{4\lambda_i + 2\delta_i} \hat{x}_{i,i} \mu_{a(i)} \right) s_i^r \\ & = 1 - \frac{\delta_i}{4\lambda_i + 2\delta_i} \hat{x}_{i,i} \mu_{a(i)} s_i^r < 0, \end{aligned} \quad (8.50)$$

for $r \geq r^*$, and where, in the third inequality we have used (8.16)–(8.18) along with (8.42), and in the last inequality we have used (8.19). It then follows from (8.40) that

$$\beta_{i,n}^r = \tau_{i,2n}^r - \tau_{i,2n-1}^r < s_i^r \text{ on } \Upsilon_i^{r,n} \text{ for } r \geq r^*. \quad (8.51)$$

Furthermore, on $\Upsilon_i^{r,n}$ for $0 \leq s \leq s_i^r$,

$$1 + A_{i,i}^{r,n}(s) - \sum_{j \in \mathcal{J}_i} S_{i,j}^{r,n}(T_{i,j}^{r,n}(s)) \leq 1 + (\lambda_i^r + \epsilon_i) s_i^r = L_i^r - |\underline{\mathcal{J}}_i| - 1, \quad (8.52)$$

by (8.7). Hence by (8.40) we have that on $\Upsilon_i^{r,n}$, $R_i^r(s) < L_i^r - |\underline{\mathcal{J}}_i|$ for $s \in (\tau_{i,2n-1}^r, \tau_{i,2n}^r)$, whenever $r \geq r^*$. Thus the last probability in (8.49) is zero for all such r .

It remains to estimate

$$\begin{aligned} & \mathbf{P} \left(\tau_{i,2n-1}^r \leq r^2 t, (\Upsilon_i^{r,n})^c, \rho_i^{r,n} \leq \beta_{i,n}^r \right) \\ & \leq \mathbf{P} \left(\rho_i^{r,n} < s_i^r, \tau_{i,2n-1}^r \leq r^2 t \right) \\ & \quad + \mathbf{P} \left(A_{i,i}^{r,n}(s_i^r) > (\lambda_i^r + \epsilon_i) s_i^r, \tau_{i,2n-1}^r \leq r^2 t \right) \\ & \quad + \sum_{j \in \mathcal{J}_i} \mathbf{P} \left(S_{i,j}^{r,n}(x_j^* - \epsilon_i) s_i^r < (\mu_j^r - \epsilon_i)(x_j^* - \epsilon_i) s_i^r, \tau_{i,2n-1}^r \leq r^2 t \right) \\ & \quad + \mathbf{P} \left(S_{i,a(i)}^{r,n}((\hat{x}_{i,i} - \epsilon_i) s_i^r) < (\mu_{a(i)}^r - \epsilon_i)(\hat{x}_{i,i} - \epsilon_i) s_i^r, \tau_{i,2n-1}^r \leq r^2 t \right) \\ & \quad + \sum_{j \in \mathcal{J}_i} \mathbf{P} \left(T_{i,j}^{r,n}(s_i^r) \leq (x_j^* - \epsilon_i) s_i^r, s_i^r \leq \beta_{i,n}^r, \tau_{i,2n-1}^r \leq r^2 t \right) \\ & \quad + \mathbf{P} \left(T_{i,a(i)}^{r,n}(s_i^r) \leq (\hat{x}_{i,i} - \epsilon_i) s_i^r, s_i^r \leq \beta_{i,n}^r, \tau_{i,2n-1}^r \leq r^2 t \right), \end{aligned} \quad (8.53)$$

for $r \geq r^*$ and $t > 0$ satisfying $r^2 t \geq M^r$.

Now, the set $\{\tau_{i,2n-1}^r < \infty\}$ is contained in the set $\{\mathcal{T}_{n,i}^r \in \mathbb{N}^{\mathbf{I}} \times \mathbb{N}^{\mathbf{J}}\}$ (cf. Lemma 7.5.2). Using Lemmas 7.5.2 and 7.5.3, we conclude that on $\{\mathcal{T}_{n,i}^r \in \mathbb{N}^{\mathbf{I}} \times \mathbb{N}^{\mathbf{J}}\}$, the conditional distribution of $\{u_i^r(A_i^r(\tau_{i,2n-1}^r) + m), m = 1, 2, 3, \dots\}$ given $\mathcal{F}_{\mathcal{T}_{n,i}^r}^r$ is equal to that of a sequence of strictly positive independent random variables where the members indexed by $m = 2, 3, \dots$ are identically distributed with the same distribution as $u_i^r(1)$ (for this lemma, observe that $L_0^r \geq \mathbf{J} + 1$ by (8.35) and (7.5) since $\epsilon_i \leq \hat{\epsilon} < 1/2(\mathbf{J} + 1)$). Then, as in equation (84) of [2], for $r \geq r^*$, and

$t > 0$ satisfying $r^2 t \geq M^r$,

$$\begin{aligned}
& \mathbf{P} \left(A_{i,i}^{r,n}(s_i^r) > (\lambda_i^r + \epsilon_i) s_i^r, \tau_{i,2n-1}^r \leq r^2 t \right) \\
& \leq \mathbf{E} \left(\mathbf{1}_{\{\mathcal{T}_{n,i}^r \in \mathbb{N}^{\mathbf{I}} \times \mathbb{N}^{\mathbf{J}}\}} \mathbf{P} \left(\check{A}_{i,i}^{r,n}(s_i^r) > (\lambda_i^r + \epsilon_i) s_i^r \mid \mathcal{F}_{\mathcal{T}_{n,i}^r}^r \right) \right) \\
& \leq \exp \left(- \left((\lambda_i^r + \epsilon_i) s_i^r - 1 \right) \Lambda_i^{a,*} \left(\frac{1}{\lambda_i} \left(\frac{1}{1 + \frac{\epsilon_i}{3\lambda_i}} \right) \right) \right) \\
& \leq K_2 \exp(-K_2' L_i^r), \tag{8.54}
\end{aligned}$$

by Lemma 7.6.1 (since $s_i^r > 2/\epsilon_i$ by (8.9), (8.12), and (8.35)), and where $K_2 = \exp((|\underline{\mathcal{J}}_i| + 3)\Lambda_i^{a,*}((\lambda_i(1 + \epsilon_i/3\lambda_i))^{-1})) > 0$ and $K_2' = \Lambda_i^{a,*}((\lambda_i(1 + \epsilon_i/3\lambda_i))^{-1}) > 0$ do not depend on t , n , or r . Similarly, for $r \geq r^*$, $t > 0$ satisfying $r^2 t \geq M^r$,

$$\begin{aligned}
& \mathbf{P}(\rho_i^{r,n} < s_i^r, \tau_{i,2n-1}^r \leq r^2 t) \\
& \leq \mathbf{E} \left(\mathbf{1}_{\{\mathcal{T}_{n,i}^r \in \mathbb{N}^{\mathbf{I}} \times \mathbb{N}^{\mathbf{J}}\}} \mathbf{P} \left(\rho_i^{r,n} < s_i^r \mid \mathcal{F}_{\mathcal{T}_{n,i}^r}^r \right) \right) \\
& \leq \mathbf{E} \left(\mathbf{1}_{\{\mathcal{T}_{n,i}^r \in \mathbb{N}^{\mathbf{I}} \times \mathbb{N}^{\mathbf{J}}\}} \mathbf{P} \left(\check{A}_{i,i}^{r,n}(s_i^r) \geq L_i^r - |\underline{\mathcal{J}}_i| - 1 \mid \mathcal{F}_{\mathcal{T}_{n,i}^r}^r \right) \right) \\
& \leq \mathbf{E} \left(\mathbf{1}_{\{\mathcal{T}_{n,i}^r \in \mathbb{N}^{\mathbf{I}} \times \mathbb{N}^{\mathbf{J}}\}} \mathbf{P} \left(\check{A}_{i,i}^{r,n}(s_i^r) > (\lambda_i^r + \epsilon_i) s_i^r \mid \mathcal{F}_{\mathcal{T}_{n,i}^r}^r \right) \right) \\
& \leq K_2 \exp(-K_2' L_i^r), \tag{8.55}
\end{aligned}$$

where the third inequality follows by the definition of s_i^r (cf. (8.7)).

In a similar manner (cf. (85)–(86) of [2]), for all $r \geq r^*$, $t > 0$ satisfying $r^2 t \geq M^r$, and $j \in \mathcal{J}_i$,

$$\begin{aligned}
& \mathbf{P} \left(S_{i,j}^{r,n}((\tilde{x}_j^* - \epsilon_i) s_i^r) < (\mu_j^r - \epsilon_i)(\tilde{x}_j^* - \epsilon_i) s_i^r, \tau_{i,2n-1}^r \leq r^2 t \right) \\
& \leq \mathbf{E} \left(\mathbf{1}_{\{\mathcal{T}_{n,i}^r \in \mathbb{N}^{\mathbf{I}} \times \mathbb{N}^{\mathbf{J}}\}} \mathbf{P} \left(\check{S}_{i,j}^{r,n}((\tilde{x}_j^* - \epsilon_i) s_i^r) < (\mu_j^r - \epsilon_i)(\tilde{x}_j^* - \epsilon_i) s_i^r, \right. \right. \\
& \quad \left. \left. \tau_{i,2n-1}^r \leq r^2 t \mid \mathcal{F}_{\mathcal{T}_{n,i}^r}^r \right) \right) \\
& \leq \exp \left(- (\mu_j - 2\epsilon_i)(\tilde{x}_j^* - \epsilon_i) s_i^r \Lambda_j^{s,*} \left(\frac{1}{\mu_j} \left(1 + \frac{\epsilon_i}{2\mu_j} \right) \right) \right) \\
& \quad + \mathbf{E} \left(\mathbf{1}_{\{\mathcal{T}_{n,i}^r \in \mathbb{N}^{\mathbf{I}} \times \mathbb{N}^{\mathbf{J}}\}} \mathbf{P} \left(v_{i,j}^{r,n} > \frac{\epsilon_i}{2\mu_j^r} (\tilde{x}_j^* - \epsilon_i) s_i^r, \right. \right. \\
& \quad \left. \left. \tau_{i,2n-1}^r \leq r^2 t \mid \mathcal{F}_{\mathcal{T}_{n,i}^r}^r \right) \right), \tag{8.56}
\end{aligned}$$

by Lemma 7.6.1, where $\tilde{x}_j^* = x_j^*$ if $j \in \underline{\mathcal{J}}_i$, $\tilde{x}_{a(i)}^* = \hat{x}_{i,i}$, $v_{i,j}^{r,n} = v_j^r(S_j^r(T_j^r(\tau_{i,2n-1}^r)) + 1)$, and $\Lambda_j^{s,*}$, $j \in \mathcal{J}_i$ is the Legendre-Fenchel transform of the logarithmic moment

generating function Λ_j^s of $v_j(1)$, $j \in \mathcal{J}_i$ (cf. (3.3)–(3.4)). (Note here that when we use (7.24) of Lemma 7.6.1, we do not need the condition that $s > 2/\epsilon$ required in (7.23), e.g., in the case above, we do not require that $(\tilde{x}_j^* - \epsilon_i)s_i^r > 2/\epsilon_i$.) Now, in a similar manner to that for (87) in [2], using (7.26), we have for all $r \geq r^*$, $t > 0$ satisfying $r^2t \geq M^r$,

$$\begin{aligned}
& \mathbf{P} \left(v_{i,j}^{r,n} > \frac{\epsilon_i}{2\mu_j^r} (\tilde{x}_j^* - \epsilon_i) s_i^r, \tau_{i,2n-1}^r \leq r^2t \right) \\
&= \mathbf{P} \left(v_j^r(S_j^r(T_j^r(\tau_{i,2n-1}^r)) + 1) > \frac{\epsilon_i}{2\mu_j^r} (\tilde{x}_j^* - \epsilon_i) s_i^r, \tau_{i,2n-1}^r \leq r^2t \right) \\
&\leq \mathbf{P} \left(\max_{m=1}^{S_j^r(r^2t)+1} v_j^r(m) > \frac{\epsilon_i}{2\mu_j^r} (\tilde{x}_j^* - \epsilon_i) s_i^r \right) \\
&\leq \mathbf{P} \left(\max_{m=1}^{\lfloor (\mu_j^r + \epsilon_i)r^2t \rfloor + 1} v_j^r(m) > \frac{\epsilon_i}{2\mu_j^r} (\tilde{x}_j^* - \epsilon_i) s_i^r \right) \\
&\quad + \mathbf{P} (S_j^r(r^2t) > (\mu_j^r + \epsilon_i)r^2t) \\
&\leq (\lfloor (\mu_j^r + \epsilon_i)r^2t \rfloor + 1) K_4 \exp(-K_4' s_i^r) + K_5 \exp(-K_5' r^2t), \tag{8.57}
\end{aligned}$$

where $K_4 = \max\{\exp(\Lambda_j^s(l_0)) : j \in \mathcal{J}_i\} > 0$, $K_4' = \min\{(l_0\epsilon_i/2\mu_j)(\tilde{x}_j^* - \epsilon_i) : j \in \mathcal{J}_i\} > 0$, where $0 < l_0 \in \mathcal{O}_0$ (cf. (3.3)–(3.4)), and where we have used (7.26) to obtain the estimate involving K_4 and K_4' in (8.57). Also, $K_5 = \max\{\exp(\Lambda_j^{s,*}((\mu_j(1 + \epsilon_i/3\mu_j))^{-1})) : j \in \mathcal{J}_i\} > 0$ and $K_5' = \min\{\mu_j\Lambda_j^{s,*}((\mu_j(1 + \epsilon_i/3\mu_j))^{-1}) : j \in \mathcal{J}_i\} > 0$. Here we have used (7.23), and for this we note that since $M^r > 2/\epsilon_i$, we have that $r^2t > 2/\epsilon_i$. Note that K_4, K_4', K_5, K_5' do not depend on t, n , or r . It follows that the last term in (8.56) is bounded by the expression in (8.57) for all $r \geq r^*$, $t > 0$ satisfying $r^2t \geq M^r$, and $j \in \mathcal{J}_i$.

Combining all of the above (from (8.44) onwards), we have for all $r \geq r^*$, $t > 0$

satisfying $r^2t \geq M^r$,

$$\begin{aligned}
& \mathbf{P}\left(R_i^r(s) \geq L_i^r - |\underline{\mathcal{J}}_i| \text{ some } s \in [0, r^2t]\right) \\
& \leq K_1 \exp(-K_1' r^2t) \\
& \quad + (n_i^r - 1) \left\{ 2K_2 \exp(-K_2' L_i^r) + |\mathcal{J}_i| K_3 \exp(-K_3' s_i^r) \right. \\
& \quad + \sum_{j \in \mathcal{J}_i} \left(\lfloor (\mu_j^r + \epsilon_i) r^2t \rfloor + 1 \right) K_4 \exp(-K_4' s_i^r) + |\mathcal{J}_i| K_5 \exp(-K_5' r^2t) \\
& \quad \left. + \sum_{j \in \mathcal{J}_i} \sup_{n < n_i^r} \mathbf{P}(T_{i,j}^{r,n}(s_i) \leq (\tilde{x}_j^* - \epsilon_i) s_i^r, s_i^r \leq \beta_{i,n}^r, \tau_{i,2n-1}^r \leq r^2t) \right\}, \quad (8.58)
\end{aligned}$$

where $K_3 = 1$, $K_3' = \min\{(\mu_j - 2\epsilon_i)(\tilde{x}_j^* - \epsilon_i)\Lambda_j^{s,*}(\mu_j^{-1}(1 + \epsilon_i/2\mu_j)) : j \in \mathcal{J}_i\} > 0$ from (8.56).

Now we consider the “down” excursions of R_i^r . For this, we assume that $\underline{\mathcal{J}}_i \neq \emptyset$, i.e., we assume that class i is a transition class, since (8.58) suffices for the proof (cf. (8.77)) that (I.1) holds when class i is a non-transition class. The treatment of the down excursions is very similar to that for the up excursions.

For each $n \geq 1$, $r \geq 1$, $j \in \underline{\mathcal{J}}_i$, and $s \geq 0$, on $\{d_{\tau_{i,2n-1}}^r < \infty\}$, define shifted processes

$${}^d A_{i,i}^{r,n}(s) = A_i^r(d_{\tau_{i,2n-1}}^r + s) - A_i^r(d_{\tau_{i,2n-1}}^r), \quad (8.59)$$

$${}^d S_{i,j}^{r,n}(s) = S_j^r(T_j^r(d_{\tau_{i,2n-1}}^r) + s) - S_j^r(T_j^r(d_{\tau_{i,2n-1}}^r)), \quad (8.60)$$

$${}^d \check{A}_{i,i}^{r,n}(s) = \sup\{m \geq 0 : \xi_i^r(A_i^r(d_{\tau_{i,2n-1}}^r) + m) - \xi_i^r(A_i^r(d_{\tau_{i,2n-1}}^r)) \leq s\}, \quad (8.61)$$

$$\begin{aligned}
{}^d \check{S}_{i,j}^{r,n}(t) &= \sup\{m \geq 0 : \eta_j^r(S_j^r(T_j^r(d_{\tau_{i,2n-1}}^r)) + m + 1) \\
&\quad - \eta_j^r(S_j^r(T_j^r(d_{\tau_{i,2n-1}}^r)) + 1) \leq s\}, \quad (8.62)
\end{aligned}$$

and, for concreteness, on $\{d_{\tau_{i,2n-1}}^r = \infty\}$ define ${}^d A_{i,i}^{r,n}$, ${}^d S_{i,j}^{r,n}$, ${}^d \check{A}_{i,i}^{r,n}$, ${}^d \check{S}_{i,j}^{r,n}$ to be identically zero. Recall the definition of ${}^d T_{i,j}^{r,n}$ from (7.13). Consider the n^{th} “down” excursion interval for R_i^r . We have that on $\{d_{\tau_{i,2n-1}}^r < \infty\}$,

$${}^d R_i^r(d_{\tau_{i,2n-1}}^r + s) = 1 + \sum_{j \in \underline{\mathcal{J}}_i} {}^d S_{i,j}^{r,n}({}^d T_{i,j}^{r,n}(s)) - {}^d A_{i,i}^{r,n}(s), \quad 0 \leq s \leq d_{\beta_{i,n}^r}^r, \quad (8.63)$$

since $T_{a(i)}^r(d_{T_{i,2n-1}}^r + s) = T_{a(i)}^r(d_{T_{i,2n-1}}^r)$, for $0 \leq s \leq d\beta_{i,n}^r$, and in a similar manner to that for (8.41) we have that,

$$dA_{i,i}^{r,n}(s) \geq d\check{A}_{i,i}^{r,n}(s), \text{ and } dS_{i,j}^{r,n}(s) \leq d\check{S}_{i,j}^{r,n}(s) + 1, \text{ for all } j \in \underline{\mathcal{J}}_i. \quad (8.64)$$

By (7.5), we have that

$$\epsilon_i < \min \left\{ \frac{\lambda_i}{2}, \frac{\mu_j}{2}, x_j^*, \left(\frac{4\lambda_i}{4\lambda_i - d\delta_i} - 1 \right) x_j^*, : j \in \underline{\mathcal{J}}_i \right\}, \quad (8.65)$$

where $d\delta_i$ is defined in (8.20).

In the following, we assume that $r \geq r^*$ and $t > 0$ satisfies $r^2t \geq Mr$. Then $r^2t > 2/\epsilon_i$ by (8.9)–(8.12), and (8.35). Define $d_n^r = \left\lfloor \sum_{j \in \underline{\mathcal{J}}_i} (\mu_j^r + \epsilon_i)r^2t \right\rfloor + |\underline{\mathcal{J}}_i|$. Since each “down” excursion of R_i^r is initiated by a completion of a class i job by some activity $j \in \underline{\mathcal{J}}_i$, using the large deviations bounds for renewal processes given in Lemma 7.6.1, we have the following estimate of the probability that there are at least $d_n^r - 1$ complete down excursions in $[0, r^2t]$:

$$\begin{aligned} \mathbf{P}(d_{T_{i,2}^r}^{r, d_n^r - 1} \leq r^2t) &\leq \mathbf{P}\left(\sum_{j \in \underline{\mathcal{J}}_i} S_j^r(T_j^r(r^2t)) \geq d_n^r\right) \\ &\leq \mathbf{P}\left(\sum_{j \in \underline{\mathcal{J}}_i} S_j^r(r^2t) \geq d_n^r\right) \\ &\leq \sum_{j \in \underline{\mathcal{J}}_i} \mathbf{P}(S_j^r(r^2t) > (\mu_j^r + \epsilon_i)r^2t) \\ &\leq K_6 \exp(-K'_6 r^2t), \end{aligned} \quad (8.66)$$

where $K_6 = |\underline{\mathcal{J}}_i| \max\{\exp(\Lambda_j^{s,*}((\mu_j(1 + \epsilon_i/3\mu_j))^{-1})) : j \in \underline{\mathcal{J}}_i\} > 0$, $K'_6 = \min\{\mu_j \Lambda_j^{s,*}((\mu_j(1 + \epsilon_i/3\mu_j))^{-1}) : j \in \underline{\mathcal{J}}_i\} > 0$ are independent of t and r .

Now (cf. (8.45)),

$$\begin{aligned}
& \mathbf{P}\left(R_i^r(s) \leq |\underline{\mathcal{J}}_i| - L_i^r \text{ some } s \in [\tau_{i,0}^r, r^2t]\right) \\
& \leq \mathbf{P}\left({}^dR_i^r(s) \geq L_i^r - |\underline{\mathcal{J}}_i| \text{ some } s \in [\tau_{i,0}^r, r^2t]\right) \\
& \leq \mathbf{P}\left({}^d\tau_{i,2}^r \leq r^2t\right) \\
& \quad + \sum_{n=1}^{d_{n_i}^r-1} \mathbf{P}\left({}^dR_i^r(s) \geq L_i^r - |\underline{\mathcal{J}}_i| \text{ some } s \in ({}^d\tau_{i,2n-1}^r, {}^d\tau_{i,2n}^r), \right. \\
& \quad \left. {}^d\tau_{i,2n-1}^r \leq r^2t\right). \tag{8.67}
\end{aligned}$$

For each positive integer n , let

$$\begin{aligned}
{}^d\Upsilon_i^{r,n} &= \left\{ {}^dA_{i,i}^{r,n}({}^dS_i^r) \geq (\lambda_i^r - \epsilon_i) {}^dS_i^r; \right. \\
& \quad \left. {}^dS_{i,j}^{r,n}((x_j^* + \epsilon_i) {}^dS_i^r) \leq (\mu_j^r + \epsilon_i)(x_j^* + \epsilon_i) {}^dS_i^r, j \in \underline{\mathcal{J}}_i; \right. \\
& \quad \left. {}^d\Upsilon_{i,j}^{r,n}({}^dS_i^r) < (x_j^* + \epsilon_i) {}^dS_i^r, j \in \underline{\mathcal{J}}_i; {}^d\tau_{i,2n-1}^r \leq r^2t \right\}. \tag{8.68}
\end{aligned}$$

Then,

$$\begin{aligned}
& \mathbf{P}\left({}^dR_i^r(s) \geq L_i^r - |\underline{\mathcal{J}}_i| \text{ some } s \in ({}^d\tau_{i,2n-1}^r, {}^d\tau_{i,2n}^r), {}^d\tau_{i,2n-1}^r \leq r^2t\right) \\
& \leq \mathbf{P}\left({}^d\tau_{i,2n-1}^r \leq r^2t, ({}^d\Upsilon_i^{r,n})^c\right) \\
& \quad + \mathbf{P}\left({}^dR_i^r(s) \geq L_i^r - |\underline{\mathcal{J}}_i| \text{ some } s \in ({}^d\tau_{i,2n-1}^r, {}^d\tau_{i,2n}^r), {}^d\Upsilon_i^{r,n}\right). \tag{8.69}
\end{aligned}$$

Now, on $\mathcal{A}\Upsilon_i^{r,n}$, using (8.65), (8.20)–(8.22), and (8.23), we have for $r \geq r^*$,

$$\begin{aligned}
& 1 + \sum_{j \in \underline{\mathcal{J}}_i} dS_{i,j}^{r,n}(\mathcal{A}T_{i,j}^{r,n}(d_{S_i}^r)) - dA_{i,i}^{r,n}(d_{S_i}^r) \\
& \leq 1 + \sum_{j \in \underline{\mathcal{J}}_i} dS_{i,j}^{r,n}((x_j^* + \epsilon_i) d_{S_i}^r) - dA_{i,i}^{r,n}(d_{S_i}^r) \\
& \leq 1 + \sum_{j \in \underline{\mathcal{J}}_i} (x_j^* + \epsilon_i)(\mu_j^r + \epsilon_i) d_{S_i}^r - (\lambda_i^r - \epsilon_i) d_{S_i}^r \\
& \leq 1 + \left(\sum_{j \in \underline{\mathcal{J}}_i} \frac{4\lambda_i}{4\lambda_i - d\delta_i} x_j^* \cdot \frac{8\lambda_i - 3d\delta_i}{8\lambda_i - 4d\delta_i} \mu_j^r - \frac{2\lambda_i}{2\lambda_i - d\delta_i} \sum_{j \in \underline{\mathcal{J}}_i} x_j^* \mu_j^r \right) d_{S_i}^r \\
& = 1 + \left(\frac{2\lambda_i}{2\lambda_i - d\delta_i} \sum_{j \in \underline{\mathcal{J}}_i} x_j^* \mu_j^r - \frac{\lambda_i d\delta_i}{(4\lambda_i - d\delta_i)(2\lambda_i - d\delta_i)} \sum_{j \in \underline{\mathcal{J}}_i} x_j^* \mu_j^r \right. \\
& \quad \left. - \frac{2\lambda_i}{2\lambda_i - d\delta_i} \sum_{j \in \underline{\mathcal{J}}_i} x_j^* \mu_j^r \right) d_{S_i}^r \\
& = 1 - \left(\frac{\lambda_i d\delta_i}{(4\lambda_i - d\delta_i)(2\lambda_i - d\delta_i)} \sum_{j \in \underline{\mathcal{J}}_i} x_j^* \mu_j^r \right) d_{S_i}^r < 0. \tag{8.70}
\end{aligned}$$

It then follows from (8.63) that for $r \geq r^*$,

$$d\beta_{i,n}^r = d_{\tau_{i,2n}}^r - d_{\tau_{i,2n-1}}^r < d_{S_i}^r \text{ on } \mathcal{A}\Upsilon_i^{r,n}. \tag{8.71}$$

Furthermore, on $\mathcal{A}\Upsilon_i^{r,n}$ for $0 \leq s \leq d_{S_i}^r$,

$$\begin{aligned}
& 1 + \sum_{j \in \underline{\mathcal{J}}_i} dS_{i,j}^{r,n}(\mathcal{A}T_{i,j}^{r,n}(s)) - dA_{i,i}^{r,n}(s) \\
& \leq 1 + \sum_{j \in \underline{\mathcal{J}}_i} (\mu_j^r + \epsilon_i)(x_j^* + \epsilon_i) d_{S_i}^r = L_i^r - (|\underline{\mathcal{J}}_i| + 1) \tag{8.72}
\end{aligned}$$

by (8.10), and hence we have that on $\mathcal{A}\Upsilon_i^{r,n}$, $dR_i^r(s) < L_i^r - |\underline{\mathcal{J}}_i|$ for $s \in (d_{\tau_{i,2n-1}}^r, d_{\tau_{i,2n}}^r)$, whenever $r \geq r^*$. Thus the last probability in (8.69) is zero for all such r .

It remains to estimate

$$\begin{aligned}
& \mathbf{P}\left(d_{\mathcal{T}_{i,2n-1}}^r \leq r^2 t, (d\Upsilon_i^{r,n})^c\right) \\
\leq & \mathbf{P}\left(dA_{i,i}^{r,n}(d_s^r) < (\lambda_i^r - \epsilon_i) d_s^r, d_{\mathcal{T}_{i,2n-1}}^r \leq r^2 t\right) \\
& + \sum_{j \in \underline{\mathcal{J}}_i} \mathbf{P}\left(dS_{i,j}^{r,n}((x_j^* + \epsilon_i) d_s^r) > (\mu_j^r + \epsilon_i)(x_j^* + \epsilon_i) d_s^r, d_{\mathcal{T}_{i,2n-1}}^r \leq r^2 t\right) \\
& + \sum_{j \in \underline{\mathcal{J}}_i} \mathbf{P}\left(dT_{i,j}^{r,n}(d_s^r) \geq (x_j^* + \epsilon_i) d_s^r, d_{\mathcal{T}_{i,2n-1}}^r \leq r^2 t\right), \tag{8.73}
\end{aligned}$$

for $r \geq r^*$ and $t > 0$ satisfying $r^2 t \geq M^r$.

As in (8.56)–(8.57), letting $u_{i,i}^{r,n} = u_i^r(A_i^r(d_{\mathcal{T}_{i,2n-1}}^r) + 1)$, we have for $r \geq r^*$ and $t > 0$ satisfying $r^2 t \geq M^r$ (which implies that $r^2 t > 2/\epsilon_i$),

$$\begin{aligned}
& \mathbf{P}\left(dA_{i,i}^{r,n}(d_s^r) < (\lambda_i^r - \epsilon_i) d_s^r, d_{\mathcal{T}_{i,2n-1}}^r \leq r^2 t\right) \\
\leq & K_7 \exp(-K_7' d_s^r) \\
& + \mathbf{E}\left(1_{\{d_{\mathcal{T}_{n,i}}^r \in \mathbb{N}^I \times \mathbb{N}^J\}} \mathbf{P}\left(u_{i,i}^{r,n} > \frac{\epsilon_i}{2\lambda_i^r} d_s^r, d_{\mathcal{T}_{i,2n-1}}^r \leq r^2 t \mid \mathcal{F}_{d_{\mathcal{T}_{n,i}}^r}^r\right)\right) \\
\leq & K_7 \exp(-K_7' d_s^r) + (\lfloor (\lambda_i^r + \epsilon_i) r^2 t \rfloor + 1) K_8 \exp(-K_8' d_s^r) \\
& + K_9 \exp(-K_9' r^2 t), \tag{8.74}
\end{aligned}$$

by Lemmas 7.5.2–7.6.1 and (7.26), where $K_7 = 1$, $K_7' = (\lambda_i - 2\epsilon_i)\Lambda_i^{a,*}(\lambda_i^{-1}(1 + \epsilon_i/2\lambda_i)) > 0$, $K_8 = \exp(\Lambda_i^a(l_0)) > 0$, $K_8' = l_0\epsilon_i/2\lambda_i > 0$, $0 < l_0 \in \mathcal{O}_0$, $K_9 = \exp(\Lambda_i^{a,*}((\lambda_i(1 + \epsilon_i/3\lambda_i))^{-1})) > 0$, and $K_9' = \lambda_i\Lambda_i^{a,*}((\lambda_i(1 + \epsilon_i/3\lambda_i))^{-1}) > 0$, which are all independent of t , n , and r .

