Dynamic Scheduling of Open Multiclass Queueing Networks in a Slowly Changing Environment

A Thesis
Presented to
The Academic Faculty

by

Junxia Chang

In Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy in Operations Research

School of Industrial and Systems Engineering
Georgia Institute of Technology
September 2004

Dynamic Scheduling of Open Multiclass Queueing Networks in a Slowly Changing Environment

Approved by:

Dr. Hayriye Ayhan, Committee Chair

Dr. Jim Dai, Advisor

Dr. Cathy Xia

Dr. Anton Kleywedgt

Dr. Spiridon Reveliotis

Date Approved: November 8th, 2004

To My Parents.

ACKNOWLEDGEMENTS

I am grateful to all those who have contributed towards shaping this thesis.

I wish to thank my advisors Dr. Jim Dai and Dr. Hayriye Ayhan for their guidance and supervision over the past several years. With the direction and guidance of Dr. Jim Dai, I am able to define, propose and defend this dissertation. And the supreme patience, support, encouragement from Dr. Hayriye Ayhan have always been accompanying me throughout this long journey. I extremely appreciate the time Dr. Hayriye Ayhan has spent on reading this thesis and helping me to refine it to the very detail. The insights, advice, and comments from Dr. Jim Dai and Dr. Hayriye Ayhan make the working with them very rewarding.

I am also indebted to Dr. Cathy Xia and Dr. Mark Squillante for leading me into studying the first problem of this dissertation. I am thankful to them for their guidance and help when I was working at IBM T.J. Watson Research Center in the summer of year 2001. I am grateful to Dr. Cathy Xia for her continuing mentoring and advising since then. I am thankful to Dr. Anton Kleywegt and Dr. Spirion Reveliotis for spending time on reading this thesis. Their comments and feedback have greatly improved this manuscript. I also benefited from the discussions with my colleagues, Wuqin Lin, Zhaosong Lu, Jiheng Zhang, Josh Reed and many others.

I would like to express my sincere gratitude to the Georgia Institute of Technology for the supportive environment. It is the guidance of Dr. Jia-an Yan and Dr. Xuan-Da Hu that has made my graduate study in the Georgia Institute of Technology possible. There are too many friends to name who have helped me adjust myself to the life and culture here during the past several years. Their friendship, support, and prayer have always been the sustaining sunshine of my life.

Above all, I owe all that I have achieved to my parents, whose forever love is the source of all my strength.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

SUMMARY

In this thesis we investigated the dynamic scheduling of computer communication networks that can be periodically overloaded. Such networks are modelled as mutliclass queueing networks in a slowly changing environment. A hierarchical framework is established to search for a suitable scheduling policy for such networks through its connection with stochastic fluid models. In this work, the dynamic scheduling of a specific multiclass stochastic fluid model is studied first. Then, a bridge between the scheduling of stochastic fluid models and that of the queueing networks in a changing environment is established.

In the multiclass stochastic fluid model, the focus is on a system with two fluid classes and a single server whose capacity can be shared arbitrarily among these two classes. The server may be overloaded transiently and it is under a quality of service contract which is indicated by a threshold value of each class. Whenever the fluid level of a certain class is above the designated threshold value, the penalty cost is incurred to the server. The optimal and asymptotically optimal scheduling policies are specified for such a stochastic fluid model.

Afterwards, a connection between the optimization of the queueing networks and that of the stochastic fluid models is established. This connection involves two steps. The first step is to approximate such networks by their corresponding stochastic fluid models with a proper scaling method. The second step is to construct a suitable policy for the queueing network through a successful interpretation of the stochastic fluid model solution, and a successful interpretation method is provided in this study.

CHAPTER 1

INTRODUCTION

1.1 Motivation

The Internet has been growing rapidly as a medium to store, process and deliver information since its birth. The Internet was first introduced when ARPANET adopted TCP/IP in the late twentieth century. With only 213 hosts in 1981, now the Internet has more than 200 million hosts and more than 840 million users as of September of 2002 (from the information released by Netsizer.com).

Accompanying the growth of the Internet, various Internet applications have been developed. These applications range from text-based utilities such as file transfer and remote login to the integrated advent such as the World Wide Web and multimedia streaming. Companies and costumers are increasingly reliant on these applications, especially the World Wide Web service, which can provide dynamic content, integrate with databases and offer secure commercial transactions. More and more people around the world tend to seek information and services from the Web, such as looking for driving directions, checking flight information, booking hotels, banking, and stock trading.

An important factor in the growth of the World Wide Web is the deployment of the electronic business (e-business). As a new communications medium, the Web becomes an electronic market for companies or organizations to advertise and sell products or services to consumers. With the trust in the provisioning of the Web sites, consumers also seek the information, buy products or services, and complete the business transactions through the on line services offered by those companies or organizations participating in this electronic market.

In a general e-business environment, most companies or organizations that sell products or services actually buy Internet services from a common Internet service provider such as IBM, HP, Intel. When a customer visits the Web sites of a company and requests a Web page, the request is actually directly served by the Internet service provider of that company. If the quality of service (QoS) provided by the Internet service provider is unsatisfactory, then the company lose potential online customers. Therefore, as part of the contract between each company (or organization) and the Internet service provider, the service level agreement (SLA) is specified. In the service level agreement, the Internet service provider guarantees to meet certain quality of service performance for each company. Each company or organization will also pay the Internet service provider according to the quality of service provided by the Internet service provider. The quality of service levels specified between the Internet service provider and companies (or organizations) are different based on the price negotiated between those companies (or organizations) and the Internet service provider. A critical issue for the Internet service provider is how to allocate its resources to meet the service level agreements and maximize its profits or minimize its costs.

We herein focus on determining the optimal decision for the Internet service provider to maximize its profits or minimize its costs with regard to what is specified in the service level agreement contracts. It is difficult to make the optimal decision for the Internet service provider due to the complexity of the computer networks it is facing. First, the quality of service levels as well as the prices for different companies or organizations are all different. Second, it is extremely difficult to predict the online behaviors of customers who come from all over the world. For example, when the customers will visit the Web sites, what Web pages they will request, how long they will stay at these Web sites, and what is the next Web page they will request are highly variable. Above all, the advancements of new computer technologies continue to bring in new Internet applications and services, and thus the complexity of the Internet also continues to grow rapidly.

We propose an analytical approach to investigate the decision problems concerning the service performance for the Internet service providers, such as resource allocation, performance prediction and quality of service provisioning. The analytical results can provide us with a better understanding of the fundamental issues and tradeoffs at the core of performance problems in the design and implementation of complex computer systems, networks and applications. The mathematical model we consider is very general in the sense that it is a stochastic network model, where we allow that the time between consecutive customer requests follows an unknown probabilistic distribution and the service time of each customer request can also be random. Details of the model are provided in Section 1.2.

1.2 The mathematical model, challenges, and objective

In this section, we will provide more details of the mathematical models we investigate and discuss the difficulties and challenges of analyzing such models.

1.2.1 The mathematical model

We consider optimizing the scheduling discipline of a multiclass queueing network model, where multiple classes of jobs wait in buffers before being served by an available server. Differentiation of different classes may result from different e-businesses with different quality of service contracts. Jobs of each class represent the requests of Web pages submitted by the customers visiting Web sites. It is important to note, however, that while our mathematical model and analysis are motivated by the e-business and Internet environments, they also apply to more general networks, including manufacturing networks with high volume production of small items.

We refer to the service requests submitted to a Web server as jobs. The workload characterization study of Web servers reveals that jobs arrive at Web servers in a bursty fashion. Not only are the inter-arrival times between jobs are random, but also the *average* inter-arrival time can change significantly over time. There are some sustainable periods when the Web site observes higher customer demand. For example, from the study of Arlitt and Jin in [1], the minimum number of requests received by the the 1998 World Cup Web site per hour from 16:30 to 21:00 on June 29th, 1998 is more than 4 millions, while the maximum number of requests per hour from 01:00am to 1:00pm on June 30th, 1998 is around only 1 million. The average number of requests per minute during 11:30pm to 11:45pm on June 30th, 1998 is 19 times more than the average number of requests per minute from June 7th to the July 18th of 1998. This type of bursty Web traffic is also observed by other commercial Web sites and Web sites of research institutions, as reported in Arlitt and Williamson [2].

This non-homogeneous behavior of system parameters motivated us to model the queueing networks as those operating in a changing environment. The change of the environment state triggers the change of the arrival rates, service rates, and routing probabilities. In particular, we consider that the state of the environment takes only discrete values. At each state of the environment and within each minute, the high speed and the large scale of contemporary Internet makes it possible that there are thousands of jobs, i.e. service requests, arriving to the network and thousands of jobs completed by the network. However, the time scale for a change in the environment state is larger than minutes. Therefore, the state of the queueing network changes much faster than the state of the environment does. In general, we consider that the network is operating in a slowly changing environment. By slow, we mean that the number of environment transitions is much fewer than the number of changes of the network state.

Our objective is to improve the performance of a mutliclass queueing network operating in a slowly changing environment. The set of parameters to describe the multiclass queueing network, such as the arrival rates, service rates, and the routing matrix, will change whenever the state of the environment changes. At each state of the operating environment, for a multiclass queueing network, there might be more than one class of jobs for a server to process. Whenever a server is available, one needs to determine which job to be processed next, (i.e the scheduling policy). In this mathematical model, we assume that violating the quality of service level agreement will result in the profit loss or cost increase for the Internet service provider. Our objective is to find an optimal or near optimal scheduling policy to maximize the profits or minimize the cost for such a mutliclass queueing network operating in a changing environment.

1.2.2 The objective

It is well known that finding the optimal scheduling policy for a multiclass queueing network is difficult even when the environment state attains only one value. For the queueing network in a changing environment, the scheduling problem is even more challenging since the environment process is a stochastic process. We adopt a relatively modest objective and plan to establish a hierarchical frame work to search for an asymptotically optimal scheduling policy. In this frame work, we first study a stochastic fluid model which has a simpler structure than the original queueing network in a changing environment. Then we derive a suitable policy for the discrete queueing network based on the stochastic fluid model solution.

We plan to establish this hierarchical frame work by providing a general method to derive an asymptotically optimal scheduling policy for the queueing network if the optimal policy of the stochastic fluid model is given.

We give a brief review of the related literature in Section 1.3 and discuss our results and contributions in Section 1.4

1.3 Literature review

As mentioned earlier, in the Web traffic characterization study, it is observed that there exist non-stationary effects and high peak-to-mean ratios in the Web traffic. It is reported in Arlitt and Jin [1] that the traffic of the 1998 World Cup Web site is quite bursty. The capacity of the system can hardly maintain the immediate responsiveness to all users' requests during the peak hours. Similar observation is also reported in Iyenger, Squillante, and Zhang [22]. Arlitt and Williamson [2] study the traffic pattern of six different Web sites, including research institute Web sites and commercial Web sites. They find out that generally the number of requests received per unit time during the peak hours is significantly larger than the other hours, and point out the failure of modelling the system by a time homogeneous network. The high peak-to-mean ratio of the demand pattern implies the Web server can be potentially overloaded during the peak hours if the capacity planning is made according to the mean value. The Web server being overloaded will result in longer response time to Web page requests, and therefore the quality of service level agreements might be violated. This consequently results in profit loss or cost increase of the Internet service providers. The significance of sustainable peak hours and the need to optimize the profit creates the necessity to model the Internet as a network in a changing environment.

As we have pointed out in Section 1.2, it is difficult to optimize such a network. However, certain connections between the standard queuing network (i.e the case that the environment stays at a single state) and its corresponding fluid model have been established. Scaling the time and space properly, Chen and Mandelbaum [13] show that a general standard class of queueing networks converges to deterministic fluid networks. In [15], Dai further reveals that the queueing network is stable if its corresponding deterministic fluid model is stable. In [12], Chen and Meyn suggest using the value function of the fluid model to initialize the value iteration algorithm for the queueing network and show through some numerical examples that such a choice may lead to faster convergence to an optimal policy.

With the hope that there is a connection between the optimal policy of the fluid model and its corresponding queueing network, more studies in optimizing deterministic fluid models have been conducted. Avram, Bertsimas, and Ricard have provided optimal solutions in [4] for various deterministic fluid models. In [11], Chen and Yao provide the conditions under which the index policy is optimal for mutliclass fluid networks. Weiss [37] provides a general algorithm to search for the optimal solution of the deterministic fluid models.

However, even if the optimal solutions of the fluid models are provided, how to derive a good policy for the original discrete queueing network is still difficult. Intuitively, one would consider to employ the solution of the fluid model for the queuing network. However, if the solution of fluid model is employed in an unmodified way, the derived scheduling policy of the queueing network may end in poor performance. This is indicated by the examples in Yeh, Dai, and Zhou [39]. Examples in Meyn [30] and Maglaras [27] also show that the derived policy may not even possess the *fluid scale asymptotic optimality* if the fluid policy is not modified properly when it is applied to the discrete queueing network.

The *fluid scale asymptotic optimality* criterion is proposed by Meyn in [29] to measure the goodness of a policy for the queueing network. If under the fluid scaling, the performance of the queueing network under a policy converges to an optimal solution of the fluid model, this policy is called an asymptotically optimal policy for this queueing network in the fluid scale, or this policy possesses the *fluid scale asymptotic optimality.* Despite the modest objective of *fluid scale asymptotic optimality*, the meaning of the fluid model solution is still subtle for the queueing network. In [30], Meyn suggests that the fluid model policy can be translated by an *affine shift* method, i.e shifting the origin to a constant value. But Meyn [30] does not provide the proof that this method will be effective in general. Maglaras [27] proposes a general method to translate the optimal fluid model solution to get a fluid scale asymptotically optimal scheduling policy for the queueing networks, although proofs of his results are not mathematically rigorous. Bäuerle [5] studies the asymptotic optimality of tracking policies for stochastic networks. However, the result in [5] relies on the assumption of the piecewise constant structure of fluid model solutions and the exponential type of distribution of inter-arrival times and service times.

Note that the above research activities concentrate on the setting where the network operates at a single environment state, (a special case of the model we consider here).

The relation between the stochastic fluid models and the queueing networks in a random environment is touched in Choudhury, Mandelbaum, Reiman, and Whitt [14]. In [14], Choudhury et al show that queueing systems in a random environment can be approximated by a stochastic fluid model, but this queueing system is mainly a single class queueing system and the mathematical model of this queueing system is not completely and rigorously built. Although the connection between optimizing the stochastic fluid model and optimizing the queueing network in a changing environment is not well established, there are already exsiting results on stochastic fluid models. In [6], Bäuerle and Rieder show that the index policy is optimal for a multiclass fluid network where the external arrival process of fluid is driven by a continuous time Markov chain with finite state space. The performance measure in [6] is to maximize the expected total discounted rewards or the expected total discounted costs. Note that the index type policies may not be optimal for general type of cost functions as indicated by this study in Chapter 2. Harrison and Zeevi [19] study a call center staffing problem using a stochastic fluid model.

This dissertation provides a bridge between the results of stochastic fluid models and the scheduling policies of queueing networks in a changing environment.

1.4 The approach, results, and contributions

In this section, we briefly describe our approach to search for the asymptotically optimal policy for the queueing network model introduced in Section 1.2. We will also present the results of this study and discuss its contributions.

1.4.1 The approach

We plan to take a two step approach to search for an asymptotically optimal policy of a mutliclass queueing network in a slowly changing environment. The first step is to investigate a stochastic fluid model which has a structure less complex than but is similar to that of the queueing network in a slowly changing environment. The second step is to derive a scheduling policy for the original queueing network from the stochastic fluid model solution.

Even though the stochastic fluid model is simpler than the original queueing network, it still keeps certain structure of the queueing network model as well as the stochastic pattern of the changing environment. In fact, the stochastic fluid model is an approximation of the queueing network model. In Chapter 3, we provide the relative result on how to approximate a queueing network in a slowly changing environment by its corresponding stochastic fluid model. Rigorous description of the stochastic fluid model is provided in Theorem 14 in Section 3.3.2.

Next, we assume that a solution of the corresponding stochastic fluid model is given, then we derive a scheduling policy for the queueing network model. This step is referred to as the *translation* or the *tracking* of the stochastic fluid model solution. If with proper scaling, the performance of the queueing network operating under the derived policy converges to the performance of the stochastic fluid model solution, then we say the translation of the stochastic fluid model solution is successful with respect to this performance measure.

When translating the fluid model policy back to get a scheduling policy for queueing networks, the caution is needed at the boundary of the fluid model. What we do is to keep a certain number of jobs at each buffer to be above certain value, which is referred to as the safety stock level. If the queue length of each buffer is above its designated level, we implement the policy suggested by the fluid model solution; otherwise, we implement a special policy such that the state of network is adjusted as quickly as possible to reach that level. We choose the safety stock level to be negligible compared to the network processing speed, i.e the network can be emptied in a very short time if the queue length is at or below the safety stock level. Essentially, we try to maintain the state of the network to be always away from the boundary to avoid potential adverse consequences. But we do not want to move the boundary too far, otherwise it may result in the profit loss or the cost increase. So we need to be cautious when choosing how much to shift the boundary. The detailed description of the translation method is provided in Section 3.4.

1.4.2 The results and contributions

In this dissertation, we build a mathematical model for multiclass queueing networks operating in a slowly changing environment. With the fluid scaling method, we show that the mutliclass queueing network in a slowly changing environment can be approximated by a stochastic fluid model. Then we provide a general translation method to derive a scheduling policy for the queueing network from a given stochastic fluid model solution. We also prove that the provided translation method is successful, i.e the derived scheduling policy is asymptotically optimal in the fluid scale if the given stochastic fluid model solution is optimal.

We also investigate the policy for a Web server through a stochastic fluid model, where the Web server could be overloaded periodically. In this stochastic fluid model, we address the service level agreement by adopting a threshold type cost function. This work shows that the optimal policy of stochastic fluid models under service level contracts seriously depend on the service level contract and the traffic pattern of the network. The simple structured policies such as the index policy may not be optimal even for fluid models with a simple cost structure to address certain quality of service contracts.

The contributions of our work are as follows.

- A complete mathematical model for multiclass queueing networks in a slowly changing environment is built. Note that in the related work [14], Choudhury, Mandelbaum, Reiman, and Whitt simply provide some suggestions to build such a network, but do not actually build the mathematical model rigorously.

- The multiclass queueing network in a slowly changing environment generalizes the standard mutliclass queueing network. In the model considered in this study, the state of the environment takes values from a discrete set, while in the standard mutliclass queueing network model, the state of environment takes only one value.

- We generalize the result provided in Choudhury et al [14] and we present the result in a more rigorous way. The result of [14] shows that a queueing system in a random environment can be approximated by a stochastic fluid model, but it concentrates on single class queueing systems only. And the result itself is not mathematically and rigorously presented. Our result is for a general multiclass queueing networks, and we state our result rigorously in Theorem 14.

- We provide a general method to translate the stochastic fluid model solution and this method is easy to implement. This method is described in Section 3.4.

- We prove that the translation method is successful under moderate conditions, i.e the derived scheduling policy for the queueing network is *good* by the *fluid scale asymptotic optimality* criterion proposed in Meyn [29]. Note that even for the standard queueing network case where the operating environment is not changing, it is difficult to provide a general and successful translation method as we have discussed in Section 1.3. Although Maglaras [27] provides a translation method for standard multiclass queueing networks, his proof lacks mathematical rigor.

- We establish a hierarchical frame work to facilitate the search for the fluid-scale asymptotically optimal scheduling policy for multiclass queueing networks in a slowly changing environment. Our approach involves three steps. The first step is to approximate the original network by a stochastic fluid network model. The approximation is provided in Section 3.3.2. The second step is to find the optimal scheduling policy for the stochastic fluid model. The third step is to apply the translation method we provide in Section 3.4 to obtain the scheduling policy for the original network.

- We provide an asymptotically optimal scheduling policy for a stochastic fluid model where Web servers are under quality of service contracts and can be overloaded periodically.

1.5 Outline of the dissertation

The rest of the thesis is organized as follows. In Chapter 2, we provide the asymptotically optimal scheduling policies for Web servers that are under quality of service contracts and can be overloaded periodically. Then we bridge the gap between the solutions of general stochastic fluid models and their corresponding queueing networks in Chapter 3. In particular, we establish a frame work in order to search for an asymptotically optimal scheduling policy for mutliclass queueing networks in a slowly changing environment. In Chapter 4, we conclude this work.

Throughout the manuscript, we use $\mathbb{R}$ to denote the real line, and $\mathbb{R}_+$ to denote the nonnegative real numbers, i.e. $[0, \infty)$. We use $'$ to denote the transpose operation on a vector or a matrix. Operations taken on vectors are interpreted as operations taken on each corresponding component. For example, for K-dimensional vectors $a = (a_1, \ldots, a_K)'$ and $b = (b_1, \ldots, b_K)'$, $a + b = (a_1 + b_1, \ldots, a_K + b_K)'$, $a \leq b$ means that $a_i \leq b_i$ for all $1 \leq i \leq K$, and $a \ngeq b$ means that there exists an i, $1 \leq i \leq K$ such that $a_i < b_i$. We also use $a = (a_i, 1 \leq i \leq K)'$ to denote a K-dimensional vector a. We use $|a|$ to denote $\max\{a, -a\}$, where a is a real number or a vector of real numbers.

CHAPTER 2

A STOCHASTIC FLUID MODEL WITH TRANSIENT OVERLOAD AND QUALITY OF SERVICE AGREEMENTS

In this chapter, we study a specific multiclass stochastic fluid model for a Web server with two classes of jobs. The Web server is under quality of service contract and can be overloaded periodically. After an introduction and the model description, we present our results. For the rest of the manuscript, jobs or customers all mean service requests.