Finally, for all $r \geq r^*$, $t > 0$ satisfying $r^2 t \geq M^r$, by Lemmas 7.5.2–7.6.1 (since $(x_j^* + \epsilon_i) d_s^r > 4/\epsilon_i$ by (8.9), (8.12), and (8.35)), we have

$$\begin{aligned}
& \sum_{j \in \underline{\mathcal{J}}_i} \mathbf{P}\left(dS_{i,j}^{r,n}((x_j^* + \epsilon_i) d_s^r) > (\mu_j^r + \epsilon_i)(x_j^* + \epsilon_i) d_s^r, d_{\mathcal{T}_{i,2n-1}}^r \leq r^2 t\right) \\
\leq & \sum_{j \in \underline{\mathcal{J}}_i} \mathbf{E}\left(1_{\{d_{\mathcal{T}_{n,i}}^r \in \mathbb{N}^I \times \mathbb{N}^J\}} \mathbf{P}\left(d\check{S}_{i,j}^{r,n}((x_j^* + \epsilon_i) d_s^r) \right. \right. \\
& \quad \left. \left. > (\mu_j^r + \epsilon_i/2)(x_j^* + \epsilon_i) d_s^r, d_{\mathcal{T}_{i,2n-1}}^r \leq r^2 t \mid \mathcal{F}_{d_{\mathcal{T}_{n,i}}^r}^r\right)\right) \\
\leq & K_{10} \exp(-K_{10}' d_s^r), \tag{8.75}
\end{aligned}$$

where $K_{10} = |\underline{\mathcal{J}}_i| \max\{\exp(\Lambda_j^{s,*}((\mu_j(1 + \epsilon_i/6\mu_j)^{-1})) : j \in \underline{\mathcal{J}}_i)\} > 0$, $K'_{10} = \min\{\mu_j(x_j^* + \epsilon_i)\Lambda_j^{s,*}((\mu_j(1 + \epsilon_i/6\mu_j)^{-1}) : j \in \underline{\mathcal{J}}_i)\} > 0$, which are independent of t , n , and r , and where we have used (8.64) together with the fact that $(x_j^* + \epsilon_i)^{d_{s_i^r}} \geq 2/\epsilon_i$ (cf. the statement preceding (8.75)) in the first inequality.

Combining all of the above (from (8.66) onwards), we have for all $r \geq r^*$, $t > 0$ satisfying $r^2t \geq M^r$,

$$\begin{aligned}
& \mathbf{P}\left(R_i^r(s) \leq |\underline{\mathcal{J}}_i| - L_i^r \text{ some } s \in [0, r^2t]\right) \\
&= \mathbf{P}\left({}^dR_i^r(s) \geq L_i^r - |\underline{\mathcal{J}}_i| \text{ some } s \in [0, r^2t]\right) \\
&\leq K_6 \exp(-K'_6 r^2t) + (q_{n_i^r} - 1) \left\{ K_7 \exp(-K'_7 d_{s_i^r}^r) \right. \\
&\quad \left. + (\lfloor (\lambda_i^r + \epsilon_i)r^2t \rfloor + 1) K_8 \exp(-K'_8 d_{s_i^r}^r) + K_9 \exp(-K'_9 r^2t) \right. \\
&\quad \left. + K_{10} \exp(-K'_{10} d_{s_i^r}^r) \right. \\
&\quad \left. + \sum_{j \in \underline{\mathcal{J}}_i} \sup_{n < q_{n_i^r}} \mathbf{P}\left({}^dT_{i,j}^{r,n}(d_{s_i^r}^r) \geq (x_j^* + \epsilon_i)^{d_{s_i^r}^r}, d_{T_{i,2n-1}}^r \leq r^2t\right)\right\}. \quad (8.76)
\end{aligned}$$

On combining the results (8.58) and (8.76) for the up and down excursions, assumptions (i), (ii), and (iv) of Lemma 8.3.1, the definitions of n_i^r , $q_{n_i^r}$, s_i^r , $d_{s_i^r}^r$, L_i^r , and the fact that $s_i^r \geq L_0^r$ and $d_{s_i^r}^r \geq L_0^r$ for all $i \in \mathcal{I}$ (by (8.9) and (8.12)), it follows that for $r \geq r^*$ and $t > 0$ satisfying $r^2t \geq M^r$, we have

$$\begin{aligned}
& \mathbf{P}\left(\sup_{\tau_{i,0}^r \leq s \leq r^2t} |R_i^r(s)| \geq L_i^r - |\underline{\mathcal{J}}_i|\right) \\
&\leq p_{1,i}(r^2t) \left(C_{1,i}^{(1)} \exp(-C_{1,i}^{(2)} L_0^r) + C_{1,i}^{(3)} \exp(-C_{1,i}^{(4)} r^2t) \right), \quad (8.77)
\end{aligned}$$

where $C_{1,i}^{(m)} > 0$, $m = 1, 2, 3, 4$, and $p_{1,i}$ is a polynomial (of degree at most $i + 1$) with non-negative coefficients, and where the constants and the polynomial are independent of r and t . Note that since $i \geq 1$, terms involving $(r^2t)^2$ are absorbed in the polynomial $p_{1,i}(r^2t)$.

Proof of (I.2). Since $\underline{\mathcal{K}}_i = \emptyset$, if i is a non-transition class, (I.2) trivially holds in this case. So it suffices to consider the case when i is a transition class. Note that

for $0 \leq s \leq \tau_{i,0}^r$,

$$Q_i^r(s) = A_i^r(s) - \sum_{j \in \underline{\mathcal{J}}_i} S_j^r(T_j^r(s)), \quad (8.78)$$

since activity $a(i)$ will be turned off for such s . By (7.5) and (8.24), we have

$$0 < \tilde{\epsilon}_i < \min \left\{ 1, \frac{\lambda_i - \sum_{j \in \underline{\mathcal{J}}_i} x_j^* \mu_j}{4(|\underline{\mathcal{J}}_i| + 2)}, \frac{\lambda_i}{2}, \frac{\mu_j}{2} : j \in \underline{\mathcal{J}}_i \right\}, \quad (8.79)$$

since

$$\epsilon_i < \min \left\{ \frac{\lambda_i - \sum_{j \in \underline{\mathcal{J}}_i} \mu_j}{2048(\mathbf{J} + 2) \sum_{j \in \underline{\mathcal{J}}} \mu_j}, \frac{1}{2 \sum_{j \in \underline{\mathcal{J}}_i} \mu_j}, \frac{\lambda_i}{4 \sum_{j \in \underline{\mathcal{J}}_i} \mu_j}, \frac{\mu_j}{4 \sum_{j \in \underline{\mathcal{J}}_i} \mu_j} : j \in \underline{\mathcal{J}}_i \right\}, \quad (8.80)$$

$x_j^* \leq 1$ for all $j \in \underline{\mathcal{J}}_i$, and $1 > 1 - 2|\underline{\mathcal{J}}_i|\epsilon_i > 1/2$.

We have,

$$\begin{aligned} & \mathbf{P}(I_k^r(\tau_{i,0}^r) \geq t_i^r) \\ & \leq \mathbf{P}(\tau_{i,0}^r \geq t_i^r) \\ & \leq \mathbf{P}\left(A_i^r(t_i^r) - \sum_{j \in \underline{\mathcal{J}}_i} S_j^r(T_j^r(t_i^r)) \leq L_i^r\right), \quad \text{by (8.78),} \\ & \leq \mathbf{P}\left(A_i^r(t_i^r) \geq (\lambda_i^r - \tilde{\epsilon}_i)t_i^r, T_j^r(t_i^r) \leq (x_j^* + \epsilon_i)t_i^r \text{ for all } j \in \underline{\mathcal{J}}_i, \right. \\ & \quad \left. S_j^r((x_j^* + \epsilon_i)t_i^r) \leq (\mu_j^r + \tilde{\epsilon}_i)(x_j^* + \epsilon_i)t_i^r \text{ for all } j \in \underline{\mathcal{J}}_i, \right. \\ & \quad \left. ((\lambda_i^r - \tilde{\epsilon}_i) - \sum_{j \in \underline{\mathcal{J}}_i} (x_j^* + \epsilon_i)(\mu_j^r + \tilde{\epsilon}_i))t_i^r \leq L_i^r\right) \\ & + \mathbf{P}\left(A_i^r(t_i^r) < (\lambda_i^r - \tilde{\epsilon}_i)t_i^r\right) \\ & + \sum_{j \in \underline{\mathcal{J}}_i} \mathbf{P}\left(S_j^r((x_j^* + \epsilon_i)t_i^r) > (\mu_j^r + \tilde{\epsilon}_i)(x_j^* + \epsilon_i)t_i^r\right) \\ & + \sum_{j \in \underline{\mathcal{J}}_i} \mathbf{P}\left(T_j^r(t_i^r) > (x_j^* + \epsilon_i)t_i^r\right). \end{aligned} \quad (8.81)$$

Now, for $r \geq r^*$, using (8.25), (8.26), (8.24), (8.79), (8.8), and the fact that $x_j^* < 1$

for all $j \in \underline{\mathcal{J}}_i$, we have

$$\begin{aligned}
& \left(\lambda_i^r - \sum_{j \in \underline{\mathcal{J}}_i} x_j^* \mu_j^r - \tilde{\epsilon}_i - \tilde{\epsilon}_i \sum_{j \in \underline{\mathcal{J}}_i} x_j^* - \epsilon_i \sum_{j \in \underline{\mathcal{J}}_i} (\mu_j^r + \tilde{\epsilon}_i) \right) t_i^r \\
& \geq \left(\frac{1}{2} \left(\lambda_i - \sum_{j \in \underline{\mathcal{J}}_i} x_j^* \mu_j \right) - (|\underline{\mathcal{J}}_i| + 1) \tilde{\epsilon}_i - \epsilon_i \sum_{j \in \underline{\mathcal{J}}_i} (\mu_j + 2\tilde{\epsilon}_i) \right) t_i^r \\
& \geq \left(\frac{1}{2} \left(\lambda_i - \sum_{j \in \underline{\mathcal{J}}_i} x_j^* \mu_j \right) - (|\underline{\mathcal{J}}_i| + 2) \tilde{\epsilon}_i \right) t_i^r > L_i^r. \tag{8.82}
\end{aligned}$$

Hence, the first probability in the last expression of (8.81) is zero. From Lemma 7.6.1 (since $(x_j^* + \epsilon_i)t_i^r > 2/\tilde{\epsilon}_i$, for all $j \in \underline{\mathcal{J}}_i$, by (8.9), (8.12), and (8.35)) we have for all $j \in \underline{\mathcal{J}}_i$, $r \geq r^*$,

$$\begin{aligned}
& \mathbf{P}(S_j^r((x_j^* + \epsilon_i)t_i^r) > (\mu_j^r + \tilde{\epsilon}_i)(x_j^* + \epsilon_i)t_i^r) \\
& \leq \exp \left(-(\mu_j(x_j^* + \epsilon_i)t_i^r - 1) \Lambda_j^{s,*} \left(\frac{1}{\mu_j} \left(\frac{1}{1 + \frac{\tilde{\epsilon}_i}{3\mu_j}} \right) \right) \right) \\
& \leq K_{11} \exp(-K'_{11}t_i^r), \tag{8.83}
\end{aligned}$$

where $K_{11} = \max\{\exp(\Lambda_j^{s,*}((\mu_j(1 + \tilde{\epsilon}_i/3\mu_j))^{-1})) : j \in \underline{\mathcal{J}}_i\} > 0$, and $K'_{11} = \min\{\mu_j(x_j^* + \epsilon_i)\Lambda_j^{s,*}((\mu_j(1 + \tilde{\epsilon}_i/3\mu_j))^{-1}) : j \in \underline{\mathcal{J}}_i\} > 0$. Using (7.26) in conjunction with Lemma 7.6.1, we have for $r \geq r^*$, $0 < l_0 \in \mathcal{O}_0$,

$$\begin{aligned}
\mathbf{P}(A_i^r(t_i^r) < (\lambda_i^r - \tilde{\epsilon}_i)t_i^r) & \leq \exp \left(-(\lambda_i - 2\tilde{\epsilon}_i)t_i^r \Lambda_i^{a,*} \left(\frac{1}{\lambda_i} \left(1 + \frac{\tilde{\epsilon}_i}{2\lambda_i} \right) \right) \right) \\
& \quad + \exp \left(-\frac{l_0 \tilde{\epsilon}_i t_i^r}{2\lambda_i} \right) \exp(\Lambda_i^a(l_0)) \\
& \leq K_{12} \exp(-K'_{12}t_i^r), \tag{8.84}
\end{aligned}$$

where $K_{12} = \max\{1, \exp(\Lambda_i^a(l_0))\} > 0$, $K'_{12} = \min\{(\lambda_i - 2\tilde{\epsilon}_i)\Lambda_i^{a,*}(\lambda_i^{-1}(1 + \tilde{\epsilon}_i/2\lambda_i)), l_0 \tilde{\epsilon}_i/2\lambda_i\} > 0$. Note that K_{11} , K'_{11} , K_{12} , K'_{12} do not depend on r .

Then, combining (8.81)–(8.84), we have for each $k \in \underline{\mathcal{K}}_i$, for all $r \geq r^*$,

$$\begin{aligned}
\mathbf{P}(I_k^r(\tau_{i,0}^r \geq t_i^r) & \leq K_{11} \exp(-K'_{11}t_i^r) + K_{12} \exp(-K'_{12}t_i^r) \\
& \quad + \sum_{j \in \underline{\mathcal{J}}_i} \mathbf{P}(T_j^r(t_i^r) > (x_j^* + \epsilon_i)t_i^r). \tag{8.85}
\end{aligned}$$

By assumption (iii) of Lemma 8.3.1, the definitions of t_i^r and L_i^r , and the fact that $t_i^r \geq L_0^r$ (by (8.9) and (8.12)) for $r \geq r^*$, it follows that for all $k \in \underline{\mathcal{K}}_i$, $r \geq r^*$, and $t > 0$ satisfying $r^2t \geq M^r$, we have

$$\begin{aligned} & \mathbf{P} \left(I_k^r(\tau_{i,0}^r) \geq t_i^r \right) \\ & \leq p_{2,i}(r^2t) \left(C_{2,i}^{(1)} \exp(-C_{2,i}^{(2)}L_0^r) + C_{2,i}^{(3)} \exp(-C_{2,i}^{(4)}r^2t) \right), \end{aligned} \quad (8.86)$$

where $p_{2,i}$ is a polynomial (of degree at most i) with non-negative coefficients, and $C_{2,i}^{(m)} > 0$, for $m = 1, 2, 3, 4$, and where the constants and the polynomial do not depend on t or r .

Proof of (I.3). Since $\underline{\mathcal{K}}_i = \emptyset$ if i is not a transition class, (I.3) holds trivially in this case. So suppose i is a transition class. Note that under the threshold policy, since class i (being above server k) is the lowest priority class for server $k \in \underline{\mathcal{K}}_i$, I_k^r can increase only when $Q_i^r \leq |\underline{\mathcal{K}}_i| = |\underline{\mathcal{J}}_i|$. The bound of $|\underline{\mathcal{J}}_i|$ occurs here because there may be a class i job in service or in suspension at each of the other $|\underline{\mathcal{J}}_i|$ servers ($|\underline{\mathcal{J}}_i| - 1$ servers below and one server above i) that can serve class i . In particular, if server $k \in \underline{\mathcal{K}}_i$ incurs some idletime in $[\tau_{i,0}^r, r^2t]$, i.e., $\tau_{i,0}^r < r^2t$ and $I_k^r(r^2t) - I_k^r(\tau_{i,0}^r) > 0$, then $R_i^r(s) \leq -L_i^r + |\underline{\mathcal{J}}_i|$ for some $s \in [\tau_{i,0}^r, r^2t]$. Thus, for $r \geq r^*$, $t > 0$ satisfying $r^2t \geq M^r$, and $k \in \underline{\mathcal{K}}_i$,

$$\begin{aligned} & \mathbf{P} \left(I_k^r(r^2t) - I_k^r(\tau_{i,0}^r) > 0, \tau_{i,0}^r < r^2t \right) \\ & \leq \mathbf{P} \left(\inf_{\tau_{i,0}^r \leq s \leq r^2t} R_i^r(s) \leq -L_i^r + |\underline{\mathcal{J}}_i| \right) \\ & \leq p_{3,i}(r^2t) \left(C_{3,i}^{(1)} \exp(-C_{3,i}^{(2)}L_0^r) + C_{3,i}^{(3)} \exp(-C_{3,i}^{(4)}r^2t) \right), \end{aligned} \quad (8.87)$$

where $p_{3,i}$ is a polynomial (of degree at most $i + 1$) with non-negative coefficients, and $C_{3,i}^{(m)} > 0$, for $m = 1, 2, 3, 4$, by (8.77), where the constants and the polynomial do not depend on t or r .

□

8.5 Proof of Lemma 8.3.2: Estimates on Allocations for Activities Immediately Below Each Buffer

Proof of Lemma 8.3.2. Fix $i \in \mathcal{I} \setminus \{i^*\}$. Suppose that i is a transition buffer, so that $\underline{\mathcal{J}}_i \neq \emptyset$, and assume that (I) and (II) hold with i' in place of i , for all $i' < i$ (for $i = 1$ this is a vacuous assumption). In the following, recall (cf. Section 1.1) that a sum over an empty set is defined to equal zero. In particular, the results below hold even if $\underline{\mathcal{I}}_k = \emptyset$.

Proof of (i). For $r \geq r^*$, $t > 0$ satisfying $r^2 t \geq M^r$, $j \in \underline{\mathcal{J}}_i$, and $k = k(j)$, we have

$$\begin{aligned}
& \mathbf{P}(T_{i,j}^{r,n}(s_i^r) \leq (x_j^* - \epsilon_i)s_i^r, s_i^r \leq \beta_{i,n}^r, \tau_{i,2n-1}^r \leq r^2 t) \\
&= \mathbf{P}\left(s_i^r - \sum_{i' \in \underline{\mathcal{I}}_k} T_{i,a(i')}^{r,n}(s_i^r) \leq (x_j^* - \epsilon_i)s_i^r, s_i^r \leq \beta_{i,n}^r, \tau_{i,2n-1}^r \leq r^2 t\right) \\
&= \mathbf{P}\left(\sum_{i' \in \underline{\mathcal{I}}_k} T_{i,a(i')}^{r,n}(s_i^r) \geq \sum_{i' \in \underline{\mathcal{I}}_k} x_{a(i')}^* s_i^r + \epsilon_i s_i^r, s_i^r \leq \beta_{i,n}^r, \tau_{i,2n-1}^r \leq r^2 t\right) \\
&\leq \sum_{i' \in \underline{\mathcal{I}}_k} \mathbf{P}\left(T_{i,a(i')}^{r,n}(s_i^r) \geq (x_{a(i')}^* + \epsilon_{i'})s_i^r, s_i^r \leq \beta_{i,n}^r, \tau_{i,2n-1}^r \leq r^2 t\right). \quad (8.88)
\end{aligned}$$

The first equality holds since server k , $k \in \underline{\mathcal{K}}_i$, does not incur any idle time in the n^{th} (up) excursion interval for R_i^r which is of length $\beta_{i,n}^r$, and the second equality holds since $\sum_{i' \in \underline{\mathcal{I}}_k} x_{a(i')}^* + x_j^* = 1$. For the last inequality, we have used the fact that $\epsilon_i \geq \mathbf{I}\epsilon_{i'} \geq |\underline{\mathcal{I}}_k| \epsilon_{i'}$, by (7.7) (since $\gamma_l \leq 1$ for all $l \in \mathcal{I}$, and the fact that $i' \in \underline{\mathcal{I}}_k$ satisfies $i' < i$ by the ordering assumed for the buffer numbering). Hence, for $r \geq r^*$, $t > 0$ satisfying $r^2 t \geq M^r$,

$$\begin{aligned}
& \sup_{n \geq 1} \mathbf{P}\left(T_{i,j}^{r,n}(s_i^r) \leq (x_j^* - \epsilon_i)s_i^r, s_i^r \leq \beta_{i,n}^r, \tau_{i,2n-1}^r \leq r^2 t\right) \\
&\leq p_{7,i}(r^2 t) \left(C_{7,i}^{(1)} \exp(-C_{7,i}^{(2)} L_0^r) + C_{7,i}^{(3)} \exp(-C_{7,i}^{(4)} r^2 t) \right), \quad (8.89)
\end{aligned}$$

since (II.1) was assumed to hold with i' in place of i for all $i' < i$ and $|\underline{\mathcal{I}}_k| < \infty$, and where $p_{7,i}$ is a polynomial (of degree at most i) with positive coefficients, $C_{7,i}^{(m)} > 0$

for $m = 1, 2, 3, 4$; the polynomial and the constants do not depend on t , n , or r . Thus, (i) of Lemma 8.3.1 holds.

Proof of (ii). For $r \geq r^*$, $t > 0$ satisfying $r^2 t \geq M^r$, $j \in \underline{\mathcal{J}}_i$, and $k = k(j)$, we have

$$\begin{aligned}
& \mathbf{P}(dT_{i,j}^{r,n}(d_{\mathcal{S}_i}^r) \geq (x_j^* + \epsilon_i) d_{\mathcal{S}_i}^r, d_{\mathcal{T}_{i,2n-1}}^r \leq r^2 t) \\
&= \mathbf{P}\left(d_{\mathcal{S}_i}^r - \sum_{i' \in \underline{\mathcal{I}}_k} dT_{i,a(i')}^{r,n}(d_{\mathcal{S}_i}^r) - (I_k^r(d_{\mathcal{T}_{i,2n-1}}^r + d_{\mathcal{S}_i}^r) - I_k^r(d_{\mathcal{T}_{i,2n-1}}^r))\right. \\
&\quad \left. \geq (x_j^* + \epsilon_i) d_{\mathcal{S}_i}^r, d_{\mathcal{T}_{i,2n-1}}^r \leq r^2 t\right) \\
&\leq \mathbf{P}\left(\sum_{i' \in \underline{\mathcal{I}}_k} dT_{i,a(i')}^{r,n}(d_{\mathcal{S}_i}^r) \leq \sum_{i' \in \underline{\mathcal{I}}_k} x_{a(i')}^* d_{\mathcal{S}_i}^r - \epsilon_i d_{\mathcal{S}_i}^r, d_{\mathcal{T}_{i,2n-1}}^r \leq r^2 t\right) \\
&\leq \sum_{i' \in \underline{\mathcal{I}}_k} \mathbf{P}\left(dT_{i,a(i')}^{r,n}(d_{\mathcal{S}_i}^r) \leq (x_{a(i')}^* - \epsilon_i) d_{\mathcal{S}_i}^r, d_{\mathcal{T}_{i,2n-1}}^r \leq r^2 t\right). \tag{8.90}
\end{aligned}$$

In the above, the equality follows from (2.9); the first inequality holds since the idletime process, I_k^r , is non-decreasing and non-negative, and since $\sum_{i' \in \underline{\mathcal{I}}_{k'}} x_{a(i')}^* + x_j^* = 1$. Note also for the last inequality that $\epsilon_i \geq \mathbf{I}\epsilon_{i'} \geq |\underline{\mathcal{I}}_k| \epsilon_{i'}$. Hence, for $r \geq r^*$, $t > 0$ satisfying $r^2 t \geq M^r$,

$$\begin{aligned}
& \sup_{n \geq 1} \mathbf{P}\left(dT_{i,j}^{r,n}(d_{\mathcal{S}_i}^r) \geq (x_j^* + \epsilon_i) d_{\mathcal{S}_i}^r, d_{\mathcal{T}_{i,2n-1}}^r \leq r^2 t\right) \\
&\leq p_{8,i}(r^2 t) \left(C_{8,i}^{(1)} \exp(-C_{8,i}^{(2)} L_0^r) + C_{8,i}^{(3)} \exp(-C_{8,i}^{(4)} r^2 t) \right), \tag{8.91}
\end{aligned}$$

since (II.2) was assumed to hold with i' in place of i for all $i' < i$, and $|\underline{\mathcal{I}}_k| < \infty$, and where $p_{8,i}$ is a polynomial (of degree at most i), $C_{8,i}^{(m)} > 0$ for $m = 1, 2, 3, 4$. The polynomial and the constants do not depend on t , n , or r . Thus, (ii) of Lemma 8.3.1 holds.

Proof of (iii). For $j \in \underline{\mathcal{J}}_i$, and $k = k(j)$, we have in a similar manner to that for

the proof of (ii) above that for $r \geq r^*$,

$$\begin{aligned}
& \mathbf{P}\left(T_j^r(t_i^r) \geq (x_j^* + \epsilon_i)t_i^r\right) \\
& \leq \mathbf{P}\left(t_i^r - \sum_{i' \in \mathcal{I}_k} T_{a(i')}^r(t_i^r) - I_k^r(t_i^r) \geq (x_j^* + \epsilon_i)t_i^r\right) \\
& \leq \mathbf{P}\left(\sum_{i' \in \mathcal{I}_k} T_{a(i')}^r(t_i^r) \leq \sum_{i' \in \mathcal{I}_k} (x_{a(i')}^* - \epsilon_{i'})t_i^r\right) \\
& \leq \sum_{i' \in \mathcal{I}_k} \mathbf{P}\left(T_{a(i')}^r(t_i^r) \leq (x_{a(i')}^* - \epsilon_{i'})t_i^r\right). \tag{8.92}
\end{aligned}$$

Hence, for $r \geq r^*$, $t > 0$ satisfying $r^2 t \geq M^r$,

$$\begin{aligned}
& \mathbf{P}\left(T_j^r(t_i^r) \geq (x_j^* + \epsilon_i)t_i^r\right) \\
& \leq p_{9,i}(r^2 t) \left(C_{9,i}^{(1)} \exp(-C_{9,i}^{(2)} L_0^r) + C_{9,i}^{(3)} \exp(-C_{9,i}^{(4)} r^2 t) \right), \tag{8.93}
\end{aligned}$$

since (II.3) was assumed to hold with i' in place of i , and $\iota > i'$, for all $i' < i$, and where $p_{9,i}$ is a polynomial (of degree at most i), $C_{9,i}^{(m)} > 0$ for $m = 1, 2, 3, 4$. The polynomial and the constants do not depend on t or r . Thus, (iii) of Lemma 8.3.1 holds. \square

8.6 Proof of Lemma 8.3.3: Estimate on Allocation for the Activity Immediately Above Each Buffer

Proof of Lemma 8.3.3. Fix $i \in \mathcal{I} \setminus \{i^*\}$, and let $k = k(a(i))$. Assume that (I) and (II) hold with i' in place of i for all $i' < i$. Note that by the priorities assigned to buffers by server k , we have that

$$\sum_{\substack{i' \in \mathcal{I}_k \\ i' \leq i}} T_{i,a(i')}^{r,n}(s) = s, \quad \text{for } 0 \leq s \leq \beta_{i,n}^r, \tag{8.94}$$

since activities that have lower priority than activity $a(i)$ will not be “on” and server k will not idle in the n^{th} up excursion interval for R_i^r . We then have for

$r \geq r^*$, $t > 0$ satisfying $r^2 t \geq M^r$,

$$\begin{aligned}
& \mathbf{P} \left(T_{i,a(i)}^{r,n}(s_i^r) \leq (\hat{x}_{i,i} - \epsilon_i) s_i^r, s_i^r \leq \beta_{i,n}^r, \tau_{i,2n-1}^r \leq r^2 t \right) \\
&= \mathbf{P} \left(s_i^r - \sum_{\substack{i' \in \underline{\mathcal{I}}_k \\ i' < i}} T_{i,a(i')}^{r,n}(s_i^r) \leq (\hat{x}_{i,i} - \epsilon_i) s_i^r, s_i^r \leq \beta_{i,n}^r, \tau_{i,2n-1}^r \leq r^2 t \right) \\
&= \mathbf{P} \left(\sum_{\substack{i' \in \underline{\mathcal{I}}_k \\ i' < i}} T_{i,a(i')}^{r,n}(s_i^r) \geq \sum_{\substack{i' \in \underline{\mathcal{I}}_k \\ i' < i}} \hat{x}_{i,i'} s_i^r + \epsilon_i s_i^r, s_i^r \leq \beta_{i,n}^r, \tau_{i,2n-1}^r \leq r^2 t \right) \\
&\leq \mathbf{P} \left(\sum_{\substack{i' \in \underline{\mathcal{I}}_k \\ i' < i}} T_{i,a(i')}^{r,n}(s_i^r) \geq \sum_{\substack{i' \in \underline{\mathcal{I}}_k \\ i' < i}} x_{a(i')}^* s_i^r + \epsilon_i s_i^r, s_i^r \leq \beta_{i,n}^r, \tau_{i,2n-1}^r \leq r^2 t \right) \\
&\leq \sum_{\substack{i' \in \underline{\mathcal{I}}_k \\ i' < i}} \mathbf{P} \left(T_{i,a(i')}^{r,n}(s_i^r) \geq (x_{a(i')}^* + \epsilon_{i'}) s_i^r, s_i^r \leq \beta_{i,n}^r, \tau_{i,2n-1}^r \leq r^2 t \right). \quad (8.95)
\end{aligned}$$

The first equality follows by (8.94) and the second uses (7.3). The second to last inequality holds since $\hat{x}_{i,i'} \geq x_{a(i')}^*$ for all $i' \in \underline{\mathcal{I}}_k$, $i' < i$, and the last inequality follows since $\epsilon_i \geq \mathbf{I}\epsilon_{i'} \geq |\{i' \in \underline{\mathcal{I}}_k : i' < i\}| \epsilon_{i'}$. (Notice that if buffer i is the highest priority buffer for server k , i.e., $\{i' \in \underline{\mathcal{I}}_k : i' < i\} = \emptyset$, then the first probability in (8.95) is zero since $T_{i,a(i)}^{r,n}(s_i^r) = s_i^r$ for $s_i^r \leq \beta_{i,n}^r$, by (8.94).) Hence, for $r \geq r^*$, $t > 0$ satisfying $r^2 t \geq M^r$,

$$\begin{aligned}
& \sup_{n \geq 1} \mathbf{P} \left(T_{i,a(i)}^{r,n}(s_i^r) \leq (\hat{x}_{i,i} - \epsilon_i) s_i^r, s_i^r \leq \beta_{i,n}^r, \tau_{i,2n-1}^r \leq r^2 t \right) \\
&\leq p_{10,i}(r^2 t) \left(C_{10,i}^{(1)} \exp(-C_{10,i}^{(2)} L_0^r) + C_{10,i}^{(3)} \exp(-C_{10,i}^{(4)} r^2 t) \right), \quad (8.96)
\end{aligned}$$

since (II.1) was assumed to hold with i' in place of i , for all $i' < i$, and where $p_{10,i}$ is a polynomial (of degree at most i), $C_{10,i}^{(m)} > 0$ for $m = 1, 2, 3, 4$. The polynomial and the constants do not depend on t , n , or r . Thus (iv) of Lemma 8.3.1 holds. \square

8.7 Proof of Lemma 8.3.4: Transition Between Layers in the Server-Buffer Tree

We will show below that (II) holds for $i \in \mathcal{I} \setminus \{i^*\}$ given that (I) and (II) hold with i' in place of i , for all $i' < i$.

Proof of Lemma 8.3.4. Fix $i \in \mathcal{I} \setminus \{i^*\}$. Assume that (I) and (II) hold with i' in place of i , for all $i' < i$. Then, from Lemmas 8.3.1–8.3.3, we have that (I) holds for i as well. By (7.5) and (7.7), we have, for any $\iota \in \mathcal{I}$ satisfying $\iota > i$,

$$\epsilon_i < \min \left\{ \frac{\prod_{m=1}^{\mathbf{I}} \gamma_m}{\mathbf{I}}, \frac{\mu_{a(i)}}{1024\lambda_\iota}, \frac{\mu_{a(i)}}{1024 \sum_{j \in \mathcal{J}_i} (\mu_j + \epsilon_\iota)}, \frac{\lambda_\iota - \sum_{j \in \underline{\mathcal{J}}_i} x_j^* \mu_j}{18\lambda_\iota}, \frac{\mu_{a(i)}}{16(\lambda_\iota - \sum_{j \in \underline{\mathcal{J}}_i} x_j^* \mu_j)} \right\}, \quad (8.97)$$

since $\epsilon_l < \hat{\epsilon}$, for all $l \in \mathcal{I}$. To validate the denominator in the third term in (8.97), we note that $1024 \sum_{j \in \mathcal{J}_i} (\mu_j + \epsilon_\iota) \leq 1024(\mu_{\text{sum}} + |\mathcal{J}_i| \hat{\epsilon}) \leq 2048\mu_{\text{sum}}$, which implies that the inequality holds for the third term. For the fourth and fifth term recall that $\lambda_\iota - \sum_{j \in \underline{\mathcal{J}}_i} x_j^* \mu_j = x_{a(\iota)}^* \mu_{a(\iota)} > 0$ for all $\iota \in \mathcal{I}$.

Proof of (II.1). For $r \geq 1$, $n \geq 1$, $s \geq 0$, $\iota \in \mathcal{I}$ such that $\iota > i$, $j \in \mathcal{J}_i$, $k = k(j)$, on $\{\tau_{\iota, 2n-1}^r < \infty\}$, define

$$A_{\iota, i}^{r, n}(s) = A_i^r(\tau_{\iota, 2n-1}^r + s) - A_i^r(\tau_{\iota, 2n-1}^r), \quad (8.98)$$

$$S_{\iota, j}^{r, n}(s) = S_j^r(T_j^r(\tau_{\iota, 2n-1}^r) + s) - S_j^r(T_j^r(\tau_{\iota, 2n-1}^r)). \quad (8.99)$$

$$\check{S}_{\iota, j}^{r, n}(s) = \sup\{m \geq 0 : \eta_j^r(S_j^r(T_j^r(\tau_{\iota, 2n-1}^r)) + m) - \eta_j^r(S_j^r(T_j^r(\tau_{\iota, 2n-1}^r))) \leq s\}, \quad (8.100)$$

$$\check{A}_{\iota, i}^{r, n}(s) = \sup\{m \geq 0 : \xi_i^r(A_i^r(\tau_{\iota, 2n-1}^r) + m + 1) - \xi_i^r(A_i^r(\tau_{\iota, 2n-1}^r) + 1) \leq s\}, \quad (8.101)$$

$$I_{\iota, k}^{r, n}(s) = I_k^r(\tau_{\iota, 2n-1}^r + s) - I_k^r(\tau_{\iota, 2n-1}^r), \quad (8.102)$$

and for concreteness on $\{\tau_{\iota, 2n-1}^r = \infty\}$, we define $A_{\iota, i}^{r, n}$, $S_{\iota, j}^{r, n}$, $\check{S}_{\iota, j}^{r, n}$, $\check{A}_{\iota, i}^{r, n}$, $I_{\iota, k}^{r, n}$, to be

identically zero. Then we have, for all $s \geq 0$,

$$\check{A}_{\iota,i}^{r,n}(s) \geq A_{\iota,i}^{r,n}(s) - 1, \quad S_{\iota,j}^{r,n}(s) \geq \check{S}_{\iota,j}^{r,n}(s). \quad (8.103)$$

Now, for $r \geq r^*$, $t > 0$ satisfying $r^2t \geq M^r$, $\iota \in \mathcal{I}$ satisfying $\iota > i$, and $n \geq 1$,

$$\begin{aligned} & \mathbf{P} \left(T_{\iota,a(i)}^{r,n}(s_\iota^r) \geq (x_{a(i)}^* + \epsilon_i)s_\iota^r, s_\iota^r \leq \beta_{\iota,n}^r, \tau_{\iota,2n-1}^r \leq r^2t \right) \\ & \leq \mathbf{P} \left(S_{\iota,a(i)}^{r,n}(T_{\iota,a(i)}^{r,n}(s_\iota^r)) \geq S_{\iota,a(i)}^{r,n}((x_{a(i)}^* + \epsilon_i)s_\iota^r), s_\iota^r \leq \beta_{\iota,n}^r, \tau_{\iota,2n-1}^r \leq r^2t \right) \\ & \leq \mathbf{P} \left(S_{\iota,a(i)}^{r,n}((x_{a(i)}^* + \epsilon_i)s_\iota^r) \leq Q_i^r(\tau_{\iota,2n-1}^r) + A_{\iota,i}^{r,n}(s_\iota^r) \right. \\ & \quad \left. - \sum_{j \in \underline{\mathcal{J}}_i} S_{\iota,j}^{r,n}(T_{\iota,j}^{r,n}(s_\iota^r)), s_\iota^r \leq \beta_{\iota,n}^r, \tau_{\iota,2n-1}^r \leq r^2t \right) \\ & \leq \mathbf{P} \left(S_{\iota,a(i)}^{r,n}((x_{a(i)}^* + \epsilon_i)s_\iota^r) < (\mu_{a(i)}^r - \epsilon_{2,i})(x_{a(i)}^* + \epsilon_i)s_\iota^r, \tau_{\iota,2n-1}^r \leq r^2t \right) \\ & \quad + \mathbf{P} \left(Q_i^r(\tau_{\iota,2n-1}^r) + A_{\iota,i}^{r,n}(s_\iota^r) - \sum_{j \in \underline{\mathcal{J}}_i} S_{\iota,j}^{r,n}(T_{\iota,j}^{r,n}(s_\iota^r)) \right. \\ & \quad \left. \geq (x_{a(i)}^* + \epsilon_i)(\mu_{a(i)}^r - \epsilon_{2,i})s_\iota^r, s_\iota^r \leq \beta_{\iota,n}^r, \tau_{\iota,2n-1}^r \leq r^2t \right). \end{aligned} \quad (8.104)$$

For the second inequality, we have used the fact that the number of class i jobs processed by activity $a(i)$, between time $\tau_{\iota,2n-1}^r$ and time $\tau_{\iota,2n-1}^r + s_\iota^r$, namely, $S_{\iota,a(i)}^{r,n}(T_{\iota,a(i)}^{r,n}(s_\iota^r))$, is less than or equal to the number of class i jobs present at time $\tau_{\iota,2n-1}^r$, namely, $Q_i^r(\tau_{\iota,2n-1}^r)$, plus the number of arrivals between time $\tau_{\iota,2n-1}^r$ and time $\tau_{\iota,2n-1}^r + s_\iota^r$, namely $A_{\iota,i}^{r,n}(s_\iota^r)$, less the number of class i jobs processed during this time interval by the activities indexed by $\underline{\mathcal{J}}_i$, namely, $\sum_{j \in \underline{\mathcal{J}}_i} S_{\iota,j}^{r,n}(T_{\iota,j}^{r,n}(s_\iota^r))$ (cf. (2.10)).

Letting $v_{\iota,a(i)}^{r,n} = v_{a(i)}^r(S_{a(i)}^r(T_{a(i)}^r(\tau_{\iota,2n-1}^r)) + 1)$, by Lemmas 7.5.2–7.6.1 together with (7.26) and (8.33), in a similar manner to (8.56)–(8.57), for $r \geq r^*$, $n \geq 1$, and $t > 0$ satisfying $r^2t \geq M^r$, we have

$$\begin{aligned} & \mathbf{P} \left(S_{\iota,a(i)}^{r,n}((x_{a(i)}^* + \epsilon_i)s_\iota^r) < (\mu_{a(i)}^r - \epsilon_{2,i})(x_{a(i)}^* + \epsilon_i)s_\iota^r, \tau_{\iota,2n-1}^r \leq r^2t \right) \\ & \leq K_{13} \exp(-K'_{13}s_\iota^r) \\ & \quad + \mathbf{E} \left(1_{\{\tau_{n,\iota}^r \in \mathbb{N}^I \times \mathbb{N}^J\}} \mathbf{P} \left(v_{\iota,a(i)}^{r,n} > \frac{\epsilon_{2,i}}{2\mu_{a(i)}^r} (x_{a(i)}^* + \epsilon_i)s_\iota^r, \tau_{\iota,2n-1}^r \leq r^2t \mid \mathcal{F}_{\tau_{n,\iota}^r}^r \right) \right) \\ & \leq K_{13} \exp(-K'_{13}s_\iota^r) + (\lfloor (\mu_{a(i)}^r + \epsilon_i)r^2t \rfloor + 1) K_{14} \exp(-K'_{14}s_\iota^r) \\ & \quad + K_{15} \exp(-K'_{15}r^2t), \end{aligned} \quad (8.105)$$

where $K_{13} = 1$, $K'_{13} = (\mu_{a(i)} - 2\epsilon_{2,i})(x_{a(i)}^* + \epsilon_i)\Lambda_{a(i)}^{s,*}(\mu_{a(i)}^{-1}(1 + \epsilon_{2,i}/2\mu_{a(i)})) > 0$, $K_{14} = \exp(\Lambda_{a(i)}^s(l_0)) > 0$, $K'_{14} = (l_0\epsilon_{2,i}/2\mu_{a(i)})(x_{a(i)}^* + \epsilon_i) > 0$, $0 < l_0 \in \mathcal{O}_0$, $K_{15} = \exp(\Lambda_{a(i)}^{s,*}((\mu_{a(i)}(1 + \epsilon_i/3\mu_{a(i)}))^{-1})) > 0$, and $K'_{15} = \mu_{a(i)}\Lambda_{a(i)}^{s,*}((\mu_{a(i)}(1 + \epsilon_i/3\mu_{a(i)}))^{-1}) > 0$, which are all independent of t , n , and r .

For $r \geq 1$, $n \geq 1$, $i \in \mathcal{I}$ such that $\iota > i$, let

$$\begin{aligned} \Upsilon_{\iota,i}^{r,n} &= \{Q_i^r(\tau_{\iota,2n-1}^r) \leq (\mu_{a(i)}^r \epsilon_i / 16) s_\iota^r; A_{\iota,i}^{r,n}(s_\iota^r) \leq (\lambda_i^r + \mu_{a(i)}^r \epsilon_i / 16) s_\iota^r + 1; \\ &S_{\iota,j}^{r,n}((x_j^* - \epsilon_{1,i}) s_\iota^r) \geq \left(\mu_j^r - \frac{\mu_{a(i)}^r \epsilon_i}{16|\underline{\mathcal{J}}_i|}\right) (x_j^* - \epsilon_{1,i}) s_\iota^r, j \in \underline{\mathcal{J}}_i; \\ &T_{\iota,j}^{r,n}(s_\iota^r) > (x_j^* - \epsilon_{1,i}) s_\iota^r, j \in \underline{\mathcal{J}}_i; \tau_{\iota,2n-1}^r \leq r^2 t\}, \end{aligned} \quad (8.106)$$

where $\epsilon_{1,i}$ is defined in (8.32), and if $\underline{\mathcal{J}}_i = \emptyset$, we omit the terms involving $j \in \underline{\mathcal{J}}_i$ from the definition of $\Upsilon_{\iota,i}^{r,n}$.

For the last probability in (8.104) we have,

$$\begin{aligned} &\mathbf{P} \left(Q_i^r(\tau_{\iota,2n-1}^r) + A_{\iota,i}^{r,n}(s_\iota^r) - \sum_{j \in \underline{\mathcal{J}}_i} S_{\iota,j}^{r,n}(T_{\iota,j}^{r,n}(s_\iota^r)) \right. \\ &\quad \left. \geq (x_{a(i)}^* + \epsilon_i)(\mu_{a(i)}^r - \epsilon_{2,i}) s_\iota^r, s_\iota^r \leq \beta_{\iota,n}^r, \tau_{\iota,2n-1}^r \leq r^2 t \right) \\ &\leq \mathbf{P} \left(Q_i^r(\tau_{\iota,2n-1}^r) + A_{\iota,i}^{r,n}(s_\iota^r) - \sum_{j \in \underline{\mathcal{J}}_i} S_{\iota,j}^{r,n}(T_{\iota,j}^{r,n}(s_\iota^r)) \right. \\ &\quad \left. \geq (x_{a(i)}^* + \epsilon_i)(\mu_{a(i)}^r - \epsilon_{2,i}) s_\iota^r, \Upsilon_{\iota,i}^{r,n}, \tau_{\iota,2n-1}^r \leq r^2 t \right) \\ &+ \mathbf{P} \left(Q_i^r(\tau_{\iota,2n-1}^r) > (\mu_{a(i)}^r \epsilon_i / 16) s_\iota^r, \tau_{\iota,2n-1}^r \leq r^2 t \right) \\ &+ \mathbf{P} \left(A_{\iota,i}^{r,n}(s_\iota^r) > \left(\lambda_i^r + \frac{\mu_{a(i)}^r \epsilon_i}{16} \right) s_\iota^r + 1, \tau_{\iota,2n-1}^r \leq r^2 t \right) \\ &+ \sum_{j \in \underline{\mathcal{J}}_i} \mathbf{P} \left(S_{\iota,j}^{r,n}((x_j^* - \epsilon_{1,i}) s_\iota^r) < \left(\mu_j^r - \frac{\mu_{a(i)}^r \epsilon_i}{16|\underline{\mathcal{J}}_i|} \right) (x_j^* - \epsilon_{1,i}) s_\iota^r, \tau_{\iota,2n-1}^r \leq r^2 t \right) \\ &+ \sum_{j \in \underline{\mathcal{J}}_i} \mathbf{P} \left(T_{\iota,j}^{r,n}(s_\iota^r) \leq (x_j^* - \epsilon_{1,i}) s_\iota^r, s_\iota^r \leq \beta_{\iota,n}^r, \tau_{\iota,2n-1}^r \leq r^2 t \right). \end{aligned} \quad (8.107)$$

For the first term in the right side of the inequality in (8.107) we have on $\Upsilon_{\iota,i}^{r,n}$ that

for $r \geq r^*$, $n \geq 1$,

$$\begin{aligned}
& Q_i^r(\tau_{l,2n-1}^r) + A_{l,i}^{r,n}(s_l^r) - \sum_{j \in \underline{\mathcal{J}}_i} S_{l,j}^{r,n}(T_{l,j}^{r,n}(s_l^r)) \\
& \leq \frac{\mu_{a(i)}^r \epsilon_i}{16} s_l^r + \lambda_i^r s_l^r + \frac{\mu_{a(i)}^r \epsilon_i}{16} s_l^r + 1 - \sum_{j \in \underline{\mathcal{J}}_i} \left(\mu_j^r - \frac{\mu_{a(i)}^r \epsilon_i}{16 |\underline{\mathcal{J}}_i|} \right) (x_j^* - \epsilon_{1,i}) s_l^r \\
& = \frac{\mu_{a(i)}^r \epsilon_i}{8} s_l^r + \lambda_i^r s_l^r - \sum_{j \in \underline{\mathcal{J}}_i} x_j^* \mu_j^r s_l^r + 1 \\
& \quad + \sum_{j \in \underline{\mathcal{J}}_i} \frac{\mu_{a(i)}^r \epsilon_i}{16 |\underline{\mathcal{J}}_i|} x_j^* s_l^r + \sum_{j \in \underline{\mathcal{J}}_i} \left(\mu_j^r - \frac{\mu_{a(i)}^r \epsilon_i}{16 |\underline{\mathcal{J}}_i|} \right) \epsilon_{1,i} s_l^r \\
& \leq \frac{\mu_{a(i)}^r \epsilon_i}{8} s_l^r + \lambda_i^r s_l^r - \sum_{j \in \underline{\mathcal{J}}_i} x_j^* \mu_j^r s_l^r + 1 + \frac{\mu_{a(i)}^r \epsilon_i}{16} s_l^r + \sum_{j \in \underline{\mathcal{J}}_i} \mu_j^r \epsilon_{1,i} s_l^r \\
& \leq \frac{\mu_{a(i)}^r \epsilon_i}{4} s_l^r + \lambda_i^r s_l^r - \sum_{j \in \underline{\mathcal{J}}_i} x_j^* \mu_j^r s_l^r + 1 \\
& \leq \left(\lambda_i^r - \sum_{j \in \underline{\mathcal{J}}_i} x_j^* \mu_j^r + \frac{\mu_{a(i)}^r \epsilon_i}{2} \right) s_l^r \\
& \leq \left(x_{a(i)}^* \mu_{a(i)}^r + \frac{\mu_{a(i)}^r \epsilon_i}{32} + \frac{\mu_{a(i)}^r \epsilon_i}{2} \right) s_l^r \\
& < (x_{a(i)}^* + \epsilon_i) (\mu_{a(i)}^r - \epsilon_{2,i}) s_l^r, \tag{8.108}
\end{aligned}$$

where in the second inequality we have used the fact that $\sum_{j \in \underline{\mathcal{J}}_i} x_j^* \leq 1$. For the third inequality we have, using (8.32) together with (8.15), if $\underline{\mathcal{J}}_i \neq \emptyset$,

$$\sum_{j \in \underline{\mathcal{J}}_i} \mu_j^r \epsilon_{1,i} = \frac{\sum_{j \in \underline{\mathcal{J}}_i} \mu_j^r}{\sum_{j \in \underline{\mathcal{J}}_i} (\mu_j + \epsilon_i)} \frac{\epsilon_i}{16} \left(\mu_{a(i)} - \frac{\mu_{a(i)} \epsilon_i}{16} \right) \leq \frac{\mu_{a(i)}^r \epsilon_i}{16}. \tag{8.109}$$

For the fourth inequality we have used (8.27), together with (8.15). The fifth inequality follows by (8.28). The final inequality follows by (8.33), together with (8.27). Hence, the first probability in the second expression of (8.107) is zero, for all $n \geq 1$, $r \geq r^*$, $t > 0$ satisfying $r^2 t \geq M^r$.

For the second term in the right side of the inequality in (8.107), using the result established at the beginning of this proof that (I.1) holds for i , we have for

$r \geq r^*$, $t > 0$ satisfying $r^2 t \geq M^r$,

$$\begin{aligned}
& \mathbf{P} \left(Q_i^r(\tau_{l,2n-1}^r) > (\mu_{a(i)}^r \epsilon_i / 16) s_l^r, \tau_{l,2n-1}^r \leq r^2 t \right) \\
& \leq \mathbf{P} \left(\sup_{\tau_{i,0}^r \leq s \leq \tau_{l,2n-1}^r} R_i^r(s) \geq L_i^r - |\underline{\mathcal{J}}_i|, \tau_{l,2n-1}^r \leq r^2 t \right) \\
& \leq \mathbf{P} \left(\sup_{\tau_{i,0}^r \leq s \leq r^2 t} R_i^r(s) \geq L_i^r - |\underline{\mathcal{J}}_i| \right) \\
& \leq p_{1,i}(r^2 t) \left(C_{1,i}^{(1)} \exp(-C_{1,i}^{(2)} L_0^r) + C_{1,i}^{(3)} \exp(-C_{1,i}^{(4)} r^2 t) \right), \quad (8.110)
\end{aligned}$$

where $p_{1,i}$ is a polynomial (of degree at most $i + 1$) with non-negative coefficients, and $C_{1,i}^{(m)} > 0$, for $m = 1, 2, 3, 4$, and where the polynomial and constants are independent of t and r . In (8.110), the first inequality holds since s_l^r (for $l > i$) is considerably larger than L_i^r . Specifically,

$$\begin{aligned}
\frac{\mu_{a(i)}^r \epsilon_i}{16} s_l^r &= \frac{\mu_{a(i)}^r \epsilon_i}{16} \cdot \frac{L_l^r - (|\underline{\mathcal{J}}_l| + 2)}{\lambda_l^r + \epsilon_l} \\
&\geq \frac{\mu_{a(i)}^r}{\mu_{a(i)}} \cdot \frac{\lambda_l}{\lambda_l^r + \epsilon_l} \left[\frac{\mu_{a(i)}}{16\lambda_l} \cdot \frac{L_l^r}{\epsilon_i} - \frac{\mu_{a(i)}}{16\lambda_l} \epsilon_i (|\underline{\mathcal{J}}_l| + 2) \right] \\
&\geq \frac{1}{4} \left[\frac{\mu_{a(i)}}{16\lambda_l} \cdot \frac{L_l^r}{\epsilon_i} - \frac{\mu_{a(i)}}{16\lambda_l} \epsilon_i (|\underline{\mathcal{J}}_l| + 2) \right] \\
&\geq 2L_i^r + \frac{1}{4} \left[8L_i^r - \frac{\mu_{a(i)}}{16\lambda_l} \epsilon_i (|\underline{\mathcal{J}}_l| + 2) \right] \\
&\geq 2L_i^r, \quad (8.111)
\end{aligned}$$

for $r \geq r^*$. In the first inequality of (8.111), we have used (7.4), and the fact that $\epsilon_l \leq 1$ for all $l \in \mathcal{I}$. In the second inequality, we have used (8.27) together with the fact that $\epsilon_l < \lambda_l/2$ to obtain $\lambda_l/(\lambda_l^r + \epsilon_l) > 1/2$. The third inequality follows since $\epsilon_i < \hat{\epsilon} < \mu_{a(i)}/256\lambda_l$ (cf. (8.97)), and the fourth inequality follows from (8.29).

For the third term in the right side of the inequality in (8.107), using (8.14), (8.31), together with the fact that $s_l^r > 2/\check{\epsilon}_i$ for all $l > i$ (by (8.9), (8.12) and

(8.35)), we have as in (8.54), for $r \geq r^*$, $n \geq 1$, $t > 0$ satisfying $r^2 t \geq M^r$,

$$\begin{aligned}
& \mathbf{P} \left(A_{l,i}^{r,n}(s_l^r) > \left(\lambda_i^r + \frac{\mu_{a(i)}^r \epsilon_i}{16} \right) s_l^r + 1, \tau_{l,2n-1}^r \leq r^2 t \right) \\
& \leq \mathbf{P} \left(A_{l,i}^{r,n}(s_l^r) > (\lambda_i^r + \check{\epsilon}_i) s_l^r + 1, \tau_{l,2n-1}^r \leq r^2 t \right) \\
& \leq \mathbf{E} \left(\mathbf{1}_{\{\mathcal{T}_{n,i}^r \in \mathbb{N}^I \times \mathbb{N}^J\}} \mathbf{P} \left(\check{A}_{l,i}^{r,n}(s_l^r) > (\lambda_i^r + \check{\epsilon}_i) s_l^r \mid \mathcal{F}_{\mathcal{T}_{n,i}^r}^r \right) \right) \\
& \leq K_{16} \exp(-K'_{16} s_l^r), \tag{8.112}
\end{aligned}$$

by Lemmas 7.5.2–7.6.1, where $K_{16} = \exp(\Lambda_i^{a,*}((\lambda_i(1 + \check{\epsilon}_i/3\lambda_i))^{-1})) > 0$ and $K'_{16} = \lambda_i \Lambda_i^{a,*}((\lambda_i(1 + \check{\epsilon}_i/3\lambda_i))^{-1}) > 0$ do not depend on t , n , or r .

Letting $v_{l,j}^{r,n} = v_j^r(S_j^r(T_j^r(\tau_{l,2n-1}^r)) + 1)$, for the fourth term in the right side of the inequality in (8.107), using (8.14) and (8.31), we have, by a similar argument to that for (8.56), that for $r \geq r^*$, $n \geq 1$, $t > 0$ satisfying $r^2 t \geq M^r$, $j \in \underline{\mathcal{J}}_i \neq \emptyset$,

$$\begin{aligned}
& \mathbf{P} \left(S_{l,j}^{r,n}((x_j^* - \epsilon_{1,i})s_l^r) < \left(\mu_j^r - \frac{\mu_{a(i)}^r \epsilon_i}{16|\underline{\mathcal{J}}_i|} \right) (x_j^* - \epsilon_{1,i})s_l^r, \tau_{l,2n-1}^r \leq r^2 t \right) \\
& \leq \mathbf{P} \left(S_{l,j}^{r,n}((x_j^* - \epsilon_{1,i})s_l^r) < (\mu_j^r - \check{\epsilon}_i) (x_j^* - \epsilon_{1,i})s_l^r, \tau_{l,2n-1}^r \leq r^2 t \right) \\
& \leq K_{17} \exp(-K'_{17} s_l^r) \\
& \quad + \mathbf{E} \left(\mathbf{1}_{\{\mathcal{T}_{n,l}^r \in \mathbb{N}^I \times \mathbb{N}^J\}} \mathbf{P} \left(v_{l,j}^{r,n} > \frac{\check{\epsilon}_i}{2\mu_j^r} (x_j^* - \epsilon_{1,i})s_l^r, \tau_{l,2n-1}^r \leq r^2 t \mid \mathcal{F}_{\mathcal{T}_{n,l}^r}^r \right) \right) \\
& \leq K_{17} \exp(-K'_{17} s_l^r) \\
& \quad + ([(\mu_j^r + \epsilon_{1,i})r^2 t] + 1) K_{18} \exp(-K'_{18} s_l^r) + K_{19} \exp(-K'_{19} r^2 t), \tag{8.113}
\end{aligned}$$

where $K_{17} = 1$, $K'_{17} = \min\{(\mu_j - 2\check{\epsilon}_i)(x_j^* - \epsilon_{1,i})\Lambda_j^{s,*}(\mu_j^{-1}(1 + \check{\epsilon}_i/2\mu_j)) : j \in \underline{\mathcal{J}}_i\} > 0$, $K_{18} = \max\{\exp(\Lambda_j^s(l_0)) > 0 : j \in \underline{\mathcal{J}}_i\}$, $K'_{18} = \min\{(l_0 \check{\epsilon}_i/2\mu_j)(x_j^* - \epsilon_{1,i}) : j \in \underline{\mathcal{J}}_i\} > 0$, $0 < l_0 \in \mathcal{O}$, $K_{19} = \max\{\exp(\Lambda_j^{s,*}((\mu_j(1 + \epsilon_i/3\mu_j))^{-1})) : j \in \underline{\mathcal{J}}_i\} > 0$, and $K'_{19} = \min\{\mu_j \Lambda_j^{s,*}((\mu_j(1 + \epsilon_i/3\mu_j))^{-1}) : j \in \underline{\mathcal{J}}_i\} > 0$, which are all independent of t , n , and r .