2.1 Introduction

Recent advances in Internet services and other emerging applications have created new computing and networking paradigms in which a set of e-commerce businesses contract with a common hosting provider of Internet applications and services for their respective customers. In such an environment, the hosting service provider needs to meet a diverse set of requirements of the various e-commerce businesses and customers. To address these diverse requirements and leverage potential economies of scale, the hosting service provider will often deploy a cluster of servers to effectively share the computing and networking resources required to support the desired Internet applications and services. A number of computer industry companies such as HP, IBM and Intel are already providing such hosting services and it appears that more companies will be doing so in the future.

To differentiate the diverse requirements of e-commerce businesses and customers, it is necessary to introduce the notion of different service classes. These service classes typically have distinct levels of importance to the hosting service provider, the businesses and their customers. Moreover, many of these service classes require specific Quality-of-Service (QoS) performance guarantees; failures to deliver such levels of QoS can have a significant impact on the e-commerce businesses and customers. For example, customers may easily lose patience and discontinue using the service if its responsiveness is perceived to be too long. Hence, as part of the contract between the service provider and each business, the hosting service provider agrees to guarantee a certain level of QoS for each class of service, and in return each e-commerce business agrees to pay the service provider for satisfying these QoS performance guarantees. Such Service-Level-Agreements (SLA) are included in service contracts between each business and the service provider, and they specify both performance targets or QoS guarantees, and financial consequences for meeting or failing to meet these targets. A service level agreement may also depend on the anticipated level of per-class workload from the customers of the business.

Thus, it is critical for the hosting service provider to dynamically allocate its server resources to optimize performance and profit measures in cluster-based computing environments with SLA contracts containing QoS performance guarantees. This is also an important issue for the continued growth and success of Internet services and applications. Therefore, in this chapter we focus on a particularly important class of dynamic scheduling problems that arise in these computing environments. However, it is important to note that while our analysis and results are motivated by such environments, they apply more generally to a wide variety of emerging computing environments with SLA-based QoS performance guarantees.

Previous studies that address QoS performance guarantees have focused mostly on throughput or mean response time measures. However, a crucial issue for Internet applications and services concerns the per-request efficiency with which the differentiated services are handled, since delays experienced by customers can result in lost revenue and customers for a business as described above. Furthermore, more standard performance metrics such as throughput and mean response time may not fully capture such QoS performance guarantees. In order to address these issues, we consider a general class of SLAs in which a threshold is defined for each class of service such that the hosting service provider gains revenues when the QoS level experienced by the class stays at or below the threshold, but the service provider pays penalties to the corresponding businesses when this threshold is exceeded. Then the optimal control problem focuses on allocating server resources in order to maximize the profit of hosting the collection of e-commerce sites under these SLA constraints.

Another big challenge of the problem concerns the diverse workloads of different e-commerce businesses and their variation over time. It is common in the computing environments of interest to have the workload of certain classes in each e-commerce site alternate between a period during which the arriving workload exceeds the allocated capacity, and a period during which the arriving workload is less than this capacity, even though the average load is within the allocated capacity; e.g., see [7]. These periods of transient overload can have a significant effect on the performance experienced by the different classes of service. This in turn can have a critical impact on the penalties that the hosting service provider is required to pay each e-commerce business according to the SLA contract between them. Hence, it is crucial to include these important workload characteristics in the analysis of the optimal control problem.

This problem falls within the general class of optimal resource control problems with the foregoing non-conventional performance metrics and workload characteristics. Several researchers have studied the issue of workloads with transient overload, but their studies have focused on single-class workloads and specific scheduling strategies, such as admission control (e.g., [21]) and direct modifications to the Internet server scheduling mechanism (e.g., [7, 10]). On the contrary, the focus in the present study is on the optimal dynamic scheduling of a multiclass system with transient overload. Furthermore, little has been done to consider the issue of maximizing profit in these computing paradigms under non-conventional performance metrics. The primary exception is the study in [25], which develops queueing-theoretic bounds and approximations to formulate the resource control optimization problem and then develops efficient algorithms to compute the optimal solution. This study is the one that is most relevant to this research, but it differs from the present study in several important aspects. The present focus is on computing the optimal dynamic scheduling policy and gaining insights into its fundamental properties, as opposed to computing the steady-state solution, and to do so under a workload with transient overload, which is not considered in [25].

The primary concern in this chapter is to investigate the preceding optimal server resource control problem as a dynamic scheduling problem. The motivation behind considering a fixed time horizon is that in reality many web sites exhibit regular daily access patterns (see [24]), typically there is one single peak period each day, the low period load is far below the system capacity so that the system usually starts empty the next day. Distributed architectures with separate machines for different geographical locations are also common in practice in order to improve the response time for accessing data over the Internet. This again validates the single period model. Hence, the traffic from the previous period does not have an effect on the next period. The approach is based on formulating the problem as a multiclass stochastic fluid model and employing optimal control theory [31, 32] to search for the optimal control policy that maximizes the total revenue over a fixed time horizon. Even though recent studies of a similar spirit for different dynamic scheduling problems include [4, 6, 17, 37], to the best our knowledge, no optimal scheduling policy is known for the general problem considered herein. As mentioned above, the present focus is on minimizing the penalty of the hosting service provider by dynamically scheduling its server resources among the fluid classes in a system that can be overloaded for a transient period. In order to capture the QoS performance guarantees in the SLA contracts, we introduce a threshold value for each fluid class such that a holding cost is incurred only if the amount of fluid of a certain class exceeds its threshold value. In this study, we consider the specific case of two fluid classes and a single server whose capacity can be shared arbitrarily among the two classes. We assume that the class 1 arrival rate changes with time and the class 1 fluid can more efficiently reduce the holding cost and develop the optimal server resource allocation policy that minimizes the holding cost in the corresponding fluid model when the arrival rate function for class 1 is known. We then study the stochastic fluid system when the arrival rate function for class 1 is random and propose various policies that are optimal or near optimal under various conditions. In particular, we consider two different types of heavy traffic regimes and prove that our proposed policies are *strongly asymptotically optimal* in the following sense: the difference between its performance and the optimality is *bounded* from above by a constant even as the optimal value itself goes to infinity. This notion of strong asymptotic optimality is used throughout this chapter and it has also been considered in [35, 38], as a measure to evaluate the closeness to optimality of approximating control policies. Numerical examples are also provided to demonstrate further that these policies yield good results in terms of minimizing the expected holding cost.

The outline of this chapter is as follows. We define our multiclass fluid model in Section 2.2. Deterministic instance of the model is analyzed in Section 2.4 where we provide the optimal control policy. Section 2.5 and Section 2.6 consider the stochastic instance of the model. In Section 2.5, we present a discrete review policy and show that it is asymptotically optimal as the expected length of the high period tends to infinity. Other policies that are asymptotically optimal are further discussed in Section 2.6. Our concluding remarks are provided in Section 2.9. Throughout, proofs are relegated to Section 2.8.

2.2 The stochastic fluid model

This chapter focuses on the following stochastic fluid system that serves two classes of fluid. Each class fluid continuously arrives at its buffer whose capacity is assumed to be infinite. Both classes are served by a single server whose service capacity can be shared arbitrarily among the two classes. When the server devotes full effort to class i, it processes class i fluid at rate μ_i, $i = 1, 2$.

Class 2 fluid arrives at a constant rate λ_2 throughout the time horizon under consideration. Class 1 fluid has a high arrival rate λ_1^h during the first part of the time interval and a low arrival rate λ_1^l in the rest of the time interval. Naturally, $\lambda_1^l \leq \lambda_1^h$. The durations of the first and second time intervals are denoted by H and L, respectively. Both H and L are random. Some of their statistics like mean remaining life times are assumed to be known. These assumptions will be spelled in more precise terms later. We call the time interval $[0, H)$ the high load period and the time interval $[H, H + L)$ the low load period.

We use $Z_i(t)$ to denote the fluid level in class i at time t, and $T_i(t)$ to denote the cumulative amount of time in $[0, t]$ that the server spends on class i fluid, $i = 1, 2$. The dynamics of the fluid model is given by the following equations

$$Z_i(t) = Z_i(0) + \int_0^t \lambda_i(s)\, ds - \mu_i T_i(t), \quad t \in [0, H+L), \tag{1}$$

$$T_i(0) = 0, \quad T_i(t) \text{ is a nondecreasing function of } t, \tag{2}$$

$$t - (T_1(t) + T_2(t)) \text{ is a nondecreasing function of } t, \tag{3}$$

where $\lambda_i(s)$ is the arrival rate to class i at time s. Since the class 1 arrival rate function $\lambda_1(\cdot)$ is random, the fluid level process Z is random as well. The allocation process $T = \{(T_1(t), T_2(t)), t \geq 0\}$ reflects how the server spends its service capacity among two classes and it is called a scheduling or a service policy.

Let $h_i > 0$ and $\theta_i \geq 0$ be constants, $i = 1, 2$. For a real number x, define $x^+ = \max(x, 0)$. Consider the integral

$$\int_0^{H+L} \sum_{i=1}^{2} h_i \left(Z_i(t) - \theta_i\right)^+ \, dt \tag{4}$$

which is called the total cost of the system. Then one interprets h_i as the holding cost per unit time when the fluid level in class i exceeds θ_i. If the fluid level in class i is below θ_i, the fluid does not accumulate cost for the system. Clearly, the cost depends on initial fluid level $z = Z(0)$, and allocation T employed. Since H and L are random variables, the cost is also random. The focus of this chapter is to find an allocation T to minimize the expected total cost for each initial point z. We assume that working on class 1 can more efficiently reduce holding costs. Namely,

$$h_1 \mu_1 > h_2 \mu_2. \tag{5}$$

If the assumption in (5) is violated, the optimal policy is a generalization of the well-known $c\mu$ rule (see for example, Smith [36], Klimov [23] and Green and Stidham [18]). Details of such an optimal policy are presented in Section 2.3.

When $\theta_i = 0$ for $i = 1, 2$, the optimal policy is again given by the $c\mu$-rule. That is the server gives priority to class i with highest $h_i \mu_i$. To the best of our knowledge, the optimal policy for our general problem is not known. In the special case when H and L are deterministic and are known at the beginning of the time window, we will present an optimal policy. Using this policy, we will construct heuristic policies, known as discrete review policies, for controlling the system. We will present numerical experiments showing that these policies perform well. We will establish asymptotic results guaranteeing good performance of these policies in certain parameter regions. We will also identify other policies that are asymptotically optimal in certain parameter regions.

For any feasible allocation T, it follows that $T(t)$ is Lipschitz continuous in t. Thus, T is absolutely continuous and has derivatives almost everywhere. Therefore, specifying an allocation T is equivalent to specifying its derivative $\dot{T}(t)$ for almost every t in $(0, H+L)$. (For a function f, $\dot{f}(t)$ denotes the derivative of f at time t. Whenever $\dot{f}(t)$ is used, the derivative of f at time t is assumed to exist.) Clearly, any feasible allocation T should be non-anticipating. Namely, $\dot{T}(t)$ depends only on the information available up to time t.

For future reference, we also define the traffic intensities of the system. The system load per unit of time contributed by class 1 fluid is $\rho_1^h = \lambda_1^h/\mu_1$ for the high load period and $\rho_1^l = \lambda_1^l/\mu_1$ for the low load period. The system load per unit of time contributed by class 2 fluid is constant and given by $\rho_2 = \lambda_2/\mu_2 > 0$. The overall system load is $\rho^h = \rho_1^h + \rho_2$ for the high load period and $\rho^l = \rho_1^l + \rho_2$ for the low load period. When $\rho^h > 1$ and $\rho^l < 1$, the total system work increases in the high load period and decreases in the low load period. In this case, the high load period is also called the overload period. Thus, when $\rho^h > 1$ and $\rho^l < 1$ the system experiences an overload period followed by an under-load period, a phenomenon known as transient overload in literature; see, for example, [7]. Although understanding transient overload is the primary motivation of this chapter, except explicitly stated otherwise, we do not assume $\rho^h > 1$.

2.3 Optimal policies if the consistent customer class is more expensive

In this section, we provide the optimal policies for the case if the assumption (5) is violated, i.e $h_1\mu_1 \le h_2\mu_2$. Under the assumption that the class 2 has constant arrival rate λ_2 and $\rho_2 < 1$, if $h_1\mu_1 \le h_2\mu_2$, then the optimal policy is a generalization of the $c\mu$ rule. Such an optimal policy is given below. The optimality of this policy can be proven using the techniques in Section 2.8.1 as is done when the assumption in (5) holds and thus omitted.

- If $Z_2(t) > \theta_2$, full capacity is given to class 2, i.e. $\dot{T}_1(t) = 0$, $\dot{T}_2(t) = 1$.

- If $Z_2(t) = \theta_2$ and $Z_1(t) > \theta_1$, enough capacity is given to class 2 such that class 2 fluid level is kept at θ_2 and the remaining capacity is used to serve class 1, i.e. $\dot{T}_1(t) = 1 - \rho_2$, $\dot{T}_2(t) = \rho_2$.

- If $Z_2(t) < \theta_2$ and $Z_1(t) \geq \theta_1$, full capacity is given to class 1, i.e. $\dot{T}_1(t) = 1$, $\dot{T}_2(t) = 0$.

- If $Z_2(t) < \theta_2$ and $Z_1(t) < \theta_1$, and the system is in the high load period ($t < H$), full capacity is given to class 1, i.e. $\dot{T}_1(t) = 1$, $\dot{T}_2(t) = 0$.

- If $Z_2(t) \leq \theta_2$ and $Z_1(t) \leq \theta_1$, and the system is in the low period ($H < t < H + L$), enough capacity is given to each class such that the fluid levels of both classes are kept below their threshold values. We have multiple choices in this case, one is to let $\dot{T}_1(t) \geq \rho_1^l$, $\dot{T}_2(t) \geq \rho_2$ such that $\dot{T}_1(t) + \dot{T}_2(t) \leq 1$.

Throughout the rest of the chapter, we assume (5).

2.4 Optimal policies in the deterministic case

In this section, we present the optimal policy when the lengths of the high period and the low period are known. Thus, H and L are deterministic quantities. The optimality of this policy is proven in Section 2.8. For the sake of notational convenience, we first define the following policy.

Definition 1 (Low-Period-Policy). The following policy referred to as the *Low-Period-Policy* is implemented in the low period, i.e, when $H < t \leq H + L$.

- If $Z_1(t) > \theta_1$, full capacity is given to class 1, i.e. $\dot{T}_1(t) = 1, \dot{T}_2(t) = 0$.

- If $Z_1(t) = \theta_1,\ Z_2(t) > \theta_2$, class 1 fluid is kept at its threshold value θ_1, while the remaining capacity is used to serve class 2, i.e. $\dot{T}_1(t) = \rho_1^l, \dot{T}_2(t) = 1 - \rho_1^l$.

- If $Z_1(t) < \theta_1,\ Z_2(t) > \theta_2$, then full capacity is given to class 2, i.e. $\dot{T}_1(t) = 0, \dot{T}_2(t) = 1$.

- If $Z_1(t) \leq \theta_1,\ Z_2(t) \leq \theta_2$, then the policy is not unique and $\dot{T}_1(t)$ and $\dot{T}_2(t)$ can be chosen from any solution satisfying $\dot{T}_1(t) \geq \rho_1^l,\ \dot{T}_2(t) \geq \rho_2$ and $\dot{T}_1(t) + \dot{T}_2(t) \leq 1$.

The optimal policy depends on the system load. In the next three sections, we will describe the optimal policy under various load conditions. In the first case, $\rho_1^h > 1, \rho^l \leq 1$, and we refer to this case as highly overloaded case; in the second case, $\rho^h > 1,\ \rho_1^h \leq 1,\ \rho^l \leq 1$, and we refer to this case as the overloaded case; and in the last case, $\rho^h \leq 1, \rho^l \leq 1$, and we refer to this case as the lightly loaded case.

2.4.1 The highly overloaded case

In this section, we assume $\rho_1^h > 1$ and $\rho^l \leq 1$ and provide the optimal policy when the duration of the high period H and the duration of the low period L is known.

Suppose that $\rho_1^h > 1$ and $\rho^l \leq 1$, then the optimal policy has the following structure:

(OPT)

$$\begin{aligned}
\forall t \in (0,\ s_1): &\quad \dot{T}_2(t) = 1,\ \dot{T}_1(t) = 0;\\
\forall t \in (s_1,\ s_2): &\quad \dot{T}_2(t) = u_2,\ \dot{T}_1(t) = u_1, u_1 + u_2 = 1;\\
\forall t \in (s_2,\ H): &\quad \dot{T}_2(t) = 0,\ \dot{T}_1(t) = 1;\\
\forall t \in (H,\ H+L): &\quad \text{Low-period-policy.}
\end{aligned}$$

Thus, the optimal policy gives fixed priority to class 2 in the interval 0 to s_1, employs processor sharing in the interval s_1 to s_2 and gives fixed priority to class 1 in the interval s_2 to H. Specific values of s_1, s_2, u_1, and u_2 depend on the initial fluid levels and the length of the high and the low periods. Before discussing the computation of s_1, s_2, u_1 and u_2 for all possible cases, we introduce the notation used in our developments:

$$d_1 = \theta_1 - Z_1(0), \quad \psi_1 = \tfrac{d_1/\mu_1}{\rho_1^h - 1}, \quad \tilde{\psi}_1 = \frac{d_1/\mu_1}{\rho_1^h}, \tag{6}$$

$$d_2 = \theta_2 - Z_2(0), \quad \psi_2 = \tfrac{d_2/\mu_2}{\rho_2}, \quad \tilde{\psi}_2 = \frac{-d_2/\mu_2}{1-\rho_2}. \tag{7}$$

The quantities ψ_1, ψ_2, $\tilde{\psi}_1$ and $\tilde{\psi}_2$ have the following interpretations. Quantity ψ_1 is the time that class 1 increases to its threshold θ_1 under the policy that gives fixed priority to class 1 if the initial fluid level of class 1 is below θ_1 and if the high period is long enough. Quantity $\tilde{\psi}_1$ is the time class 1 increases to its threshold θ_1 under the policy that gives fixed priority to class 2 if the initial fluid level of class 1 is below θ_1 and if the high period is long enough. Quantity ψ_2 is the time class 2 increases to its threshold θ_2 under the policy that gives fixed priority to class 1 if the initial fluid level of class 2 is below θ_2. Finally, $\tilde{\psi}_2$ is the time class 2 decreases to its threshold θ_2 under the policy that gives fixed priority to class 2 if the initial fluid level of class 2 is above θ_2. Clearly, d_1 and d_2 denote the initial deviation of the fluid levels from the desired thresholds for classes 1 and 2, respectively.

We also define

$$a_1 = \frac{d_1/\mu_1 + d_2/\mu_2}{\rho_1^h + \rho_2 - 1}, \quad a_2 = \frac{1-\eta\xi}{1-\eta}\psi_1^+ - \frac{\eta(1-\xi)}{1-\eta}\psi_2^+, \tag{8}$$

$$B = \frac{1-\eta\xi}{1-\eta}\psi_1^+ - \frac{(1-\rho_1^l)[1+\eta(\rho_1^h-1)] + (1-\eta)(\rho_1^h-1)}{(\rho_1^h-1)(\rho_1^h-\rho_1^l)(1-\eta)}\tilde{\psi}_2^+, \tag{9}$$

where

$$\xi = \frac{(\rho_1^h - 1)}{(\rho_1^h - \rho_1^l)} \text{ and } \eta = \frac{h_2\mu_2}{h_1\mu_1}.$$

Quantities a_1, a_2 and B have the following interpretations. Quantity a_1 is the critical value such that if the high period is longer than a_1 then under any policy either class 1 fluid level will exceed its threshold θ_1 or class 2 fluid level will exceed its threshold θ_2. Quantity a_2 is the critical value such that if the high period is longer than a_2 and the low period is long enough to reduce the fluid level of class 1 to its threshold θ_1 then fixed priority to class 1 is the optimal policy in the high period. Finally, B is the critical value such that if the high period is longer than B and the low period is long enough to reduce the fluid level of class 1 to its threshold θ_1 then the optimal policy never uses processor sharing in the high period. Finally, for the sake of simplicity, we define

$$\gamma_1 = \frac{\eta(\rho_1^h-1)(\rho_1^h+\rho_2-1)}{(1-\rho_1^l)[\rho_2+\eta(\rho_1^h-1)]+(1-\eta)\rho_2(\rho_1^h-1)}, \quad \gamma_2 = \frac{\eta\rho_1^h(\rho_1^h-1)}{(1-\rho_1^l)[1+\eta(\rho_1^h-1)]+(1-\eta)(\rho_1^h-1)}, \quad \gamma_3 = \frac{\rho_1^h-1}{1-\rho_1^l}.$$

We now provide a more detailed description of the optimal policy by considering all possible cases of the initial load. As can be seen below, Cases 1 and 3 are simple and have no subcases (i.e the policy is independent of the length of H and L). However, Cases 2 and 4 have many subcases. Hence, for the sake of clarity, we provide pictorial representations of Cases 2 and 4 in Figures 1 to 3. In particular, we present the corresponding case for each value of H and L and demonstrate that we consider all possible values for the length of the high and low periods. Depending on the relationship between $\tilde{\psi}_1$ and $\tilde{\psi}_2$, we provide the corresponding pictorial representation of Case 2, respectively in Figures 1 and 2. Figure 3 is the pictorial representation of Case 4.

- Case 1: $Z_1(0) \geq \theta_1$. In this case, the optimal policy is given by (OPT) with $s_1 = s_2 = 0$. Note that when setting $s_1 = s_2 = 0$, the (OPT) policy gives fixed priority to class 1 throughout the high period.

- Case 2: $Z_1(0) < \theta_1$, $Z_2(0) > \theta_2$. Computation of s_1, s_2, u_1 and u_2 depends on the length of the high and the low periods.

  – Case 2.1: If

$$a_1 \leq H \leq B, \quad L \geq \gamma_1(H - a_1), \tag{10}$$

  then s_1, s_2, u_1 and u_2 are computed by solving

$$Z_2(0) + (\lambda_2 - \mu_2)s_1 = \theta_2, \tag{11}$$

$$Z_1(0) + \lambda_1^h s_1 = Z_1(s_1), \tag{12}$$

$$Z_2(s_1) + (\lambda_2 - \mu_2 u_2)(s_2 - s_1) = \theta_2, \tag{13}$$

$$Z_1(s_1) + (\lambda_1^h - \mu_1 u_1)(s_2 - s_1) = Z_1(s_2), \tag{14}$$

$$u_1 + u_2 = 1, \tag{15}$$

$$Z_1(s_2) + (\lambda_1^h - \mu_1)(t_1 - s_2) = \theta_1, \tag{16}$$

$$Z_1(t_1) + (\lambda_1^h - \mu_1)(H - t_1) = Z_1(H), \tag{17}$$

$$Z_1(H) + (\lambda_1^l - \mu_1)(t_2 - H) = \theta_1, \tag{18}$$

$$\mu_1 h_1(t_2 - t_1) = \mu_2 h_2(t_2 - s_2). \tag{19}$$

  Note that equations (11) to (18) describe the evolution of the fluid levels of class 1 and class 2 from time 0 to t_2 under the optimal policy, where t_2 represents the time epoch at which the class 1 fluid level in the low period reaches its threshold value as indicated in equation (18). In particular, equations (11) and (12) describe the evolution of fluid levels from time 0 to s_1 when higher priority is given to class 2. At s_1, class 2 fluid level is reduced to its threshold θ_2 from above. Equations (13) to (15) describe the evolution of the fluid levels from s_1 to s_2 under the processor sharing policy. In $[s_1, s_2]$, class 2 fluid level remains at its threshold θ_2. Equations (16) to (18) describe the evolution of class 1 fluid level from s_2 to t_2 under the policy that gives higher priority to class 1. Equation (16) implies that at time t_1, class 1 fluid level increases to its threshold θ_1. Equation (17) records the class 1 fluid level at the end of the high period. Equation (19) ensures that the profit gained by serving class 1 is equal to the profit lost by not serving class 2. Under the conditions given in (10), it will be shown in Appendix 2.8.1 that equations (11) to (19) have a unique solution with $0 \le s_1 \le s_2 \le t_1 \le H \le t_2 \le H + L$ and $u_1,\ u_2 \ge 0$.

– Case 2.2: If

$$L \le \gamma_1(H - a_1), \quad a_1 \le H, \quad H + L \le \tilde{\psi}_1 + \frac{1 + \eta(\rho_1^h - 1)}{(1-\eta)(\rho_1^h - 1)}(\tilde{\psi}_1 - \tilde{\psi}_2),$$

then we set $t_2 = H + L$ and compute s_1, s_2, u_1, u_2 and t_1 by solving equations (11)-(17) and (19).

– Case 2.3: If

$$\max\{B, \tilde{\psi}_1\} \le H \le a_2, \quad L \ge \gamma_2(H - \tilde{\psi}_1),$$

then we set $s_1 = s_2$ and solve the equations (12) and (16)–(19) for s_2, t_1 and t_2.

– Case 2.4: If

$$L \le \gamma_2(H - \tilde{\psi}_1), \quad \max\left\{\tilde{\psi}_1, \tilde{\psi}_1 + \frac{1 + \eta(\rho_1^h - 1)}{(1-\eta)(\rho_1^h - 1)}(\tilde{\psi}_1 - \tilde{\psi}_2)\right\} \le H + L \le \frac{\psi_1}{1-\eta},$$

then we set $s_1 = s_2$ and $t_2 = H + L$ and compute s_2 and t_1, by solving equations (12), (16)-(17) and (19).

– Case 2.5: If $H \le \max\{a_1, \tilde{\psi}_1\}$, then the optimal policy is given by (OPT) with $s_1 = \min\{\tilde{\psi}_2,\ H\}$, $s_2 = H$, $u_2 = \rho_2$, and $u_1 = 1 - \rho_2$.

– Case 2.6: If $H \ge a_2$ and $H + L \ge (1-\eta)^{-1}\psi_1$, then the optimal policy is given by (OPT) with $s_1 = s_2 = 0$.

- Case 3: $Z_1(0) < \theta_1$, $Z_2(0) \le \theta_2$, $\psi_1 \le \psi_2$. In this case, the optimal policy is given by (OPT) with $s_1 = s_2 = 0$.

- Case 4: $Z_1(0) < \theta_1$, $Z_2(0) \le \theta_2$, $\psi_1 \ge \psi_2$. In this case, $s_1 = 0$. However, the computation of s_2, u_1 and u_2 depends on the lengths of the high and the low periods as discussed below.

  - Case 4.1: If $a_1 \le H \le a_2$, $L \ge \gamma_1(H - a_1)$, then s_2, u_1, u_2, t_1 and t_2 are computed by solving equations (13)–(19) with $s_1 = 0$.

  - Case 4.2: If

$$H \ge a_1,\ H + L \le \psi_1 + \frac{\eta}{1-\eta}(\psi_1 - \psi_2),\ L \le \gamma_1(H - a_1),$$

    then we set $t_2 = H + L$, and solve the equations (13)-(17) and (19) with $s_1 = 0$ to compute s_2, u_1, u_2 and t_1.

  - Case 4.3: If $H \le a_1$, then the optimal policy is given by (OPT) upon setting $s_1 = 0$, $s_2 = H$, selecting u_2 as any value in the interval $[(\rho_2 - d_2(\mu_2 H)^{-1})^+,\ d_1(\mu_1 H)^{-1} - (\rho_1^h - 1)]$ and setting $u_1 = 1 - u_2$.

  - Case 4.4: If $H \ge a_2$, $H + L > \psi_1 + \eta(1-\eta)^{-1}(\psi_1 - \psi_2)$, then the optimal policy is given by (OPT) with $s_1 = s_2 = 0$.

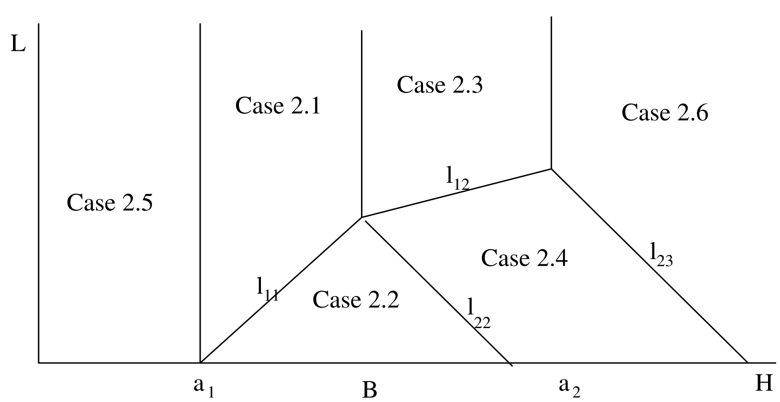

Figure 1: Optimal policies in the deterministic case for the first type initial condition.

We also provide a pictorial representation for the optimal policies corresponding to each value of H and L through three figures. Figure 1 is for the case where the parameters

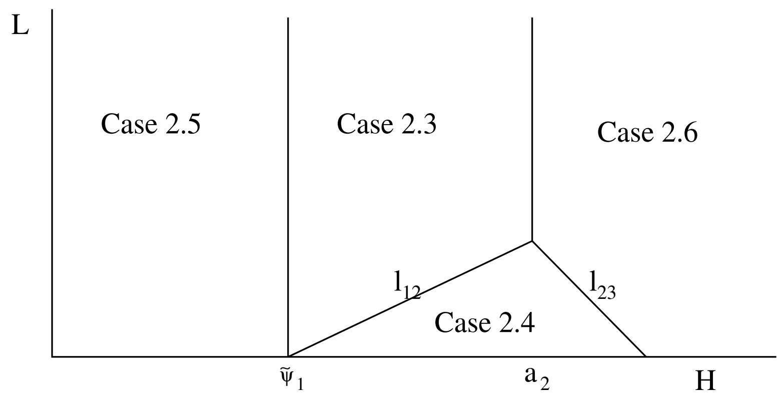

Figure 2: Optimal policies in the deterministic case for the second type initial condition.

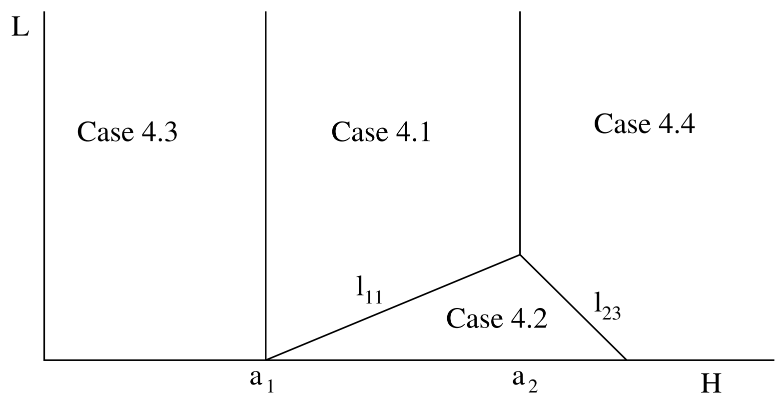

Figure 3: Optimal policies in the deterministic case for the third type initial condition.

and initial condition satisfies $Z_1(0) \leq \theta_1, Z_2(0) \geq \theta_2$ and $\tilde{\psi}_1 \geq \tilde{\psi}_2$. Figure 2 is for the case that $Z_1(0) \leq \theta_1, Z_2(0) \geq \theta_2$ and $\tilde{\psi}_1 \leq \tilde{\psi}_2$, and Figure 3 is for the case that $Z_1(0) \leq \theta_1, Z_2(0) \leq \theta_2$ and $\psi_2 \leq \psi_1$. In all these three figures, the line l_{11} satisfies $L = \gamma_1(H - a_1)$; the line l_{12} satisfies $L = \gamma_2(H - \tilde{\psi}_1)$; the line l_{13} satisfies $L = \gamma_3(H - \psi_1)$; the line l_{22} satisfies $H + L = \tilde{\psi}_1 + ((1-\eta)(\rho_1^h - 1))^{-1}(1 + \eta(\rho_1^h - 1))(\tilde{\psi}_1 - \tilde{\psi}_2)$; and line l_{23} satisfies $H + L = (1-\eta)^{-1}(\psi_1 - \eta\psi_2^+)$.

As mentioned above, we prove the optimality of this policy in Section 2.8.1. However, in order to give the reader an intuitive explanation, we consider one of the cases above, for example Case 3. We claim that if $Z_1(0) < \theta_1$, $Z_2(0) \leq \theta_2$, $\psi_1 \leq \psi_2$, then the optimal policy is given by (OPT) with $s_1 = s_2 = 0$. In order to see this, first consider the case $H \geq \psi_1$. Under the policy with $s_1 = s_2 = 0$, class 1 fluid level reaches its threshold θ_1 at time ψ_1, and class 2 fluid level reaches its threshold θ_2 at time ψ_2. Note that for any $t \geq \psi_1$, we have

$$\mu_1 h_1(t - \psi_1) \geq \mu_2 h_2(t - \psi_2),$$

since $\psi_2 \geq \psi_1 \geq 0$ and $\mu_1 h_1 > \mu_2 h_2$. Thus, it is more profitable to give fixed priority to class 1 until the class 1 fluid level decreases to its threshold in the low period. If $H < \psi_1$, then again the optimal policy is given by (OPT) upon setting $s_1 = s_2 = 0$ (i.e. giving fixed priority to class 1 in the high period), which yields a total cost of 0.

The following corollary follows from the description of the optimal policy.

Corollary 2. *If*

(i) $Z_1(0) \geq \theta_1$ *or,*

(ii) $Z_1(0) \leq \theta_1$, $Z_1(0) \leq \theta_2$ *and* $0 \leq \psi_1 \leq \psi_2$,

then the policy with

$$\begin{aligned} \forall t \in (0, H) &\qquad \dot{T}_1(t) = 1,\ \dot{T}_2(t) = 0; \\ \forall t \in (H, H+L) &\qquad \textit{Low-Period-Policy} \end{aligned}$$

is optimal for all $H \geq 0$ *and* $L \geq 0$.

Note that if the initial fluid levels satisfy the conditions in (i) or (ii), the policy described in Corollary 2 is optimal even when the length of the high period and the length of the low period are random variables.

2.4.2 The overloaded case

We assume $\rho^h > 1,\ \rho_1^h \leq 1,\ \rho^l \leq 1$ in this section and provide the optimal policy when the value of the high period duration H and the value of the low period duration L are deterministic.

When $\rho^h > 1,\ \rho_1^h \leq 1,\ \rho^l \leq 1$, the optimal policy has the following structure:

$$\begin{aligned} \forall t \in (0,\ s_1): &\qquad \dot{T}_2(t) = 1,\ \dot{T}_1(t) = 0; \\ \forall t \in (s_1,\ s_2): &\qquad \dot{T}_2(t) = \rho_2 - \frac{(\theta_2 - Z_2(s_1))/\mu_2}{a_1(s_1)},\ \dot{T}_1(t) = 1 - \dot{T}_2(t); \\ \forall t \in (s_2,\ s_3): &\qquad \dot{T}_2(t) = 0,\ \dot{T}_1(t) = 1; \end{aligned}$$

$$\forall t \in (s_3,\ H): \qquad \dot{T}_2(t) = 1 - \rho_1^h,\ \dot{T}_1(t) = \rho_1^h;$$

$$\forall t \in (H,\ H+L): \qquad \text{Low-period-policy;}$$

where

$$a_1(s_1) = \frac{(\theta_1 - Z_1(s_1))/\mu_1 + (\theta_2 - Z_2(s_1))/\mu_2}{\rho_1^h + \rho_2 - 1},$$

and s_1, s_2, s_3 are given as

$$\begin{aligned} s_1 &= \max\{t:\ 0 \le t \le H, Z_2(t) \ge \theta_2,\ Z_1(t) \le \theta_1\}, \\ s_2 &= \max\{t:\ s_1 \le t \le H, Z_1(t) \le \theta_1\}, \\ s_3 &= \max\{t:\ s_2 \le t \le H, Z_1(t) \ge \theta_1\}. \end{aligned}$$

with the convention that $\max\{t : x \le t \le y, t \in A\} = x$ if $A = \emptyset$.

2.4.3 The lightly loaded case

We assume $\rho^h \le 1,\ \rho^l \le 1$, then the optimal policy has the following structure:

$$\forall t \in (0, H) \qquad \text{Low-Period-Policy except replace } \rho_1^l \text{ by } \rho_1^h;$$

$$\forall t \in (H, H+L) \qquad \text{Low-Period-Policy.}$$

Remark 3. The policies described in Sections 2.4.2 and 2.4.3 can be implemented without knowing the length of the high and the low periods. Hence, these policies are also optimal when the length of the high period and the length of the low period are random variables.

2.5 Discrete review policies in the stochastic case

Throughout the rest of this chapter, we shall consider the stochastic instance of the fluid model described in Section 2.2. Recall that the system starts with a high period, followed by a low period. The duration of the high period H, and the duration of the low period L are independent random variables. For this stochastic fluid control problem, the optimal policy when $\rho^h > 1,\ \rho_1^h \le 1,\ \rho^l \le 1$ is given in Section 2.4.2 and the optimal policy when $\rho^h \le 1$, $\rho^l \le 1$ is given in Section 2.4.3 (see Remark 3). We therefore focus only on the case when

$$\rho_1^h > 1,\ \rho^l \le 1.$$

To specify the control policy in this case, we shall always consider the following four subcases which were first introduced in Section 2.4 and are summarized below:

$$\text{Case 1:} \qquad Z_1(0) \geq \theta_1, \tag{20}$$

$$\text{Case 2:} \qquad Z_1(0) < \theta_1,\ Z_2(0) > \theta_2, \tag{21}$$

$$\text{Case 3:} \qquad Z_1(0) < \theta_1,\ Z_2(0) \leq \theta_2,\ \psi_1 \leq \psi_2, \tag{22}$$

$$\text{Case 4:} \qquad Z_1(0) < \theta_1,\ Z_2(0) \leq \theta_2,\ \psi_1 \geq \psi_2. \tag{23}$$

In this section, we present a discrete review policy that is asymptotically optimal as the expected length of the high period tends to infinity. Under our discrete review policy, the state of the system is observed at intervals of length τ which is a predetermined positive number. Note that no assumptions are imposed on τ. Given τ, the distribution of the high period and the mean of the low period, the discrete review policy is implemented as follows. Let H_0 and L_0 denote the actual values of the high period and the low period respectively. The state of the system is observed at times $t = 0, \tau, 2\tau, \ldots, M\tau$, where

$$M = \min\{n \in \mathbb{N} : n\tau \geq H_0\}.$$

Note that we do not assume that we know H_0 initially. We assume that the system can detect the end of the high period by observing a sudden drop in the arrival rate of class 1 fluid. At each time t, we observe the fluid level of both classes, i.e., $Z_1(t)$ and $Z_2(t)$. We then predict the remaining high period $\tilde{H}(t)$ and the low period $\tilde{L}(t)$ using one of the methods described below. If $t < M\tau$, we implement the policy described in Section 2.4 from t to $t + \tau$ using $\tilde{H}(t)$ as the length of the high period, $\tilde{L}(t)$ as the length of the low period, and $Z_1(t)$ and $Z_2(t)$ as the initial fluid levels. If $t = M\tau$, we implement the Low-period-policy from t until the end of the low period.

At time t, we either set

$$\tilde{H}(t) = \mathbf{E}[H|H > t] - t, \tag{24}$$

or

$$\tilde{H}(t) = \min\{x \geq 0 : \mathbb{P}(H > x + t|H > t) = p\}, \tag{25}$$

where p will be specified later. Note that in (24) remaining high period is estimated by its expected value, and in (25) remaining high period is set equal to x which guarantees that the probability that the remaining high period is larger than x is p. While implementing the discrete review policy in the numerical examples of Section 2.7, we use both of these methods to estimate the remaining high period and we set $p = 0.25, 0.5$ and 0.75. On the other hand, the remaining low period is always set equal to its mean. Hence, $\tilde{L}(t) = \mathbf{E}[L]$.

We now show in Proposition 4 that our discrete review policy is asymptotically optimal as the expected length of the high period tends to infinity and the proof is provided in Section 2.8.2. Given the actual values of the high and low periods, let $c(H_0, L_0)$ be the holding cost under the optimal policy described in Section 2.4. The closed form expression for $c(H_0, L_0)$ is given in Appendix A. Similarly, let $c^{\mathrm{DR}}(H_0, L_0)$ denote the holding cost under our discrete review policy when the length of the high period is H_0 and the length of the low period is L_0.

Proposition 4. *There exist $D > 0$ and $\beta_1 \geq 0$ (which depend on the arrival rates, service rates, initial fluid levels, threshold values and holding costs per unit time) such that if*

$$\tilde{H}(0) \geq D,$$

then the discrete review policy is equivalent to giving fixed priority to class 1 in the high period, and we have

$$c^{\mathrm{DR}}(H_0, L_0) - c(H_0, L_0) \leq \beta_1 \tag{26}$$

for all $H_0 \geq 0$ and $L_0 \geq 0$.

In the next section, we provide various policies for different parameter sets and initial conditions. We also show that they are asymptotically optimal in certain regime.

2.6 Other policies that are asymptotically optimal

Throughout this section, we assume that $\rho_1^h > 1$ and $\rho^l \leq 1$. We are interested in two heavy traffic regimes. In the first one, the expected length of the high period tends to infinity. In the second one, traffic intensity of class 2 (i.e. ρ_2) tends to $1 - \rho_1^l$ when ρ_1^l is fixed and the low period is infinitely long. Under both these regimes, we are interested in finding the asymptotically optimal policies.

Consider the policy that gives fixed priority to class 1 in the high period and uses the Low-Period-Policy in the low period. For the rest of the paper, we will refer to this policy as FP1. We shall use $c^{FP1}(H_0, L_0)$ to denote the holding cost of the FP1 policy when the length of the high period is H_0 and the length of the low period is L_0. Recall that $c(H_0, L_0)$ denotes the holding cost of the optimal control policy (as specified in Section 2.4) when the lengths of the high and the low periods are known and equal to H_0 and L_0, respectively. Holding cost expressions for all possible values of the high and low periods under the FP1 policy and the optimal policy (as well as other policies considered in this chapter) are given in Appendix A. We have the following proposition and its proof is provided in Section 2.8.2.

Proposition 5. *There exists $\beta_2 \geq 0$, which does not depend on the duration of the high period and low period, such that*

$$c^{\mathrm{FP1}}(H_0, L_0) - c(H_0, L_0) \leq \beta_2.$$

for all $H_0 \geq 0$ and $L_0 \geq 0$.

We next consider the case that the traffic intensify of class 2 tends to $1 - \rho_1^l$ (i.e. the system is always heavily loaded) and the expected length of the low period tends to infinity. Again we consider Cases 1 to 4 given in (20) to (23), separately. We know from Corollary 2 that in Case 1 and Case 3, FP1 policy is optimal. Hence, we only consider Case 2 and Case 4. We start with Case 4.