For the fifth term in the right side of the inequality in (8.107), we have for

$r \geq r^*$, $n \geq 1$, $t > 0$ satisfying $r^2t \geq M^r$, $j \in \underline{\mathcal{J}}_i \neq \emptyset$, $k = k(j)$,

$$\begin{aligned}
& \mathbf{P}\left(T_{l,j}^{r,n}(s_l^r) \leq (x_j^* - \epsilon_{1,i})s_l^r, s_l^r \leq \beta_{l,n}^r, \tau_{l,2n-1}^r \leq r^2t\right) \\
&= \mathbf{P}\left(s_l^r - \sum_{i' \in \underline{\mathcal{I}}_k} T_{l,a(i')}^{r,n}(s_l^r) - I_{l,k}^{r,n}(s_l^r) \leq (x_j^* - \epsilon_{1,i})s_l^r, s_l^r \leq \beta_{l,n}^r, \tau_{l,2n-1}^r \leq r^2t\right) \\
&\leq \mathbf{P}\left(\sum_{i' \in \underline{\mathcal{I}}_k} T_{l,a(i')}^{r,n}(s_l^r) + I_{l,k}^{r,n}(s_l^r) \geq \sum_{i' \in \underline{\mathcal{I}}_k} x_{a(i')}^* s_l^r + \epsilon_{1,i} s_l^r, s_l^r \leq \beta_{l,n}^r, \tau_{l,2n-1}^r \leq r^2t\right) \\
&\leq \sum_{i' \in \underline{\mathcal{I}}_k} \mathbf{P}\left(T_{l,a(i')}^{r,n}(s_l^r) \geq (x_{a(i')}^* + \epsilon_{i'}) s_l^r, s_l^r \leq \beta_{l,n}^r, \tau_{l,2n-1}^r \leq r^2t\right) \\
&\quad + \mathbf{P}\left(I_{l,k}^{r,n}(s_l^r) \geq \frac{\epsilon_{1,i}}{\mathbf{I}} s_l^r, \tau_{l,2n-1}^r \leq r^2t\right). \tag{8.114}
\end{aligned}$$

For the first inequality in (8.114) we use the fact that $\sum_{i' \in \underline{\mathcal{I}}_k} x_{a(i')}^* + x_j^* = 1$. The last inequality in (8.114) follows from the fact that when $\underline{\mathcal{J}}_i \neq \emptyset$, for all $i' \in \underline{\mathcal{I}}_k$,

$$\begin{aligned}
\epsilon_{1,i} &= \frac{1}{\sum_{j \in \underline{\mathcal{J}}_i} (\mu_j + \epsilon_i)} \frac{\epsilon_i}{16} \left(\mu_{a(i)} - \frac{\mu_{a(i)} \epsilon_i}{16} \right) \\
&\geq \frac{\mu_{a(i)} \epsilon_i}{64 \sum_{j \in \underline{\mathcal{J}}_i} \mu_j} \\
&\geq \gamma_i \epsilon_i \\
&\geq \mathbf{I} \epsilon_{i'}, \tag{8.115}
\end{aligned}$$

by (7.7) (since $i' < i$), and where $\mathbf{I} \geq (|\underline{\mathcal{I}}_k| + 1)$. The first inequality in (8.115) holds since $\epsilon_i \leq \min\{\mu_{\min}, 1\}$, by (7.5), and the second inequality follows by the definition of γ_i (cf. (7.6)).

For the second term in the last expression in (8.114), we have that for $r \geq r^*$, $n \geq 1$, $t > 0$ satisfying $r^2t \geq M^r$,

$$\begin{aligned}
& \sup_{n \geq 1} \mathbf{P}\left(I_{l,k}^{r,n}(s_l^r) \geq \frac{\epsilon_{1,i}}{\mathbf{I}} s_l^r, \tau_{l,2n-1}^r \leq r^2t\right) \\
&\leq \mathbf{P}(I_k^r(2r^2t) - I_k^r(\tau_{i,0}^r) > 0, \tau_{i,0}^r \leq 2r^2t) \\
&\quad + \mathbf{P}(I_k^r(\tau_{i,0}^r) \geq t_i^r) \\
&\leq p_{3,i}(2r^2t) \left(C_{3,i}^{(1)} \exp(-C_{3,i}^{(2)} L_0^r) + C_{3,i}^{(3)} \exp(-2C_{3,i}^{(4)} r^2t) \right) \\
&\quad + p_{2,i}(r^2t) \left(C_{2,i}^{(1)} \exp(-C_{2,i}^{(2)} L_0^r) + C_{2,i}^{(3)} \exp(-C_{2,i}^{(4)} r^2t) \right), \tag{8.116}
\end{aligned}$$

where $p_{2,i}$ and $p_{3,i}$ are polynomials (of degree at most i and $i+1$, respectively) with non-negative coefficients, and $C_{l,i}^{(m)} > 0$, for $l = 2, 3$, $m = 1, 2, 3, 4$, since (I) holds for i (with $2t$ in place of t), and $k \in \underline{\mathcal{K}}_i$. The polynomials and the constants do not depend on t , n , or r . The first inequality in (8.116) holds since on $\{\tau_{l,2n-1}^r \leq r^2 t\}$, $I_{l,k}^{r,n}(s_l^r) \leq I_k^r(\tau_{l,2n-1}^r + s_l^r) \leq I_k^r(r^2 t + s_l^r) \leq I_k^r(2r^2 t)$, as $s_l^r \leq M^r \leq r^2 t$, and since by (8.115),

$$\begin{aligned}
\frac{\epsilon_{1,i}}{\mathbf{I}} s_l^r &\geq \frac{\gamma_i \epsilon_i s_l^r}{\mathbf{I}} = \frac{\gamma_i \epsilon_i (L_i^r - (|\underline{\mathcal{J}}_l| + 2))}{\mathbf{I}(\lambda_l^r + \epsilon_i)} \\
&\geq \frac{\epsilon_i^2 (L_i^r - (|\underline{\mathcal{J}}_l| + 2))}{(\lambda_l^r + \epsilon_i)} \\
&\geq \frac{\lambda_l}{\lambda_l^r + \epsilon_i} \left(\frac{L_i^r}{\lambda_l \epsilon_i} - \frac{\epsilon_i^2 (|\underline{\mathcal{J}}_l| + 2)}{\lambda_l} \right) \\
&\geq \frac{\lambda_l}{\lambda_l^r + \epsilon_i} \left(\frac{18L_i^r}{\lambda_l - \sum_{j \in \underline{\mathcal{J}}_i} x_j^* \mu_j} - \frac{\epsilon_i (|\underline{\mathcal{J}}_l| + 2)}{\lambda_l} \right) \\
&\geq \frac{1}{2} \left(\frac{17L_i^r}{\lambda_l - \sum_{j \in \underline{\mathcal{J}}_i} x_j^* \mu_j} \right) \geq t_i^r. \tag{8.117}
\end{aligned}$$

In the second inequality in (8.117), we have used the fact that $\epsilon_i < (\prod_{m=1}^{\mathbf{I}} \gamma_m) / \mathbf{I} \leq \gamma_i / \mathbf{I}$, as $\gamma_m \leq 1$ for all m (cf. (8.97)). In the third inequality of (8.117), we have used (7.4). In the fourth inequality we have used the fact that $\epsilon_i < (\lambda_l - \sum_{j \in \underline{\mathcal{J}}_i} x_j^* \mu_j) / (18\lambda_l)$ (cf. (8.97)) and $\epsilon_i < 1$. In the fifth inequality we have used (8.29) along with the estimate $\lambda_l / (\lambda_l^r + \epsilon_i) > 1/2$ used in proving (8.111).

Combining all of the above (from (8.104) onwards), we have for all $r \geq r^*$,

$t > 0$ satisfying $r^2 t \geq M^r$,

$$\begin{aligned}
& \sup_{n \geq 1} \mathbf{P} \left(T_{\iota, a(i)}^{r, n}(s_l^r) \geq (x_{a(i)}^* + \epsilon_i) s_l^r, s_l^r \leq \beta_{\iota, n}^r, \tau_{\iota, 2n-1}^r \leq r^2 t \right) \\
\leq & K_{13} \exp(-K'_{13} s_l^r) \\
& + \left(\lfloor (\mu_{a(i)}^r + \epsilon_i) r^2 t \rfloor + 1 \right) K_{14} \exp(-K'_{14} s_l^r) + K_{15} \exp(-K'_{15} r^2 t) \\
& + p_{1, i}(r^2 t) \left(C_{1, i}^{(1)} \exp(-C_{1, i}^{(2)} L_0^r) + C_{1, i}^{(3)} \exp(-C_{1, i}^{(4)} r^2 t) \right) \\
& + K_{16} \exp(-K'_{16} s_l^r) + |\underline{\mathcal{J}}_i| K_{17} \exp(-K'_{17} s_l^r) \\
& + \sum_{j \in \underline{\mathcal{J}}_i} \left(\lfloor (\mu_j^r + \epsilon_{1, i}) r^2 t \rfloor + 1 \right) K_{18} \exp(-K'_{18} s_l^r) + |\underline{\mathcal{J}}_i| K_{19} \exp(-K'_{19} r^2 t) \\
& + \sum_{j \in \underline{\mathcal{J}}_i} \left\{ \sum_{i' \in \underline{\mathcal{I}}_{k(j)}} \sup_{n \geq 1} \mathbf{P} \left(T_{\iota, a(i')}^{r, n}(s_l^r) \geq (x_{a(i')}^* + \epsilon_{i'}) s_l^r, \right. \right. \\
& \qquad \qquad \qquad \left. \left. s_l^r \leq \beta_{\iota, n}^r, \tau_{\iota, 2n-1}^r \leq r^2 t \right) \right. \\
& \left. + p_{3, i}(2r^2 t) \left(C_{3, i}^{(1)} \exp(-C_{3, i}^{(2)} L_0^r) + C_{3, i}^{(3)} \exp(-2C_{3, i}^{(4)} r^2 t) \right) \right. \\
& \left. + p_{2, i}(r^2 t) \left(C_{2, i}^{(1)} \exp(-C_{2, i}^{(2)} L_0^r) + C_{2, i}^{(3)} \exp(-C_{2, i}^{(4)} r^2 t) \right) \right\}. \quad (8.118)
\end{aligned}$$

By the assumption that (II.1) holds with i' in place of i , for all $i' < i$, the definition of s_l^r , and the fact that $s_l^r \geq L_0^r$, for all $l \in \mathcal{I}$, it follows that for $r \geq r^*$, $t > 0$ satisfying $r^2 t \geq M^r$,

$$\begin{aligned}
& \mathbf{P} \left(T_{\iota, a(i)}^{r, n}(s_l^r) \geq (x_{a(i)}^* + \epsilon_i) s_l^r, s_l^r \leq \beta_{\iota, n}^r, \tau_{\iota, 2n-1}^r \leq r^2 t \right) \\
\leq & p_{4, i}(r^2 t) \left(C_{4, i}^{(1)} \exp(-C_{4, i}^{(2)} L_0^r) + C_{4, i}^{(3)} \exp(-C_{4, i}^{(4)} r^2 t) \right), \quad (8.119)
\end{aligned}$$

where $p_{4, i}$ is a polynomial (of degree at most $i + 1$) with positive coefficients, and $C_{4, i}^{(m)} > 0$, for $m = 1, 2, 3, 4$. The polynomial and the constants do not depend on t or r . This completes the proof that (II.1) holds for i .

Proof of (II.2). Fix a transition class $\iota > i$. For each $r \geq 1$, $n \geq 1$, $s \geq 0$, $j \in \mathcal{J}_i$,

on $\{d_{\tau_{l,2n-1}}^r < \infty\}$ define

$$\begin{aligned} dA_{l,i}^{r,n}(s) &= A_i^r(d_{\tau_{l,2n-1}}^r + s) - A_i^r(d_{\tau_{l,2n-1}}^r), \\ dS_{l,j}^{r,n}(s) &= S_j^r(T_j^r(d_{\tau_{l,2n-1}}^r) + s) - S_j^r(T_j^r(d_{\tau_{l,2n-1}}^r)), \\ d\check{A}_{l,i}^{r,n}(s) &= \sup\{m \geq 0 : \xi_i^r(A_i^r(d_{\tau_{l,2n-1}}^r + m)) - \xi_i^r(A_i^r(d_{\tau_{l,2n-1}}^r)) \leq s\}, \\ d\check{S}_{l,j}^{r,n}(s) &= \sup\{m \geq 0 : \eta_j^r(S_j^r(T_j^r(d_{\tau_{l,2n-1}}^r)) + m + 1) \\ &\quad - \eta_j^r(S_j^r(T_j^r(d_{\tau_{l,2n-1}}^r)) + 1) \leq s\}, \end{aligned}$$

and for concreteness on $\{d_{\tau_{l,2n-1}}^r = \infty\}$, we define $dA_{l,i}^{r,n}$, $dS_{l,j}^{r,n}$, $d\check{A}_{l,i}^{r,n}$, $d\check{S}_{l,j}^{r,n}$, to be identically zero. Then, for each $s \geq 0$,

$$dA_{l,i}^{r,n}(s) \geq d\check{A}_{l,i}^{r,n}(s), \quad d\check{S}_{l,j}^{r,n}(s) \geq dS_{l,j}^{r,n}(s) - 1. \quad (8.120)$$

We have for $r \geq r^*$, $t > 0$ satisfying $r^2 t \geq M^r$,

$$\begin{aligned} & \mathbf{P} \left(dT_{l,a(i)}^{r,n}(d_s^r) \leq (x_{a(i)}^* - \epsilon_i) d_s^r, d_{\tau_{l,2n-1}}^r \leq r^2 t \right) \\ & \leq \mathbf{P} \left(dS_{l,a(i)}^{r,n}(dT_{l,a(i)}^{r,n}(d_s^r)) \leq dS_{l,a(i)}^{r,n}((x_{a(i)}^* - \epsilon_i) d_s^r), d_{\tau_{l,2n-1}}^r \leq r^2 t \right) \\ & \leq \mathbf{P} \left(dA_{l,i}^{r,n}(d_s^r) - dS_{l,a(i)}^{r,n}(dT_{l,a(i)}^{r,n}(d_s^r)) - \sum_{j \in \mathcal{I}_i} dS_{l,j}^{r,n}(dT_{l,j}^{r,n}(d_s^r)) \right. \\ & \quad \left. \geq dA_{l,i}^{r,n}(d_s^r) - \sum_{j \in \mathcal{I}_i} dS_{l,j}^{r,n}(dT_{l,j}^{r,n}(d_s^r)) - dS_{l,a(i)}^{r,n}((x_{a(i)}^* - \epsilon_i) d_s^r), \right. \\ & \quad \left. d_{\tau_{l,2n-1}}^r \leq r^2 t \right) \\ & \leq \mathbf{P} \left(Q_i^r(d_{\tau_{l,2n-1}}^r + d_s^r) - Q_i^r(d_{\tau_{l,2n-1}}^r) \geq (\mu_{a(i)}^r \epsilon_i / 32) d_s^r, d_{\tau_{l,2n-1}}^r \leq r^2 t \right) \\ & \quad + \mathbf{P} \left(dA_{l,i}^{r,n}(d_s^r) - \sum_{j \in \mathcal{I}_i} dS_{l,j}^{r,n}(dT_{l,j}^{r,n}(d_s^r)) \right. \\ & \quad \left. - dS_{l,a(i)}^{r,n}((x_{a(i)}^* - \epsilon_i) d_s^r) < (\mu_{a(i)}^r \epsilon_i / 32) d_s^r, d_{\tau_{l,2n-1}}^r \leq r^2 t \right), \end{aligned} \quad (8.121)$$

where the last inequality in (8.121) follows by (2.10), as in (8.104).

For each $r \geq 1$ and $n \geq 1$, let

$$\begin{aligned} \mathcal{A}_{l,i}^{r,n} &= \{dA_{l,i}^{r,n}(d_s^r) \geq (\lambda_i^r - (\mu_{a(i)}^r \epsilon_i / 32)) d_s^r; \\ & \quad dS_{l,a(i)}^{r,n}((x_{a(i)}^* - \epsilon_i) d_s^r) \leq (\mu_{a(i)}^r + \epsilon_{2,i})(x_{a(i)}^* - \epsilon_i) d_s^r + 1; \\ & \quad dS_{l,j}^{r,n}((x_j^* + \epsilon_{1,i}) d_s^r) \leq (\mu_j^r + \epsilon_{3,j})(x_j^* + \epsilon_{1,i}) d_s^r + 1, j \in \mathcal{I}_i; \\ & \quad dT_{l,j}^{r,n}(d_s^r) < (x_j^* + \epsilon_{1,i}) d_s^r, d_{\tau_{l,2n-1}}^r \leq r^2 t\}, \end{aligned} \quad (8.122)$$

where $\epsilon_{2,i}$ and $\epsilon_{3,j}$ are defined in (8.33) and (8.34), respectively.

Then, for the last term in (8.121), we have

$$\begin{aligned}
& \mathbf{P} \left(dA_{l,i}^{r,n}(d_{S_l}^r) - \sum_{j \in \underline{\mathcal{J}}_i} dS_{l,j}^{r,n}(dT_{l,j}^{r,n}(d_{S_l}^r)) \right. \\
& \quad \left. - dS_{l,a(i)}^{r,n}((x_{a(i)}^* - \epsilon_i) d_{S_l}^r) < (\mu_{a(i)}^r \epsilon_i / 32) d_{S_l}^r, d_{\tau_{l,2n-1}}^r \leq r^2 t \right) \\
\leq & \mathbf{P} \left(dA_{l,i}^{r,n}(d_{S_l}^r) - \sum_{j \in \underline{\mathcal{J}}_i} dS_{l,j}^{r,n}(dT_{l,j}^{r,n}(d_{S_l}^r)) \right. \\
& \quad \left. - dS_{l,a(i)}^{r,n}((x_{a(i)}^* - \epsilon_i) d_{S_l}^r) < (\mu_{a(i)}^r \epsilon_i / 32) d_{S_l}^r, d\Upsilon_{l,i}^{r,n}, d_{\tau_{l,2n-1}}^r \leq r^2 t \right) \\
& + \mathbf{P} \left(dA_{l,i}^{r,n}(d_{S_l}^r) < (\lambda_i^r - (\mu_{a(i)}^r \epsilon_i / 32)) d_{S_l}^r, d_{\tau_{l,2n-1}}^r \leq r^2 t \right) \\
& + \mathbf{P} \left(dS_{l,a(i)}^{r,n}((x_{a(i)}^* - \epsilon_i) d_{S_l}^r) > (\mu_{a(i)}^r + \epsilon_{2,i})(x_{a(i)}^* - \epsilon_i) d_{S_l}^r + 1, d_{\tau_{l,2n-1}}^r \leq r^2 t \right) \\
& + \sum_{j \in \underline{\mathcal{J}}_i} \mathbf{P} \left(dS_{l,j}^{r,n}((x_j^* + \epsilon_{1,i}) d_{S_l}^r) > (\mu_j^r + \epsilon_{3,j})(x_j^* + \epsilon_{1,i}) d_{S_l}^r + 1, d_{\tau_{l,2n-1}}^r \leq r^2 t \right) \\
& + \sum_{j \in \underline{\mathcal{J}}_i} \mathbf{P} \left(dT_{l,j}^{r,n}(d_{S_l}^r) \geq (x_j^* + \epsilon_{1,i}) d_{S_l}^r, d_{\tau_{l,2n-1}}^r \leq r^2 t \right). \tag{8.123}
\end{aligned}$$

On $d\Upsilon_{l,i}^{r,n}$ we have for $r \geq r^*$, $t > 0$ satisfying $r^2 t \geq Mr$,

$$\begin{aligned}
& dA_{l,i}^{r,n}(d_{S_l}^r) - dS_{l,a(i)}^{r,n}((x_{a(i)}^* - \epsilon_i) d_{S_l}^r) - \sum_{j \in \underline{\mathcal{J}}_i} dS_{l,j}^{r,n}(dT_{l,j}^{r,n}(d_{S_l}^r)) \\
\geq & \lambda_i^r d_{S_l}^r - \frac{\mu_{a(i)}^r \epsilon_i}{32} d_{S_l}^r - (\mu_{a(i)}^r + \epsilon_{2,i})(x_{a(i)}^* - \epsilon_i) d_{S_l}^r \\
& - \sum_{j \in \underline{\mathcal{J}}_i} (\mu_j^r + \epsilon_{3,j})(x_j^* + \epsilon_{1,i}) d_{S_l}^r - 2 \\
= & \left(\lambda_i^r - x_{a(i)}^* \mu_{a(i)}^r - \sum_{j \in \underline{\mathcal{J}}_i} x_j^* \mu_j^r \right) d_{S_l}^r - \frac{\mu_{a(i)}^r \epsilon_i}{32} d_{S_l}^r + \mu_{a(i)}^r \epsilon_i d_{S_l}^r - 2 \\
& - \sum_{j \in \underline{\mathcal{J}}_i} \mu_j^r \epsilon_{1,i} d_{S_l}^r - \epsilon_{2,i} (x_{a(i)}^* - \epsilon_i) d_{S_l}^r - \sum_{j \in \underline{\mathcal{J}}_i} \epsilon_{3,j} (x_j^* + \epsilon_{1,i}) d_{S_l}^r \\
\geq & \left(-\frac{\mu_{a(i)}^r \epsilon_i}{32} - \frac{\mu_{a(i)}^r \epsilon_i}{32} + \mu_{a(i)}^r \epsilon_i - \frac{\mu_{a(i)}^r \epsilon_i}{2} - \frac{\mu_{a(i)}^r \epsilon_i}{16} - \frac{\mu_{a(i)}^r \epsilon_i}{16} - \frac{\mu_{a(i)}^r \epsilon_i}{16} \right) d_{S_l}^r \\
> & \frac{\mu_{a(i)}^r \epsilon_i}{32} d_{S_l}^r, \tag{8.124}
\end{aligned}$$

where in the second to the last inequality we have used (8.28), (8.30), (8.109),

(8.33), and (8.34), and in the last inequality we have used the fact that $\mu_{a(i)} < 2\mu_{a(i)}^r$ by (8.27). Hence, the first probability in the right side of (8.123) is zero.

For the second term in the right side of (8.123), letting $u_{l,i}^{r,n} = u_i^r(A_i^r(d_{\tau_{l,2n-1}}^r) + 1)$, we have as in (8.74), using Lemmas 7.5.2–7.6.1 and (7.26), for $r \geq r^*$, $t > 0$ satisfying $r^2t \geq M^r$,

$$\begin{aligned}
& \mathbf{P}\left(d_{l,i}^{r,n}(d_{s_l}^r) < (\lambda_i^r - (\mu_{a(i)}^r \epsilon_i / 32))d_{s_l}^r, d_{\tau_{l,2n-1}}^r \leq r^2t\right) \\
& \leq \mathbf{P}\left(d_{l,i}^{r,n}(d_{s_l}^r) < (\lambda_i^r - \check{\epsilon}_i / 2)d_{s_l}^r, d_{\tau_{l,2n-1}}^r \leq r^2t\right) \\
& \leq K_{20} \exp(-K'_{20} d_{s_l}^r) \\
& \quad + \mathbf{E}\left(1_{\{d_{\tau_{n,t}}^r \in \mathbb{N}^I \times \mathbb{N}^J\}} \mathbf{P}\left(u_{l,i}^{r,n} > \frac{\check{\epsilon}_i}{4\lambda_i^r} d_{s_l}^r, d_{\tau_{l,2n-1}}^r \leq r^2t \mid \mathcal{F}_{d_{\tau_{n,t}}^r}^r\right)\right) \\
& \leq K_{20} \exp(-K'_{20} d_{s_l}^r) + (\lfloor (\lambda_i^r + \epsilon_i)r^2t \rfloor + 1) K_{21} \exp(-K'_{21} d_{s_l}^r) \\
& \quad + K_{22} \exp(-K'_{22} r^2t), \tag{8.125}
\end{aligned}$$

where $K_{20} = 1$, $K'_{20} = (\lambda_i - \check{\epsilon}_i)\Lambda_i^{a,*}(\lambda_i^{-1}(1 + \check{\epsilon}_i/4\lambda_i)) > 0$, $K_{21} = \exp(\Lambda_i^a(l_0)) > 0$, $K'_{21} = l_0\check{\epsilon}_i/4\lambda_i > 0$, $0 < l_0 \in \mathcal{O}_0$, $K_{22} = \exp(\Lambda_i^{a,*}((\lambda_i(1 + \epsilon_i/3\lambda_i))^{-1})) > 0$, and $K'_{22} = \lambda_i\Lambda_i^{a,*}((\lambda_i(1 + \epsilon_i/3\lambda_i))^{-1}) > 0$, which are all independent of t , n , and r .

For the third and fourth terms in the right side of (8.123), we have in a similar manner to (8.112), using Lemmas 7.5.2–7.6.1 (since $(x_{a(i)}^* - \epsilon_i)d_{s_l}^r > 2/\epsilon_{2,i}$, and $(x_j^* + \epsilon_{1,i})d_{s_l}^r > 2/\epsilon_{3,j}$, for all $j \in \underline{\mathcal{J}}_i$, by (8.9), (8.12), and (8.35)) for $r \geq r^*$, $t > 0$ satisfying $r^2t \geq M^r$,

$$\begin{aligned}
& \mathbf{P}\left(d_{l,a(i)}^{r,n}((x_{a(i)}^* - \epsilon_i)d_{s_l}^r) > (\mu_{a(i)}^r + \epsilon_{2,i})(x_{a(i)}^* - \epsilon_i)d_{s_l}^r + 1, d_{\tau_{l,2n-1}}^r \leq r^2t\right) \\
& \quad + \sum_{j \in \underline{\mathcal{J}}_i} \mathbf{P}\left(d_{l,j}^{r,n}((x_j^* + \epsilon_{1,i})d_{s_l}^r) > (\mu_j^r + \epsilon_{3,j})(x_j^* + \epsilon_{1,i})d_{s_l}^r + 1, d_{\tau_{l,2n-1}}^r \leq r^2t\right) \\
& \leq \mathbf{P}\left(d_{l,a(i)}^{\check{r},n}((x_{a(i)}^* - \epsilon_i)d_{s_l}^r) > (\mu_{a(i)}^r + \epsilon_{2,i})(x_{a(i)}^* - \epsilon_i)d_{s_l}^r, d_{\tau_{l,2n-1}}^r \leq r^2t\right) \\
& \quad + \sum_{j \in \underline{\mathcal{J}}_i} \mathbf{P}\left(d_{l,j}^{\check{r},n}((x_j^* + \epsilon_{1,i})d_{s_l}^r) > (\mu_j^r + \epsilon_{3,j})(x_j^* + \epsilon_{1,i})d_{s_l}^r, d_{\tau_{l,2n-1}}^r \leq r^2t\right) \\
& \leq K_{23} \exp(-K'_{23} d_{s_l}^r), \text{ where} \tag{8.126}
\end{aligned}$$

$K_{23} = |\underline{\mathcal{J}}_i| \max\{\exp(\Lambda_{a(i)}^{s,*}((\mu_{a(i)}(1 + \epsilon_{2,i}/3\mu_{a(i)})^{-1})); \exp(\Lambda_j^{s,*}((\mu_j(1 + \epsilon_{3,j}/3\mu_j)^{-1})) : j \in \underline{\mathcal{J}}_i\} > 0$, $K'_{23} = \min\{\mu_{a(i)}(x_{a(i)}^* - \epsilon_i)\Lambda_{a(i)}^{s,*}((\mu_{a(i)}(1 + \epsilon_{2,i}/3\mu_{a(i)})^{-1}); \mu_j(x_j^* +$

$\epsilon_{1,j})\Lambda_j^{s,*}((\mu_j(1 + \epsilon_{3,j}/3\mu_j)^{-1}) : j \in \underline{\mathcal{J}}_i\} > 0$, which are independent of t , n , and r .

For the last probability in (8.123), we have for $r \geq r^*$, $t > 0$ satisfying $r^2t \geq M^r$, $j \in \underline{\mathcal{J}}_i$, $k = k(j)$, as in (8.90),

$$\begin{aligned} & \mathbf{P}\left(d_{\mathcal{I}_i,j}^{T^r,n}(d_{S_\ell}^r) \geq (x_j^* + \epsilon_{1,i})d_{S_\ell}^r, d_{\mathcal{I}_i,2n-1}^r \leq r^2t\right) \\ & \leq \sum_{i' \in \underline{\mathcal{I}}_k} \mathbf{P}\left(d_{\mathcal{I}_i,a(i')}^{T^r,n}(d_{S_\ell}^r) \leq (x_{a(i')}^* - \epsilon_{i'})d_{S_\ell}^r, d_{\mathcal{I}_i,2n-1}^r \leq r^2t\right). \end{aligned} \quad (8.127)$$

Finally, for the first term in the last expression of (8.121) we have that for $r \geq r^*$, $t > 0$ satisfying $r^2t \geq M^r$,

$$\begin{aligned} & \mathbf{P}\left(Q_i^r(d_{\mathcal{I}_i,2n-1}^r + d_{S_\ell}^r) - Q_i^r(d_{\mathcal{I}_i,2n-1}^r) \geq (\mu_{a(i)}^r \epsilon_i / 32) d_{S_\ell}^r, d_{\mathcal{I}_i,2n-1}^r \leq r^2t\right) \\ & \leq \mathbf{P}\left(\sup_{0 \leq s \leq 2r^2t} Q_i^r(s) \geq (\mu_{a(i)}^r \epsilon_i / 32) d_{S_\ell}^r\right) \\ & \leq \mathbf{P}\left(\sup_{\tau_{i,0}^r \leq s \leq 2r^2t} R_i^r(s) \geq L_i^r - |\underline{\mathcal{J}}_i|\right) \\ & \leq p_{1,i}(2r^2t) \left(C_{1,i}^{(1)} \exp(-C_{1,i}^{(2)} L_i^r) + C_{1,i}^{(3)} \exp(-2C_{1,i}^{(4)} r^2t) \right), \end{aligned} \quad (8.128)$$

where $p_{1,i}$ is a polynomial (of degree at most $i + 1$) with non-negative coefficients, and $C_{1,i}^{(m)} > 0$, for $m = 1, 2, 3, 4$, since (I.1) was already proved to hold for i (with $2t$ in place of t). The polynomial and the constants do not depend on t or r . The first inequality above uses the fact that $M^r \geq d_{S_\ell}^r$, and the second inequality holds since for $r \geq r^*$,

$$\begin{aligned} & \frac{\epsilon_i \mu_{a(i)}^r}{32} \cdot d_{S_\ell}^r \\ & = \frac{\epsilon_i \mu_{a(i)}^r}{32} \cdot \frac{L_i^r - (|\underline{\mathcal{J}}_i| + 2)}{\sum_{j \in \underline{\mathcal{J}}_i} (\mu_j^r + \epsilon_i)(x_j^* + \epsilon_i)} \\ & = \frac{\epsilon_i \mu_{a(i)}^r}{32 \mu_{a(i)}} \cdot \frac{\sum_{j \in \underline{\mathcal{J}}_i} (\mu_j + \epsilon_i)(x_j^* + \epsilon_i)}{\sum_{j \in \underline{\mathcal{J}}_i} (\mu_j^r + \epsilon_i)(x_j^* + \epsilon_i)} \cdot \frac{\mu_{a(i)}}{\sum_{j \in \underline{\mathcal{J}}_i} (\mu_j + \epsilon_i)(x_j^* + \epsilon_i)} (L_i^r - (|\underline{\mathcal{J}}_i| + 2)) \\ & \geq \frac{\epsilon_i \mu_{a(i)}}{256 \sum_{j \in \underline{\mathcal{J}}_i} (\mu_j + \epsilon_i)(x_j^* + \epsilon_i)} \left(\frac{L_i^r}{\epsilon_i^3} - (|\underline{\mathcal{J}}_i| + 2) \right) \\ & \geq 2L_i^r, \end{aligned} \quad (8.129)$$

by (8.27), (8.30), (7.4), (8.29), and since $\epsilon_i < \min \left\{ \frac{\mu_{a(i)}}{1024 \sum_{j \in \underline{\mathcal{J}}_i} (\mu_j + \epsilon_i)}, 1 \right\}$ (cf. (8.97)). (We have used the fact that $x_j^* + \epsilon_i \leq 2$, since $x_j^* \leq 1$ and $\epsilon_i \leq x_j^*$ by (7.5).)

Combining all of the above from (8.121) onwards, we have for $r \geq r^*$, $t > 0$ satisfying $r^2 t \geq M^r$,

$$\begin{aligned}
& \sup_{n \geq 1} \mathbf{P} \left({}^d T_{l,a(i)}^{r,n} ({}^d s_l^r) \leq (x_{a(i)}^* - \epsilon_i) {}^d s_l^r, {}^d_{\mathcal{T}_{l,2n-1}} s_l^r \leq r^2 t \right) \\
\leq & p_{1,i}(2r^2 t) \left(C_{1,i}^{(1)} \exp(-C_{1,i}^{(2)} L_0^r) + C_{1,i}^{(3)} \exp(-2C_{1,i}^{(4)} r^2 t) \right) \\
& + K_{20} \exp(-K'_{20} {}^d s_l^r) + (\lfloor (\lambda_i^r + \epsilon_i) r^2 t \rfloor + 1) K_{21} \exp(-K'_{21} {}^d s_l^r) \\
& + K_{22} \exp(-K'_{22} r^2 t) + K_{23} \exp(-K'_{23} {}^d s_l^r) \\
& + \sum_{j \in \mathcal{I}_i} \sum_{i' \in \mathcal{I}_{k(j)}} \sup_{n \geq 1} \mathbf{P} \left({}^d T_{l,a(i')}^{r,n} ({}^d s_l^r) \leq (x_{a(i')}^* - \epsilon_{i'}) {}^d s_l^r, \right. \\
& \left. {}^d_{\mathcal{T}_{l,2n-1}} s_l^r \leq r^2 t \right). \tag{8.130}
\end{aligned}$$

By the induction assumption that (II.2) holds with i' in place of i for all $i' < i$, the definition of ${}^d s_l^r$, and the fact that ${}^d s_l^r \geq L_0^r$ for all $l \in \mathcal{I}$, it follows that for $r \geq r^*$, $t > 0$ satisfying $r^2 t \geq M^r$,

$$\begin{aligned}
& \sup_{n \geq 1} \mathbf{P} \left({}^d T_{l,a(i)}^{r,n} ({}^d s_l^r) \leq (x_{a(i)}^* - \epsilon_i) {}^d s_l^r, {}^d_{\mathcal{T}_{l,2n-1}} s_l^r \leq r^2 t \right) \\
\leq & p_{5,i}(r^2 t) \left(C_{5,i}^{(1)} \exp(-C_{5,i}^{(2)} L_0^r) + C_{5,i}^{(3)} \exp(-C_{5,i}^{(4)} r^2 t) \right), \tag{8.131}
\end{aligned}$$

where $p_{5,i}$ is a polynomial (of degree at most $i + 1$) with non-negative coefficients, and $C_{5,i}^{(m)} > 0$, for $m = 1, 2, 3, 4$. The polynomial and the constants do not depend on r or t . Thus, (II.2) holds for i .

Proof of (II.3). For $r \geq r^*$ and $\iota > i$ fixed,

$$\begin{aligned}
& \mathbf{P} \left(T_{a(i)}^r(t_\iota^r) \leq (x_{a(i)}^* - \epsilon_i)t_\iota^r \right) \\
& \leq \mathbf{P} \left(S_{a(i)}^r(T_{a(i)}^r(t_\iota^r)) \leq S_{a(i)}^r((x_{a(i)}^* - \epsilon_i)t_\iota^r) \right) \\
& \leq \mathbf{P} \left(A_i^r(t_\iota^r) - \sum_{j \in \underline{\mathcal{I}}_i} S_j^r(T_j^r(t_\iota^r)) - S_{a(i)}^r(T_{a(i)}^r(t_\iota^r)) \right. \\
& \quad \left. \geq A_i^r(t_\iota^r) - \sum_{j \in \underline{\mathcal{I}}_i} S_j^r(T_j^r(t_\iota^r)) - S_{a(i)}^r((x_{a(i)}^* - \epsilon_i)t_\iota^r) \right) \\
& \leq \mathbf{P} \left(Q_i^r(t_\iota^r) \geq (\mu_{a(i)}^r \epsilon_i / 32)t_\iota^r \right) \\
& \quad + \mathbf{P} \left(A_i^r(t_\iota^r) - \sum_{j \in \underline{\mathcal{I}}_i} S_j^r(T_j^r(t_\iota^r)) - S_{a(i)}^r((x_{a(i)}^* - \epsilon_i)t_\iota^r) \right. \\
& \quad \left. < (\mu_{a(i)}^r \epsilon_i / 32)t_\iota^r \right). \tag{8.132}
\end{aligned}$$

Let,

$$\begin{aligned}
\Upsilon_{\iota,i}^r = & \left\{ A_i^r(t_\iota^r) \geq (\lambda_i^r - (\mu_{a(i)}^r \epsilon_i / 32))t_\iota^r; \right. \\
& S_{a(i)}^r((x_{a(i)}^* - \epsilon_i)t_\iota^r) \leq (\mu_{a(i)}^r + \epsilon_{2,i})(x_{a(i)}^* - \epsilon_i)t_\iota^r; \\
& S_j^r((x_j^* + \epsilon_{1,i})t_\iota^r) \leq (\mu_j^r + \epsilon_{3,j})(x_j^* + \epsilon_{1,i})t_\iota^r, j \in \underline{\mathcal{I}}_i; \\
& \left. T_j^r(t_\iota^r) < (x_j^* + \epsilon_{1,i})t_\iota^r, j \in \underline{\mathcal{I}}_i \right\}. \tag{8.133}
\end{aligned}$$

Now, for the last term in (8.132), we have

$$\begin{aligned}
& \mathbf{P} \left(A_i^r(t_\iota^r) - \sum_{j \in \underline{\mathcal{I}}_i} S_j^r(T_j^r(t_\iota^r)) - S_{a(i)}^r((x_{a(i)}^* - \epsilon_i)t_\iota^r) < (\mu_{a(i)}^r \epsilon_i / 32)t_\iota^r \right) \\
& \leq \mathbf{P} \left(A_i^r(t_\iota^r) - \sum_{j \in \underline{\mathcal{I}}_i} S_j^r(T_j^r(t_\iota^r)) \right. \\
& \quad \left. - S_{a(i)}^r((x_{a(i)}^* - \epsilon_i)t_\iota^r) < (\mu_{a(i)}^r \epsilon_i / 32)t_\iota^r, \Upsilon_{\iota,i}^r \right) \\
& \quad + \mathbf{P} \left(A_i^r(t_\iota^r) < (\lambda_i^r - (\mu_{a(i)}^r \epsilon_i / 32))t_\iota^r \right), \\
& \quad + \mathbf{P} \left(S_{a(i)}^r((x_{a(i)}^* - \epsilon_i)t_\iota^r) > (\mu_{a(i)}^r + \epsilon_{2,i})(x_{a(i)}^* - \epsilon_i)t_\iota^r \right) \\
& \quad + \sum_{j \in \underline{\mathcal{I}}_i} \mathbf{P} \left(S_j^r((x_j^* + \epsilon_{1,i})t_\iota^r) > (\mu_j^r + \epsilon_{3,j})(x_j^* + \epsilon_{1,i})t_\iota^r \right) \\
& \quad + \sum_{j \in \underline{\mathcal{I}}_i} \mathbf{P} \left(T_j^r(t_\iota^r) \geq (x_j^* + \epsilon_{1,i})t_\iota^r \right). \tag{8.134}
\end{aligned}$$

For the first term in the right side of (8.134) we have that on $\Upsilon_{l,i}^r$, for $r \geq r^*$,

$$A_i^r(t_l^r) - \sum_{j \in \underline{\mathcal{J}}_i} S_j^r(T_j^r(t_l^r)) - S_{a(i)}^r((x_{a(i)}^* - \epsilon_i)t_l^r) \geq \frac{\mu_{a(i)}^r \epsilon_i}{32} t_l^r, \quad (8.135)$$

in a similar manner to that in (8.124). Hence the first probability in the right side of (8.134) is zero.

For the second term in the right side of (8.134) we have, using (8.31), for $r \geq r^*$,

$$\begin{aligned} & \mathbf{P}\left(A_i^r(t_l^r) < (\lambda_i^r - (\mu_{a(i)}^r \epsilon_i / 32))t_l^r\right) \\ & \leq \mathbf{P}\left(A_i^r(t_l^r) < (\lambda_i^r - \check{\epsilon}_i / 2)t_l^r\right) \\ & \leq K_{24} \exp(-K'_{24} t_l^r) + K_{25} \exp(-K'_{25} t_l^r), \end{aligned} \quad (8.136)$$

by Lemma 7.6.1 and (7.26), where $K_{24} = 1$, $K'_{24} = (\lambda_i - \check{\epsilon}_i) \Lambda_i^{a,*} (\lambda_i^{-1} (1 + \check{\epsilon}_i / 4 \lambda_i)) > 0$, $K_{25} = \exp(\Lambda_i^a(l_0)) > 0$, $K'_{25} = l_0 \check{\epsilon}_i / 4 \lambda_i > 0$, for $0 < l_0 \in \mathcal{O}_0$.

For the third and fourth terms in the right side of (8.134) we have, using Lemma 7.6.1 (since $(x_{a(i)}^* - \epsilon_i)t_l^r > 2/\epsilon_{2,i}$ and $(x_j^* + \epsilon_{1,i})t_l^r > 2/\epsilon_{3,j}$, for all $j \in \underline{\mathcal{J}}_i$, by (8.9), (8.12), and (8.35)), for $r \geq r^*$,

$$\begin{aligned} & \mathbf{P}\left(S_{a(i)}^r((x_{a(i)}^* - \epsilon_i)t_l^r) > (\mu_{a(i)}^r + \epsilon_{2,i})(x_{a(i)}^* - \epsilon_i)t_l^r\right) \\ & \quad + \sum_{j \in \underline{\mathcal{J}}_i} \mathbf{P}\left(S_j^r((x_j^* + \epsilon_{1,i})t_l^r) > (\mu_j^r + \epsilon_{3,j})(x_j^* + \epsilon_{1,i})t_l^r\right) \\ & \leq K_{26} \exp(-K'_{26} t_l^r), \text{ where} \end{aligned} \quad (8.137)$$

$K_{26} = |\mathcal{J}_i| \max\{\exp(\Lambda_{a(i)}^{s,*} ((\mu_{a(i)}(1 + \epsilon_{2,i}/3\mu_{a(i)}))^{-1})); \exp(\Lambda_j^{s,*} ((\mu_j(1 + \epsilon_{3,j}/3\mu_j))^{-1})) : j \in \underline{\mathcal{J}}_i\} > 0$, $K'_{26} = \min\{\mu_{a(i)}(x_{a(i)}^* - \epsilon_i) \Lambda_{a(i)}^{s,*} ((\mu_{a(i)}(1 + \epsilon_{2,i}/3\mu_{a(i)}))^{-1}); \mu_j(x_j^* + \epsilon_{1,i}) \Lambda_j^{s,*} ((\mu_j(1 + \epsilon_{3,j}/3\mu_j))^{-1}) : j \in \underline{\mathcal{J}}_i\} > 0$.

For the fifth term in the right side of (8.134) we have for $j \in \underline{\mathcal{J}}_i$, $k = k(j)$, in a similar manner to that in (8.92), for $r \geq r^*$,

$$\mathbf{P}(T_j^r(t_l^r) \geq (x_j^* + \epsilon_{1,i})t_l^r) \leq \sum_{i' \in \underline{\mathcal{I}}_k} \mathbf{P}(T_{a(i')}^r(t_l^r) < (x_{a(i')}^* - \epsilon_{i'})t_l^r). \quad (8.138)$$

For the first term in the last expression of (8.132) we have that for $r \geq r^*$, $t > 0$ satisfying $r^2 t \geq M^r$,

$$\begin{aligned}
& \mathbf{P} \left(Q_i^r(t_l^r) \geq (\mu_{a(i)}^r \epsilon_i / 32) t_l^r \right) \\
& \leq \mathbf{P} \left(\sup_{\tau_{i,0}^r \leq s \leq r^2 t} R_i^r(s) \geq L_i^r - |\mathcal{J}_i| \right) \\
& \leq p_{1,i}(r^2 t) \left(C_{1,i}^{(1)} \exp(-C_{1,i}^{(2)} L_0^r) + C_{1,i}^{(3)} \exp(-C_{1,i}^{(4)} r^2 t) \right), \quad (8.139)
\end{aligned}$$

where $p_{1,i}$ is a polynomial (of degree at most $i + 1$) with non-negative coefficients, and $C_{1,i}^{(m)} > 0$, for $m = 1, 2, 3, 4$, since (I.1) was proved to hold for i . The polynomial and the constants do not depend on t or r . For the first inequality in (8.139), we have used the fact that

$$\begin{aligned}
\frac{\mu_{a(i)}^r \epsilon_i}{32} t_l^r &= \frac{\mu_{a(i)}^r \epsilon_i}{32} \cdot \frac{8L_l^r}{\lambda_l - \sum_{j \in \mathcal{J}_l} x_j^* \mu_j} \\
&\geq \frac{\mu_{a(i)}^r}{\mu_{a(i)}} \cdot \frac{\mu_{a(i)} \epsilon_i}{4(\lambda_l - \sum_{j \in \mathcal{J}_l} x_j^* \mu_j)} \cdot \frac{L_l^r}{\epsilon_i^3} \\
&\geq \frac{\mu_{a(i)}}{8(\lambda_l - \sum_{j \in \mathcal{J}_l} x_j^* \mu_j)} \cdot \frac{L_l^r}{\epsilon_i^2} \\
&\geq 2L_l^r, \quad (8.140)
\end{aligned}$$

by (7.4), (8.27), and the fact that $\epsilon_i < \min \left\{ \frac{\mu_{a(i)}}{16(\lambda_l - \sum_{j \in \mathcal{J}_l} x_j^* \mu_j)}, 1 \right\}$ (cf. (8.97)).

Combining all of the above (from (8.132) onwards), we have for all $r \geq r^*$, $t > 0$ satisfying $r^2 t \geq M^r$,

$$\begin{aligned}
& \mathbf{P} \left(T_{a(i)}^r(t_l^r) \leq (x_{a(i)}^* - \epsilon_i) t_l^r \right) \\
& \leq p_{1,i}(r^2 t) \left(C_{1,i}^{(1)} \exp(-C_{1,i}^{(2)} L_0^r) + C_{1,i}^{(3)} \exp(-C_{1,i}^{(4)} r^2 t) \right) \\
& \quad + K_{24} \exp(-K'_{24} t_l^r) + K_{25} \exp(-K'_{25} t_l^r) + K_{26} \exp(-K'_{26} t_l^r) \\
& \quad + \sum_{j \in \mathcal{J}_i} \sum_{i' \in \mathcal{I}_{k(j)}} \mathbf{P} \left(T_{a(i')}^r(t_l^r) < (x_{a(i')}^* - \epsilon_{i'}) t_l^r \right). \quad (8.141)
\end{aligned}$$

By the assumption that (II.3) holds with i' in place of i , for all $i' < i$, and the

definition of t_ι^r , it follows that for $r \geq r^*$, $t > 0$ satisfying $r^2t \geq M^r$,

$$\begin{aligned} & \mathbf{P} \left(T_{a(i)}^r(t_\iota^r) \leq (x_{a(i)}^* - \epsilon_i)t_\iota^r \right) \\ & \leq p_{6,i}(r^2t) \left(C_{6,i}^{(1)} \exp(-C_{6,i}^{(2)}L_0^r) + C_{6,i}^{(3)} \exp(-C_{6,i}^{(4)}r^2t) \right), \end{aligned} \quad (8.142)$$

where $p_{6,i}$ is a polynomial (of degree at most $i+1$) with non-negative coefficients, and $C_{6,i}^{(m)} > 0$ for $m = 1, 2, 3, 4$, since (I.1) was already proved to hold for i . Thus, (II.3) holds for i . \square

8.8 Proofs of Theorems 7.3.1, 8.1.1, and 8.3.5

We will now establish the state space collapse result described in Section 7.3.

Proof of Theorem 8.3.5. Fix $i \in \mathcal{I} \setminus \{i^*\}$ and assume that (I) and (II) hold for all $i' < i$. Then, by Lemmas 8.3.2–8.3.3, (i)–(iv) in Lemma 8.3.1 hold for $r \geq r^*$, $t > 0$ satisfying $r^2t \geq M^r$, for i and each $j \in \underline{\mathcal{J}}_i$, and hence (I) holds for i by Lemma 8.3.1. By Lemma 8.3.4, we have that (II) also holds for i . The conclusion of the first statement in Theorem 8.3.5 then follows by the induction principle.

For (III), suppose that i^* is a transition class, and let $k \in \underline{\mathcal{K}}_{i^*}$. By the above, for $r \geq r^*$, $t > 0$ satisfying $r^2t \geq M^r$, we have that (II) holds with i^* in place of ι , $i' \in \underline{\mathcal{I}}_k$ in place of i .

For $r \geq r^*$, by the same proof as for Lemma 8.3.2, (ii) of Lemma 8.3.1 holds with i^* in place of i (cf. (8.90)), and then by the same proof as for Lemma 8.3.1 (cf. the second half of (I.1), (8.59)–(8.77), for the (down) excursions of $R_{i^*}^r$), with i^* in place of i there, we have that for $r \geq r^*$

$$\begin{aligned} & \mathbf{P} \left(\inf_{\tau_{i^*,0}^r \leq s \leq r^2t} R_{i^*}^r(s) \leq -L_{i^*}^r + |\underline{\mathcal{J}}_{i^*}| \right) \\ & \leq p_{1,i^*}(r^2t) \left(C_{1,i^*}^{(1)} \exp(-C_{1,i^*}^{(2)}L_0^r) + C_{1,i^*}^{(3)} \exp(-C_{1,i^*}^{(4)}r^2t) \right), \end{aligned} \quad (8.143)$$

where p_{1,i^*} is a polynomial (of degree at most $\mathbf{I}+1$) with non-negative coefficients, and $C_{1,i^*}^{(m)} > 0$, $m = 1, 2, 3, 4$ are independent of r and t . Thus (III.1) holds.

To establish (III.2), we note that by the same proof as for Lemma 8.3.2, for $r \geq r^*$, (iii) of Lemma 8.3.1 holds with i^* in place of i (i.e., for $j \in \underline{\mathcal{J}}_{i^*}$), since (II.3) holds with i^* in place of ι and $i' \in \underline{\mathcal{I}}_k$ (where $k = k(j)$) in place of i (cf. (8.92)), for $r \geq r^*$. Then, by the same proof as for (I.2) in Lemma 8.3.1 (with i^* in place of i there, cf. (8.78)–(8.86)), we have,

$$\begin{aligned} & \mathbf{P} \left(I_k^r(\tau_{i^*,0}^r) \geq t_{i^*}^r \right) \\ & \leq p_{2,i^*}(r^2t) \left(C_{2,i^*}^{(1)} \exp \left(- C_{2,i^*}^{(2)} L_0^r \right) + C_{2,i^*}^{(3)} \exp \left(- C_{2,i^*}^{(4)} r^2t \right) \right), \end{aligned} \quad (8.144)$$

where p_{2,i^*} is a polynomial (of degree at most **I**) with non-negative coefficients, and $C_{2,i^*}^{(m)} > 0$, $m = 1, 2, 3, 4$ are independent of r and t .

Finally, for (III.3) we argue as in the proof of (I.3) in Lemma 8.3.1 (cf. (8.87)), that I_k^r , $k \in \underline{\mathcal{K}}_{i^*}$, can increase only at times $s \geq 0$ such that $Q_{i^*}^r(s) \leq |\underline{\mathcal{J}}_{i^*}|$, so that for $r \geq r^*$, $t > 0$ satisfying $r^2t \geq M^r$,

$$\begin{aligned} & \mathbf{P} \left(I_k^r(r^2t) - I_k^r(\tau_{i^*,0}^r) > 0, \tau_{i^*,0}^r < r^2t \right) \\ & \leq \mathbf{P} \left(\inf_{\tau_{i^*,0}^r \leq s \leq r^2t} Q_{i^*}^r(s) \leq |\underline{\mathcal{J}}_{i^*}| \right) \\ & \leq p_{1,i^*}(r^2t) \left(C_{1,i^*}^{(1)} \exp \left(- C_{1,i^*}^{(2)} L_0^r \right) + C_{1,i^*}^{(3)} \exp \left(- C_{1,i^*}^{(4)} r^2t \right) \right), \end{aligned} \quad (8.145)$$

by (8.143). \square

Proof of Theorem 8.1.1. Note that (I) and (II) hold for all $i \in \mathcal{I} \setminus \{i^*\}$ by Theorem 8.3.5. Fix $i \in \mathcal{I} \setminus \{i^*\}$, $k \in \underline{\mathcal{K}}_i$, and $\epsilon > 0$. For $t = 0$, (8.1) and (8.3) hold trivially since $R_i^r(0) = 0$ if $\tau_{i,0}^r = 0$. So we assume that $t > 0$ is fixed. Since $M^r = O(\log r)$, there exists an $r_t \geq r^*$ such that for all $r \geq r_t$, $r^2t \geq M^r$. Then for $r \geq r_t$,

$$\begin{aligned} & \mathbf{P} \left(\sup_{\tau_{i,0}^r \leq s \leq r^2t} |R_i^r(s)| \geq L_i^r - |\underline{\mathcal{J}}_i| \right) \\ & \leq p_{1,i}(r^2t) \left(C_{1,i}^{(1)} \exp \left(- C_{1,i}^{(2)} L_0^r \right) + C_{1,i}^{(3)} \exp \left(- C_{1,i}^{(4)} r^2t \right) \right), \end{aligned} \quad (8.146)$$

by (I.1), where $p_{1,i}$ is a polynomial of degree at most $i + 1$ and $C_{1,i}^{(m)}$, $m = 1, 2, 3, 4$ are positive constants. The polynomial and the constants are independent of t and r .

Since $t_i^r = o(r)$, there is an $r' \geq r_t$, such that $t_i^r \leq r\epsilon$ for all $r \geq r'$. Then, by (I.2), for $r \geq r'$

$$\begin{aligned} & \mathbf{P} \left(I_k^r(\tau_{i,0}^r) \geq r\epsilon \right) \\ & \leq \mathbf{P} \left(I_k^r(\tau_{i,0}^r) \geq t_i^r \right) \\ & \leq p_{2,i}(r^2t) \left(C_{2,i}^{(1)} \exp(-C_{2,i}^{(2)}L_0^r) + C_{2,i}^{(3)} \exp(-C_{2,i}^{(4)}r^2t) \right), \end{aligned} \quad (8.147)$$

where $p_{2,i}$ is a polynomial of degree at most i and $C_{2,i}^{(m)}$, $m = 1, 2, 3, 4$, are positive constants. The polynomial and constants do not depend on t or r .

Finally, we have by (I.3) that for $r \geq r_t$,

$$\begin{aligned} & \mathbf{P} \left(I_k^r(r^2t) - I_k^r(\tau_{i,0}^r) > 0, \tau_{i,0}^r < r^2t \right) \\ & \leq p_{3,i}(r^2t) \left(C_{3,i}^{(1)} \exp(-C_{3,i}^{(2)}L_0^r) + C_{3,i}^{(3)} \exp(-C_{3,i}^{(4)}r^2t) \right), \end{aligned} \quad (8.148)$$

where $p_{3,i}$ is a polynomial of degree at most $i + 1$ and $C_{3,i}^{(m)}$, $m = 1, 2, 3, 4$ are positive constants. The polynomial and constants do not depend on t or r .

Since $L_0^r = \lceil c \log r \rceil$ (and hence for $d > 0$, $\exp(-dL_0^r) \leq r^{-cd}$), it follows from (8.146)–(8.148) that there is a constant $c_0 > 0$ (not depending on t or r) such that if $c \geq c_0$, the expressions in (8.146)–(8.148) tend to zero as $r \rightarrow \infty$ (for each fixed $t > 0$).

For (8.4)–(8.6), suppose that i^* is a transition class. Let $k \in \underline{\mathcal{K}}_{i^*}$, and $\epsilon > 0$. For $t = 0$, (8.4) and (8.6) hold trivially since $\tau_{i^*,0}^r > 0$ for a transition class, so we fix $t > 0$. Then, as above, for $r \geq r_t$, we have that $r^2t \geq M^r$. We then have, for $r \geq r_t$,

$$\begin{aligned} & \mathbf{P} \left(\inf_{\tau_{i^*,0}^r \leq s \leq r^2t} Q_{i^*}^r(s) \leq |\underline{\mathcal{J}}_{i^*}| \right) \\ & = \mathbf{P} \left(\inf_{\tau_{i^*,0}^r \leq s \leq r^2t} R_{i^*}^r(s) \leq -L_{i^*}^r + |\underline{\mathcal{J}}_{i^*}| \right) \\ & \leq p_{1,i^*}(r^2t) \left(C_{1,i^*}^{(1)} \exp(-C_{1,i^*}^{(2)}L_0^r) + C_{1,i^*}^{(3)} \exp(-C_{1,i^*}^{(4)}r^2t) \right), \end{aligned} \quad (8.149)$$

by (III.1), where p_{1,i^*} is a polynomial (of degree at most $\mathbf{I} + 1$) with non-negative coefficients, and $C_{1,i^*}^{(m)} > 0$, $m = 1, 2, 3, 4$ are independent of r and t .

Since $t_{i^*}^r = o(r)$, there exist $r'' \geq r_t$, such that $t_{i^*}^r \leq r\epsilon$, for all $r \geq r''$. Then by (III.2), we have for $r \geq r''$,

$$\begin{aligned} & \mathbf{P} \left(I_k^r(\tau_{i^*,0}^r) \geq r\epsilon \right) \\ & \leq \mathbf{P} \left(I_k^r(\tau_{i^*,0}^r) \geq t_{i^*}^r \right) \\ & \leq p_{2,i^*}(r^2t) \left(C_{2,i^*}^{(1)} \exp(-C_{2,i^*}^{(2)}L_0^r) + C_{2,i^*}^{(3)} \exp(-C_{2,i^*}^{(4)}r^2t) \right), \end{aligned} \quad (8.150)$$

where p_{2,i^*} is a polynomial (of degree at most \mathbf{I}) with non-negative coefficients, and $C_{2,i^*}^{(m)} > 0$, $m = 1, 2, 3, 4$ are independent of r and t .

Finally, we have by (III.3) that for $r \geq r_t$,

$$\begin{aligned} & \mathbf{P}(I_k^r(r^2t) - I_k^r(\tau_{i^*,0}^r) > 0, \tau_{i^*,0}^r < r^2t) \\ & \leq p_{3,i^*}(r^2t) \left(C_{3,i^*}^{(1)} \exp(-C_{3,i^*}^{(2)}L_0^r) + C_{3,i^*}^{(3)} \exp(-C_{3,i^*}^{(4)}r^2t) \right), \end{aligned} \quad (8.151)$$

where p_{3,i^*} is a polynomial of degree at most $\mathbf{I} + 1$ and $C_{3,i^*}^{(m)}$, $m = 1, 2, 3, 4$ are positive constants. The polynomial and the constants are independent of t and r .

Since $L_0^r = \lceil c \log r \rceil$, it follows that there is a constant $c_1 \geq c_0$ (not depending on t or r), such that if $c \geq c_1$, then the expressions in (8.149)–(8.151) tend to zero as $r \rightarrow \infty$ (for each fixed $t > 0$). \square

We now show that state space collapse occurs for our parallel server system operating under the threshold policy, $\{T^{r,*}\}$.

Proof of Theorem 7.3.1. It suffices to show that as $r \rightarrow \infty$, $(\hat{Q}_i^r : i \in \mathcal{I} \setminus \{i^*\}; \hat{I}_k^r : k \in \underline{\mathcal{K}}_i, i \in \mathcal{I}) \Rightarrow \mathbf{0}$. Note for this that $k \in \mathcal{K} \setminus \{k^*\}$ is either in $\underline{\mathcal{K}}_i$ for some $i \in \mathcal{I} \setminus \{i^*\}$ or in $\underline{\mathcal{K}}_{i^*}$ if i^* is a transition class.