Definition 6. Assume conditions of Case 4. We define the π^{a_1} policy as follows:

$$\begin{aligned}
&\forall t \in (0, a_1 \wedge H), && \dot{T}_2(t) = \rho_2 - \frac{\theta_2 - Z_2(0)}{a_1 \mu_1}, \quad \dot{T}_1(t) = 1 - \dot{T}_2(t);\\
&\forall t \in (a_1 \wedge H, H), && \dot{T}_2(t) = 0, \quad \dot{T}_1(t) = 1;\\
&\forall t \in (H, H+L), && \text{Low-period-policy.}
\end{aligned}$$

Under Case 4, since initially both class 1 and class 2 fluid levels are below their threshold values, π^{a_1} policy starts with processor sharing. In the processor sharing serving scheme,

$\dot{T}_1(t)$ and $\dot{T}_2(t)$ are chosen such that the time that class 2 fluid level reaches its threshold is delayed while ensuring that the cost accumulated from class 1 in the high period is not too high. Moreover, this choice of $\dot{T}_1(t)$ and $\dot{T}_2(t)$ guarantees that class 1 and class 2 reach their thresholds from below at the same time if H is long enough to do so. Thus, during the processor sharing period, the π^{a_1} policy gives as much proportion of service as possible to class 2 while maintaining class 1 below its threshold. Note that if the traffic intensity in the low period is close to 1 and the low period is long, the holding cost for class 2 fluid in the low period can be high. Hence, it is important to reduce the amount of class 2 fluid at the beginning of the low period without incurring too much cost from class 1 fluid. We will show in Proposition 9 that when $\rho_2 \to 1-\rho_1^l$ and $\mathbf{E}[L] \to \infty$, π^{a_1} is strongly asymptotically optimal under the assumptions of Case 4. We use a notion of strongly asymptotically optimal (as introduced in [35]) in the following sense:

Definition 7. Consider a control problem where the performance measure $J(u, \alpha)$ is a function of the control policy u and parameter α. Let the optimal control policy be $u^*(\alpha)$, and suppose $J(u^*(\alpha), \alpha) \to \infty$ as $\alpha \to \alpha_0$. A control policy $\hat{u}$ is called *strongly asymptotically optimal* if there exists $K < \infty$ such that

$$J(\hat{u}(\alpha), \alpha) - J(u^*(\alpha), \alpha) \leq K, \text{ as } \alpha \to \alpha_0.$$

We will also use the following notation.

Definition 8. For $f : \mathbb{R} \to \mathbb{R}$, we write

$$f(r) \quad = \quad \mathcal{O}(1) \text{ as } r \to r_0$$

to mean that there exists a constant $M > 0$ such that $|f(r)| < M$ as $r \to r_0$.

Let $c^{a_1}(H, L)$ denote the holding cost under policy π^{a_1} when the length of the high period is H and the length of the low period is L. The closed form expression for $c^{a_1}(H, L)$ is given in Appendix A.

Proposition 9. *Assume conditions of Case 4. Suppose H and L are random variables with $\mathbf{E}[H^2] < \infty$. If $\mathbf{E}[L] \to \infty$ and $\rho_2 \to (1-\rho_1^l)$ (where ρ_1^l is fixed), then*

$$\mathbf{E}[c^{a_1}(H, L) - c(H, L)] \quad = \quad \mathcal{O}(1),$$

and π^{a_1} is strongly asymptotically optimal.

The proof of Proposition 9 is provided in Section 2.8.2.

We next consider Case 2 given in (21) and define the following policy.

Definition 10. Assume conditions of Case 2. We define the FP2-FP1 policy as follows:

$$\begin{array}{rccl} \forall t \in (0,H), & \text{if } Z_2(t) > \theta_2,\, Z_1(t) < \theta_1 & \text{then } \dot{T}_2(t) = 1,\, \dot{T}_1(t) = 0; \\ \forall t \in (0,H) & \text{if } Z_2(t) = \theta_2,\, Z_1(t) < \theta_1 & \text{then } \dot{T}_2(t) = \rho_2,\, \dot{T}_1(t) = 1-\rho_2; \\ \forall t \in (0,H) & \text{if } Z_1(t) \geq \theta_1 & \text{then } \dot{T}_2(t) = 0,\, \dot{T}_1(t) = 1; \\ \forall t \in (H,H+L) & & \text{Low-Period-Policy.} \end{array}$$

Note that FP2-FP1 policy is similar to the π^{a_1} policy. However, since initially class 2 fluid is above its threshold level, FP2-FP1 policy starts with giving fixed priority to class 2. Let $c^{\mathrm{FP2-FP1}}(H,L)$ denote the holding cost under the FP2-FP1 policy when the length of the high period is H and the length of the low period is L. The closed form expression for $c^{\mathrm{FP2-FP1}}(H,L)$ is given in Appendix A.

Proposition 11. *Assume conditions of Case 2. Suppose H and L are random variables with $\mathbf{E}[H^2] < \infty$. If $\mathbf{E}[L] \to \infty$ and $\rho_2 \to 1-\rho_1^l$ (where ρ_1^l is fixed), then*

$$\mathbf{E}[c^{\mathrm{FP2-FP1}}(H,L) - c(H,L)] \;=\; \mathcal{O}(1),$$

and FP2-FP1 policy is strongly asymptotically optimal.

The proof of Proposition 11 is provided in Section 2.8.2.

2.7 Numerical results

In this section, we provide numerical examples to demonstrate the performance of the discrete review policy described in Section 2.5 in systems with random high and low periods. Ideally, once the exact lengths of the high and low periods (H and L) are known, one can follow the optimal policy in the deterministic case described in Section 2.4. Recall that $c(H,L)$ denotes the total holding cost under the optimal policy when the lengths of the high and low periods are known. Since one can not observe the true lengths of the either periods until they end, such a policy is not implementable. However, the quantity $\mathbf{E}[c(H, L)]$ can be used as a lower bound of the cost function since no other policy can outperform such a policy with perfect knowledge of H and L. We will use this lower bound (which will be referred as LB) as a guideline to evaluate the performance of other implementable policies.

While implementing the discrete review policy, we use both of the methods given in (24) and (25) to estimate the remaining high period and set $p = 0.25, 0.5$ and 0.75. Recall that the remaining low period is always set equal to its mean. The discrete review policy implemented with the method in (24) (i.e. the remaining high period is set equal to its expected value) will be called DR1, and the discrete review policies implemented with the method given in (25) with $p = 0.25, 0.5$ and 0.75 will be called DR2, DR3, and DR4, respectively. We compare the expected holding cost of these four policies with the lower bound LB, the expected holding cost of the FP1 policy and the expected holding cost of the π^{a_1} policy.

Even though we have considered several systems, in the interest of space we report our findings from two sets of examples referred to as System I and System II respectively. In System I, parameters are set as follow: $\theta_1 = 50$, $\theta_2 = 100$, $h_1 = 2$, $h_2 = 1$, $Z_1(0) = 0$, $Z_2(0) = 90$, $\rho_1^h = 2$, $\rho_1^l = 0.1$ and $\rho_2 = 0.4$. In System II, $\rho_1^l = 0$ and $\rho_2 = 0.95$ and the remaining parameters remain the same. We consider four different distributions (referred to as Case A, Case B, Case C and Case D respectively) for the length of the high (H) and the low (L) periods, In Case A, both H and L are Erlang-2 random variables. In Case B, both H and L are exponential random variables. In Cases C and D, both H and L are hyper-exponential random variables with squared coefficient of variation 2 and 10, respectively. Note that the squared coefficient of variation of the distributions in Case A and Case B are 1/2 and 1, respectively. In our experiments, $\mathbf{E}[H]$ attains the values: 5, 12.5, 25, 37.5 and 50 and $\mathbf{E}[L]$ attains the values: 12.5, 25, 50 and 1000.

Under a specified distribution with fixed values of $\mathbf{E}[H]$ and $\mathbf{E}[L]$, we generate 500,000 sets of H and L values. For each set of H and L values, we compute $c(H, L)$ (lower bound), $c^{FP1}(H, L)$, $c^{a_1}(H, L)$ and the holding costs of the four discrete review policies. We then compute the average holding costs over 500,000 replications. In all our numerical experiments, while implementing the discrete review policies, we set τ equal to 0.1. The value of τ is determined by simulating the systems that we consider under the discrete review policies with different τ values and eventually picking the τ value which yields a good holding cost performance while keeping the run times reasonably short. Tables 1 through 4 display the average value of the lower bound on holding cost and the percentage difference off the lower bound of the average holding cost of the FP1, π^{a_1}, DR1, DR2, DR3 and DR4 policies.

As Tables 1 through 4 show, discrete review policies have a good holding cost performance. The largest percentage difference between the holding cost of discrete review policies and the lower bound on the holding cost is approximately 21%. Moreover, the discrete review policies are more robust than the FP1 and the π^{a_1} policies. Note that the average holding cost under the discrete review policies is much less than the average holding cost under the FP1 policy in Cases A and B when $\mathbf{E}[H]$ is small to moderate. The same result also holds for Case C when $\rho_1^l = 0$ and $\rho_2 = 0.95$. However, as the variability increases, FP1 policy outperforms all other policies. In particular, in Case D the holding cost under the FP1 policy is less than the holding cost under all discrete review policies except when $E[L]$ is large (see Table 4). Discrete review policies outperform π^{a_1} policy in Cases C and D. When the system variability is low, for systems with $\rho_1^l = 0.1$ and $\rho_2 = 0.4$, the discrete review policies outperform the π^{a_1} policy. For systems with $\rho_1^l = 0$ and $\rho_2 = 0.95$, the same observation holds for the DSview1, DSview3 and DSview4 policies. If $\rho_1^l = 0$ and $\rho_2 = 0.95$, DSview2 has higher holding cost than π^{a_1} policy in Cases A and B when $\mathbf{E}[H]$ is small and $\mathbf{E}[L]$ is not large or when $\mathbf{E}[L]$ is large.

In systems with $\rho_1^l = 0.1$ and $\rho_2 = 0.4$, in general DSview4 policy has a poor performance compared to the other discrete review policies. It performs well only for small values of $\mathbf{E}[H]$ in Case A. On the other hand, DSview2 significantly outperforms DSview1 and DSview3 policies in Cases A and B and in Case C when $\mathbf{E}[H]$ is not large. In Case C, as $\mathbf{E}[H]$ increases, DSview1 policy starts dominating the other discrete review policies. On the other hand, in Case D, DSview1 policy always outperforms the other discrete review policies in systems with $\rho_1^l = 0.1$ and $\rho_2 = 0.4$. The same assertion holds for systems with $\rho_1^l = 0$ and

$\rho_2 = 0.95$ except when $E[L]$ and $E[H]$ are both large (see Table 4).

In systems with $\rho_1^l = 0$ and $\rho_2 = 0.95$, the performances of DR2 and DR4 policies depend on the expected length of the low period. Even though the DR4 policy shows poor performance (compared to the other discrete review policies) when $\mathbf{E}[L]$ is small, its performance improves (in particular in Cases A and B) as $\mathbf{E}[L]$ gets large. On the other hand, even though DR2 policy has one of the best performances among the discrete review policies when $\mathbf{E}[L]$ is small, its performance deteriorates in Cases A and B as $\mathbf{E}[L]$ gets large. However, in Cases C and D, DR1 and DR2 policies always have better holding cost performance than the other discrete review policies.

In conclusion, discrete review policies yield good holding cost performance and they are robust with respect to the system parameters. Among the discrete review policies, one can employ the DR2 policy (in order to reduce the total holding cost) if class 2 is not heavily loaded and the coefficient of variation of the high and the low periods is not large. However, if the coefficient of variation of the high and the low periods is large, DR1 policy seems to outperform the other discrete review policies. On the other hand, if class 2 is heavily loaded, DR1 policy has a good overall policy.

2.8 Proof of the desired results

In this section, we provide the proof of the desired results in the earlier sections. We first prove in Section 2.8.1 the optimality of policies provide in Section 2.4. In Section 2.8.2, we show the asymptotic optimality of the discrete review policies provided in Section 2.5 and we also provide enough details for the proof of the asymptotic optimality proposed policies in Section 2.6.

2.8.1 Proof of the optimality of the policies in the deterministic case

In this section, we provide the detailed proof of the optimality of the policies proposed in Section 2.4. We first develop the lemmas needed and then we prove the desired result.

To prove the optimality of the policies given in Section 2.4, we first provide a lemma related to the Pontryagin maximum principle. Originally, this lemma was given in Seierstad and Sydsaeter [32] but the version stated here was tailored for our problem. For completeness, we also provide the proof of the lemma.

Consider an optimal control problem as follows,

$$\max \int_{B_0}^{B_1} f^0(x(t), u(t), t)\, dt \tag{27}$$

such that

$$\dot{x}(t) = f(x(t), u(t), t), \tag{28}$$

$$x(B_0) = x_0, \tag{29}$$

$$x(B_1) \geq x_1, \tag{30}$$

$$u(t) \in \mathsf{U} \text{ where } \mathsf{U} \subset \mathbb{R}^r \text{ and } (x(t), u(t)) \in \mathbb{R}^n \times \mathbb{R}^r, \tag{31}$$

where $f^0(x(t), u(t), t)$, and $f(x(t), u(t), t)$ are continuous functions of t over $[B_0,\ B_1]$ except at finite number of points.

We say that $(x(t), u(t))$ is an *admissible pair* if $x(t)$ is absolutely continuous, and $(x(t), u(t))$ satisfies (28) to (31). We want to find an optimal admissible pair $(x(t),\ u(t))$ that maximizes integral in (27). In the following lemma, for vectors a and b, $a \cdot b$ denotes the usual inner product of a and b.

Lemma 12. *Let $(\bar{x}(t), \bar{u}(t))$ be an admissible pair for the problem given in (27) to (31). Suppose there exists a continuous function $p(t) = (p_1(t),\ p_2(t), \ldots, p_n(t))$ on $[B_0,\ B_1]$ such that it has a piecewise continuous derivative $\dot{p}(t)$, the continuity of $\dot{p}(t)$ is violated only at finite number of points, and $p(t)$ satisfies*

$$p_i(B_1) \geq 0, \ \textit{and}\ p_i(B_1) = 0 \ \textit{if}\ \bar{x}_i(B_1) > x_1^i,\ \forall i = 1, \ldots, n. \tag{32}$$

In addition, the Hamiltonian function

$$H(x(t), u(t), p(t), t) = f^0(x(t), u(t), t) + p(t) \cdot f(x(t), u(t), t) \tag{33}$$

satisfies the following

$$H(\bar{x}(t), \bar{u}(t), p(t), t) - H(x(t), u(t), p(t), t) \geq \dot{p}(t) \cdot (x(t) - \bar{x}(t)) \tag{34}$$

for all admissible pairs $(x(t), u(t))$, for all $t \in [B_0, B_1]$ except at finite number of points. Then $(\bar{x}(t), \bar{u}(t))$ is an optimal pair for problem (27) to (31).

Proof of Lemma 12. We use Δ to denote the following

$$\Delta = \int_{B_0}^{B_1} f^0(\bar{x}(t), \bar{u}(t), t)\mathrm{d}t - \int_{B_0}^{B_1} f^0(x(t), u(t), t)\, dt.$$

Then the optimality of $(\bar{x}(t), \bar{u}(t))$ is equivalent to $\Delta \geq 0$ for all admissible pairs $(x(t), u(t))$.

According to (33) we have

$$\begin{aligned}\Delta \ &= \int_{B_0}^{B_1} [H(\bar{x}(t), \bar{u}(t), p(t), t) - H(x(t), u(t), p(t), t)]\, dt \\ &\quad + \int_{B_0}^{B_1} p(t) \cdot [f(x(t), u(t), t) - f(\bar{x}(t), \bar{u}(t), t)]\, dt.\end{aligned}$$

It then follows from (28) and (34) that

$$\Delta \ \geq \int_{B_0}^{B_1} \dot{p}(t) \cdot [x(t) - \bar{x}(t)]\, dt + \int_{B_0}^{B_1} p(t) \cdot [\dot{x}(t) - \dot{\bar{x}}(t)]\, dt.$$

Assume that $B_0 = \xi_0 < \xi_1 < \cdots \xi_k < \xi_{k+1} = B_1$, are all the possible discontinuity points of $\dot{p}(t),\ \dot{x}(t)$ and $\dot{\bar{x}}(t)$. So the right hand side of the above inequality can be written as

$$\sum_{i=0}^{k} \Big\{ \int_{\xi_i}^{\xi_{i+1}} \dot{p}(t) \cdot [x(t) - \bar{x}(t)]\mathrm{d}t + \int_{\xi_i}^{\xi_{i+1}} p(t) \cdot [\dot{x}(t) - \dot{\bar{x}}(t)\, dt] \Big\}$$

$$
\begin{aligned}
&= \sum_{i=0}^{k} \int_{\xi_i}^{\xi_{i+1}} \frac{d}{dt}[p(t) \cdot (x(t) - \bar{x}(t))] \\
&= \sum_{i=0}^{k} \Big[p(\xi_{i+1}) \cdot (x(\xi_{i+1}) - \bar{x}(\xi_{i+1})) - p(\xi_i) \cdot (x(\xi_i) - \bar{x}(\xi_i)) \Big] \\
&= p(B_1) \cdot (x(B_1) - \bar{x}(B_1)) \\
&\geq 0,
\end{aligned}
$$

where the last equality is due to the continuity of $p(t), x(t), \bar{x}(t)$ and (29), and the last inequality is based on (30) and (32). Hence, $\Delta \geq 0$, and the optimality of $(\bar{x}(t), \bar{u}(t))$ is proven. $\square$

We next prove that the policy specified in Section 3 is optimal for our original problem described in Section 2.2 with deterministic high and low periods. First, replacing $\dot{T}_i(t)$ by $u_i(t)$, notice that our original control problem is equivalent to

$$
\begin{aligned}
\max \quad & \int_0^{H+L} \sum_{i=1}^{2} -h_i \left(Z_i(t) - \theta_i\right)^+ \, dt. && (35) \\
\text{such that} \quad & \dot{Z}_i(t) = \lambda_i(t) - \mu_i u_i(t) \qquad i = 1, 2 && (36) \\
& Z_i(t) \geq 0 \qquad \forall t \in [0, H+L], \quad i = 1, 2 && (37) \\
& u_i(t) \geq 0 \qquad \forall t \in [0, H+L], \quad i = 1, 2 && (38) \\
& u_1(t) + u_2(t) \leq 1 \qquad \forall t \in [0, H+L], && (39)
\end{aligned}
$$

where $\lambda_1(t) = \lambda_1^h, \ \forall t \in (0, H)$, and $\lambda_1(t) = \lambda_1^l, \ \forall t \in (H, H+L)$, and $\lambda_2(t) = \lambda_2, \ \forall t \in (0, H+L)$.

Hereafter, we are going to use $u^*(t)$ to denote the proposed policy given in Section 3, and $Z^*(t)$ to denote the fluid level under this policy.

Based on Lemma 12, in order to prove the optimality of (Z^*, u^*), it suffices to construct continuous functions $p_i(t), \ i = 1, 2$, with piecewise continuous derivatives such that $(Z^*(t), u^*(t), p(t))$ satisfies (32) and (34). In what follows, we illustrate the basic idea of the construction and proof by focusing on only one special case in Section 3. Notice that other cases can be proved similarly.

2.8.1.1 *Proof of the optimality for the highly overloaded case*

Before introducing our construction of p's, we first describe the fluid level evolution of both classes under the policy u^* specified in Section 3.1.

Notice that under the policy u^*, class 1 will have higher priority starting from time s_2 until time t in the low period such that $Z_1^*(t) \leq \theta_1$. Corresponding to this policy, we define two critical time instances for class 1 as follow

$$t_1 = \max\{t : s_2 \leq t \leq H,\, Z_1^*(t) \leq \theta_1\}, \tag{40}$$

$$t_2 = \max\{t : H \leq t \leq H + L,\, Z_1^*(t) \geq \theta_1\}, \tag{41}$$

where t_1 is the time that class 1 increases to its threshold from below in the high period if the duration of high period is long enough and t_2 is the time that class 1 decreases to its threshold from above in the low period if the duration of the low period is long enough.

Similarly, we define two critical time instances for class 2

$$\tilde{s}_2 = \max\{t : s_2 \leq t \leq t_2,\, Z_2^*(t) \leq \theta_2\}, \tag{42}$$

$$\tilde{t}_2 = \max\{t : t_2 \leq t \leq H + L,\, Z_2^*(t) \geq \theta_2\}, \tag{43}$$

where $\tilde{s}_2$ is the time that class 2 increases to its threshold from below during the time interval that class 1 has higher priority, i.e. during interval $[s_2, t_2]$ and $\tilde{t}_2$ is the time that class 2 decreases to its threshold from above in the low period if the duration of the low period is long enough. Note that after class 1 decreases to its threshold from above in the low period at t_2, the Low-period-policy gives enough capacity to class 2 to decrease class 2 fluid level.

Based on the definition of $s_1,\ s_2$ (described in Section 2.4) and the definition of $t_1,\ t_2, \tilde{s}_2,\ \tilde{t}_2$, we claim the following holds:

Claim 1:

$$s_1 \leq s_2 \leq t_1 \leq H \leq t_2 \leq H + L,$$

$$s_1 \leq s_2 \leq \tilde{s}_2 \leq t_2 \leq \tilde{t}_2 \leq H + L,$$

Claim 2:

$$\forall t \in (0,\, s_1) \qquad Z_1^*(t) < \theta_1,\, Z_2^*(t) > \theta_2,$$

$$\forall t \in (s_1,\, s_2) \qquad Z_1^*(t) < \theta_1,\; Z_2^*(t) \le \theta_2,$$

$$\begin{aligned}
\forall t \in (s_2,\, t_1) &\qquad Z_1^*(t) < \theta_1,\\
\forall t \in (t_1,\, t_2) &\qquad Z_1^*(t) > \theta_1,\\
\forall t \in (t_2,\, H+L) &\qquad Z_1^*(t) \le \theta_1,
\end{aligned}$$

$$\begin{aligned}
\forall t \in (s_2,\, \tilde{s}_2) &\qquad Z_2^*(t) < \theta_2,\\
\forall t \in (\tilde{s}_2,\, \tilde{t}_2) &\qquad Z_2^*(t) > \theta_2,\\
\forall t \in (\tilde{t}_2,\, H+L) &\qquad Z_2^*(t) \le \theta_2.
\end{aligned}$$

For ease of readability, we defer the proof of the claims to the end and next show how to construct the auxiliary functions $p(t)$.