Fix $t > 0$. By Theorem 8.1.1 and the properties of $\{L_i^r\}_{i \in \mathcal{I}}$, for each $\epsilon > 0$, there is an $r'' = r''(\epsilon, t) \geq 1$ such that whenever $r \geq r''$ for all $i \in \mathcal{I} \setminus \{i^*\}$, $k \in \underline{\mathcal{K}}_i$, $2L_i^r/r < \epsilon$, and

$$\mathbf{P} \left(\sup_{\tau_{i,0}^r \leq s \leq r^2t} |R_i^r(s)| \geq L_i^r - |\underline{\mathcal{J}}_i| \right) < \epsilon, \quad (8.152)$$

$$\mathbf{P} \left(I_k^r(\tau_{i,0}^r) \geq r\epsilon \right) < \epsilon, \quad (8.153)$$

$$\mathbf{P} \left(I_k^r(r^2t) - I_k^r(\tau_{i,0}^r) > 0, \tau_{i,0}^r < r^2t \right) < \epsilon, \quad (8.154)$$

and if i^* is a transition buffer, for all $k \in \underline{\mathcal{K}}_{i^*}$,

$$\mathbf{P}\left(I_k^r(r^2t) - I_k^r(\tau_{i^*,0}^r) > 0, \tau_{i^*,0}^r < r^2t\right) < \varepsilon. \quad (8.155)$$

Recalling the definition of $\|\cdot\|_t$ from Section 1.1, we then have for all $r \geq r''$,

$$\begin{aligned} & \mathbf{P}\left(\|\hat{Q}_i^r\|_t \geq \varepsilon \text{ for some } i \in \mathcal{I} \setminus \{i^*\}, \text{ or } \|\hat{I}_k^r\|_t \geq \varepsilon \text{ for some } k \in \mathcal{K} \setminus \{k^*\}\right) \\ = & \mathbf{P}\left(\|Q_i^r\|_{r^2t} \geq r\varepsilon \text{ for some } i \in \mathcal{I} \setminus \{i^*\}, \text{ or } \|I_k^r\|_{r^2t} \geq r\varepsilon \text{ for some } k \in \mathcal{K} \setminus \{k^*\}\right) \\ \leq & \mathbf{P}\left(\sup_{\tau_{i,0}^r \leq s \leq r^2t} Q_i^r(s) \geq 2L_i^r \text{ for some } i \in \mathcal{I} \setminus \{i^*\}, \text{ or} \right. \\ & \left. I_k^r(\tau_{i,0}^r) \geq r\varepsilon \text{ for some } k \in \underline{\mathcal{K}}_i \text{ and } i \in \mathcal{I}, \text{ or} \right. \\ & \left. \tau_{i,0}^r < r^2t \text{ and } I_k^r(r^2t) - I_k^r(\tau_{i,0}^r) > 0 \text{ for some } k \in \underline{\mathcal{K}}_i \text{ and } i \in \mathcal{I}\right) \\ \leq & \sum_{i \in \mathcal{I} \setminus \{i^*\}} \mathbf{P}\left(\sup_{\tau_{i,0}^r \leq s \leq r^2t} |R_i^r(s)| \geq L_i^r - |\underline{\mathcal{J}}_i|\right) \\ & + \sum_{i \in \mathcal{I}} \sum_{k \in \underline{\mathcal{K}}_i} \left\{ \mathbf{P}(I_k^r(\tau_{i,0}^r) \geq r\varepsilon) + \mathbf{P}(I_k^r(r^2t) - I_k^r(\tau_{i,0}^r) > 0, \tau_{i,0}^r < r^2t) \right\} \\ < & (\mathbf{I} + 2\mathbf{K} - 3)\varepsilon. \quad (8.156) \end{aligned}$$

□

9

Weak Convergence under the Threshold Policy

This chapter is devoted to the proof of Theorem 6.2.1. Throughout this chapter we assume that the allocation processes $T^{r,*}$ and thresholds for our threshold policy as described in Chapters 6 and 7 are used in the r^{th} parallel server system. We append a superscript $*$ to the queue length and idle time processes $(Q^{r,*}, I^{r,*})$ to indicate that the threshold policy is employed. Recall that $T_j^{r,*} \equiv \mathbf{0}$ for the non-basic activities $j = \mathbf{B} + 1, \dots, \mathbf{J}$, and for all $r \geq 1$.

In addition to the scaled processes defined in Chapter 3, we define the following fluid scaled processes. For each $t \geq 0$, let

$$\bar{A}^r(t) = r^{-2}A^r(r^2t), \tag{9.1}$$

$$\bar{S}^r(t) = r^{-2}S^r(r^2t), \tag{9.2}$$

$$\bar{I}^r(t) = r^{-2}I^r(r^2t), \tag{9.3}$$

$$\bar{Q}^r(t) = r^{-2}Q^r(r^2t). \tag{9.4}$$

9.1 Fluid Limits for Allocation Processes

Recall the definitions from Section 2.2 of the functions,

$$i : \mathcal{J} \rightarrow \mathcal{I}, \quad k : \mathcal{J} \rightarrow \mathcal{K}, \quad (9.5)$$

where for $j \in \mathcal{J}$, $i(j)$ is the buffer processed by activity j and $k(j)$ is the server which processes activity j . Also, recall from Assumption 3.1.1 that

$$\lambda^r \rightarrow \lambda, \quad \mu^r \rightarrow \mu, \quad \text{as } r \rightarrow \infty. \quad (9.6)$$

Recall from Chapters 6 and 7 that $a(i)$ is the activity above buffer i in the server-buffer tree, and k^* is the server at the root of the server-buffer tree in layer $l = l^*$. For the top layer $l = l^*$ and $i \in \mathcal{I}^{l^*}$, define

$$z_{k^*}^r = z_{k^*}, \quad y_i^r = \frac{z_{k^*}^r}{\mu_{a(i)}^r}, \quad (9.7)$$

and for layer $l = l^* - 1, \dots, 1$, define by backwards induction on l ,

$$\begin{aligned} z_k^r &= y_{i(b(k))}^r \mu_{b(k)}^r, \quad \text{for each } k \in \mathcal{K}^l, \\ y_i^r &= \frac{z_{k(a(i))}^r}{\mu_{a(i)}^r}, \quad \text{for each } i \in \mathcal{I}^l, \end{aligned} \quad (9.8)$$

where $b(k)$ is the activity immediately above server k which links it to a buffer in the next highest layer (see Figure 9.1). Here $z_{k^*}^*$ is the variable, determined from one component of the unique solution (y^*, z^*) of the dual problem (5.2) defined in Chapter 5. Then, since each basic activity either links a buffer in some layer $l + 1$ to a server below it in layer l , or links a server in some layer l to a buffer immediately below it in layer l , we have for each basic activity j ,

$$y_{i(j)}^r \mu_j^r = z_{k(j)}^r, \quad j = 1, \dots, \mathbf{B}, \quad (9.9)$$

i.e., for any basic activity j ,

$$\left((y^r)' \mathbf{R}^r \right)_j = \left((z^r)' \mathbf{A} \right)_j, \quad (9.10)$$




Figure 9.1: Activity $b(k)$ above a server k that is in layer l

where $y^r = (y_i^r : i \in \mathcal{I})$, $z^r = (z_k^r : k \in \mathcal{K})$,

$$\mathbf{R}^r = \mathbf{C} \text{diag}(\mu^r), \quad \mathbf{C}_{ij} = \delta_{i,i(j)}, \quad \mathbf{A}_{kj} = \delta_{k,k(j)}, \quad \text{for all } j \in \mathcal{J}, \quad (9.11)$$

and $\delta_{x,y} = 1$ if $x = y$ and $\delta_{x,y} = 0$ otherwise. The dimensions of \mathbf{C} and \mathbf{A} are $\mathbf{I} \times \mathbf{J}$ and $\mathbf{K} \times \mathbf{J}$, respectively (cf. (2.1), (2.2), and below (3.2.5)).

Lemmas 9.1.1, 9.1.2, and 9.1.4, which are proved below, will be used in Chapter 10 for the proof of asymptotic optimality (Theorem 6.2.2) as well as in the proof of Lemma 9.1.5, which shows that the fluid scaled allocation processes, $\{\bar{T}^{r,*}\}$, associated with the threshold policy, converge in distribution (as r goes to infinity) to the nominal allocation processes \bar{T}^* defined in (4.1).

Lemma 9.1.1 *We have that $(y^r, z^r) \rightarrow (y^*, z^*)$ as $r \rightarrow \infty$, where (y^*, z^*) is the optimal solution to the dual linear program specified in (5.2).*

Proof. We prove that $z_k^r \rightarrow z_k^*$, $y_i^r \rightarrow y_i^*$, as $r \rightarrow \infty$, for all $k \in \mathcal{K}^l$, $i \in \mathcal{I}^l$, for $l = 1, 2, \dots, l^*$, by backward induction on l . The result of the lemma then follows, since each server belongs to some layer and similarly for each buffer. From (5.5) and (9.9), we have that $y_{i(j)}^* \mu_j = z_{k(j)}^*$ and $y_{i(j)}^r \mu_j = z_{k(j)}^r$ for all basic activities $j = 1, 2, \dots, \mathbf{B}$.

Note that layer l^* has one server k^* and the buffers in that layer are indexed by $\underline{\mathcal{I}}_{k^*}$. We have that (by definition) $z_{k^*}^r = z_{k^*}$. Then for $i \in \underline{\mathcal{I}}_{k^*}$, we have by (9.6)

and (9.7), that

$$y_i^r = \frac{z_{k^*}^*}{\mu_{a(i)}^r} \rightarrow \frac{z_{k^*}^*}{\mu_{a(i)}} = \frac{z_{k(a(i))}^*}{\mu_{a(i)}} = y_{i(a(i))}^* = y_i^*, \quad \text{as } r \rightarrow \infty. \quad (9.12)$$

Now, for the induction step, suppose that $(y_i^r, z_k^r) \rightarrow (y_i^*, z_k^*)$ as $r \rightarrow \infty$, for all $i \in \mathcal{I}^l$, $k \in \mathcal{K}^l$, some $l \geq 2$. For $k \in \mathcal{K}^{(l-1)}$, $i \in \underline{\mathcal{I}}_k$, we have by (9.6), (9.9), and (5.5) that as $r \rightarrow \infty$,

$$z_k^r = y_{i(b(k))}^r \mu_{b(k)}^r \rightarrow y_{i(b(k))}^* \mu_{b(k)} = z_{k(b(k))}^* = z_k^*, \quad (9.13)$$

since $i(b(k)) \in \mathcal{I}^l$, and

$$y_i^r = \frac{z_{k(a(i))}^r}{\mu_{a(i)}^r} \rightarrow \frac{z_{k(a(i))}^*}{\mu_{a(i)}} = y_{a(i)}^* = y_i^*, \quad (9.14)$$

since $k(a(i)) = k$. This completes the induction step and the conclusion of the lemma then follows. \square

We have by (3.17), (3.19), (3.20), Assumption 3.2.2, and (9.11) that

$$\bar{Q}^r(t) = r^{-1} \hat{A}^r(t) - r^{-1} \mathbf{C} \hat{S}^r(\bar{T}^r(t)) + (\lambda^r - \mathbf{R}^r x^*)t + \mathbf{R}^r(x^*t - \bar{T}^r(t)) \quad (9.15)$$

$$\bar{I}^r(t) = \mathbf{1}t - \mathbf{A} \bar{T}^r(t), \quad (9.16)$$

where $\mathbf{1}$ is the \mathbf{K} -dimensional vector of all ones. Recall from Assumption 3.2.5 that

$$r(\lambda^r - \mathbf{R}^r x^*) \rightarrow \theta \quad \text{as } r \rightarrow \infty, \quad (9.17)$$

where $\theta \in \mathbb{R}^{\mathbf{I}}$.

Lemma 9.1.2 *For (y^r, z^r) as defined in (9.7)–(9.8), we have that*

$$(y^r)' \mathbf{R}^r(x^*t - \bar{T}^{r,*}(t)) = (z^r)' \bar{I}^{r,*}(t), \quad \text{for all } t \geq 0.$$

Proof. Since $x_j^* = 0$ and $\bar{T}_j^{r,*} \equiv \mathbf{0}$ for all non-basic activities j , we have from Assumption 3.2.2, (9.10) and (9.16) that for all $t \geq 0$,

$$\begin{aligned} (y^r)' \mathbf{R}^r (x^* t - \bar{T}^{r,*}(t)) &= (z^r)' (\mathbf{A} x^* t - \mathbf{A} \bar{T}^{r,*}(t)) \\ &= (z^r)' (\mathbf{1} t - \mathbf{A} \bar{T}^{r,*}(t)) \\ &= (z^r)' \bar{T}^{r,*}(t). \end{aligned} \tag{9.18}$$

□

Definition 9.1.3 *A sequence of processes with paths in D^m for some $m \geq 1$ is called C-tight if it is tight in D^m and any weak limit point of the sequence (obtained as a weak limit along a subsequence) has continuous paths almost surely.*

For Lemma 9.1.4 below, note that the result holds for *any* sequence of scheduling policies, not just for the threshold policy.

Lemma 9.1.4 *Let $\{T^r\}$ be any sequence of scheduling control policies (one for each member of the sequence of parallel server systems). Then*

$$\left\{ (\bar{Q}^r(\cdot), \bar{A}^r(\cdot), \bar{S}^r(\cdot), \bar{T}^r(\cdot), \bar{I}^r(\cdot)) \right\} \text{ is C-tight.}$$

Proof. It follows from (3.23) that

$$(\bar{A}^r(\cdot), \bar{S}^r(\cdot)) \Rightarrow (\lambda(\cdot), \mu(\cdot)) \quad \text{as } r \rightarrow \infty, \tag{9.19}$$

where $\lambda(t) = \lambda t$ and $\mu(t) = \mu t$ for all $t \geq 0$. In addition, since they correspond to cumulative allocations of time, each of the components of T^r is uniformly Lipschitz continuous with a Lipschitz constant less than or equal to one and this property is preserved by the fluid scaled processes \bar{T}^r . It follows immediately from this and (9.19) that $\{(\bar{A}^r(\cdot), \bar{S}^r(\cdot), \bar{T}^r(\cdot))\}$ is C-tight, cf. Theorem 15.5 in [3]. From the equations (2.16)–(2.17) for queue length and idletime we have that for each $t \geq 0$,

$$\bar{Q}^r(t) = \bar{A}^r(t) - \mathbf{C} \bar{S}^r(\bar{T}^r(t)), \tag{9.20}$$

$$\bar{I}^r(t) = \mathbf{1} t - \mathbf{A} \bar{T}^r(t). \tag{9.21}$$

Combining these with the C-tightness established above and a random time change theorem (cf. [3], p. 145), we obtain the desired result. \square

Lemma 9.1.5 below is needed in the proof of Theorem 6.2.1 in combining the functional central limit theorem, (3.23), with a random time change theorem.

Lemma 9.1.5 *For the fluid scaled allocation processes, $\bar{T}_j^{r,*}$, $j \in \mathcal{J}$, we have,*

$$\bar{T}^{r,*} \Rightarrow \bar{T}^* \quad \text{as } r \rightarrow \infty, \quad (9.22)$$

where $\bar{T}^*(t) = x^*t$, for all $t \geq 0$.

Proof. We note first that since $\bar{T}_j^{r,*} \equiv \mathbf{0}$ and $\bar{T}_j^* \equiv 0$ for $j = \mathbf{B} + 1, \dots, \mathbf{J}$, for all $r \geq 1$, we have that (trivially)

$$\bar{T}_j^{r,*}(\cdot) \Rightarrow \bar{T}_j^*(\cdot) \quad \text{as } r \rightarrow \infty, \quad \text{for } j = \mathbf{B} + 1, \dots, \mathbf{J}. \quad (9.23)$$

Now from (3.23) and the fact that $\bar{T}_j^{r,*}(t) \leq t$, $j \in \mathcal{J}$, for all $t \geq 0$, we have that

$$(r^{-1}\hat{A}_i^r(\cdot), r^{-1}\hat{S}_j^r(\bar{T}_j^{r,*}(\cdot))) : i \in \mathcal{I}, j \in \mathcal{J} \Rightarrow \mathbf{0}, \quad \text{as } r \rightarrow \infty, \quad (9.24)$$

where $\mathbf{0}$ is the identically zero function in $D^{\mathbf{I}+\mathbf{J}}$.

Using (9.15) and Lemma 9.1.2 we have that for each $t \geq 0$,

$$y^r \cdot \bar{Q}^{r,*}(t) = y^r \cdot \bar{X}^{r,*}(t) + z^r \cdot \bar{I}^{r,*}(t), \quad (9.25)$$

where

$$\bar{X}^{r,*}(t) = r^{-1}\hat{A}^r(t) - r^{-1}\mathbf{C}\hat{S}^r(\bar{T}^{r,*}(t)) + (\lambda^r - \mathbf{R}^r x^*)t, \quad (9.26)$$

and by (9.24), (9.6) and (9.17),

$$\bar{X}^{r,*}(\cdot) \Rightarrow \mathbf{0} \quad \text{as } r \rightarrow \infty. \quad (9.27)$$

Thus, for $t \geq 0$,

$$y_{i^*}^r \bar{Q}_{i^*}^{r,*}(t) = \bar{\zeta}^{r,*}(t) + z_{k^*}^r \bar{I}_{k^*}^{r,*}(t), \quad (9.28)$$

where, by (9.25), (9.27), Lemma 9.1.1, and Theorem 7.3.1 we have

$$\bar{\zeta}^{r,*}(\cdot) \equiv y^r \cdot \bar{X}^{r,*}(\cdot) - \sum_{i \in \mathcal{I} \setminus \{i^*\}} y_i^r \bar{Q}_i^{r,*}(\cdot) + \sum_{k \in \mathcal{K} \setminus \{k^*\}} z_k^r \bar{I}_k^{r,*}(\cdot) \Rightarrow \mathbf{0} \quad (9.29)$$

as $r \rightarrow \infty$, where $\mathbf{0}$ is the one-dimensional zero process. Here $y^r > 0$, $z^r > 0$, $\bar{Q}_{i^*}^{r,*} \geq 0$, $\bar{\zeta}^{r,*}(0) = 0$, $\bar{I}_{k^*}^{r,*}(0) = 0$, $\bar{I}_{k^*}^{r,*}(\cdot)$ is continuous and non-decreasing. Furthermore, from the definition of the threshold policy in Chapter 6, for r sufficiently large, we note that, in the r^{th} system operating under the threshold policy, the idletime at server k^* can increase only if the queue length at buffer i^* is at or below its associated threshold level. Hence $\bar{I}_{k^*}^{r,*}$ can only increase if $\bar{Q}_{i^*}^{r,*}$ is at or below the level $L_{i^*}^r/r^2$. (Since $L_{i^*}^r \rightarrow \infty$ as $r \rightarrow \infty$, we have that $L_{i^*}^r > |\underline{\mathcal{J}}_{i^*}|$ for sufficiently large r . If this condition were violated, then the idletime at server k^* might increase if there were between $L_{i^*}^r$ and $|\underline{\mathcal{J}}_{i^*}|$ jobs either in service or in suspension at the servers in $\underline{\mathcal{K}}_{i^*}$.) Note that the above even applies to the case when i^* is a non-transition class since in this case server k^* is busy whenever buffer i^* is nonempty, and in particular whenever $Q_{i^*}^r > L_{i^*}^r > 0$.

It then follows from an oscillation inequality for solutions of a perturbed Skorokhod problem (cf. the proof of Theorem 5.1 in [37]) that

$$z_{k^*}^r \bar{I}_{k^*}^{r,*}(t) \leq - \inf_{0 \leq s \leq t} (\bar{\zeta}^{r,*}(s)) + y_{i^*}^r L_{i^*}^r r^{-2}. \quad (9.30)$$

Hence, it follows by (9.29), the continuous mapping theorem, Lemma 9.1.1, and the fact that $z_{k^*} > 0$, that

$$\bar{I}_{k^*}^{r,*} \Rightarrow \mathbf{0} \text{ as } r \rightarrow \infty. \quad (9.31)$$

By (9.28) and (9.29), we then have that

$$\bar{Q}_{i^*}^{r,*} \Rightarrow \mathbf{0} \text{ as } r \rightarrow \infty, \quad (9.32)$$

since $y_{i^*}^* > 0$. To obtain the conclusion of the lemma, we appeal to Lemma 9.1.4 and assume that $(\bar{Q}(\cdot), \bar{A}(\cdot), \bar{S}(\cdot), \bar{T}(\cdot), \bar{I}(\cdot))$ is obtained as a weak limit point of $\{(\bar{Q}^{r,*}(\cdot), \bar{A}^r(\cdot), \bar{S}^r(\cdot), \bar{T}^{r,*}(\cdot), \bar{I}^{r,*}(\cdot))\}$ along a subsequence. Since $(\bar{Q}(\cdot), \bar{I}(\cdot)) \equiv \mathbf{0}$

(by Theorem 7.3.1, (9.31), and (9.32)), and $(\bar{A}(\cdot), \bar{S}(\cdot)) = (\lambda(\cdot), \mu(\cdot))$ (by (9.19)) we have by passing to the limit in (9.20)–(9.21) that \bar{T} satisfies

$$\mathbf{0} = \lambda t - \mathbf{R}\bar{T}(t), \quad (9.33)$$

$$\mathbf{0} = \mathbf{1}t - \mathbf{A}\bar{T}(t), \quad (9.34)$$

for all $t \geq 0$, where \bar{T} inherits the properties (3.7) from $\bar{T}^{r,*}$. Thus, the fluid system is balanced and incurs no idleness (cf. Chapter 3) under \bar{T} . Hence, by Definition 3.2.1, Assumption 3.2.2, and Lemma 3.2.3, we have that $\bar{T} = \bar{T}^*$. It follows that $(\bar{Q}^{r,*}(\cdot), \bar{A}^r(\cdot), \bar{S}^r(\cdot), \bar{T}^{r,*}(\cdot), \bar{I}^{r,*}(\cdot)) \Rightarrow (\mathbf{0}, \lambda(\cdot), \mu(\cdot), \bar{T}^*(\cdot), \mathbf{0})$, as $r \rightarrow \infty$. The conclusion of the lemma then follows. \square

9.2 Proof of Theorem 6.2.1: Convergence of Diffusion Scaled Performance Measures under the Threshold Policy

We now prove that the diffusion scaled performance measures for our sequence of parallel server systems operating under the allocation processes $\{T^{r,*}\}$ converge in distribution to the processes given in (5.13)–(5.14).

Proof of Theorem 6.2.1. For each $t \geq 0$, we have by multiplying (9.25) through by r that

$$\hat{W}^{r,*}(t) \equiv y^r \cdot \hat{Q}^{r,*}(t) = y^r \cdot \hat{X}^{r,*}(t) + z^r \cdot \hat{I}^{r,*}(t) \quad (9.35)$$

where $\hat{X}^{r,*}$ is defined by (4.2) (with \bar{T}^r replaced by $\bar{T}^{r,*}$ there). By the functional central limit theorem (3.23), Lemma 9.1.5, and a random time change theorem (cf. [3] §17), we have that as $r \rightarrow \infty$,

$$(\hat{A}^r(\cdot), \hat{S}^r(\bar{T}^{r,*}(\cdot))) \Rightarrow (\tilde{A}(\cdot), \tilde{S}(\bar{T}^*(\cdot))). \quad (9.36)$$

It then follows from the definition of $\hat{X}^{r,*}$, (9.6), and (9.17) that

$$\hat{X}^{r,*} \Rightarrow \tilde{X}, \quad \text{as } r \rightarrow \infty, \quad (9.37)$$

where

$$\tilde{X}(t) = \tilde{A}(t) - \mathbf{C}\tilde{S}(\bar{T}^*(t)) + \theta t, \quad t \geq 0, \quad (9.38)$$

is an \mathbf{I} -dimensional Brownian motion with drift θ and a diagonal covariance matrix whose i^{th} diagonal entry is $\lambda_i a_i^2 + \sum_{j=1}^{\mathbf{J}} \mathbf{C}_{ij} \mu_j b_j^2 x_j^*$.

Rearranging (9.35), we have that for $t \geq 0$,

$$y_{i^*}^r \hat{Q}_{i^*}^{r,*}(t) = \hat{\zeta}^{r,*}(t) + z_{k^*}^* \hat{I}_{k^*}^{r,*}(t), \quad (9.39)$$

where

$$\hat{\zeta}^{r,*}(\cdot) \equiv y^r \cdot \hat{X}^{r,*}(\cdot) - \sum_{i \in \mathcal{I} \setminus \{i^*\}} y_i^r \hat{Q}_i^{r,*}(\cdot) + \sum_{k \in \mathcal{K} \setminus \{k^*\}} z_k^r \hat{I}_k^{r,*}(\cdot) \Rightarrow y^* \cdot \tilde{X}(\cdot), \quad (9.40)$$

as $r \rightarrow \infty$, by (9.37), Lemma 9.1.1, and Theorem 7.3.1. Here $y^r > 0$, $z^r > 0$, $\hat{Q}_{i^*}^{r,*} \geq 0$, $\hat{\zeta}^{r,*}(0) = 0$, $\hat{I}_{k^*}^{r,*}(0) = 0$, $\hat{I}_{k^*}^{r,*}(\cdot)$ is continuous and non-decreasing. Furthermore, from the definition of the threshold policy in Chapter 6, for sufficiently large r , $\hat{I}_{k^*}^{r,*}$ can increase only if $\hat{Q}_{i^*}^{r,*}$ is at or below level $L_{i^*}^r/r$ (since $L_{i^*}^r > |\underline{\mathcal{J}}_{i^*}|$, for sufficiently large r).

By Theorem 4.1 in [37], since $L_{i^*}^r/r \rightarrow 0$ as $r \rightarrow \infty$, we then have that

$$(\hat{Q}_{i^*}^{r,*}, \hat{I}_{k^*}^{r,*}) \Rightarrow (\tilde{Q}_{i^*}^*, \tilde{I}_{k^*}^*), \quad \text{as } r \rightarrow \infty, \quad (9.41)$$

where $\tilde{Q}_{i^*}^*, \tilde{I}_{k^*}^*$ are given by (5.13)-(5.14). Combining this with Theorem 7.3.1 and Lemma 9.1.1 yields,

$$(\hat{Q}^{r,*}, \hat{I}^{r,*}, \hat{W}^{r,*}) \Rightarrow (\tilde{Q}^*, \tilde{I}^*, \tilde{W}^*), \quad \text{as } r \rightarrow \infty, \quad (9.42)$$

where \tilde{W}^* is defined in (5.11)-(5.12), and \tilde{Q}^*, \tilde{I}^* are defined by (5.13)-(5.14).

10

Asymptotic Optimality of the Threshold Policy

In this chapter we prove Theorem 6.2.2. We follow a similar development to that in Section 9 of [2]. Before proceeding with the proof, we first establish some preliminary results concerning fluid scaled processes under a sequence of scheduling control policies, $T = \{T^r\}$ (one for each member of the sequence of parallel server systems). The associated queue length and idletime processes will be denoted by Q^r, I^r , and the fluid and diffusion scaled versions of these processes will be denoted by \bar{Q}^r, \bar{I}^r and \hat{Q}^r, \hat{I}^r , respectively. We also let

$$\underline{J}(T) = \liminf_{r \rightarrow \infty} \hat{J}^r(T^r), \quad (10.1)$$

where $\hat{J}^r(T^r)$ is defined by (3.15). When our sequence of threshold policies $\{T^{r,*}\}$ is used, we append a superscript $*$ to the queue length, idletime etc. processes, i.e., we use $Q^{r,*}, I^{r,*}$, etc.

10.1 A Necessary Condition for Asymptotic Optimality

The next lemma implies that, when searching for an asymptotically optimal policy, we may restrict to those policies whose associated fluid scaled allocation processes converge (along a subsequence) to those given by \bar{T}^* .

Lemma 10.1.1 *Let $T = \{T^r\}$ be a sequence of scheduling control policies such that $\underline{J}(T) < \infty$. Consider a subsequence $\{T^{r'}\}$ of $\{T^r\}$ along which the \liminf in the definition of $\underline{J}(T)$ is achieved, i.e.,*

$$\lim_{r' \rightarrow \infty} \hat{J}^{r'}(T^{r'}) = \underline{J}(T). \quad (10.2)$$

Then,

$$(\bar{Q}^{r'}(\cdot), \bar{A}^{r'}(\cdot), \bar{S}^{r'}(\cdot), \bar{T}^{r'}(\cdot), \bar{I}^{r'}(\cdot)) \Rightarrow (\mathbf{0}, \lambda(\cdot), \mu(\cdot), \bar{T}^*(\cdot), \mathbf{0}) \quad \text{as } r' \rightarrow \infty, \quad (10.3)$$

where $\bar{T}^r(t) = T^r(r^2t)/r^2$ and $\bar{T}^*(t) = x^*t$, for all $r \geq 1$ and $t \geq 0$, x^* is given by the heavy traffic Assumption 3.2.2, $\mathbf{0}$ denotes the constant process that stays at the origin for all time, and $\lambda(t) = \lambda t$, $\mu(t) = \mu t$ for all $t \geq 0$.

Proof. It follows from Lemma 9.1.4 that

$$\left\{ (\bar{Q}^{r'}(\cdot), \bar{A}^{r'}(\cdot), \bar{S}^{r'}(\cdot), \bar{T}^{r'}(\cdot), \bar{I}^{r'}(\cdot)) \right\} \quad (10.4)$$

is C-tight. Thus, it suffices to show that all weak limit points of this sequence are given by the right member of (10.3). For this, suppose that

$$(\bar{Q}(\cdot), \bar{A}(\cdot), \bar{S}(\cdot), \bar{T}(\cdot), \bar{I}(\cdot)), \quad (10.5)$$

is obtained as a weak limit of (10.4) along a subsequence indexed by r'' . Without loss of generality, by appealing to the Skorokhod representation theorem (cf. [8], Theorem 3.1.8), we may choose an equivalent distributional representation (for which we use the same symbols) such that all of the random processes in (10.4)

indexed by r'' in place of r' , as well as the limit (10.5), are defined on the same probability space and the convergence in distribution is replaced by almost sure convergence uniformly on compact time intervals, so that a.s., as $r'' \rightarrow \infty$.

$$(\bar{Q}^{r''}(\cdot), \bar{A}^{r''}(\cdot), \bar{S}^{r''}(\cdot), \bar{T}^{r''}(\cdot), \bar{I}^{r''}(\cdot)) \rightarrow (\bar{Q}(\cdot), \bar{A}(\cdot), \bar{S}(\cdot), \bar{T}(\cdot), \bar{I}(\cdot)) \quad \text{u.o.c.} \quad (10.6)$$

From (9.19) we have that a.s., $\bar{A}(\cdot) = \lambda(\cdot)$ and $\bar{S}(\cdot) = \mu(\cdot)$. We next show that a.s., $\bar{Q}(\cdot) \equiv \mathbf{0}$. Combining the fact that $\lim_{r'' \rightarrow \infty} \hat{J}^{r''}(T^{r''}) = \underline{J}(T) < \infty$ with (10.6) and Fatou's lemma, we have

$$\begin{aligned} 0 = \lim_{r'' \rightarrow \infty} \frac{1}{r''} \hat{J}^{r''}(T^{r''}) &\geq \mathbf{E} \left(\int_0^\infty e^{-\gamma t} \liminf_{r'' \rightarrow \infty} (h \cdot \bar{Q}^{r''}(t)) dt \right) \\ &= \mathbf{E} \left(\int_0^\infty e^{-\gamma t} h \cdot \bar{Q}(t) dt \right). \end{aligned} \quad (10.7)$$

Since $h_i > 0$, for all $i \in \mathcal{I}$, and a.s., \bar{Q} has continuous paths in $\mathbb{R}_+^{\mathbf{I}}$, it follows from the above that a.s., $\bar{Q}(\cdot) \equiv \mathbf{0}$. Then, by letting $r = r'' \rightarrow \infty$ in (9.20)–(9.21) and using (9.19), (10.6), and the definition of \mathbf{R} , we have a.s., for each $t \geq 0$,

$$\mathbf{0} = \lambda t - \mathbf{R}\bar{T}(t), \quad (10.8)$$

$$\bar{I}(t) = \mathbf{1}t - \mathbf{A}\bar{T}(t). \quad (10.9)$$

Multiplying (10.8) by $(y^*)'$ while recalling that $y^* \cdot \lambda = 1 = z^* \cdot \mathbf{1}$ and $(y^*)'\mathbf{R} = (z^*)'\mathbf{A} - [0' \ (u^*)']$, where $u^* > 0$ (cf. Theorem 5.1.1 and Lemma 5.1.7), we obtain

$$0 = z^* \cdot \bar{I}(t) + [0' \ (u^*)']\bar{T}(t). \quad (10.10)$$

Since a.s., the components of $\bar{I}(\cdot)$ and $\bar{T}(\cdot)$ inherit the property from $\bar{I}^{r''}(\cdot)$, $\bar{T}^{r''}(\cdot)$ that they are all non-negative for all time, it follows from (10.10), and the fact that $z_k^* > 0$ for all $k \in \mathcal{K}$, $u_j^* > 0$ for all $j = 1, 2, \dots, \mathbf{J} - \mathbf{B}$, that a.s., $\bar{I}(\cdot) = \mathbf{0}$, and $\bar{T}_j(\cdot) = \mathbf{0}$ for all $j = \mathbf{B} + 1, \dots, \mathbf{J}$. We then observe that \bar{T} is a fluid control under which the fluid system in (3.5)–(3.6) is balanced and incurs no idleness (cf. Chapter 3). Hence, by Definition 3.2.1, Assumption 3.2.2, and Lemma 3.2.3 we have that $\bar{T}(\cdot) = \bar{T}^*(\cdot)$. \square

10.2 Proof of Asymptotic Optimality

We now prove that our threshold policy is asymptotically optimal.

Proof of Theorem 6.2.2. We first concentrate on proving the inequality on the left side of (6.2). For this, let $T \equiv \{T^r\}$ be a sequence of scheduling control policies. If $\underline{J}(T) = \infty$, then the inequality holds trivially and so we assume that $\underline{J}(T) < \infty$. Recall the definitions of (y^r, z^r) from (9.7)–(9.8). For each $j = \mathbf{B} + 1, \dots, \mathbf{J}$, let $u_{j-\mathbf{B}}^r = ((z^r)' \mathbf{A} - (y^r)' \mathbf{R}^r)_j$. Then by (9.10), we have that

$$(y^r)' \mathbf{R}^r = (z^r)' \mathbf{A} - [0' (u^r)'], \quad (10.11)$$

where for sufficiently large r , $u^r > 0$ by Lemma 5.1.7, (9.6), and Lemma 9.1.1.

For each $i \in \mathcal{I}$, let

$$h_i^r = \frac{h_i y_i^r}{y_i^*}, \quad (10.12)$$

where h is defined in the paragraph following (3.15). Since i^* is, by assumption, the “cheapest” buffer (cf. Chapter 6), we have that

$$\frac{h_{i^*}}{y_{i^*}^*} \leq \frac{h_i}{y_i^*}, \quad \text{for all } i \in \mathcal{I}. \quad (10.13)$$

Then, using (10.11), (3.19)–(3.20) and (10.12)–(10.13), we have for all $t \geq 0$,

$$\begin{aligned} h^r \cdot \hat{Q}^r(t) &= \sum_{i=1}^{\mathbf{I}} h_i^r \hat{Q}_i^r(t) \\ &\geq \frac{h_{i^*}}{y_{i^*}^*} \sum_{i=1}^{\mathbf{I}} y_i^r \hat{Q}_i^r(t) \\ &= \frac{h_{i^*}}{y_{i^*}^*} y^r \cdot \hat{Q}^r(t) \\ &= \frac{h_{i^*}}{y_{i^*}^*} \left(y^r \cdot \hat{X}^r(t) + \hat{V}^r(t) \right), \end{aligned} \quad (10.14)$$

where \hat{X}^r is given in (4.2),

$$\hat{V}^r(t) = ((z^r)' \mathbf{A} - [0' (u^r)'])' \hat{Y}^r(t) = z^r \cdot \hat{I}^r(t) - u^r \cdot \hat{Y}_N^r(t), \quad t \geq 0, \quad (10.15)$$

where \hat{Y}^r is defined by (3.17) and \hat{Y}_N^r is the $(\mathbf{J} - \mathbf{B})$ -dimensional process whose components are \hat{Y}_j^r , $j = \mathbf{B} + 1, \dots, \mathbf{J}$.

Now, since $h^r \cdot \hat{Q}^r(t) \geq 0$ for all $t \geq 0$, $y^r \cdot \hat{X}^r$ starts from zero, and \hat{V}^r is non-decreasing (for sufficiently large r) and starts from zero, it follows from the well known minimality of the solution of the Skorokhod problem (cf. e.g. Appendix B in [2]) that for all r sufficiently large,

$$\hat{V}^r(t) \geq \sup_{0 \leq s \leq t} \left(-y^r \cdot \hat{X}^r(s) \right) \quad \text{for all } t \geq 0, \quad (10.16)$$

and hence

$$h^r \cdot \hat{Q}^r(t) \geq \frac{h_{i^*}}{y_{i^*}^*} \varphi \left(y^r \cdot \hat{X}^r \right) (t) \quad \text{for all } t \geq 0, \quad (10.17)$$

where $\varphi(x)(t) \equiv x(t) + \sup_{0 \leq s \leq t} (-x(s))$ for all $t \geq 0$ and $x \in \mathbf{D}$ satisfying $x(0) = 0$.

Now, let $\{T^{r'}\}$ be a subsequence of $\{T^r\}$ such that $\lim_{r' \rightarrow \infty} \hat{J}^{r'}(T^{r'}) = \underline{J}(T)$. By Lemma 10.1.1, the fact that the limit there is deterministic, and (3.23), we have that as $r' \rightarrow \infty$,

$$\left(\hat{A}^{r'}(\cdot), \hat{S}^{r'}(\cdot), \bar{T}^{r'}(\cdot) \right) \Rightarrow \left(\tilde{A}(\cdot), \tilde{S}(\cdot), \bar{T}^*(\cdot) \right). \quad (10.18)$$

By invoking the Skorokhod representation theorem, we may assume without loss of generality that the convergence above is almost surely uniform on compact time intervals (u.o.c.) and then for \hat{X}^r given by (4.2), using Assumption 3.2.5, we have that a.s. as $r' \rightarrow \infty$,

$$\left(\hat{A}^{r'}(\cdot), \hat{S}^{r'}(\cdot), \bar{T}^{r'}(\cdot), \hat{X}^{r'}(\cdot) \right) \rightarrow \left(\tilde{A}(\cdot), \tilde{S}(\cdot), \bar{T}^*(\cdot), \tilde{X}(\cdot) \right) \quad \text{u.o.c.}, \quad (10.19)$$

where $\tilde{X}(t) = \tilde{A}(t) - \mathbf{C}\tilde{S}(\bar{T}^*(t)) + \theta t$, for $t \geq 0$, defines a Brownian motion as described in Definition 4.1.1. Then by Fatou's lemma, we have

$$\underline{J}(T) = \lim_{r' \rightarrow \infty} \hat{J}^{r'}(T^{r'}) \geq \mathbf{E} \left(\int_0^\infty e^{-\gamma t} \liminf_{r' \rightarrow \infty} (h \cdot \hat{Q}^{r'}(t)) dt \right). \quad (10.20)$$

Now we claim that a.s., for each $t \geq 0$,

$$\liminf_{r' \rightarrow \infty} (h \cdot \hat{Q}^{r'}(t)) \geq h \cdot \tilde{Q}^*(t), \quad (10.21)$$

where \tilde{Q}^* is given by (5.11)–(5.14). To see this, fix $\omega \in \Omega$ such that ω is in the set of probability one where the convergence in (10.19) holds u.o.c., and the limits have continuous paths. Fix $t \geq 0$. If the left member of the inequality (10.21) is infinite at ω , then the inequality clearly holds. On the other hand, if the left member is finite at ω , then there is a further subsequence indexed by r'' (possibly depending on ω and t) such that

$$\lim_{r'' \rightarrow \infty} (h \cdot \hat{Q}^{r''}(t, \omega)) = \liminf_{r' \rightarrow \infty} (h \cdot \hat{Q}^{r'}(t, \omega)) < \infty. \quad (10.22)$$

Since $h_i > 0$ and $\hat{Q}_i^{r''}(t, \omega) \geq 0$, for all $i \in \mathcal{I}$, it follows that $\hat{Q}_i^{r''}(t, \omega)$ is bounded as $r'' \rightarrow \infty$, for all $i \in \mathcal{I}$, and then using the fact that $h_i^r \rightarrow h_i$, for all $i \in \mathcal{I}$ (cf. Lemma 9.1.1), we have

$$\lim_{r'' \rightarrow \infty} (h - h^{r''}) \cdot \hat{Q}^{r''}(t, \omega) = 0. \quad (10.23)$$

Then, using (10.17), (10.19), the continuity of φ on D , the fact that $\varphi(y^* \cdot \tilde{X})(\cdot, \omega)$ is continuous, and (5.12)–(5.13), we have

$$\begin{aligned} \lim_{r'' \rightarrow \infty} h \cdot \hat{Q}^{r''}(t, \omega) &= \lim_{r'' \rightarrow \infty} \left(h^{r''} \cdot \hat{Q}^{r''}(t, \omega) + (h - h^{r''}) \cdot \hat{Q}^{r''}(t, \omega) \right) \\ &\geq \liminf_{r'' \rightarrow \infty} \frac{h_{i^*}}{y_{i^*}^*} \varphi \left(y^{r''} \cdot \tilde{X}^{r''} \right) (t, \omega) \\ &= \frac{h_{i^*}}{y_{i^*}^*} \varphi \left(y^* \cdot \tilde{X} \right) (t, \omega) = \frac{h_{i^*}}{y_{i^*}^*} \tilde{W}^*(t, \omega) = h \cdot \tilde{Q}^*(t, \omega). \end{aligned}$$

Thus, (10.21) holds. Now, substituting this in (10.20), we conclude that

$$\underline{J}(T) \geq \mathbf{E} \left(\int_0^\infty e^{-\gamma t} h \cdot \tilde{Q}^*(t) dt \right) \equiv J^*. \quad (10.24)$$

This completes the proof of the inequality in the left side of (6.2).

Suppose now that the threshold policy $T^{r,*}$ is used in the r^{th} parallel server system. For the purpose of establishing the finiteness of J^* and the equality in the right side of (6.2), by appealing to Theorem 6.2.1 and the Skorokhod representation theorem, we may assume that a.s.,

$$\hat{Q}^{r,*} \rightarrow \tilde{Q}^* \quad \text{u.o.c. as } r \rightarrow \infty, \quad (10.25)$$

where \tilde{Q}^* is given by (5.11)–(5.13). Then for

$$\hat{H}^{r,*} \equiv h \cdot \hat{Q}^{r,*} \quad \text{and} \quad \tilde{H}^* \equiv h \cdot \tilde{Q}^* \quad (10.26)$$

we have

$$\hat{H}^{r,*} \rightarrow \tilde{H}^* \quad (m \times \mathbf{P})\text{-a.e. on } \mathbb{R}_+ \times \Omega, \quad (10.27)$$

where $dm = \gamma e^{-\gamma t} dt$ on $(\mathbb{R}_+, \mathcal{B}_+)$ and \mathcal{B}_+ denotes the Borel σ -algebra on \mathbb{R}_+ . Then, since $(\mathbb{R}_+ \times \Omega, \mathcal{B}_+ \times \mathcal{F}, m \times \mathbf{P})$ is a probability space, to establish

$$\hat{J}^r(T^{r,*}) \equiv \mathbf{E} \left(\int_0^\infty e^{-\gamma t} \hat{H}^{r,*}(t) dt \right) \rightarrow J^* < \infty \text{ as } r \rightarrow \infty, \quad (10.28)$$

it suffices to show that

$$\limsup_{r \rightarrow \infty} \mathbf{E} \left(\int_0^\infty e^{-\gamma t} (\hat{H}^{r,*}(t))^2 dt \right) < \infty \quad (10.29)$$

which implies the required uniform integrability. From (9.35) we have

$$\hat{H}^{r,*} = h \cdot \hat{Q}^{r,*} \leq \left(\sum_{i \in \mathcal{I}} \frac{h_i}{y_i^r} \right) \hat{W}^{r,*} \quad (10.30)$$

where

$$\hat{W}^{r,*} = y^r \cdot \hat{Q}^{r,*} = y^r \cdot \hat{X}^{r,*} + \hat{V}^{r,*} \quad (10.31)$$

and for each $t \geq 0$,

$$\hat{X}^{r,*}(t) = \hat{A}^r(t) - \mathbf{C} \hat{S}^r(\bar{T}^{r,*}(t)) + r(\lambda^r - \mathbf{R}^r x^*)t, \quad (10.32)$$

$$\hat{V}^{r,*}(t) = z^r \cdot \hat{I}^{r,*}(t). \quad (10.33)$$

For each $r \geq 1$ and $t > 0$, let

$$G^{r,t} = \left\{ Q_i^{r,*}(s) \leq 2L_i^r \text{ for all } s \in [0, r^2 t], i \in \mathcal{I} \setminus \underline{\mathcal{I}}_{k^*} \right\}. \quad (10.34)$$

By the definition of $\bar{T}^{r,*}$, on $G^{r,t}$ we have that (for r large enough so that $L_i^r \geq |\underline{\mathcal{I}}_i|$ for all $i \in \underline{\mathcal{I}}_{k^*}$), $\hat{I}_{k^*}^{r,*}$ can have a point of increase at $s \in [0, t]$ only if

$$\hat{W}^{r,*}(s) \leq \left(\sum_{i \in \mathcal{I} \setminus \underline{\mathcal{I}}_{k^*}} 2y_i^r L_i^r + \sum_{i \in \underline{\mathcal{I}}_{k^*}} y_i^r L_i^r \right) / r, \quad (10.35)$$

i.e., for which $\hat{Q}_i^r(s)$ is at or below the level L_i^r/r for all $i \in \underline{\mathcal{I}}_{k^*}$ (this applies to both the case when $i \in \underline{\mathcal{I}}_{k^*}$ is a transition class as well as when i is a non-transition class, since for a non-transition class i , server k^* is busy whenever buffer i is nonempty, in particular, whenever $\hat{Q}_i^r > L_i^r/r > 0$).

Thus, on $G^{r,t}$, it follows from an oscillation inequality for solutions of a perturbed Skorokhod problem (cf. the proof of Theorem 5.1 in [37]) that

$$\begin{aligned}
z_{k^*}^r \hat{I}_{k^*}^{r,*}(t) &\leq \sup_{0 \leq s \leq t} \left(-y^r \cdot \hat{X}^{r,*}(s) - \sum_{k \neq k^*} z_k^r \hat{I}_k^{r,*}(s) \right) \\
&\quad + \left(\sum_{i \in \mathcal{I} \setminus \underline{\mathcal{I}}_{k^*}} 2y_i^r L_i^r + \sum_{i \in \underline{\mathcal{I}}_{k^*}} y_i^r L_i^r \right) r^{-1} \\
&\leq \sup_{0 \leq s \leq t} |y^r \cdot \hat{X}^{r,*}(s)| + \sum_{k \neq k^*} z_k^r \hat{I}_k^{r,*}(t) \\
&\quad + \left(\sum_{i \in \mathcal{I} \setminus \underline{\mathcal{I}}_{k^*}} 2y_i^r L_i^r + \sum_{i \in \underline{\mathcal{I}}_{k^*}} y_i^r L_i^r \right) r^{-1}, \tag{10.36}
\end{aligned}$$

where we have used the fact that $\hat{I}_k^{r,*}$, $k \in \mathcal{K} \setminus \{k^*\}$, is non-decreasing, to obtain the last inequality.

Since $I_k^r(r^2t) \leq r^2t$ for all $k \in \mathcal{K}$, we have for t satisfying $r^2t < M^r$,

$$\mathbf{E} \left((\hat{I}_k^{r,*}(t))^2 \right) \leq \left(\frac{M^r}{r} \right)^2 \text{ for all } k \in \mathcal{K}. \tag{10.37}$$

On the other hand, for $r \geq r^*$, for each $t > 0$ satisfying $r^2t \geq M^r$,

$$\begin{aligned}
&\mathbf{E} \left((\hat{I}_{k^*}^{r,*}(t))^2; \Omega \setminus G^{r,t} \right) \\
&\leq r^2t^2 \mathbf{P} \left(\Omega \setminus G^{r,t} \right) \\
&\leq r^2t^2 \sum_{i \in \mathcal{I} \setminus \underline{\mathcal{I}}_{k^*}} \mathbf{P} \left(Q_i^{r,*}(s) > 2L_i^r \text{ for some } s \in [0, r^2t] \right) \\
&\leq r^2t^2 \sum_{i \in \mathcal{I} \setminus \underline{\mathcal{I}}_{k^*}} \mathbf{P} \left(\sup_{\tau_{i,0}^r \leq s \leq r^2t} |R_i^r(s)| \geq L_i^r - |\underline{\mathcal{J}}_i| \right) \\
&\leq r^2t^2 \sum_{i \in \mathcal{I} \setminus \underline{\mathcal{I}}_{k^*}} p_{1,i}(r^2t) \left(C_{1,i}^{(1)} \exp(-C_{1,i}^{(2)} L_0^r) + C_{1,i}^{(3)} \exp(-C_{3,i}^{(4)} r^2t) \right) \\
&\leq t p(r^2t) \left(C^{(1)} \exp(-C^{(2)} L_0^r) + C^{(3)} \exp(-C^{(4)} r^2t) \right), \tag{10.38}
\end{aligned}$$

by Theorem 8.3.5, where the $p_{1,i}$, $C_{1,i}^{(m)}$, $m = 1, 2, 3, 4$, are as in (I.1), $p(s) = s \sum_{i \in \mathcal{I} \setminus \underline{\mathcal{I}}_{k^*}} p_{1,i}(s)$, for all $s \geq 0$, $C^{(1)} = \max\{C_{1,i}^{(1)} : i \in \mathcal{I} \setminus \underline{\mathcal{I}}_{k^*}\}$, $C^{(3)} = \max\{C_{1,i}^{(3)} : i \in \mathcal{I} \setminus \underline{\mathcal{I}}_{k^*}\}$, $C^{(2)} = \min\{C_{1,i}^{(2)} : i \in \mathcal{I} \setminus \underline{\mathcal{I}}_{k^*}\}$, $C^{(4)} = \min\{C_{1,i}^{(4)} : i \in \mathcal{I} \setminus \underline{\mathcal{I}}_{k^*}\}$. Hence p is a polynomial (of degree at most $\mathbf{I} + 1$) with non-negative coefficients and $C^{(m)} > 0$, $m = 1, 2, 3, 4$; the polynomial and constants do not depend on t or r .

By (9.17) and the fact that L_i^r is of order $\log r$ for all $i \in \mathcal{I}$, there is $r' \geq r^*$, such that for all $r \geq r'$, $y^r \cdot (\lambda^r - \mathbf{R}^r x^*)r \leq y^* \cdot |\theta| + 1$ and $\left(\sum_{i \in \mathcal{I} \setminus \underline{\mathcal{I}}_{k^*}} 2y_i^r L_i^r + \sum_{i \in \underline{\mathcal{I}}_{k^*}} y_i^r L_i^r\right)r^{-1} \leq 1$. Then we have for $r \geq r'$, and $t > 0$ satisfying $r^2 t \geq M^r$, using the inequality $\left(\sum_{i=1}^n x_i\right)^2 \leq n \sum_{i=1}^n x_i^2$ repeatedly, that

$$\begin{aligned}
& \mathbf{E}\left(\left(\hat{I}_{k^*}^{r,*}(t)\right)^2\right) \\
&= \mathbf{E}\left(\left(\hat{I}_{k^*}^{r,*}(t)\right)^2; G^{r,t}\right) + \mathbf{E}\left(\left(\hat{I}_{k^*}^{r,*}(t)\right)^2; \Omega \setminus G^{r,t}\right) \\
&\leq \frac{5}{(z_{k^*}^r)^2} \left\{ \mathbf{I} \sum_{i \in \mathcal{I}} (y_i^r)^2 \mathbf{E}\left(\sup_{0 \leq s \leq t} \left(\hat{A}_i^r(s)\right)^2\right) \right. \\
&\quad \left. + \mathbf{J} \sum_{j \in \mathcal{J}} (y_{i(j)}^r)^2 \mathbf{E}\left(\sup_{0 \leq s \leq t} \left(\hat{S}_j^r(T_j^{r,*}(s))\right)^2\right) + ((y^* \cdot |\theta| + 1)t)^2 \right. \\
&\quad \left. + (\mathbf{K} - 1) \sum_{k \neq k^*} (z_k^r)^2 \mathbf{E}\left(\left(\hat{I}_k^{r,*}(t)\right)^2\right) + 1 \right\} \\
&\quad + t p(r^2 t) \left(C^{(1)} \exp(-C^{(2)} L_0^r) + C^{(3)} \exp(-C^{(4)} r^2 t)\right), \quad (10.39)
\end{aligned}$$

where the first term to the right of the equality sign is controlled via (10.36), and the second term is controlled via (10.38). Using the fact that $\exp(-C^{(2)} L_0^r) \leq r^{-C^{(2)}c}$ (since $L_0^r = \lceil c \log r \rceil$), and the fact that for any polynomial q and $d > 0$, $q(x)e^{-dx}$ is bounded for $x \in \mathbb{R}_+$, we have that there is a constant $c_2 \geq c_1$ (independent of t and r) and $r'' \geq r'$ such that for each fixed $c \geq c_2$ and all $r \geq r''$, the last term in (10.39) is bounded by a polynomial in t , independent of r .

Then, using (10.30)–(10.33), (10.37), (10.39), and the fact that $M^r = o(r)$, for c sufficiently large (chosen independently of t and r), we see that to prove (10.29) it suffices to show that, as functions of t , the following are all in a bounded subset

of $L^1(m) \equiv L^1(\mathbb{R}_+, \mathcal{B}_+, m)$ for r sufficiently large:

$$\mathbf{E} \left(\sup_{0 \leq s \leq t} \left(\hat{A}_i^r(s) \right)^2 \right), \quad \mathbf{E} \left(\sup_{0 \leq s \leq t} \left(\hat{S}_j^r(\bar{T}_j^{r,*}(s)) \right)^2 \right), \quad \mathbf{E} \left(\left(\hat{I}_k^{r,*}(t) \right)^2 \right), \quad (10.40)$$

for all $i \in \mathcal{I}$, $j \in \mathcal{J}$, $k \in \mathcal{K} \setminus \{k^*\}$. The bounds for the first two expectations can be obtained as in Section 9 in [2]. For the third bound, fix $k \in \mathcal{K} \setminus \{k^*\}$, and let i be the transition buffer that is immediately above k . Let $\epsilon > 0$ and choose $r''' \geq r''$ such that for all $i \in \mathcal{I}$ and $r \geq r'''$, $r\epsilon \geq t_i^r$ (cf. (8.8)). Fix $r \geq r'''$. By (10.37), we need only consider $t > 0$ satisfying $r^2 t \geq M^r$. Since $I_k^{r,*}(r^2 t) \leq r^2 t$,

$$\begin{aligned} & \mathbf{E} \left(\left(\hat{I}_k^{r,*}(t) \right)^2 \right) \\ &= \int_0^\infty \mathbf{P} \left(\left(\hat{I}_k^{r,*}(t) \right)^2 > s \right) ds \\ &= \int_0^{r^2 t^2} \mathbf{P} \left(I_k^{r,*}(r^2 t) > r\sqrt{s} \right) ds \\ &\leq \int_0^{r^2 t^2} \left\{ \mathbf{P} \left(I_k^{r,*}(\tau_{i,0}^r) > r\sqrt{s} \right) \right. \\ &\quad \left. + \mathbf{P} \left(\inf_{\tau_{i,0}^r \leq u \leq r^2 t} R_i^r(u) \leq -L_i^r + |\underline{\mathcal{J}}_i| \right) \right\} ds \\ &\leq \epsilon^2 + \int_{\epsilon^2}^{r^2 t^2} \mathbf{P} \left(I_k^{r,*}(\tau_{i,0}^r) > r\epsilon \right) ds \\ &\quad + r^2 t^2 \mathbf{P} \left(\inf_{\tau_{i,0}^r \leq u \leq r^2 t} R_i^r(u) \leq -L_i^r + |\underline{\mathcal{J}}_i| \right) \\ &\leq \epsilon^2 + t \left\{ r^2 t \mathbf{P} \left(I_k^{r,*}(\tau_{i,0}^r) \geq t_i^r \right) \right. \\ &\quad \left. + r^2 t \mathbf{P} \left(\inf_{\tau_{i,0}^r \leq u \leq r^2 t} R_i^r(u) \leq -L_i^r + |\underline{\mathcal{J}}_i| \right) \right\} \\ &\leq \epsilon^2 + t \left\{ r^2 t p_{2,i}(r^2 t) \left(C_{2,i}^{(1)} \exp(-C_{2,i}^{(2)} L_0^r) \right. \right. \\ &\quad \left. \left. + C_{2,i}^{(3)} \exp(-C_{2,i}^{(4)} r^2 t) \right) \right. \\ &\quad \left. + r^2 t p_{1,i}(r^2 t) \left(C_{1,i}^{(1)} \exp(-C_{1,i}^{(2)} L_0^r) \right. \right. \\ &\quad \left. \left. + C_{1,i}^{(3)} \exp(-C_{1,i}^{(4)} r^2 t) \right) \right\}, \quad (10.41) \end{aligned}$$

by (I.1)–(I.2) which hold for all $i \neq i^*$ by Theorem 8.3.5, or by (III.1)–(III.2) (if $k \in \underline{\mathcal{K}}_{i^*}$) which also hold by Theorem 8.3.5. In the first inequality in (10.41), we

have used that fact that $I_k^{r,*}(r^2t) - I_k^{r,*}(\tau_{i,0}^r) = 0$ if $\inf_{\tau_{i,0}^r \leq u \leq r^2t} R_i^r(u) > -L_i^r + |\underline{\mathcal{J}}_i|$.

Since $\exp(-C_{m,i}^{(2)}L_0^r) \leq r^{-C_{m,i}^{(2)}c}$, $m = 1, 2$, and the fact that for any polynomial q and $d > 0$, $q(x)e^{-dx}$ is bounded for $x \in \mathbb{R}_+$, it follows that there is a constant $c_3 \geq c_2$ (independent of t and r), $r^{**} \geq r'''$, such that for each fixed $c \geq c_3$, and all $r \geq r^{**}$, the right member above can be bounded by a polynomial in t (not depending on r). Combining (10.37) with the above, we conclude that $\{\mathbf{E}((\hat{I}_k^{r,*}(\cdot))^2)\}$ is a bounded sequence of functions in $L^1(m)$, provided that c is fixed and sufficiently large. \square

Appendix A

Appendix: Proof of Multiparameter Stopping Time Property

Proof of Lemma 7.5.2. We show the case of $\mathcal{T}_{n,\iota}^r$ only, the case of ${}^d\mathcal{T}_{n,\iota}^r$ being similar. Let $\iota \in \mathcal{I}$. For the proof, we fix $n \geq 1$ and $r \geq r^*$ (so that $L_i \equiv L_i^r \geq |\underline{\mathcal{J}}_i| + 1$, for all $i \in \mathcal{I}$), and to simplify the notation, we suppress the superscript r . We first define a sequence of stopping times $\{\sigma_\ell\}_{\ell=0}^\infty$ that specifies the successive times at which an arrival occurs to, or a departure occurs from, any class $i \in \mathcal{I}$. This allows us to give a discrete-event-type description of the dynamics of our parallel server system. Let $\sigma_0 = 0$, and σ_ℓ ($\ell \geq 1$) be the time of the ℓ^{th} change in the number of arrivals to, or the number of departures from, any class $i \in \mathcal{I}$. It can be readily shown that $\sigma_\ell < \infty$ a.s. for each ℓ , and that $\sigma_\ell \rightarrow \infty$ a.s. as $\ell \rightarrow \infty$. Since $\{\mathcal{F}_{pq} : (p, q) \in \mathbb{N}_\infty^{\mathbf{I}} \times \mathbb{N}_\infty^{\mathbf{J}}\}$ contains all of the \mathbf{P} -null sets and the underlying probability space is complete, for the proof of the stopping time property, we may ignore the exceptional \mathbf{P} -null sets on which the aforementioned properties do not hold. Thus, without loss of generality, for this proof, we assume that for each $\omega \in \Omega$, $\sigma_\ell(\omega) < \infty$ for all ℓ , and $\sigma_\ell(\omega) \rightarrow \infty$ as $\ell \rightarrow \infty$. For $k \in \mathcal{K}$, we define $\underline{\mathcal{I}}_k^{\mathbf{T}} = \{i \in \underline{\mathcal{I}}_k : i \text{ is a transition class}\}$ and

$\underline{\mathcal{I}}_k^{\text{NT}} = \{i \in \underline{\mathcal{I}}_k : i \text{ is a non-transition class}\}$.