It follows from the Pontryagin maximal principle that the optimal policy has to satisfy $\dot{p}_i(t) = \frac{\partial}{\partial Z_i} H(Z(t), p(t), t)$ at the differentiable points, where the Hamiltonian function is given by

$$H(Z(t), u(t), p(t), t) = \sum_{i=1}^{2} \left(-h_i \left(Z_i(t) - \theta_i \right)^+ + p_i(t)(\lambda_i(t) - \mu_i u_i(t)) \right). \tag{44}$$

We therefore construct $p_i(t)$, $i = 1, 2$ (in a backward fashion) as follows:

$$\begin{aligned}
&\qquad p_i(H+L) = 0;\; i = 1, 2,\\
\forall t \in (\tilde{t}_2,\, H+L): &\qquad \dot{p}_1(t) = 0,\; \dot{p}_2(t) = 0,\\
\forall t \in (t_2,\, \tilde{t}_2): &\qquad \dot{p}_1(t) = \frac{\mu_2 h_2}{\mu_1},\; \dot{p}_2(t) = h_2,
\end{aligned}$$

$$\begin{aligned}
\forall t \in (t_1, t_2): &\qquad \dot{p}_1(t) = h_1,\\
\forall t \in (s_2,\, t_1): &\qquad \dot{p}_1(t) = 0,
\end{aligned}$$

$$\begin{aligned}
\forall t \in (\tilde{s}_2,\, t_2): &\qquad \dot{p}_2(t) = h_2,\\
\forall t \in (s_2,\, \tilde{s}_2): &\qquad \dot{p}_2(t) = 0,
\end{aligned}$$

$$\forall t \in (s_1,\, s_2): \qquad \dot{p}_1(t) = 0;\, \dot{p}_2(t) = 0,$$

$$\forall t \in (0,\, s_1): \qquad \dot{p}_1(t) = 0;\, \dot{p}_2(t) = h_2.$$

Based on the above construction, we have the following properties stated as Claim 3, whose proof is also deferred to the end of this section.

Claim 3:

$$\forall t \in (t_2,\, H+L): \qquad \mu_1 p_1(t) = \mu_2 p_2(t) \le 0;$$

$$\forall t \in (s_2,\, t_2): \qquad \mu_1 p_1(t) < \mu_2 p_2(t) \le 0;$$

$$\forall t \in (s_1,\, s_2): \qquad \mu_1 p_1(t) = \mu_2 p_2(t) \le 0;$$

$$\forall t \in (0,\, s_1): \qquad 0 \ge \mu_1 p_1(t) > \mu_2 p_2(t).$$

Based on Lemma 12, the optimality follows once we show that $(Z^*(t), u^*(t), p(t))$ satisfies (32) and (34). From the construction of $p_i(t)$, (32) holds immediately. It remains to show that (34) holds in each time interval throughout $(0, H+L)$ under all four cases given in (20) to (23). Here, we focus only on Case 2.1 to illustrate the basic idea. The other cases can be proved similarly.

Consider, for example, the first time interval $(0, s_1)$. The policy in this period is $u_1^*(t) = 0$, $u_2^*(t) = 1$, and from Claim 2 we have $Z_1^*(t) < \theta_1$, $Z_2^*(t) > \theta_2$. Note that no other admissible policy can reduce more class 2 fluid level than u^*, thus under any admissible policy $u_i(t)$, the fluid level will satisfy $Z_1(t) < \theta_1$ and $Z_2(t) > \theta_2$ for $t \in (0, s_1)$. Plugging this in (44), we have the left hand side of (34) equal to

$$h_2(Z_2(t) - Z_2^*(t)) + \sum_{i=1}^{2} -\mu_i p_i(t)(u_i^*(t) - u_i(t)).$$

Based on Claim 3, for all t in $(0, s_1)$, we have $-\mu_2 p_2(t) \ge -\mu_1 p_1(t) \ge 0$. Therefore,

$$\sum_{i=1}^{2} -\mu_i p_i(t)(u_i^*(t) - u_i(t)) \ge -\mu_1 p_1(t)(u_1^*(t) + u_2^*(t) - u_1(t) - u_2(t)).$$

Note that $u_1^*(t) + u_2^*(t) = 1$, and the admissible $u_i(t)$, $i = 1, 2$ satisfies $u_1(t) + u_2(t) \le 1$, so the right hand side of the above inequality is non-negative. It follows immediately that (34) holds for all time t in the interval $(0, s_1)$.

Repeating this procedure for the remaining intervals, we can similarly prove that (34) holds for all time t in $(0, H+L)$. Hence the optimality of the proposed policy is guaranteed.

We now prove the three claims we made earlier. Again, we focus only on Case 2.1 to illustrate the basic idea. The other cases can be proved similarly.

• **Proof for Claim 1 and Claim 2 in Case 2.1.** Recall that in Case 2.1, we assume that $Z_1^*(0) < \theta_1$, $Z_2^*(0) > \theta_2$, and condition (10) holds.

In this case, s_1 and s_2 are solved using the equations given in (11) to (19). Simultaneously, we also compute u_1, u_2, t_1 and t_2. They can all be expressed in terms of initial fluid levels $Z_i^*(0)$, $i = 1, 2$, durations of the high and low periods H and L, the arrival rates λ_1^h, λ_1^l, and λ_2, service rates $\mu_i, i = 1, 2$, and holding cost rates h_i, $i = 1, 2$.

Since $Z_2^*(0) > \theta_2$ and $\rho_2 < 1$ (i.e $\lambda_2 < \mu_2$), it follows from (11) that $s_1 > 0$ (s_1 is the time that class 2 decreases to its threshold when it has higher priority). Since $Z_2^*(s_1) = \theta_2$, it follows from (13) that $u_2 = \rho_2 > 0$. Hence, from (15) $u_1 = 1 - \rho_2 > 0$. One can check that the requirement $t_2 \leq H + L$ is equivalent to $L \geq \gamma_1(H - a_1)$. In addition, $t_1 \leq H \leq t_2$ is equivalent to $a_1 \leq H$, and $s_1 \leq s_2$ is equivalent to $H \leq B$. So, in Case 2.1 of Section 3.1, condition (10) guarantees that we have $0 \leq s_1 \leq s_2 \leq t_1 \leq H \leq t_2 \leq H + L$ and $u_1 > 0$, $u_2 > 0$.

Under the proposed policy, we know that $\lambda_1^h > \mu_1$. Hence, the fluid level $Z_1^*(t)$ increases in the interval $(0, H)$ and $Z_1^*(0) < \theta_1$ and $Z_1^*(t_1) = \theta_1$ (see (16)). Thus, for any $t \in (0,\, t_1)$, we know that $Z_1^*(t) < \theta_1$ and for any $t \in (t_1,\, H)$, $Z_1^*(t) > \theta_1$. Under the proposed policy, in the low period, the fluid level $Z_1^*(t)$ decreases until it hits its threshold at t_2 (see (18)). Hence, for any $t \in (t_1,\, t_2)$, $Z_1^*(t) > \theta_1$. Then we can see that t_1 and t_2 obtained from the set of equations of Case 2.1 coincide with their definitions given in (40) and (41). Hence, the first inequality of claim 1 holds. From the definition of $\tilde{s}_2$ and $\tilde{t}_2$, we can immediately see that the second inequality of claim 1 also holds.

We now prove Claim 2. While proving Claim 1, we have already shown that $Z_1^*(t)$ satisfies the inequalities in Claim 2 for all $t < t_2$. Since $\lambda_2 < \mu_2$ and $u_2 = \rho_2$, under the proposed policy, $Z_2^*(t)$ decreases in the interval $(0,\, s_1)$, until it reaches θ_2 at s_1 (see (11)).

It is kept at its threshold θ_2 in the interval (s_1, s_2) since $\lambda_2 = \mu_2 u_2$. Then it increases in the interval (s_2, H) since class 1 has higher priority. Since $Z_1^*(t) > \theta_1$ in the interval (H, t_2), under the proposed Low-period-policy, class 1 still has higher priority and class 2 fluid continues to increase until class 1 fluid decreases to its threshold at t_2. Hence,

$$\begin{aligned} \forall t \in (0, s_1), &\qquad Z_2^*(t) > \theta_2,\, Z_2^*(s_1) = \theta_2, \\ \forall t \in (s_1,\, s_2), &\qquad Z_2^*(t) = \theta_2,\, Z_2^*(s_2) = \theta_2, \\ \forall t \in (s_2,\, t_2), &\qquad Z_2^*(t) > \theta_2, Z_2^*(t_2) \geq \theta_2. \end{aligned}$$

After t_2, under the proposed Low-period-policy, if $Z_2^*(t_2) > \theta_2$, then class 1 fluid is going to be kept at its threshold by setting $u_1^*(t) = \rho_1^l$, and class 2 fluid is going to decrease by holding service capacity at $u_2^*(t) = 1 - \rho_1^l > \rho_2$ until class 2 fluid reaches its threshold from above at $\tilde{t}_2$ (see the definition of $\tilde{t}_2$ given in (43)). After $\tilde{t}_2$, $u_1^*(t) > \rho_1^l$ and $u_2^*(t) > \rho_2$. So, fluid levels of both classes are going to decrease and are maintained below their thresholds. Hence,

$$\begin{aligned} \forall t \in (t_2,\, \tilde{t}_2), &\qquad Z_2^*(t) > \theta_2,\, Z_1^*(t) = \theta_1, \\ \forall t \in (\tilde{t}_2,\, H + L), &\qquad Z_2^*(t) \leq \theta_2,\, Z_1^*(t) \leq \theta_1. \end{aligned}$$

This completes the proofs of Claims 1 and 2.

- **Proof for Claim 3 in Case 2.1.** From the proofs of Claims 1 and 2, we know that in this case $\tilde{s}_2 = s_2$.

From the construction of $p_i(t)$, $i = 1, 2$, we know that they are piecewise linear functions. To compare their values, it is sufficient to compare them at the end points of each interval. Since $p_i(H + L) = 0$ and $\dot{p}_i(t) \geq 0$ at all differentiable points, we know $p_i(t) \leq 0$, $i = 1, 2$, for all $t \in [0, H + L]$. Note that since $p_1(H + L) = p_2(H + L) = 0$ and $\mu_1 \dot{p}_1(t) = \mu_2 \dot{p}_2(t)$ for $t \in (t_2, H + L)$, we have $\mu_1 p_1(t) = \mu_2 p_2(t)$ for $t \in [t_2,\, H + L]$. Based on the derivatives, we then have

$$\forall t \in [t_1,\, t_2], \qquad \mu_i p_i(t) = \mu_i p_i(t_2) + \mu_i h_i(t - t_2),\;\; i = 1, 2.$$

Using the fact that $\mu_1 h_1 > \mu_2 h_2$, $\mu_1 p_1(t_2) = \mu_2 p_2(t_2)$ and noting $t - t_2 < 0$ for $t \in (t_1, t_2)$, we have

$$\forall t \in (t_1,\, t_2), \qquad \mu_2 p_2(t) > \mu_1 p_1(t).$$

Based on the derivatives of $p(t)$, we have

$$\begin{aligned} &\forall t \in [s_2,\, t_1], \qquad \mu_1 p_1(t) = \mu_1 p_1(t_1), \\ &\forall t \in [s_2,\, t_2], \qquad \mu_2 p_2(t) = \mu_2 p_2(t_2) + \mu_2 h_2 (t - t_2). \end{aligned}$$

From (19) and $\mu_1 p_1(t_2) = \mu_2 p_2(t_2)$, we have $\mu_1 p_1(s_2) = \mu_2 p_2(s_2)$. Combining this with $\mu_1 p_1(t_1) \leq \mu_2 p_2(t_1)$, we have

$$\forall t \in (s_2,\, t_1), \qquad \mu_1 p_1(t) \leq \mu_2 p_2(t).$$

From $\mu_1 p_1(s_2) = \mu_2 p_2(s_2)$ and $\dot{p}_i(t) = 0$, $i = 1, 2$, for $t \in (s_1,\, s_2)$, we can immediately see that

$$\forall t \in [s_1,\, s_2], \qquad \mu_1 p_1(t) = \mu_2 p_2(t) = \mu_2 p_2(s_2).$$

For $t \in (0, s_1)$, based on the derivatives of $p(t)$, we have

$$\begin{aligned} &\forall t \in [0,\, s_1], \qquad \mu_2 p_2(t) = \mu_2 p_2(s_1) + \mu_2 h_2 (t - s_1), \\ &\forall t \in [0,\, s_1], \qquad \mu_1 p_1(t) = \mu_1 p_1(s_1). \end{aligned}$$

Note that $\mu_i p_i(t)$ has the same value at s_1 for $i = 1, 2$ and for $t \in (0,\, s_1)$, $\dot{p}_2(t) = h_2 > 0 = \dot{p}_1(t)$, then we have

$$\forall t \in (0,\, s_1), \qquad \mu_1 p_1(t) > \mu_2 p_2(t).$$

This completes the proof of Claim 3.

2.8.1.2 Proof of the optimality for the overloaded case

We will only construct the auxiliary function $p_i(t)$, $i = 1, 2$. To complete the proof of (34), one only needs to go through the routine procedure as described in Section 2.8.1.1. We define t_2, $\tilde{s}_2$ and $\tilde{t}_2$ in the same way as in (41), (42) and (43) but now they are defined under the policy given in Section 3.2. According to the definition of the break points s_i, $i = 1, 2, 3$, $\tilde{s}_2$, t_2, and $\tilde{t}_2$, we can specify the fluid level evolution for each time interval, and the derivatives of $p_i(t)$, $i = 1, 2$. In the equations given below, if the right hand side of an interval is not strictly larger than the left side of the interval, then that interval does not exist but this does not affect our definition of the derivatives of $p_i(t)$ and the fluid level description $Z_i^*(t)$ for $i = 1, 2$. We have

$$\begin{aligned}
\forall t \in (0,\, s_1): &\quad Z_1^*(t) < \theta_1,\ Z_2^*(t) > \theta_2,\ \dot{p}_1(t) = 0,\ \dot{p}_2(t) = h_2,\\
\forall t \in (s_1,\, s_2): &\quad Z_1^*(t) < \theta_1,\ Z_2^*(t) \le \theta_2,\ \dot{p}_1(t) = 0,\ \dot{p}_2(t) = 0,
\end{aligned}$$

$$\begin{aligned}
\forall t \in (s_2,\, \tilde{s}_2): &\quad Z_2^*(t) < \theta_2,\ \dot{p}_2(t) = 0,\\
\forall t \in (\tilde{s}_2,\, \tilde{t}_2): &\quad Z_2^*(t) > \theta_2,\ \dot{p}_2(t) = h_2,
\end{aligned}$$

$$\begin{aligned}
\forall t \in (s_2,\, s_3): &\quad Z_1^*(t) > \theta_1,\ \dot{p}_1(t) = h_1,\\
\forall t \in (s_3,\, H): &\quad Z_1^*(t) = \theta_1,\ \dot{p}_1(t) = \mu_2 \dot{p}_2(t)/\mu_1,\\
\forall t \in (H,\, t_2): &\quad Z_1^*(t) > \theta_1,\ \dot{p}_1(t) = h_1,\\
\forall t \in (t_2,\, \tilde{t}_2): &\quad Z_1^*(t) = \theta_1,\ \dot{p}_1(t) = \mu_2 \dot{p}_2(t)/\mu_1,
\end{aligned}$$

$$\forall t \in (\tilde{t}_2,\, H + L): \quad Z_1^*(t) \le \theta_1,\ Z_2^*(t) \le \theta_2,\ \dot{p}_i(t) = 0,\ i = 1, 2,$$

and we let $p_i(H + L) = 0$, $i = 1, 2$. Thus, we can construct continuous and piecewise linear functions $p_i(t)$, $i = 1, 2$ which have the specified derivatives in each interval and satisfy (32).

2.8.1.3 Proof of the optimality for the lightly loaded case

As in the proof of the optimality of the policies given in Sections 3.1 and 3.2, the proof involves constructing the functions $p_i(t)$, $i = 1, 2$ based on the Pontryagin maximal principle and is omitted.

2.8.2 Proof of the asymptotic optimality of the policies in the stochastic case

In this section, we provide the proof of the asymptotic optimality of the polices proposed in Section 2.5 and Section 2.6. We first prove the that the discrete review policies proposed in Section 2.5 is asymptotically optimal as provided by Proposition 4. Then we prove the results of Proposition 5, Proposition 9 and Proposition 11 in Section 2.6.

We first prove that the discrete review policies are asymptotically optimal when the expected high period goes to infinity as provided in Proposition 4.

Proof of Proposition 4. We provide the proof for the discrete review policy where $\tilde{H}(t)$ is calculated based on the method given in (24). The proof for the discrete review policy implemented with the method given in (25) is similar.

With a slight abuse of notation, we use $d_i(t)$ and $\psi_i(t)$, $i = 1, 2$ to denote the quantities defined in (6) and (7) at time t when fluid levels are $Z_i(t)$, $i = 1, 2$. Similarly, let $a_i(t), i = 1, 2$ denote the corresponding quantities given in (8) at time t. Hence, $d_i(0) = d_i$, $\psi_i(0) = \psi_i$ and $a_i(0) = a_i$ for $i = 1, 2$. Let

$$D = \max\Big\{a_2(0),\ \ \psi_1(0) + \frac{\eta}{1-\eta}(\psi_1(0) - \psi_2(0))\Big\}. \tag{45}$$

We first show by induction that for all $0 \le n \le M - 1$, the discrete review policy sets $\dot{T}_1(t) = 1$, $\dot{T}_2(t) = 0$ for all $t \in [n\tau, (n + 1)\tau)$. Hence the discrete review policy is equivalent to giving fixed priority to class 1 in the high period $[0, H_0)$.

First consider $t = 0$. Note that for Case 1 and Case 3, it follows immediately from Corollary 2 that the discrete review policy gives fixed priority to class 1, i.e. $\dot{T}_1(t) = 1$, $\dot{T}_2(t) = 0$ for all $t \in [0, \tau)$.

For Case 2, note that $\psi_2 = \psi_2(0) \le 0$, then $D \ge a_2$ and $D \ge \psi_1 + \eta(1-\eta)^{-1}\psi_1 = (1-\eta)^{-1}\psi_1$. Hence, $\tilde{H}(t) \ge D$ (where D is given in (45)) which implies that the condition of Case 2.6 in Section 2.4.1 is satisfied, where the discrete review policy gives fixed priority to class 1 in the interval $[0, \tau)$.

For Case 4, $\tilde{H}(t) \ge D$ (where D is given in (45)) which implies that the condition of Case 4.4 in Section (2.4.1) is satisfied, where the discrete review policy gives fixed priority to class 1 in $[0, \tau)$.

Therefore the claim is true for $n = 0$. Now assume that under the discrete review policy fixed priority is given to class 1 until $t = n\tau$ for $1 \le n \le M - 1$. Then the fluid levels of the two classes at time $t = n\tau$ are $Z_1(n\tau) = Z_1(0) + n\tau(\lambda_1^h - \mu_1)$, and $Z_2(n\tau) = Z_2(0) + n\tau\lambda_2$, respectively. It is easily checked from (6),(7) and (8) that

$$\psi_1(n\tau) = \psi_1(0) - n\tau, \quad \psi_2(n\tau) = \psi_2(0) - n\tau, \quad a_2(n\tau) = a_2(0) - n\tau.$$

To specify the discrete review policy at time $t = n\tau$, we again consider Cases 1 to 4 given in (20) to (23) separately. Note that the conditions of these four cases should now be evaluated at time $t = n\tau$ based on $Z_i(n\tau)$ and $\psi_i(n\tau)$, $i = 1, 2$.

Again under Case 1 and Case 3, Corollary 2 applies, hence, the discrete review policy sets $\dot{T}_1(t) = 1$, $\dot{T}_2(t) = 0$ and gives fixed priority to class 1 for all $t \in [n\tau, (n+1)\tau)$.

Under Case 2, since $\tilde{H}(0) = \mathbf{E}[H]$ and

$$\tilde{H}(n\tau) \ge \mathbf{E}[H] - n\tau \ge D - n\tau = \max\{a_2(n\tau), \ \ \psi_1(n\tau) + \frac{\eta}{1-\eta}(\psi_1(n\tau) - \psi_2(n\tau))\}, \quad (46)$$

it follows from Case 2.6 in Section 2.4.1 that the discrete review policy gives fixed priority to class 1 in the interval $[n\tau, (n+1)\tau)$.

Similarly, for Case 4, (46) implies that conditions of Case 4.4 in Section (2.4.1) hold, hence the discrete review policy gives fixed priority to class 1 in the interval $[n\tau, (n+1)\tau)$.

This then completes the induction and we therefore conclude that the discrete review policy sets $\dot{T}_1(t) = 1$, $\dot{T}_2(t) = 0$ for all $0 \le t \le H_0$. The result in (26) then follows from Proposition 5 in Section 2.6. $\square$

Remark 13. The proof for other methods are the same except $\mathbf{E}[H]$ is replaced by $\tilde{H}(0)$ in (46).

Next,we prove that the FP1 policy is asymptotically optimal when the expected high period goes to infinity as stated in Proposition 5.

Proof of Proposition 5. We need to consider the holding costs under Cases 1 to 4 separately. Note that for Case 1 and Case 3, Corollary 2 applies and the optimal policy is FP1, hence we can take $\beta_2 = 0$ for these two cases.

Now consider Case 4. Note that the optimal policy (as described in Section 2.4) is the same as the FP1 policy in Case 4.4, and differs from FP1 only under Cases 4.1, 4.2 and 4.3. Thus, the two costs differ only when (H_0, L_0) belongs to the regions considered in Cases 4.1, 4.2 and 4.3. Our proof involves providing an upper bound on the difference between the holding costs of the FP1 policy and the optimal policy. In the interest of space, we only derive this upper bound when (H_0, L_0) is in the region given in Case 4.1. However, as it will become clear from our analysis below, this will lead to subcases. Since the computation of the upper bound for these subcases is similar, we only provide the analysis when (H_0, L_0) satisfies (48) below.

We start by computing the holding cost expression for the FP1 policy. Under the FP1 policy, for any $t \in (0, H_0)$, we have $Z_1(t) = Z_1(0) + (\lambda_1^h - \mu_1)t$, and $Z_1(\psi_1) = \theta_1$ if $H_0 \geq \psi_1$. We consider the sample paths such that conditions of Case 4.1 and $H_0 \geq \psi_1$ are both satisfied. Thus, we have

$$\psi_1 \leq H_0 \leq a_2,\ L_0 \geq \gamma_1(H_0 - a_1). \tag{47}$$

For (H_0, L_0) such that (47) is satisfied, we can specify the fluid level evolution under the FP1 policy. Class 1 fluid level increases to its threshold value at ψ_1 and stays above its threshold until it decreases to its threshold value in the low period. Let t_2' denote the time that the fluid level of class 1 decreases to its threshold value in the low period. Then $Z_1(\psi_1) + (\lambda_1^h - \mu_1)(H_0 - \psi_1) + (\lambda_1^l - \mu_1)(t_2' - H_0) = \theta_1$. Since $Z_1(\psi_1) = \theta_1$, we obtain

$$t_2' = H_0 + (1 - \rho_1^l)^{-1}(\rho_1^h - 1)(H_0 - \psi_1).$$

Note that the conditions of Case 4.1 imply that $t_2' \leq H_0 + L_0$. Thus, the fluid level of class 1 can decrease to its threshold in the low period. On the other hand, since before t_2' class 2 is not served, its fluid level increases at rate λ_2 and reaches its threshold value at ψ_2. Conditions of (47) imply that $H_0 \geq \psi_2$. Hence, the fluid level of class 2 is above its threshold in the interval (ψ_2, t_2'). After t_2', class 2 fluid level decreases at rate $\mu_2(1-\rho_1^l) - \lambda_2$. Let $\tilde{t}_2'$ denote the time that class 2 decreases to its threshold value in the low period. Then

$$Z_2(\psi_2) + \lambda_2(t_2' - \psi_2) + (\lambda_2 - \mu_2(1 - \rho_1^l))(\tilde{t}_2' - t_2') = \theta_2.$$

Since $Z_2(\psi_2) = \theta_2$, we get

$$\tilde{t}_2' = t_2' + (1 - \rho_2 - \rho_1^l)^{-1} \rho_2 (t_2' - \psi_2).$$

In order to have $\tilde{t}_2' \le H_0 + L_0$, we need $L_0 \ge \gamma_4 (H_0 - a_1)$, where $\gamma_4 = (1 - \rho_2 - \rho_1^l)^{-1}(\rho_1^h + \rho_2 - 1)$. Thus, we consider sample paths such that

$$\psi_1 \le H_0 \le a_2,\ L_0 \ge \gamma_4 (H_0 - a_1) \tag{48}$$

and specify the fluid level evolution of class 1 and class 2 as

$$\begin{aligned}
\text{if } t \in [0, \psi_1], &\qquad Z_1(t) = Z_1(0) + (\lambda_1^h - \mu_1)t \le \theta_1,\\
\text{if } t \in (\psi_1, H_0], &\qquad Z_1(t) = Z_1(\psi_1) + (\lambda_1^h - \mu_1)(t - \psi_1) > \theta_1,\\
\text{if } t \in [H_0, t_2'), &\qquad Z_1(t) = Z_1(H_0) + (\lambda_1^l - \mu_1)(t - H_0) > \theta_1,\\
\text{if } t \in [t_2', H_0 + L_0], &\qquad Z_1(t) \le \theta_1,
\end{aligned}$$

and

$$\begin{aligned}
\text{if } t \in [0, \psi_2], &\qquad Z_2(t) = Z_2(0) + \lambda_2 t \le \theta_2,\\
\text{if } t \in (\psi_2, t_2'], &\qquad Z_2(t) = Z_2(\psi_2) + \lambda_2 (t - \psi_1) > \theta_2,\\
\text{if } t \in (t_2', \tilde{t}_2'), &\qquad Z_2(t) = Z_1(t_2') + (\lambda_2 - \mu_2(1 - \rho_1^l))(t - t_2') > \theta_2,\\
\text{if } t \in [\tilde{t}_2', H_0 + L_0], &\qquad Z_2(t) \le \theta_2.
\end{aligned}$$

So, we can calculate the holding cost under the FP1 policy for (H_0, L_0) satisfying (48) as

$$\begin{aligned}
\int_0^{H_0+L_0} \sum_{i=1}^{2} h_i (Z_i(t) - \theta_i)^+ \, dt &= \frac{1}{2} h_1 \mu_1 \Big((\rho_1^h - 1)(H_0 - \psi_1)^2 + (1 - \rho_1^l)(t_2' - H_0)^2 \Big)\\
&\quad + \frac{1}{2} h_2 \mu_2 \Big(\rho_2 (t_2' - \psi_2)^2 + (1 - \rho_2 - \rho_1^l)(\tilde{t}_2' - t_2')^2 \Big).
\end{aligned}$$

Plugging in the expressions of t_2' and $\tilde{t}_2'$, we obtain

$$\begin{aligned}
& c^{FP1}(H_0, L_0) - c(H_0, L_0) \quad \le \quad c^{FP1}(H_0, L_0)\\
= \ & \frac{1}{2} h_2 \mu_2 \Big\{ \frac{(\rho_1^h - 1)(\rho_1^h - \rho_1^l)}{\eta (1 - \rho_1^l)} (H_0 - \psi_1)^2\\
& + \frac{(1 - \rho_2 - \rho_1^l)^2}{\rho_2} \Big[\frac{\rho_1^h + \rho_2 - 1}{1 - \rho_1^l - \rho_2} (H_0 - a_1) - \frac{(\rho_1^h - 1)}{(1 - \rho_1^l)} (H_0 - \psi_1) \Big]^2\\
& + (1 - \rho_1^l - \rho_2) \Big[\frac{\rho_1^h + \rho_2 - 1}{1 - \rho_1^l - \rho_2} (H_0 - a_1) - \frac{(\rho_1^h - 1)}{(1 - \rho_1^l)} (H_0 - \psi_1) \Big]^2 \Big\}.
\end{aligned}$$

Since $H_0 \leq a_2$, $c^{\text{FP1}}(H_0, L_0) - c(H_0, L_0)$ is bounded when (H_0, L_0) satisfies (48). Hence, the holding cost under FP1 policy differs from the holding cost of the optimal policy by a constant. This completes the proof when (H_0, L_0) satisfies (48). Expressions in Appendix A illustrate that the difference between the holding costs of the FP1 policy and the optimal policy is also bounded by a constant for other values of the high and the low periods (i.e. when (H_0, L_0) does not satisfy (48)).

The proof for Case 2 is similar and thus omitted. □

Next, we show that the π^{a_1} policy is asymptotically optimal when the traffic intensity of class 2 increases as stated in Proposition 9.

Proof of Proposition 9. Similar to the proof of Proposition 5, we obtain an upper bound on the difference between the holding costs of the π^{a_1} policy and the optimal policy for each possible value of H and L. In the interest of space, we only consider the values of H and L that satisfy the conditions of Case 4.4. However, as it will become clear from our analysis below, this will lead to subcases. Since the computation of the upper bound for these subcases is similar, we only provide the analysis when (H, L) satisfies (53) below.

If H and L belong to the region given in Case 4.4, the optimal policy is the same as the FP1 policy, which corresponds to $s_1 = s_2 = 0$ (see Section 2.4.1). We start with computing the holding cost under the optimal policy and the π^{a_1} policy when H and L belong to the region of Case 4.4. Under the optimal policy, even though class 1 receives full capacity, its fluid level increases in the high period. Let t_1 denote the time that fluid level of class 1 reaches its threshold θ_1 in the high period. Then we can solve for t_1 which is equal to ψ_1 in this case. Note that the conditions of Case 4.4, in particular $H \geq a_2$ and $\psi_1 \geq \psi_2$ imply that $H \geq \psi_1$. The fluid level of class 1 continues to increase after ψ_1 during the high period, and it is above its threshold at the beginning of the low period. Under the Low-Period-Policy, class 1 still has full service capacity. Let t_2 denote the time that the fluid level of class 1 decreases to its threshold θ_1 in the low period. Then

$$Z_1(\psi_1) + (\lambda_1^h - \mu_1)(H - \psi_1) + (\lambda_1^l - \mu_1)(t_2 - H) = \theta_1,$$

where $Z_1(\psi_1) = \theta_1$ and we can compute t_2 as

$$t_2 = (\rho_1^h - 1)(1 - \rho_1^l)^{-1}(H - \psi_1) + H. \tag{49}$$

Note that $t_2 \leq H + L$ implies that $L \geq \gamma_3(H - \psi_1)$. Thus, we consider sample paths such that (H, L) satisfies both the conditions of Case 4.4 and $L \geq \gamma_3(H - \psi_1)$, i.e

$$H \geq a_2, \quad H + L \geq \psi_1 + \frac{\eta}{1 - \eta}(\psi_1 - \psi_2), \quad L \geq \gamma_3(H - \psi_1),$$

which is equivalent to

$$H \geq a_2, \quad L \geq \gamma_3(H - \psi_1). \tag{50}$$

If H and L satisfy (50), the evolution of class 1 fluid under the optimal policy is as follows

$$\begin{aligned}
\text{if } t \in [0, \psi_1], &\qquad Z_1(t) = Z_1(0) + (\lambda_1^h - \mu_1)t \leq \theta_1, \\
\text{if } t \in (\psi_1, H), &\qquad Z_1(t) = \theta_1 + (\lambda_1^h - \mu_1)(t - \psi_1) > \theta_1, \\
\text{if } t \in [H, t_2), &\qquad Z_1(t) = Z_1(H) + (\lambda_1^h - \mu_1)(t - H) > \theta_1, \\
\text{if } t \in [t_2, H + L], &\qquad Z_1(t) \leq \theta_1,
\end{aligned}$$

where t_2 is given in (49). The holding cost incurred by class 1 is given as

$$\int_0^{H+L} h_1(Z_1(t) - \theta_1)^+ \, dt = \frac{1}{2} h_1 \mu_1 \frac{(\rho_1^h - 1)(\rho_1^h - \rho_1^l)}{(1 - \rho_1^l)}(H - \psi_1)^2. \tag{51}$$

Next we compute the holding cost incurred by class 2 under the optimal policy when H and L satisfy (50). Under the optimal policy, class 2 is not served during the high period and not served in the low period until class 1 fluid level decreases to its threshold. Hence, class 2 is not served until t_2. Therefore, the fluid level of class 2 increases until t_2. Let $\tilde{t}_1$ denote the time that the fluid level of class 2 increases to its threshold. We can compute $\tilde{t}_1$ as $\tilde{t}_1 = \psi_2$. Note that conditions in (50) imply that class 2 increases to its threshold in the high period and reaches its threshold earlier than class 1. After t_2, fluid level of class 2 begins to decrease at rate $\mu_2(1 - \rho_1^l) - \lambda_2$ under the Low-Period-Policy. Let $\tilde{t}_2$ denote the time that class 2 decreases to its threshold in the low period. Then

$$Z_2(\psi_2) + \lambda_2(t_2 - H) + (\lambda_2 - \mu_2(1 - \rho_1^l))(\tilde{t}_2 - t_2) = \theta_2,$$

where $Z_2(\psi_2) = \theta_2$ and we have

$$\tilde{t}_2 = t_2 + \frac{\rho_2}{1-\rho_2-\rho_1^l}(t_2 - H), \tag{52}$$

where t_2 is given in (49). Note that $\tilde{t}_2 \leq H + L$ implies that $L \geq \gamma_4(H - a_1)$.

Thus, we consider the sample paths such that H and L satisfy both (50) and $L \geq \gamma_4(H - a_1)$, which is equivalent to

$$H \geq a_2, \quad L \geq \gamma_4(H - a_1). \tag{53}$$

Now we can specify the evolution of class 2 fluid under the optimal policy if (H, L) satisfies (53). That is

$$\begin{aligned}
\text{if } t \in [0, \psi_2], &\qquad Z_2(t) = Z_2(0) + \lambda_2 t \leq \theta_2, \\
\text{if } t \in (\psi_2, t_2), &\qquad Z_2(t) = \theta_2 + \lambda_2(t - \psi_2) > \theta_2, \\
\text{if } t \in (t_2, \tilde{t}_2), &\qquad Z_2(t) = Z_2(t_2) + (\lambda_2 - \mu_2(1-\rho_1^l))(t - t_2) > \theta_2, \\
\text{if } t \in [\tilde{t}_2, H+L], &\qquad Z_2(t) \leq \theta_2,
\end{aligned}$$

where t_2 and $\tilde{t}_2$ are given in (49) and (52), respectively. The holding cost incurred by class 2 under the optimal policy if (H, L) satisfies (53) is given as

$$\int_0^{H+L} h_2(Z_2(t) - \theta_2)^+ \, dt = \frac{1}{2} h_2 \mu_2 \Big(\rho_2 (t_2 - \psi_2)^2 + (1 - \rho_2 - \rho_1^l)(\tilde{t}_2 - t_2)^2 \Big). \tag{54}$$

Plugging in the expressions of t_2 and $\tilde{t}_2$ and using the fact that $h_2\mu_2 = \eta h_1 \mu_1$, the sum of (51) and (54) is equal to

$$\begin{aligned}
c(H, L) \;\; = \;\; & \frac{1}{2} h_2 \mu_2 \left\{ \frac{(\rho_1^h - 1)(\rho_1^h - \rho_1^l)}{\eta(1 - \rho_1^l)} (H - \psi_1)^2 \right. \\
& + \frac{(1 - \rho_2 - \rho_1^l)^2}{\rho_2} \Big[\frac{\rho_1^h + \rho_2 - 1}{1 - \rho_1^l - \rho_2} (H - a_1) - \frac{(\rho_1^h - 1)}{(1 - \rho_1^l)} (H - \psi_1) \Big]^2 \\
& \left. + (1 - \rho_1^l - \rho_2) \Big[\frac{\rho_1^h + \rho_2 - 1}{1 - \rho_1^l - \rho_2} (H - a_1) - \frac{(\rho_1^h - 1)}{(1 - \rho_1^l)} (H - \psi_1) \Big]^2 \right\}.
\end{aligned} \tag{55}$$

The expression in (55) yields the lower bound for the holding cost if (H, L) satisfies (53).

Next we calculate the holding cost under the π^{a_1} policy when (H, L) belongs to the region in (53). Note that $H \geq a_2$ and $\psi_1 \geq \psi_2 \geq 0$ imply that $H \geq a_1$. According to the π^{a_1} policy, we know that both classes share the service capacity until a_1 as specified in Definition 6. Since $a_1 \leq H$, class 1 fluid increases before a_1. Since the service speed for class 2 is slower than its arrival rate under the π^{a_1} policy before a_1, class 2 fluid level also increases before a_1. Moreover, we can calculate that the fluid level of each class at a_1 is equal to its threshold value, i.e $Z_i(a_1) = \theta_i$ for $i = 1, 2$. From a_1 to H, class 1 has higher priority and gets the full service capacity. However, since $\rho_1^h > 1$, the fluid level of class 1 continues to increase after a_1 and reaches its highest level at the end of the high period. Afterwards, under the Low-Period-Policy, the fluid level of class 1 decreases. Let t_2' denote the time that the fluid level of class 1 decreases to its threshold value in the low period. Then

$$Z_1(a_1) + (\lambda_1^h - \mu_1)(H - a_1) + (\lambda_1^l - \mu_1)(t_2' - H) = \theta_1,$$

where $Z_1(a_1) = \theta_1$ and from the above equation we can solve for t_2' as

$$t_2' = \frac{\rho_1^h - 1}{1 - \rho_1^l}(H - a_1) + H. \tag{56}$$

Since $t_2' \leq H + L$, $L \geq (\rho_1^h - 1)(1 - \rho_1^l)^{-1}(H - a_1)$. Note that for every (H, L) that satisfies (53), this condition is satisfied. That is if (H, L) belongs to the region in (53), the fluid level of class 1 decreases to its threshold before the low period is over. Then we can specify the evolution of class 1 fluid under the π^{a_1} policy as

$$\begin{aligned}
\text{if } t \in [0, a_1], &\qquad Z_1(t) = Z_1(0) + (\lambda_1^h - \mu_1 a_1)t \leq \theta_1, \\
\text{if } t \in (a_1, H), &\qquad Z_1(t) = \theta_1 + (\lambda_1^h - \mu_1)(t - a_1) > \theta_1, \\
\text{if } t \in [H, t_2'), &\qquad Z_1(t) = Z_1(H) + (\lambda_1^h - \mu_1)(t - H) > \theta_1, \\
\text{if } t \in [t_2', H + L], &\qquad Z_1(t) \leq \theta_1,
\end{aligned}$$

where t_2' is given in (56). The holding cost incurred by class 1 under the π^{a_1} policy is equal to

$$\int_0^{H+L} h_1(Z_1(t) - \theta_1)^+ \, dt = \frac{1}{2} h_1 \mu_1 \frac{(\rho_1^h - 1)(\rho_1^h - \rho_1^l)}{1 - \rho_1^l}(H - a_1)^2. \tag{57}$$

Finally, we specify the evolution of class 2 fluid under the π^{a_1} policy. Note that the fluid level of class 2 increases to its threshold level at a_1 and class 2 is not served in the interval (a_1, t_2').

Hence, the fluid level of class 2 is above its threshold value in the interval (a_1, t'_2). After t'_2, class 2 is served at the speed $\mu_2(1-\rho_1^l)$ under the Low-Period-Policy. Since $\mu_2(1-\rho_1^l) > \lambda_2$, the fluid level of class 2 begins to decrease after t'_2 and reaches its threshold at some point in the low period denoted by $\tilde{t}'_2$. Then

$$Z_2(a_1) + \lambda_2(t'_2 - a_1) + (\lambda_2 - \mu_2(1-\rho_1^l))(\tilde{t}'_2 - t'_2) = \theta_2,$$

where $Z_2(a_1) = \theta_2$. We can solve the above equation for $\tilde{t}'_2$ and compute

$$\tilde{t}'_2 = t'_2 + \frac{\rho_2}{1-\rho_2-\rho_1^l}(t'_2 - a_1). \tag{58}$$

Since $\tilde{t}'_2 \le H + L$, $L \ge \gamma_4(H - a_1)$. For each sample path such that (H, L) satisfies (53), the fluid level of class 2 decreases to its threshold before the low period is over. Now we can specify the evolution of class 2 fluid as

$$\begin{aligned}
\text{if } t \in [0, a_1], &\qquad Z_2(t) = Z_2(0) + (\lambda_2 - \mu_2 a_1)t \le \theta_2,\\
\text{if } t \in (a_1, t'_2), &\qquad Z_2(t) = \theta_2 + \lambda_2(t - a_1) > \theta_2,\\
\text{if } t \in (t'_2, \tilde{t}'_2), &\qquad Z_2(t) = Z_2(t'_2) + (\lambda_2 - \mu_2(1-\rho_1^l))(t - t'_2) > \theta_2,\\
\text{if } t \in [\tilde{t}'_2, H+L], &\qquad Z_2(t) \le \theta_2,
\end{aligned}$$

where t'_2 and $\tilde{t}'_2$ are given in (56) and (58), respectively. The holding cost incurred by class 2 under the π^{a_1} policy when (H, L) satisfies (53) is equal to

$$\begin{aligned}
\int_0^{H+L} h_2(Z_2(t)-\theta_2)^+\,dt &= \frac{1}{2}h_2\mu_2\Big\{\rho_2(t'_2-a_1)^2 + (1-\rho_2-\rho_1^l)(\tilde{t}'_2 - t'_2)^2\Big\}\\
&= \frac{1}{2}h_2\mu_2\frac{\rho_2(\rho_1^h-\rho_1^l)^2}{(1-\rho_1^l)(1-\rho_1^l-\rho_2)}(H-a_1)^2.
\end{aligned} \tag{59}$$

Summing (57) and (59), we obtain the total holding cost under π^{a_1} when (H, L) satisfies (53) as

$$c^{a_1}(H, L) = \frac{1}{2}h_2\mu_2\left(\frac{(\rho_1^h-1)(\rho_1^h-\rho_1^l)}{\eta(1-\rho_1^l)} + \frac{\rho_2(\rho_1^h-\rho_1^l)^2}{(1-\rho_1^l)(1-\rho_1^l-\rho_2)}\right)(H-a_1)^2. \tag{60}$$

We can now compute the difference between the holding costs of the optimal policy and the π^{a_1} policy. Subtracting (55) from (60), we have

$$c^{a_1}(H, L) - c(H, L)$$

$$
\begin{aligned}
= & \frac{1}{2}\mu_2 h_2 \left\{ \left[\frac{(\rho_1^h - 1)(\rho_1^h - \rho_1^l)}{\eta(1-\rho_1^l)} \right] [(H-a_1)^2 - (H-\psi_1)^2] \right. \\
& + \frac{\rho_2(\rho_1^h - \rho_1^l)^2}{(1-\rho_1^l)(1-\rho_1^l-\rho_2)}(H-a_1)^2 - \frac{\rho_2(\rho_1^h-1)^2}{(1-\rho_1^l)(1-\rho_1^l-\rho_2)}(H-\psi_1)^2 \\
& \left. - \frac{2\rho_2(\rho_1^h-1)}{1-\rho_2-\rho_1^l}(H-\psi_2)(H-\psi_1) - \frac{\rho_2(1-\rho_1^l)}{1-\rho_2-\rho_1^l}(H-\psi_2)^2 \right\}. \qquad (61)
\end{aligned}
$$

First consider the last three terms in (61). Factoring out $\rho_2[(1-\rho_1^l)(1-\rho_1^l-\rho_2)]^{-1}$, we can combine them into

$$
-\frac{\rho_2}{(1-\rho_1^l)(1-\rho_1^l-\rho_2)}[(\rho_1^h-1)(H-\psi_1)+(1-\rho_1^l)(H-\psi_2)]^2.
$$

Adding this value to the second term in (61) and taking the common factor $\rho_2[(1-\rho_1^l)(1-\rho_1^l-\rho_2)]^{-1}$ out, we can combine all the terms with $(1-\rho_2-\rho_1^l)$ in the denominator into

$$
\begin{aligned}
& \frac{\rho_2}{(1-\rho_1^l)(1-\rho_1^l-\rho_2)} \left\{ [(\rho_1^h-\rho_1^l)(H-a_1)+(\rho_1^h-1)(H-\psi_1)+(1-\rho_1^l)(H-\psi_2)] \right. \\
& \left. \times[(\rho_1^h-\rho_1^l)(H-a_1)-(\rho_1^h-1)(H-\psi_1)-(1-\rho_1^l)(H-\psi_2)] \right\}. \qquad (62)
\end{aligned}
$$

From the definitions of a_1, ψ_1, and ψ_2, we know that $a_1 = ((\rho_1^h-1)\psi_1+\rho_2\psi_2)(\rho_1^h+\rho_2-1)^{-1}$. Plugging in this expression of a_1, we can further simplify the expression in the second line of (62) as

$$
-\frac{(\rho_1^h-1)(1-\rho_2-\rho_1^l)}{(\rho_1^h+\rho_2+1)}(\psi_1-\psi_2).
$$

Thus, we have

$$
\begin{aligned}
& c^{a_1}(H,L) - c(H,L) \\
= & \frac{1}{2}\mu_2 h_2 \left\{ \frac{(\rho_1^h-1)(\rho_1^h-\rho_1^l)}{\eta(1-\rho_1^l)}(2H-a_1-\psi_1)(\psi_1-a_1) \right. \\
& \left. - \frac{\rho_2(\rho_1^h-\rho_1^l)(\rho_1^h-1)}{(1-\rho_1^l)(\rho_1^h+\rho_2-1)}(2H-a_1-\frac{(\rho_1^h-1)\psi_1+(1-\rho_1^l)\psi_2}{\rho_1^h-\rho_1^l})(\psi_1-\psi_2) \right\} \\
\leq & \frac{1}{2}h_2\mu_2 \left\{ \frac{(\rho_1^h-1)(\rho_1^h-\rho_1^l)}{\eta(1-\rho_1^l)}(2H-a_1-\psi_1)(\psi_1-a_1) \right\},
\end{aligned}
$$

where the inequality follows from the fact that the second term is not positive since $0 \leq \psi_2 \leq a_1 \leq \psi_1 \leq a_2 \leq H$. At the same time, since $0 \leq \psi_1 - a_1 \leq H$ and $0 \leq (2H-a_1-\psi_1) \leq 2H$, we obtain

$$
\begin{aligned}
c^{a_1}(H,L) - c(H,L) \leq & \frac{1}{2}h_2\mu_2 \left(\frac{(\rho_1^h-1)(\rho_1^h-\rho_1^l)}{\eta(1-\rho_1^l)} \right) 2H^2 \\
= & \frac{1}{2}h_1\mu_1 \left(\frac{(\rho_1^h-1)(\rho_1^h-\rho_1^l)}{(1-\rho_1^l)} \right) 2H^2, \qquad (63)
\end{aligned}
$$

where the equality follows from the definition of η. Since $E[H^2] \leq \infty$, we have the desired result.