For each ℓ , we define the following auxiliary variables below: $\mathcal{C}^\ell, Q_i^\ell, \mathcal{S}_j^\ell, u_i^\ell, v_j^\ell$, $i \in \mathcal{I}$, $j \in \mathcal{J}$, and for convenience, we let $\chi^\ell = (\sigma_\ell, \mathcal{C}^\ell, Q_i^\ell, \mathcal{S}_j^\ell, u_i^\ell, v_j^\ell : i \in \mathcal{I}, j \in \mathcal{J})$. For $\ell = 0$, define $\mathcal{C}^0 = 0$, $Q_i^0 = Q_i(0) = 0$, $\mathcal{S}_j^0 = 0$, $u_i^0 = u_i(1)$, $v_j^0 = v_j(1)$, $i \in \mathcal{I}$, $j \in \mathcal{J}$. Now, consider $\ell \geq 1$. Let $Q_i^\ell = Q_i(\sigma_\ell)$, $i \in \mathcal{I}$. For each $i \in \mathcal{I}$, to keep track of whether there is a class i job in process or in suspension for activity $j \in \mathcal{J}_i$, we define the (indicator) random variables \mathcal{S}_j^ℓ , $j \in \mathcal{J}_i$, as follows. For $j \in \mathcal{J}_i$, let $\mathcal{S}_j^\ell = 1$ if there is a class i job (in process or in suspension) at server $k(j)$ at the time σ_ℓ , otherwise, let $\mathcal{S}_j^\ell = 0$. Since \mathcal{J} is the disjoint union of \mathcal{J}_i over $i \in \mathcal{I}$, this defines \mathcal{S}_j^ℓ for all $j \in \mathcal{J}$. (Note that $\sum_{j \in \mathcal{J}_i} \mathcal{S}_j^\ell$ is a lower bound on the number of class i jobs in the system at the time σ_ℓ .) To count the number of times that Q_ℓ hits $L_\ell + 1$ from below, define $\mathcal{C}^\ell = \sup\{m \geq 1 : \tau_{\ell, 2m-1} \leq \sigma_\ell\}$. Let u_i^ℓ denote the residual interarrival time for class i , $i \in \mathcal{I}$, as measured from σ_ℓ , i.e., u_i^ℓ is the amount of time remaining after σ_ℓ until the next class i arrival with the convention that if σ_ℓ is the time of a class i arrival then $u_i^\ell = u_i(A_i(\sigma_\ell) + 1)$.

For $j \in \underline{\mathcal{J}}_i$, if $Q_{i'}^{\ell-1} = 0$ for all $i' \in \underline{\mathcal{I}}_{k(j)}^{\text{NT}}$, and $Q_{i'}^{\ell-1} \leq L_{i'}$, for all $i' \in \underline{\mathcal{I}}_{k(j)}^{\text{T}}$, and if in addition $\mathcal{S}_j^{\ell-1} = 1$, let v_j^ℓ be the residual service time of the class i job being served by activity $j \in \underline{\mathcal{J}}_i$ at the time σ_ℓ with the convention that if a class i job completed service at server $k(j)$ at the time σ_ℓ , then $v_j^\ell = v_j(\mathcal{S}_j(T_j(\sigma_\ell)) + 1)$, the service time for the next class i job to be served at server $k(j)$. Otherwise, let $v_j^\ell = v_j^{\ell-1}$.

For $k \in \mathcal{K}$, $i \in \underline{\mathcal{I}}_k^{\text{T}}$, if $Q_i^{\ell-1} > L_i$ and $Q_{i'}^{\ell-1} \leq L_{i'}$ for all $i' \in \underline{\mathcal{I}}_k$, $i' < i$, let $v_{a(i)}^\ell$ be the residual service time of the class i job being served by server k at time σ_ℓ with the convention that if a class i job completed service at server k at the time σ_ℓ , then $v_{a(i)}^\ell = v_{a(i)}(S_{a(i)}(T_{a(i)}(\sigma_\ell)) + 1)$. If $Q_i^{\ell-1} \leq L_i$ or $Q_{i'}^{\ell-1} > L_{i'}$, some $i' \in \underline{\mathcal{I}}_k$, $i' < i$, server k does not allocate any service time to class i jobs in the interval $[\sigma_{\ell-1}, \sigma_\ell)$ and we let $v_{a(i)}^\ell = v_{a(i)}^{\ell-1}$. For $i \in \underline{\mathcal{I}}_k^{\text{NT}}$, if $Q_i^{\ell-1} > 0$, $Q_{i'}^{\ell-1} \leq L_{i'}$, for all $i' \in \underline{\mathcal{I}}_k$, and $Q_{i'}^{\ell-1} = 0$, for all $i' \in \underline{\mathcal{I}}_k^{\text{NT}}$, $i' < i$, let $v_{a(i)}^\ell$ be the residual service time of the class i job being served by server k at time σ_ℓ with the convention that if a class i job

completed service at server k at the time σ_ℓ , then $v_{a(i)}^\ell = v_{a(i)}(S_{a(i)}(T_{a(i)}(\sigma_\ell)) + 1)$. If $Q_i^{\ell-1} = 0$ or $Q_{i'}^{\ell-1} > L_{i'}$ for some $i' \in \underline{\mathcal{I}}_k^T$, or $Q_{i'}^{\ell-1} > 0$, for some $i' \in \underline{\mathcal{I}}_k^{\text{NT}}$, $i' < i$, server k does not allocate any service time to class i jobs in the interval $[\sigma_{\ell-1}, \sigma_\ell)$ and we let $v_{a(i)}^\ell = v_{a(i)}^{\ell-1}$. Here, we have divided the activities into those below buffers (above servers) and those below servers (above buffers).

For each $(p, q) \in \mathbb{N}^{\mathbf{I}} \times \mathbb{N}^{\mathbf{J}}$ let

$$B_{pq} \equiv \{A_i(\tau_{i,2n-1}) = p_i, S_j(T_j(\tau_{i,2n-1})) = q_j : i \in \mathcal{I}, j \in \mathcal{J}\}. \quad (\text{A.1})$$

By the convention concerning the removal of exceptional \mathbf{P} -null sets made above and the definition of A_i, S_j , $i \in \mathcal{I}$, $j \in \mathcal{J}$, at the time ∞ , we have

$$B_{\infty\infty} = \{\tau_{i,2n-1} = \infty\} = \Omega \setminus \{\tau_{i,2n-1} < \infty\} \quad (\text{A.2})$$

$$= \Omega \setminus \bigcup_{(p,q) \in \mathbb{N}^{\mathbf{I}} \times \mathbb{N}^{\mathbf{J}}} B_{pq} \quad (\text{A.3})$$

and so it suffices to show that $B_{pq} \in \mathcal{F}_{pq}$ for each $(p, q) \in \mathbb{N}^{\mathbf{I}} \times \mathbb{N}^{\mathbf{J}}$. Now for each $(p, q) \in \mathbb{N}^{\mathbf{I}} \times \mathbb{N}^{\mathbf{J}}$ and $\ell \geq 0$, let

$$B_{pq}^\ell \equiv \{A_i(\sigma_\ell) = p_i, S_j(T_j(\sigma_\ell)) = q_j : i \in \mathcal{I}, j \in \mathcal{J}\}. \quad (\text{A.4})$$

Now, when $\tau_{i,2n-1}$ is finite, it is equal to σ_ℓ for some ℓ (depending on ω), and so we have

$$\begin{aligned} B_{pq} &= \bigcup_{\ell=1}^{\infty} B_{pq}^\ell \cap \{\tau_{i,2n-1} = \sigma_\ell\} \\ &= \bigcup_{\ell=1}^{\mathbf{1}_I \cdot p + \mathbf{1}_J \cdot q} B_{pq}^\ell \cap \{Q_i^\ell = L_i + 1\} \cap \{Q_{i'}^{\ell-1} < L_{i'} + 1\} \cap \{\mathcal{C}^\ell = n\}, \end{aligned} \quad (\text{A.5})$$

where $\mathbf{1}_I$ and $\mathbf{1}_J$ are \mathbf{I} and \mathbf{J} -dimensional vectors of ones, respectively. We show by induction on ℓ that the following two properties hold for $\ell = 0, 1, \dots$, for each $(p, q) \in \mathbb{N}^{\mathbf{I}} \times \mathbb{N}^{\mathbf{J}}$:

- (i) $B_{pq}^\ell \in \mathcal{F}_{pq}$,
- (ii) $\mathbf{1}_{B_{pq}^\ell} \mathcal{H}^\ell \in \mathcal{F}_{pq}$,

where

$$\mathcal{H}^\ell = \{\chi^m, m = 0, \dots, \ell\}. \quad (\text{A.6})$$

(Note that B_{pq}^ℓ is a set whereas $1_{B_{pq}^\ell} \mathcal{H}^\ell$ is a random vector so that (ii) means that this vector is measurable with respect to the σ -algebra \mathcal{F}_{pq} .) It is easy to see that these properties (i)–(ii) and (A.5) imply that $B_{pq} \in \mathcal{F}_{pq}$ for each $(p, q) \in \mathbb{N}^{\mathbf{I}} \times \mathbb{N}^{\mathbf{J}}$.

The truth of (i)–(ii) for $\ell = 0$ and all $(p, q) \in \mathbb{N}^{\mathbf{I}} \times \mathbb{N}^{\mathbf{J}}$ is readily verified since $\sigma_0 = 0$, the arrival, service, allocation, and queue length processes all start from zero, and \mathcal{F}_{00} contains the initial interarrival and service times, $u_i(1), v_j(1), i \in \mathcal{I}, j \in \mathcal{J}$.

For the induction step, assume that for some $\ell \geq 0$, (i)–(ii) hold for all $(p, q) \in \mathbb{N}^{\mathbf{I}} \times \mathbb{N}^{\mathbf{J}}$. Now,

$$B_{pq}^{\ell+1} = \bigcup_{(n,m)} (B_{pq}^{\ell+1} \cap B_{nm}^\ell) \quad (\text{A.7})$$

where the union is over all $(n, m) \in \mathbb{N}^{\mathbf{I}} \times \mathbb{N}^{\mathbf{J}}$, such that $n_i \leq p_i, m_j \leq q_j$, for all $i \in \mathcal{I}, j \in \mathcal{J}$. By the induction assumption, for fixed $(p, q) \in \mathbb{N}^{\mathbf{I}} \times \mathbb{N}^{\mathbf{J}}$ and any $(n, m) \in \mathbb{N}^{\mathbf{I}} \times \mathbb{N}^{\mathbf{J}}$ such that $n_i \leq p_i, m_j \leq q_j, i \in \mathcal{I}, j \in \mathcal{J}$, we have

$$B_{nm}^\ell \in \mathcal{F}_{nm}, \quad 1_{B_{nm}^\ell} \mathcal{H}^\ell \in \mathcal{F}_{nm}. \quad (\text{A.8})$$

For $k \in \mathcal{K}, i \in \underline{\mathcal{I}}_k^{\mathbf{T}}, j \in \underline{\mathcal{J}}_i$, define the sets

$$\begin{aligned} \mathcal{O}_{\ell, a(i)} &= \{Q_i^\ell > L_i\} \cap \{Q_{i'}^\ell \leq L_{i'} \text{ for all } i' \in \underline{\mathcal{I}}_k : i' < i\}, \\ \mathcal{O}_{\ell, j} &= \{Q_{i'}^\ell \leq L_{i'} \text{ for all } i' \in \underline{\mathcal{I}}_{k(j)}^{\mathbf{T}}\} \cap \{Q_{i'}^\ell = 0 \text{ for all } i' \in \underline{\mathcal{I}}_{k(j)}^{\mathbf{NT}}\} \cap \{\mathcal{S}_j^\ell = 1\}. \end{aligned}$$

For $i \in \underline{\mathcal{I}}_k^{\mathbf{NT}}$, define

$$\begin{aligned} \mathcal{O}_{\ell, a(i)} &= \{Q_i^\ell > 0\} \cap \{Q_{i'}^\ell \leq L_{i'} \text{ for all } i' \in \underline{\mathcal{I}}_k^{\mathbf{T}}\} \cap \\ &\quad \{Q_{i'}^\ell = 0 \text{ for all } i' \in \underline{\mathcal{I}}_k^{\mathbf{NT}} : i' < i\}. \end{aligned}$$

Here \mathcal{O} is mnemonic for “occupied”. For example, the set $\mathcal{O}_{\ell, j}$, is the set on which, at the time σ_ℓ , server $k(j)$ is “occupied” by a class i job (busy serving a class i job) since, on this set, the number of jobs in each of the transition buffers in the layer

below buffer i served by server $k(j)$ is below its threshold, the number of jobs in each non-transition buffer served by server $k(j)$ is zero, and in addition either (a) the total number of class i jobs in the system is greater than the number of servers below buffer i or (b) the total number of class i jobs in the system, say m , is less than or equal to the number of servers below buffer i and at most $m - 1$ of these servers (other than server $k(j)$) have class i jobs that they are serving or have in suspension.

For each $i \in \mathcal{I}$, let $k_{i,1}, \dots, k_{i,|\mathcal{J}_i|}$ be a list of the servers $k(j) : j \in \mathcal{J}_i$ arranged in increasing order, and let $\{j_{i,1}, j_{i,2}, \dots, j_{i,|\mathcal{J}_i|}\}$ be the corresponding list of activities: $k(j_{i,m}) = k_{i,m}$, $m = 1, \dots, \mathcal{J}_i$. (This enumeration is done to simplify the treatment of the tie breaking rule.) For fixed $i \in \mathcal{I}$, $x \in \{0, 1\}^{|\mathcal{J}_i|}$, let

$$\mathcal{O}_{x,i}^\ell = \left(\bigcup_{k: x_k=1} \mathcal{O}_{\ell, j_{i,k}} \right) \cup \left(\bigcup_{k: x_k=0} \mathcal{O}_{\ell, j_{i,k}}^c \right), \quad v_{x,i}^\ell = \bigwedge_{k: x_k=1} v_{j_{i,k}}^\ell,$$

where we set $v_{\vec{0},i}^\ell = \infty$, and where $\vec{0} = x \in \{0, 1\}^{|\mathcal{J}_i|}$ such that $x_k = 0$ for $1 \leq k \leq |\mathcal{J}_i|$. The set $\mathcal{O}_{x,i}^\ell$ is the set on which the activities which serve buffer i at time σ_ℓ are given by the positive entries in x . For example, if $x = (0, 1, 1)$ then on $\mathcal{O}_{x,i}^\ell$ activities $j_{i,2}$ and $j_{i,3}$ are busy serving class i jobs at time σ_ℓ while activity $j_{i,1}$ is not active. In this case, $v_{x,i}^\ell$ is the minimum of the (residual) service times of the two class i jobs being served at the time σ_ℓ .

Then, since $\sigma_{\ell+1}$ is the first time after σ_ℓ that a new arrival or departure occurs, we have

$$\sigma_{\ell+1} = \sigma_\ell + \bigwedge_{i \in \mathcal{I}} \sum_{x \in \{0,1\}^{|\mathcal{J}_i|}} 1_{\mathcal{O}_{x,i}^\ell} (u_i^\ell \wedge v_{x,i}^\ell). \quad (\text{A.9})$$

For each $i \in \mathcal{I}$ and $\omega \in \Omega$, the sum in (A.9) collapses to a single term since $\omega \in \mathcal{O}_{x,i}^\ell$ for exactly one $x \in \{0, 1\}^{|\mathcal{J}_i|}$. The last term in (A.9) is the minimum (taken over all buffers i) of the (residual) interarrival time for buffer i and of all the (residual) service times of the activities which are serving class i jobs at time σ_ℓ .

Now, on B_{nm}^ℓ , for $i \in \mathcal{I}$, $j \in \mathcal{J}$,

$$\begin{aligned} A_i(\sigma_{\ell+1}) &= n_i + 1_{\{u_i^\ell = \sigma_{\ell+1} - \sigma_\ell\}}, \\ S_j(T_j(\sigma_{\ell+1})) &= m_j + 1_{\mathcal{O}_{\ell,j} \cap \{v_j^\ell = \sigma_{\ell+1} - \sigma_\ell\}}. \end{aligned}$$

Thus, on B_{nm}^ℓ , we can represent $A_i(\sigma_{\ell+1})$, $S_j(T_j(\sigma_{\ell+1}))$, $i \in \mathcal{I}$, $j \in \mathcal{J}$, as measurable functions of χ^ℓ , and so it follows from the induction assumption (A.8) that

$$B_{pq}^{\ell+1} \cap B_{nm}^\ell \in \mathcal{F}_{nm} \subset \mathcal{F}_{pq}. \quad (\text{A.10})$$

Then by (A.7), (i) holds with $\ell + 1$ in place of ℓ . It remains to verify (ii) with $\ell + 1$ in place of ℓ . By (i), (A.7), and (A.8), for this it suffices to show that

$$1_{B_{pq}^{\ell+1} \cap B_{nm}^\ell} \chi^{\ell+1} \in \mathcal{F}_{pq}. \quad (\text{A.11})$$

This follows from the induction assumption (A.8) and (A.10), and the following representations for the remaining variables contained in $\chi^{\ell+1}$:

$$\mathcal{C}^{\ell+1} = \mathcal{C}^\ell + 1_{\{Q_i^\ell = L_i\} \cap \{u_i^\ell = \sigma_{\ell+1} - \sigma_\ell\}} \left(\sum_{x \in \{0,1\}^{|\mathcal{J}_i|}} 1_{\mathcal{O}_{x,i}^\ell \cap \{u_i^\ell < v_{x,i}^\ell\}} \right). \quad (\text{A.12})$$

On $B_{pq}^{\ell+1}$, for $k \in \mathcal{K}$, $i \in \underline{\mathcal{I}}_k$,

$$\begin{aligned}
Q_i^{\ell+1} &= A_i(\sigma_{\ell+1}) - \sum_{j \in \mathcal{J}_i} S_j(T_j(\sigma_{\ell+1})) \\
&= p_i - \sum_{j \in \mathcal{J}_i} q_j, \\
u_i^{\ell+1} &= 1_{\{u_i^\ell > \sigma_{\ell+1} - \sigma_\ell\}}(u_i^\ell - (\sigma_{\ell+1} - \sigma_\ell)) + 1_{\{u_i^\ell = \sigma_{\ell+1} - \sigma_\ell\}} u_i(p_i + 1), \\
v_j^{\ell+1} &= 1_{\mathcal{O}_{\ell,j} \cap \{v_j^\ell > \sigma_{\ell+1} - \sigma_\ell\}}(v_j^\ell - (\sigma_{\ell+1} - \sigma_\ell)) \\
&\quad + 1_{\mathcal{O}_{\ell,j} \cap \{v_j^\ell = \sigma_{\ell+1} - \sigma_\ell\}} v_j(q_j + 1) + 1_{\mathcal{O}_{\ell,j}^c} v_j^\ell, \quad j \in \mathcal{J}_i, \\
\mathcal{S}_{a(i)}^{\ell+1} &= 1_{\mathcal{O}_{\ell+1,a(i)}} + 1_{\mathcal{O}_{\ell+1,a(i)}^c \cap \mathcal{O}_{\ell,a(i)} \cap \{v_{a(i)}^\ell > \sigma_{\ell+1} - \sigma_\ell\}} + 1_{\mathcal{O}_{\ell+1,a(i)}^c \cap \mathcal{O}_{\ell,a(i)}^c} \mathcal{S}_{a(i)}^\ell, \\
\mathcal{S}_j^{\ell+1} &= 1_{\{Q_i^{\ell+1} > |\underline{\mathcal{I}}_i|\} \cap \{Q_{i'}^{\ell+1} \leq L_{i'} \text{ for all } i' \in \underline{\mathcal{I}}_k^T(j)\} \cap \{Q_{i'}^{\ell+1} = 0 \text{ for all } i' \in \underline{\mathcal{I}}_k^{NT}(j)\}} \\
&\quad + 1_{\{Q_{i'}^{\ell+1} > L_{i'} \text{ for some } i' \in \underline{\mathcal{I}}_k^T(j)\} \cup \{Q_{i'}^{\ell+1} > 0 \text{ for some } i' \in \underline{\mathcal{I}}_k^{NT}(j)\}} \\
&\quad \cdot \left(1_{\mathcal{O}_{\ell,j} \cap \{v_j^\ell > \sigma_{\ell+1} - \sigma_\ell\}} + 1_{\mathcal{O}_{\ell,j}^c} \mathcal{S}_j^\ell \right) \\
&\quad + 1_{\{Q_i^{\ell+1} \leq |\underline{\mathcal{I}}_i|\} \cap \{Q_{i'}^{\ell+1} \leq L_{i'} \text{ for all } i' \in \underline{\mathcal{I}}_k^T(j)\} \cap \{Q_{i'}^{\ell+1} = 0 \text{ for all } i' \in \underline{\mathcal{I}}_k^{NT}(j)\}} \\
&\quad \cdot \left(1_{\mathcal{O}_{\ell,j} \cap \{v_j^\ell > \sigma_{\ell+1} - \sigma_\ell\}} + 1_{\mathcal{O}_{\ell,j}^c \cap \{\mathcal{S}_j^\ell = 1\}} \right) \\
&\quad + 1_{(\mathcal{O}_{\ell,j} \cap \{v_j^\ell = \sigma_{\ell+1} - \sigma_\ell\}) \cup (\mathcal{O}_{\ell,j}^c \cap \{\mathcal{S}_j^\ell = 0\})} \\
&\quad \cdot 1_{\left\{ Q_i^{\ell+1} > \sum_{m'=m+1}^{|\mathcal{J}_i|} \mathcal{S}_{j_{i,m'}}^\ell 1_{(\mathcal{O}_{\ell,j_{i,m'}} \cap \{v_{j_{i,m'}}^\ell > \sigma_{\ell+1} - \sigma_\ell\}) \cup \mathcal{O}_{\ell,j_{i,m'}}^c} + \sum_{m'=1}^{m-1} \mathcal{S}_{j_{i,m'}}^{\ell+1} \right\}} \Big),
\end{aligned}$$

for $j \in \underline{\mathcal{J}}_i$, where $m : j = j_{i,m}$. In the last term of the definition of $\mathcal{S}_j^{\ell+1}$, we have used our tie breaking rule from Chapter 6 which posits that lower numbered servers (and hence lower numbered activities by our enumeration of \mathcal{J}_i in this section) have higher priority when a job can be served by more than one server (activity). For $\mathcal{S}_j^{\ell+1}$, we use induction on m to show that $\mathcal{S}_{j_{i,m}}^{\ell+1}$ is an \mathcal{F}_{pq} measurable random variable for $m = 1, 2, \dots, |\mathcal{J}_i|$. (For $m = 1$, the result is trivial since the second sum in the last term of the definition of $\mathcal{S}_j^{\ell+1}$ is defined to be zero.) Suppose the result holds for all $m' < m$ and let $j = j_{i,m}$. Then the first sum in the last term of $\mathcal{S}_j^{\ell+1}$ is over those activities whose class i job assignments are made after those of activity j at the time $\sigma_{\ell+1}$ (i.e., they have lower priority than j). Such an activity contributes one to the sum if either it has a class i job in process at σ_ℓ

and that job remains with the activity (in process or in suspension) at $\sigma_{\ell+1}$, or it has a class i job in suspension at σ_{ℓ} , which will still remain with the activity at $\sigma_{\ell+1}$. The second sum is over those activities whose class i job assignment is made ahead of j (i.e., they have higher priority than j). Such an activity contributes one to the sum if there is a class i job in process or in suspension for the activity at $\sigma_{\ell+1}$. The random variables $\mathcal{S}_{j_i, m'}^{\ell+1}$, $m' < m$ appearing there are \mathcal{F}_{pq} measurable by the induction assumption, and it follows that $\mathcal{S}_j^{\ell+1}$ is \mathcal{F}_{pq} measurable.

For $\mathcal{S}_{a(i)}^{\ell+1}$, note that $\mathcal{O}_{\ell+1, j}$ depends on $Q_i^{\ell+1}$ which was determined earlier. \square

Bibliography

- [1] B. Ata and S. Kumar. Heavy traffic analysis of open processing networks with complete resource pooling: asymptotic optimality of discrete review policies. Submitted to *Annals of Applied Probability*.
- [2] S. L. Bell and R. J. Williams. Dynamic scheduling of a system with two parallel servers in heavy traffic with resource pooling: asymptotic optimality of a threshold policy. *Ann. Appl. Probab.*, **11** (2001), 608–649.
- [3] P. Billingsley. *Convergence of Probability Measures*. John Wiley and Sons, New York, 1968.
- [4] M. Bramson. State space collapse with applications to heavy traffic limits for multiclass queueing networks. *Queueing Systems*, **30** (1998), 89–148.
- [5] M. Bramson and R. J. Williams. Two workload properties for Brownian networks. To appear in *Queueing Systems*.
- [6] P. B. Chevalier and L. Wein. Scheduling networks of queues: heavy traffic analysis of a multistation closed network. *Oper. Res.* **41** (1993), 743–758.
- [7] K. L. Chung and R. J. Williams. *Introduction to Stochastic Integration*. 2nd edition, Birkhäuser, Boston, 1990.
- [8] S. N. Ethier and T. G. Kurtz. *Markov Processes: Characterization and Convergence*. Wiley, New York, 1986.
- [9] J. M. Harrison. *Brownian Motion and Stochastic Flow Systems*. John Wiley and Sons, New York, 1985.
- [10] J. M. Harrison. Brownian models of queueing networks with heterogeneous customer populations. In *Stochastic Differential Systems, Stochastic Control Theory and Their Applications, IMA Volume 10*, W. Fleming and P.L. Lions (eds.), Springer Verlag, New York, 1988, pp. 147-186.

- [11] J. M. Harrison. Balanced fluid models of multiclass queueing networks: a heavy traffic conjecture. In *Stochastic Networks*, F. P. Kelly and R. J. Williams (eds.), Springer-Verlag, 1995, pp. 1–20.
- [12] J. M. Harrison. The BIGSTEP approach to flow management in stochastic processing networks. In *Stochastic Networks: Theory and Applications*, F. P. Kelly, S. Zachary and I. Ziedins (eds.), Oxford University Press, 1996, pp. 57–90.
- [13] J. M. Harrison. Heavy traffic analysis of a system with parallel servers: asymptotic optimality of discrete-review policies. *Ann. Appl. Probab.*, **8** (1998), 822–848.
- [14] J. M. Harrison. Brownian models of open processing networks: canonical representation of workload. *Ann. Appl. Probab.*, **10** (2000), 75–103.
- [15] J. M. Harrison. A broader view of Brownian networks. To appear in *Annals of Applied Probability*.
- [16] J. M. Harrison and M. J. López. Heavy traffic resource pooling in parallel-server systems. *Queueing Systems*, **33** (1999), 339–368.
- [17] J. M. Harrison and J. A. Van Mieghem. Dynamic control of Brownian networks: state space collapse and equivalent workload formulations. *Ann. Appl. Probab.*, **7** (1997), 747–771.
- [18] J. M. Harrison and L. Wein. Scheduling networks of queues: heavy traffic analysis of a simple open network. *Queueing Systems*, **5** (1989), 265–280.
- [19] J. M. Harrison and L. Wein. Scheduling networks of queues: heavy traffic analysis of a two-station closed network. *Oper. Res.*, **38** (1990), 1052–1064.
- [20] D. L. Iglehart and W. Whitt. The equivalence of functional central limit theorems for counting processes and associated partial sums. *Ann. Math. Statist.*, **42** (1971), 1372–1378.
- [21] W. C. Jordan and C. Graves. Principles on the benefits of manufacturing process flexibility. *Management Science*, **41** (1995), 577–594.
- [22] F. P. Kelly and C. N. Laws. Dynamic routing in open queueing networks: Brownian models, cut constraints and resource pooling. *Queueing Systems*, **13** (1993), 47–86.

- [23] H. J. Kushner and Y. N. Chen. Optimal control of assignment of jobs to processors under heavy traffic. *Stochastics and Stochastics Rep.* **68** (2000), no. 3-4, pp. 177–228.
- [24] H. J. Kushner and P. Dupuis. *Numerical Methods for Stochastic Control Problems in Continuous Time*. Springer-Verlag, New York, 1992.
- [25] H. J. Kushner and L. F. Martins. Numerical methods for stochastic singular control problems. *SIAM J. Control and Optimization* **29** (1991), 1443–1475.
- [26] C. N. Laws. *Dynamic Routing in Queueing Networks*. Ph.D. dissertation, Statistical Laboratory, University of Cambridge, U.K., 1990.
- [27] C. N. Laws. Resource pooling in queueing networks with dynamic routing. *Adv. Appl. Probab.*, **24** (1992), 699–726.
- [28] C. N. Laws and G. M. Louth. Dynamic scheduling of a four-station queueing network. *Probab. Engrg. Inform. Sci.*, **4** (1990), 131–156.
- [29] A. Mandelbaum and A. L. Stolyar. Scheduling flexible servers with convex delay costs: heavy-traffic optimality of the generalized $c\mu$ -rule. Submitted to *Operations Research*.
- [30] W. P. Peterson. A heavy traffic limit theorem for networks of queues with multiple customer types. *Math. Oper. Res.* **16** (1991), 90–118.
- [31] A. A. Puhalskii and M. I. Reiman. A critically loaded multirate link with trunk reservation. *Queueing Systems*, **28** (1998), 157–190.
- [32] M. Squillante, C. H. Xia, D. Yao, and L. Zhang. Threshold based priority policies for parallel-server systems. Preprint, 2000.
- [33] A. L. Stolyar. Maxweight scheduling in a generalized switch: state space collapse and workload minimization in heavy traffic. Submitted to *Annals of Applied Probability*.
- [34] Y. C. Teh. *Threshold Routing Strategies for Queueing Networks*. D. Phil. thesis, University of Oxford, 1999.
- [35] Y. C. Teh and A. R. Ward. Critical thresholds for dynamic routing in queueing networks. To appear in *Queueing Systems*.
- [36] L. Wein. Scheduling networks of queues: heavy traffic analysis of a two-station network with controllable inputs. *Oper. Res.* **38** (1990), 1065–1078.

- [37] R. J. Williams. An invariance principle for semimartingale reflecting Brownian motions in an orthant. *Queueing Systems*, **30** (1998), 5–25.
- [38] R. J. Williams. Diffusion approximations for open multiclass queueing networks: sufficient conditions involving state space collapse. *Queueing Systems*, **30** (1998), 27–88.
- [39] R. J. Williams. On dynamic scheduling of a parallel server system with complete resource pooling. In *Analysis of Communication Networks: Call Centres, Traffic and Performance*, D. R. McDonald and S. R. E. Turner (eds.), American Mathematical Society, Providence, RI, 2000, 49-71.