□

Next, we provide the proof of Proposition 11.

Proof of Proposition 11. We again compare the holding cost under the optimal policy and the FP2-FP1 policy for each possible value of H and L. In particular, as in the proof of Proposition 9, we obtain upper bounds on the difference between the holding costs of the FP2-FP1 policy and the optimal policy. In the interest of space, we only consider the values of H and L that satisfy the conditions of Case 2.6. However, as it will become clear from our analysis below, this will lead to subcases. Since the computation of the upper bound for these subcases is similar, we only provide the analysis when (H, L) satisfies (67) below and $\tilde{\psi}_1 \leq \tilde{\psi}_2$.

If H and L belong to the region given in Case 2.6, the optimal policy is the same as the FP1 policy, which corresponds to $s_1 = s_2 = 0$ (see Section 2.4.1). Note that in this case $Z_2(0) \geq \theta_2$ and $Z_1(0) \leq \theta_1$, hence $\psi_2 \leq 0 \leq \psi_1$. Let t_1 again denote the time that the fluid level of class 1 increases to its threshold in the high period under the optimal policy. Then $Z_1(t_1) = Z_1(0) + (\lambda_1^h - \mu_1)t_1 = \theta_1$. Hence, $t_1 = \psi_1$, and the condition of Case 2.6, in particular $H \geq a_2$, guarantees $\psi_1 \leq H$. Similar to the analysis in the proof of Proposition 9, in the interval (t_1, H), the fluid level of class 1 continues to increase and reaches its highest level at the end of the high period and we have $Z_1(H) \geq \theta_1$. In the low period class 1 still has the higher priority and its fluid level starts to decrease. If the low period lasts long enough, the fluid level of class 1 decreases to its threshold at some point in the low period. Let t_2 denote the time that the fluid level of class 1 decreases to its threshold. Then

$$Z_1(t_1) + (\lambda_1^h - \mu_1)(H - t_1) + (\lambda_1^l - \mu_1)(t_2 - H) = \theta_1.$$

Note that since $t_1 = \psi_1$ and $Z_1(t_1) = Z_1(\psi_1) = \theta_1$, we have

$$t_2 = H + \frac{\rho_1^h - 1}{1 - \rho_1^l}(H - \psi_1) = H + \gamma_3(H - \psi_1). \tag{64}$$

In order to have $t_2 \leq H + L$, we need $L \geq \gamma_3(H - \psi_1)$. Thus, we consider sample paths such that (H, L) satisfies both the conditions of Case 2.6 and $L \geq \gamma_3(H - \psi_1)$, i.e

$$H \geq a_2, \quad H + L \geq (1 - \eta)^{-1}\psi_1, \quad L \geq \gamma_3(H - \psi_1),$$

which is equivalent to

$$H \geq a_2, \quad L \geq \gamma_3(H - \psi_1). \tag{65}$$

Now we can specify the evolution of class 1 fluid which is

$$\begin{aligned}
\text{if } t \in [0, \psi_1], &\qquad Z_1(t) = Z_1(0) + (\lambda_1^h - \mu_1)t \leq \theta_1, \\
\text{if } t \in (\psi_1, H), &\qquad Z_1(t) = \theta_1 + (\lambda_1^h - \mu_1)(t - \psi_1) > \theta_1, \\
\text{if } t \in [H, t_2), &\qquad Z_1(t) = Z_1(H) + (\lambda_1^h - \mu_1)(t - H) > \theta_1, \\
\text{if } t \in [t_2, H + L], &\qquad Z_1(t) \leq \theta_1,
\end{aligned}$$

where t_2 is given in (64). We can calculate the holding cost incurred by class 1 under the optimal policy for each sample path such that (H, L) satisfies (65) and it is in fact the same as the one given in (57).

Now we analyze the evolution of class 2 fluid under the optimal policy when (H, L) satisfies (65). From the optimal policy, class 2 is not served before class 1 decreases to its threshold in the low period. Since the initial fluid level of class 2 is above its threshold under the conditions of Case 2, it remains above its threshold until t_2. After t_2, class 2 is served at the speed of $\mu_2(1 - \rho_1^l)$ and its fluid level begins to decrease. If the low period lasts long enough, the fluid level of class 2 decreases to its threshold value at some point in the low period, denoted by $\tilde{t}_2$. Then

$$Z_2(0) + \lambda_2 t_2 + (\lambda_2 - \mu_2(1 - \rho_1^l))(\tilde{t}_2 - t_2) = \theta_2.$$

Solving the above equation for $\tilde{t}_2$ and plugging in the expression of ψ_2 given in (7), we have

$$\tilde{t}_2 = t_2 + \rho_2(1 - \rho_2 - \rho_1^l)^{-1}(t_2 - \psi_2), \tag{66}$$

where t_2 is given in (64). Since $\tilde{t}_2 \leq H + L$, $L \geq \gamma_4(H - a_1)$.

Now we consider sample paths such that (H, L) satisfies both (65) and $L \geq \gamma_4(H - a_1)$. Thus, (H, L) satisfies

$$H \geq a_2, \quad L \geq \gamma_4(H - a_1). \tag{67}$$

For each sample path such that (H, L) satisfies (67), the evolution of class 2 fluid can be specified according to the optimal policy as follows

$$\begin{aligned}
\text{if } t \in (0, t_2), &\qquad Z_2(t) = Z_2(0) + \lambda_2 t > \theta_2, \\
\text{if } t \in (t_2, \tilde{t}_2), &\qquad Z_2(t) = Z_2(t_2) + (\lambda_2 - \mu_2(1 - \rho_1^l))(t - t_2) > \theta_2, \\
\text{if } t \in [\tilde{t}_2, H + L], &\qquad Z_2(t) \leq \theta_2,
\end{aligned}$$

where t_2 and $\tilde{t}_2$ are given in (64) and (66), respectively. Then the holding cost incurred by class 2 under the optimal policy for each sample path with (H, L) satisfying (67) is equal to

$$\int_0^{H+L} h_2(Z_2(t) - \theta_2)^+ \, dt = \frac{1}{2} h_2 \mu_2 \Big\{ \rho_2 (t_2 - \psi_2)^2 - \rho_2 \psi_2^2 + (1 - \rho_2 - \rho_1^l)(\tilde{t}_2 - t_2)^2 \Big\}. \tag{68}$$

Summing (57) and (68) and plugging in the expressions of t_2 and $\tilde{t}_2$, we have

$$\begin{aligned}
c(H, L) \;=\; & \frac{1}{2} h_2 \mu_2 \left\{ \frac{(\rho_1^h - 1)(\rho_1^h - \rho_1^l)}{\eta(1 - \rho_1^l)} (H - \psi_1)^2 - \rho_2 (\psi_2^-)^2 \right. \\
& + \frac{(1 - \rho_2 - \rho_1^l)^2}{\rho_2} \Big[\frac{\rho_1^h + \rho_2 - 1}{1 - \rho_1^l - \rho_2} (H - a_1) - \frac{(\rho_1^h - 1)}{(1 - \rho_1^l)} (H - \psi_1) \Big]^2 \\
& \left. + (1 - \rho_1^l - \rho_2) \Big[\frac{\rho_1^h + \rho_2 - 1}{1 - \rho_1^l - \rho_2} (H - a_1) - \frac{(\rho_1^h - 1)}{(1 - \rho_1^l)} (H - \psi_1) \Big]^2 \right\}.
\end{aligned} \tag{69}$$

Now we analyze the fluid level evolution under FP2-FP1 policy when (H, L) satisfies (67). Recall that under the conditions of Case 2, $Z_2(0) > \theta_2$ and $Z_1(0) < \theta_1$. Under the FP2-FP1 policy, class 2 has higher priority before the fluid level of class 1 increases to θ_1 and class 2 decreases to its threshold θ_2. Let t_1' ($\tilde{t}_1'$) be the time that the fluid level of class 1 (class 2) increases (decreases) to θ_1 (θ_2) when class 2 has higher priority. Then

$$Z_1(0) + \lambda_1^h t_1' = \theta_1, \quad Z_2(0) + (\lambda_2 - \mu_2)\tilde{t}_1 = \theta_2,$$

and we have $t_1' = \tilde{\psi}_1$ and $\tilde{t}_1' = \tilde{\psi}_2$. We first consider the case that $\tilde{\psi}_1 \leq \tilde{\psi}_2$, i.e the fluid level of class 2 is still above its threshold θ_2 while the fluid level of class 1 increases to its threshold θ_1. According to the FP2-FP1 policy, class 1 has higher priority if its fluid level is above its threshold value. Note that the above equation is valid only if $H \geq t_1'$, i.e $H \geq \tilde{\psi}_1$. One can verify that if (H, L) satisfies (67), then $H \geq \tilde{\psi}_1$. Therefore, for any (H, L) that satisfies (67) and $\tilde{\psi}_1 \leq \tilde{\psi}_2$, under the FP2-FP1 policy, class 2 has higher priority before $\tilde{\psi}_1$, its fluid decreases before $\tilde{\psi}_1$, and is still above its threshold θ_2 at $\tilde{\psi}_1$. On the other hand, class 1 is not served before $\tilde{\psi}_1$, its fluid level increases before $\tilde{\psi}_1$, and reaches its threshold θ_1 at $\tilde{\psi}_1$. Note that since $\tilde{\psi}_1 < H$, the fluid level of class 1 increases even when it is served with higher priority. Under the FP2-FP1 policy, class 1 has higher priority before its fluid level decreases to its threshold θ_1 which can only happen in the low period. Let t_2' be the time that the fluid level of class 1 decreases to its threshold θ_1, then

$$Z_1(\tilde{\psi}_1) + (\lambda_1^h - \mu_1)(H - \tilde{\psi}_1) + (\lambda_1^l - \mu_1)(t_2' - H) = \theta_1.$$

Under the FP2-FP1 policy $Z_1(\tilde{\psi}_1) = \theta_1$. We can solve the above equation and obtain

$$t_2' = H + (\rho_1^h - 1)(1 - \rho_1^l)^{-1}(H - \tilde{\psi}_1) = H + \gamma_3(H - \tilde{\psi}_1). \tag{70}$$

If the fluid level of class 1 decreases to its threshold before the low period is over, $t_2' \leq H + L$. Thus, we need $L \geq \gamma_3(H - \tilde{\psi}_1)$. But for any sample path with (H, L) satisfying (67) and $\tilde{\psi}_1 \leq \tilde{\psi}_2$, $L \geq \gamma_3(H - \tilde{\psi}_1)$ holds.

If (H, L) satisfies (67) and $\tilde{\psi}_1 \leq \tilde{\psi}_2$, the evolution of class 1 fluid is given as

$$\begin{aligned}
\text{if } t \in [0, \tilde{\psi}_1], \qquad & Z_1(t) = Z_1(0) + \lambda_1^h t \leq \theta_1, \\
\text{if } t \in (\tilde{\psi}_1, H), \qquad & Z_1(t) = Z_1(\tilde{\psi}_1) + (\lambda_1^h - \mu_1)(t - a_1) = \theta_1 + (\lambda_1^h - \mu_1)(t - a_1) > \theta_1, \\
\text{if } t \in [H, t_2'), \qquad & Z_1(t) = Z_1(H) + (\lambda_1^h - \mu_1)(t - H) > \theta_1, \\
\text{if } t \in [t_2', H + L], \qquad & Z_1(t) \leq \theta_1,
\end{aligned}$$

where t_2' is in by (70). The holding cost incurred by class 1 is equal to

$$\int_0^{H+L} h_1(Z_1(t) - \theta_1)^+ \, dt = \frac{(\rho_1^h - 1)(\rho_1^h - \rho_1^l)}{1 - \rho_1^l}(H - \tilde{\psi}_1)^2. \tag{71}$$

We now consider the evolution of class 2 fluid when (H, L) satisfies (67) and $\tilde{\psi}_1 \leq \tilde{\psi}_2$. Recall that class 2 has higher priority before $\tilde{\psi}_1$ and its fluid level is still above its threshold at time $\tilde{\psi}_1$ when class 1 starts receiving higher priority. Before the fluid level of class 1 decreases to its threshold θ_1, class 2 is not served and its fluid level begins to increase until t'_2 (where t'_2 is the time that the fluid level of class 1 decreases to its threshold in the low period). After t'_2, class 2 is served at the speed of $\mu_2(1-\rho_1^l)$. If class 2 continues to be served at this speed, its fluid level decreases to its threshold at some time in the low period, denoted by $\tilde{t}'_2$. Then

$$Z_2(0) + (\lambda_2 - \mu_2)\tilde{\psi}_1 + \lambda_2(t'_2 - \tilde{\psi}_1) + (\lambda_2 - \mu_2(1-\rho_1^l))(\tilde{t}'_2 - t'_2) = \theta_2.$$

¿From this equation, we can get

$$\tilde{t}'_2 = t'_2 + \Big((1-\rho_2)(\tilde{\psi}_2 - \tilde{\psi}_1) + \rho_2(t'_2 - \tilde{\psi}_1)\Big)(1-\rho_2-\rho_1^l)^{-1}, \tag{72}$$

where t'_2 is given in (70). For class 2 fluid to decrease to its threshold level in the low period, we need to have $\tilde{t}'_2 \leq H + L$, which requires that $L \geq \gamma_4(H - a_1)$.

Then the evolution of class 2 fluid under FP2-FP1 policy with (H, L) satisfying (67) and $\tilde{\psi}_1 \leq \tilde{\psi}_2$ is given as

$$\begin{aligned}
\text{if } t \in [0, \tilde{\psi}_1], &\qquad Z_2(t) = Z_2(0) + (\lambda_2 - \mu_2)t \geq \theta_2,\\
\text{if } t \in (\tilde{\psi}_1, t'_2), &\qquad Z_2(t) = Z_2(\tilde{\psi}_1) + \lambda_2(t - \tilde{\psi}_1) > \theta_2,\\
\text{if } t \in (t'_2, \tilde{t}'_2), &\qquad Z_2(t) = Z_2(t'_2) + (\lambda_2 - \mu_2(1-\rho_1^l))(t - t'_2) > \theta_2,\\
\text{if } t \in [\tilde{t}'_2, H+L], &\qquad Z_2(t) \leq \theta_2,
\end{aligned}$$

where t'_2 and $\tilde{t}'_2$ are given in (70) and (72), respectively. The holding cost incurred by class 2 under the FP2-FP1 policy can be computed as

$$\begin{aligned}
&\int_0^{H+L} h_2(Z_2(t) - \theta_2)^+ \, dt \qquad (73)\\
= \;& \frac{1}{2}h_2\mu_2\Big\{2\frac{(1-\rho_2)(\rho_1^h - \rho_1^l)}{1-\rho_1^l}(\tilde{\psi}_2 - \tilde{\psi}_1)(H - \tilde{\psi}_1) + \rho_2\Big[\frac{\rho_1^h - \rho_1^l}{1-\rho_1^l}(H - \tilde{\psi}_1)\Big]^2\\
&\quad + (1-\rho_2-\rho_1^l)\Big[\frac{\rho_1^h + \rho_2 - 1}{1-\rho_2-\rho_1^l}(H - a_1) - \frac{\rho_1^h - 1}{1-\rho_1^l}(H - \tilde{\psi}_1)\Big]^2\\
&\quad + (1-\rho_2)(2\tilde{\psi}_2 - \tilde{\psi}_1)\tilde{\psi}_1\Big\}. \qquad (74)
\end{aligned}$$

Summing (71) and (74), we get the total holding cost under FP2-FP1 policy for each

(H, L) that satisfies (67) and $\tilde{\psi}_1 \le \tilde{\psi}_2$ as

$$\begin{aligned} c^{\mathrm{FP2-FP1}}(H,L) &= \frac{1}{2}h_2\mu_2\left\{\frac{(\rho_1^h-1)(\rho_1^h-\rho_1^l)}{\eta(1-\rho_1^l)}(H-\tilde{\psi}_1)^2+(1-\rho_2)(2\tilde{\psi}_2-\tilde{\psi}_1)\tilde{\psi}_1\right. \\ &\quad +2\frac{(1-\rho_2)(\rho_1^h-\rho_1^l)}{1-\rho_1^l}(\tilde{\psi}_2-\tilde{\psi}_1)(H-\tilde{\psi}_1)+\rho_2\Big[\frac{\rho_1^h-\rho_1^l}{1-\rho_1^l}(H-\tilde{\psi}_1)\Big]^2 \\ &\quad \left.+(1-\rho_2-\rho_1^l)\Big[\frac{\rho_1^h+\rho_2-1}{1-\rho_2-\rho_1^l}(H-a_1)-\frac{\rho_1^h-1}{1-\rho_1^l}(H-\tilde{\psi}_1)\Big]^2\right\}. \end{aligned} \tag{75}$$

Subtracting (69) from (75), with some algebra we have

$$\begin{aligned} & c^{\mathrm{FP2-FP1}}(H,L)-c(H,L) \\ = & \frac{1}{2}h_2\mu_2\Big\{\Big(\frac{(\rho_1^h-1)(\rho_1^h-\rho_1^l)}{\eta(1-\rho_1^l)}+\frac{(\rho_1^h-1)^2}{(1-\rho_1^l)}\Big)(2H-\psi_1-\tilde{\psi}_1)(\psi_1-\tilde{\psi}_1)-\rho_2(H-\psi_2)^2 \\ - & \frac{2\rho_2(\rho_1^h-1)}{(1-\rho_1^l)}(H-\psi_2)(H-\tilde{\psi}_1)-\frac{2(\rho_1^h-1)^2}{(1-\rho_1^l)}(H-\psi_1)(\psi_1-\tilde{\psi}_1)+(1-\rho_2)(2\tilde{\psi}_2-\tilde{\psi}_1)\tilde{\psi}_1 \\ & +\frac{2(1-\rho_2)(\rho_1^h-\rho_1^l)}{(1-\rho_1^l)}(\tilde{\psi}_2-\tilde{\psi}_1)(H-\tilde{\psi}_1)+\rho_2\psi_2^2 \\ \le & \frac{1}{2}h_2\mu_2\Big\{\Big(\frac{(\rho_1^h-1)(\rho_1^h-\rho_1^l)}{\eta(1-\rho_1^l)}+\frac{(\rho_1^h-1)^2}{(1-\rho_1^l)}\Big)(2H-\psi_1-\tilde{\psi}_1)(\psi_1-\tilde{\psi}_1)-\rho_2(H-\psi_2)^2 \\ - & \frac{2\rho_2(\rho_1^h-1)}{(1-\rho_1^l)}(H-\psi_2)(H-\tilde{\psi}_1)-\frac{2(\rho_1^h-1)^2}{(1-\rho_1^l)}(H-\psi_1)(\psi_1-\tilde{\psi}_1)+2(1-\rho_2)\tilde{\psi}_2\tilde{\psi}_1 \\ & +\frac{2(1-\rho_2)(\rho_1^h-\rho_1^l)}{(1-\rho_1^l)}\tilde{\psi}_2(H-\tilde{\psi}_1)+\rho_2\psi_2^2\Big\}. \end{aligned}$$

Since $H \ge a_2 \ge \psi_1 \ge \tilde{\psi}_1$ (which also implies that $(H-\psi_2)(H-\tilde{\psi}_1) \ge -\psi_2(H-\tilde{\psi}_1)$), we have

$$\begin{aligned} & \Big(\frac{(\rho_1^h-1)(\rho_1^h-\rho_1^l)}{\eta(1-\rho_1^l)}+\frac{(\rho_1^h-1)^2}{(1-\rho_1^l)}\Big)(2H-\psi_1-\tilde{\psi}_1)(\psi_1-\tilde{\psi}_1)-\rho_2(H-\psi_2)^2 \\ - & \frac{2\rho_2(\rho_1^h-1)}{(1-\rho_1^l)}(H-\psi_2)(H-\tilde{\psi}_1)-\frac{2(\rho_1^h-1)^2}{(1-\rho_1^l)}(H-\psi_1)(\psi_1-\tilde{\psi}_1) \\ \le & \Big(\frac{(\rho_1^h-1)(\rho_1^h-\rho_1^l)}{\eta(1-\rho_1^l)}+\frac{(\rho_1^h-1)^2}{(1-\rho_1^l)}\Big)(2H-\psi_1-\tilde{\psi}_1)(\psi_1-\tilde{\psi}_1) \\ & -\rho_2(H-\psi_2)^2+\frac{2\rho_2(\rho_1^h-1)}{(1-\rho_1^l)}\psi_2(H-\tilde{\psi}_1). \end{aligned}$$

Thus,

$$\begin{aligned} & c^{\mathrm{FP2-FP1}}(H,L)-c(H,L) \\ \le & \frac{1}{2}h_2\mu_2\Big\{\Big(\frac{(\rho_1^h-1)(\rho_1^h-\rho_1^l)}{\eta(1-\rho_1^l)}+\frac{(\rho_1^h-1)^2}{(1-\rho_1^l)}\Big)(2H-\psi_1-\tilde{\psi}_1)(\psi_1-\tilde{\psi}_1) \end{aligned}$$

$$
\begin{aligned}
&-\rho_2(H-\psi_2)^2+\frac{2\rho_2(\rho_1^h-1)}{(1-\rho_1^l)}\psi_2(H-\tilde{\psi}_1)+2(1-\rho_2)\tilde{\psi}_2\tilde{\psi}_1\\
&+\frac{2(1-\rho_2)(\rho_1^h-\rho_1^l)}{(1-\rho_1^l)}\tilde{\psi}_2(H-\tilde{\psi}_1)+\rho_2\psi_2^2\Big\}.
\end{aligned}
$$

Note that since $(1-\rho_2)\tilde{\psi}_2=-\rho_2\psi_2$, we can further simplify the above upper bound as

$$
\frac{1}{2}h_2\mu_2\Big\{\Big(\frac{(\rho_1^h-1)(\rho_1^h-\rho_1^l)}{\eta(1-\rho_1^l)}+\frac{(\rho_1^h-1)^2}{(1-\rho_1^l)}\Big)(2H-\psi_1-\tilde{\psi}_1)(\psi_1-\tilde{\psi}_1)-\rho_2H^2\Big\}.
$$

Then we have

$$
\begin{aligned}
& c^{\mathrm{FP2-FP1}}(H,L)-c(H,L)\\
\leq\;& \frac{1}{2}h_2\mu_2\Big\{\Big(\frac{(\rho_1^h-1)(\rho_1^h-\rho_1^l)}{\eta(1-\rho_1^l)}+\frac{(\rho_1^h-1)^2}{(1-\rho_1^l)}\Big)(2H-\psi_1-\tilde{\psi}_1)(\psi_1-\tilde{\psi}_1)\Big\}\\
\leq\;& \frac{1}{2}h_1\mu_1\Big\{\Big(\frac{(\rho_1^h-1)(\rho_1^h-\rho_1^l)}{(1-\rho_1^l)}+\frac{(\rho_1^h-1)^2}{(1-\rho_1^l)}\Big)(2H-\psi_1-\tilde{\psi}_1)(\psi_1-\tilde{\psi}_1)\Big\}, \qquad (76)
\end{aligned}
$$

where the last inequality follows from the definition of η and our assumption that $h_1\mu_1 > h_2\mu_2$. The analysis for $\tilde{\psi}_1 > \tilde{\psi}_2$ is similar and omitted. Since $E[H^2] \leq \infty$, we have the desired result. □

2.9 Summary and conclusions

We studied the dynamic scheduling of different classes of service in a fluid model of computing paradigms for Internet services that may be overloaded for a transient period. We focused on minimizing the penalty of the hosting service provider by scheduling its server resources among various e-commerce sites under Service-Level-Agreement (SLA) contracts with specific Quality-of-Service (QoS) performance guarantees for each class of service.

Our focus in this chapter was on a system with two fluid classes and a single server whose capacity can be shared arbitrarily among the two classes. To capture the QoS performance guarantees in the SLA contracts, we introduced a threshold value for each fluid class such that a holding cost is incurred only if the amount of fluid of a certain class exceeds its threshold value. We assumed that the class 1 arrival rate changes with time and the class 1 fluid can more efficiently reduce the holding cost. Under these assumptions, our objective is to specify the optimal server allocation policy that minimizes the total holding cost.

We first considered the case that the arrival rate function for class 1 fluid is known. In this deterministic setting we could completely characterize the optimal server allocation policy that minimizes the holding cost. We then studied the stochastic fluid system when the arrival rate function for class 1 is random. Using the key insights gained from the optimal policy in the deterministic setting, we developed simple (heuristic) server allocation policies. These policies called "discrete review policies" are not only easy to implement but also shown to be strongly asymptotically optimal for the two heavy traffic regimes considered in this chapter. Moreover, numerical studies have also demonstrated that discrete review policies yield good holding cost performance in general (not only in the asymptotic sense) and they are robust with respect to the system parameters such as load and class 1 arrival rate function.

In the next chapter, we will establish the connection between the stochastic fluid models and queueing networks in a slowly changing environment. We first show that the stochastic fluid models are actually an approximation to queueing networks in a stochastically and slowly changing environment. Then we provide a method to derive a scheduling policy from the solutions of the stochastic fluid models. We also prove that the derived scheduling policy with the provided method is asymptotically optimal if the solution of the stochastic fluid model is optimal.

Table 1: Average holding costs when $\mathbf{E}[L] = 12.5$.

System	Case	$\mathbf{E}[H]$	LB	Percentage Differences off the Lower Bound					
				FP1	π^{a_1}	DSv1	DSv2	DSv3	DSv4
I	A	5	0.00	0.0026*	0.00	0.00	0.00	0.00	0.00
		12.5	1.98	100.77	21.36	19.76	13.34	16.70	19.34
		25	132.16	9.48	15.59	12.66	7.09	10.52	13.35
		37.5	705.40	2.42	11.71	7.66	2.79	6.29	9.31
		50	1883.35	0.90	9.24	3.73	1.36	3.30	6.67
	B	5	0.03	332.61	22.97	16.94	15.29	18.56	21.09
		12.5	20.68	17.92	16.70	10.53	8.99	12.12	14.67
		25	407.05	2.43	10.90	4.90	3.50	6.41	8.88
		37.5	1535.52	0.75	8.01	2.14	0.93	3.59	6.01
		50	3516.83	0.33	6.32	0.61	0.35	1.95	4.33
	C	5	29.51	3.27	11.36	5.58	4.29	7.03	9.41
		12.5	523.38	0.33	5.65	2.05	1.84	3.01	4.26
		25	2962.87	0.17	3.18	0.97	1.08	1.95	2.61
		37.5	7544.71	0.11	2.34	0.45	0.54	1.32	1.88
		50	14337.9	0.07	1.89	0.13	0.17	0.92	1.46
	D	5	3230.38	0.0004	0.73	0.12	0.18	0.23	0.28
		12.5	21190.3	0.002	0.30	0.07	0.18	0.22	0.26
		25	86143.0	0.004	0.15	0.04	0.10	0.12	0.14
		37.5	194926	0.004	0.10	0.02	0.05	0.08	0.10
		50	347597	0.003	0.09	0.005	0.03	0.06	0.07
II	A	5	0.00	67296.6	14.87	13.84	15.74	11.54	12.75
		12.5	14.68	452.27	21.58	18.18	15.69	15.18	18.19
		25	344.88	46.81	22.15	15.58	12.67	13.42	17.60
		37.5	1332.03	12.86	19.44	10.78	8.82	9.61	14.29
		50	3073.30	4.96	16.66	6.54	4.75	6.14	11.05
	B	5	0.43	2061.88	18.97	13.30	14.39	13.82	16.32
		12.5	67.00	104.28	21.63	13.29	13.74	14.41	17.90
		25	756.57	15.74	18.41	8.85	8.76	10.35	14.30
		37.5	2404.51	5.11	14.96	5.19	4.84	6.79	10.82
		50	5085.14	2.27	12.41	2.71	2.31	4.27	8.29
	C	5	59.54	21.69	18.05	9.19	8.99	10.56	14.23
		12.5	747.00	4.34	11.20	5.02	4.82	6.92	9.12
		25	3870.27	2.01	6.82	2.77	2.95	4.53	5.79
		37.5	9600.58	1.06	5.12	1.54	1.67	3.07	4.20
		50	17991.4	0.62	4.19	0.81	0.88	2.18	3.28
	D	5	3908.27	0.01	1.62	0.36	0.56	0.71	0.87
		12.5	25373.9	0.05	0.67	0.23	0.49	0.58	0.63
		25	102840.0	0.05	0.35	0.12	0.23	0.29	0.32
		37.5	232504.0	0.04	0.25	0.06	0.13	0.18	0.22
		50	414419.0	0.02	0.20	0.03	0.07	0.13	0.17

Note: * indicates the actual value of the average holding cost for the FP1 policy.

Table 2: Average holding costs when $\mathbf{E}[L] = 25$.

System	Case	$\mathbf{E}[H]$	LB	Percentage Differences off the Lower Bound					
				FP1	π^{a_1}	DSv1	DSv2	DSv3	DSv4
I	A	5	0.00	0.0026*	0.00	0.00	0.00	0.00	0.00
		12.5	2.36	86.93	20.92	19.43	13.34	16.57	19.04
		25	160.28	8.03	15.49	12.69	7.13	10.62	13.35
		37.5	847.81	2.06	11.68	7.73	2.62	6.37	9.35
		50	2230.50	0.78	9.22	3.77	1.28	3.33	6.71
	B	5	0.03	294.06	22.96	17.21	15.60	18.78	21.18
		12.5	24.52	15.61	16.48	10.55	9.00	12.10	14.55
		25	478.94	2.13	10.83	4.92	3.46	6.43	8.86
		37.5	1774.28	0.67	7.98	2.13	0.85	3.60	6.01
		50	3998.12	0.30	6.30	0.58	0.31	1.95	4.33
	C	5	33.96	2.95	11.36	5.65	4.30	7.10	9.46
		12.5	578.39	0.31	5.63	2.05	1.84	3.01	4.25
		25	3174.51	0.17	3.19	0.97	1.08	1.95	2.61
		37.5	7964.52	0.11	2.35	0.45	0.54	1.33	1.90
		50	15003.8	0.07	1.92	0.13	0.17	0.94	1.48
	D	5	3270.14	0.0004	0.73	0.12	0.18	0.23	0.29
		12.5	21300.0	0.002	0.30	0.07	0.18	0.22	0.26
		25	86382.7	0.004	0.15	0.04	0.10	0.12	0.14
		37.5	195315	0.004	0.11	0.02	0.05	0.08	0.10
		50	348157	0.003	0.09	0.005	0.03	0.06	0.07
II	A	5	0.01	65776	9.40	8.74	12.53	7.46	8.04
		12.5	21.73	471.27	16.69	14.19	14.51	12.28	14.20
		25	462.35	50.90	19.88	14.39	13.51	12.77	16.03
		37.5	1687.87	14.39	18.44	10.73	9.88	9.76	13.79
		50	3756.04	5.66	16.17	6.80	5.35	6.45	10.93
	B	5	0.65	2113.81	14.14	10.65	12.36	10.62	12.25
		12.5	91.84	112.58	18.42	12.36	13.64	12.82	15.39
		25	949.06	17.91	17.16	9.07	9.56	10.13	13.49
		37.5	2879.99	5.99	14.39	5.58	5.57	6.88	10.52
		50	5914.70	2.72	12.11	3.06	2.76	4.43	8.18
	C	5	74.93	23.78	16.53	9.07	9.29	10.06	13.17
		12.5	852.74	5.35	10.81	5.10	4.93	6.78	8.83
		25	4212.98	2.57	6.73	2.93	3.08	4.51	5.72
		37.5	10251.4	1.37	5.11	1.70	1.82	3.09	4.19
		50	18997.4	0.79	4.21	0.95	1.01	2.22	3.30
	D	5	3962.43	0.014	1.63	0.37	0.56	0.71	0.87
		12.5	25524.5	0.07	0.67	0.23	0.49	0.58	0.63
		25	103166	0.07	0.35	0.13	0.23	0.29	0.32
		37.5	233036	0.05	0.25	0.07	0.13	0.19	0.22
		50	415185	0.03	0.20	0.04	0.08	0.13	0.17

The star (*) indicates the actual value of the average holding cost for the FP1 policy.

Table 3: Average holding costs when $\mathbf{E}[L] = 50$.

System	Case	$\mathbf{E}[H]$	LB	Percentage Differences off the Lower Bound					
				FP1	π^{a_1}	DSv1	DSv2	DSv3	DSv4
I	A	5	0.00	0.0027*	0.00	0.00	0.00	0.00	0.00
		12.5	2.69	76.98	20.01	18.63	12.91	15.96	18.26
		25	191.80	6.78	14.76	12.16	6.84	10.22	12.77
		37.5	1029.76	1.72	11.23	7.50	2.40	6.19	9.04
		50	2712.54	0.65	8.95	3.69	1.18	3.26	6.55
	B	5	0.04	264.86	22.08	16.66	15.12	18.14	20.41
		12.5	28.61	13.63	15.88	10.28	8.77	11.76	14.07
		25	570.02	1.82	10.47	4.80	3.33	6.28	8.61
		37.5	2106.81	0.58	7.76	2.07	0.76	3.53	5.87
		50	4708.96	0.26	6.16	0.54	0.27	1.92	4.25
	C	5	39.83	2.57	11.18	5.61	4.25	7.06	9.35
		12.5	659.90	0.28	5.55	2.03	1.82	2.98	4.20
		25	3511.68	0.15	3.16	0.96	1.08	1.94	2.60
		37.5	8655.56	0.10	2.35	0.45	0.54	1.33	1.90
		50	16120.7	0.06	1.93	0.13	0.17	0.94	1.50
	D	5	3346.52	0.0004	0.73	0.12	0.18	0.23	0.29
		12.5	21512.7	0.002	0.30	0.07	0.18	0.22	0.26
		25	86834.6	0.004	0.15	0.04	0.10	0.12	0.14
		37.5	196039	0.004	0.11	0.02	0.05	0.08	0.10
		50	349187	0.003	0.09	0.006	0.03	0.06	0.08
II	A	5	0.01	60834.20	5.81	5.40	8.63	4.67	4.97
		12.5	34.48	472.51	11.30	9.65	11.70	8.57	9.66
		25	679.44	54.28	15.35	11.35	12.94	10.41	12.50
		37.5	2336.96	16.06	15.52	9.46	10.62	8.84	11.76
		50	4988.97	6.52	14.34	6.55	5.93	6.30	9.84
	B	5	1.02	2040.44	9.61	7.63	9.41	7.37	8.36
		12.5	137.10	116.42	13.77	10.06	12.01	9.93	11.59
		25	1299.36	19.96	14.39	8.50	9.77	8.92	11.41
		37.5	3739.20	7.00	12.77	5.71	6.30	6.49	9.43
		50	7410.93	3.27	11.10	3.41	3.31	4.39	7.58
	C	5	102.12	25.17	13.99	8.35	9.05	8.88	11.25
		12.5	1044.08	6.34	9.93	4.98	4.87	6.34	8.15
		25	4836.52	3.24	6.43	3.06	3.16	4.36	5.48
		37.5	11441.8	1.77	4.98	1.90	1.97	3.06	4.10
		50	20846.7	1.04	4.16	1.14	1.19	2.23	3.26
	D	5	4605.41	0.02	1.62	0.37	0.56	0.71	0.87
		12.5	25815.6	0.11	0.66	0.24	0.50	0.58	0.63
		25	103803	0.10	0.35	0.14	0.24	0.29	0.33
		37.5	234065	0.07	0.25	0.08	0.14	0.19	0.23
		50	416648	0.04	0.21	0.05	0.08	0.14	0.18

The star (*) indicates the actual value of the average holding cost for the FP1 policy.

Table 4: Average holding costs when $\mathbf{E}[L] = 1000$.

System	Case	$\mathbf{E}[H]$	LB	Percentage Differences off the Lower Bound					
				FP1	π^{a_1}	DSv1	DSv2	DSv3	DSv4
I	A	5	0.00	0.0027∗	0.00	0.00	0.00	0.00	0.00
		12.5	3.04	68.30	18.55	17.29	12.03	14.85	16.96
		25	246.67	5.29	12.83	10.61	5.98	8.94	11.13
		37.5	1463.13	1.21	9.39	6.32	1.91	5.22	7.59
		50	4161.72	0.43	7.31	3.05	0.89	2.68	5.38
	B	5	0.04	231.28	21.14	16.13	14.67	17.51	19.60
		12.5	37.56	10.58	14.15	9.28	7.92	10.59	12.59
		25	846.81	1.25	8.86	4.11	2.80	5.38	7.32
		37.5	3386.34	0.36	6.41	1.70	0.53	2.96	4.88
		50	7981.57	0.15	5.03	0.40	0.16	1.58	3.50
	C	5	62.00	1.70	9.53	4.87	3.64	6.11	8.03
		12.5	1174.26	0.16	4.61	1.71	1.53	2.51	3.52
		25	6481.53	0.09	2.64	0.80	0.90	1.63	2.17
		37.5	15825.0	0.06	1.98	0.36	0.45	1.13	1.60
		50	28922.2	0.04	1.64	0.09	0.13	0.81	1.28
	D	5	5151.63	0.0002	0.69	0.12	0.17	0.22	0.27
		12.5	27731.1	0.001	0.29	0.07	0.17	0.22	0.25
		25	101086	0.004	0.15	0.04	0.10	0.12	0.14
		37.5	218685	0.004	0.11	0.02	0.05	0.08	0.10
		50	380521	0.003	0.09	0.006	0.03	0.06	0.08
II	A	5	0.03	42604.30	2.25	2.09	3.43	1.80	1.92
		12.5	158.84	297.99	2.56	2.19	3.30	1.99	2.19
		25	4177.56	31.79	2.77	2.08	4.00	2.03	2.27
		37.5	15194.7	9.47	2.80	1.85	4.69	1.89	2.14
		50	31848.9	4.03	2.77	1.64	2.95	1.69	1.94
	B	5	3.88	1275.31	2.75	2.33	3.15	2.16	2.40
		12.5	721.66	69.32	3.03	2.57	3.69	2.30	2.57
		25	7654.53	11.62	3.03	2.40	3.76	2.08	2.43
		37.50	21436.6	4.30	2.90	2.10	3.41	1.74	2.18
		50	40180.6	2.15	2.77	1.74	2.17	1.41	1.93
	C	5	525.66	17.23	3.56	2.79	3.88	2.51	2.91
		12.5	4768.37	4.52	3.15	1.94	1.97	2.11	2.61
		25	18169.3	2.99	2.56	1.65	1.61	1.80	2.20
		37.5	37865.5	1.91	2.25	1.39	1.33	1.45	1.87
		50	62926.3	1.25	2.05	1.13	1.10	1.18	1.62
	D	5	6885.53	0.03	1.38	0.32	0.48	0.61	0.74
		12.5	34368.7	0.29	0.63	0.28	0.47	0.55	0.60
		25	123211	0.36	0.34	0.24	0.24	0.29	0.32
		37.5	265930	0.25	0.26	0.19	0.16	0.19	0.23
		50	462363	0.17	0.21	0.15	0.12	0.15	0.18

The star (*) indicates the actual value of the average holding cost for the FP1 policy.

CHAPTER 3

SCHEDULING OF MULTICLASS OPEN QUEUEING NETWORKS IN A SLOWLY CHANGING ENVIRONMENT

In this chapter, we provide a relationship between the optimal scheduling policy for the stochastic fluid model and the asymptotically optimal policy for the corresponding queueing network in a slowly changing environment. We provide a general method to derive a fluid-scale asymptotically optimal scheduling policy for the queueing network if the optimal policy for its corresponding stochastic fluid model is given.

3.1 Introduction

The contemporary Internet is a large, complex, rapidly changing system characterized with many uncertainties, such as unpredictable user behaviors, server break downs, new technology and service advances. Mathematical modelling and analysis of such a system can augment the understanding of key issues of its performance problems. However it is impossible to model such a complex system precisely, therefore stochastic processing networks have been selected as a more realistic mathematical model for it. In this study, we particulary consider a multiclass open queueing network.

In a queueing network, customers (jobs or service requests) arrive randomly and wait in queue before being served. In a multiclass queuing network, one server might need to process more than one class of customers. When a server is available and there are more than one class of customers waiting in the queue, a scheduling policy determines which customer class to serve next. Our concern is to search for an optimal or near optimal scheduling policy for such networks. In this study, we only consider head of line policies, i.e for the customers of the same class, the earliest one has the highest priority. So the decision to make is what is the next class to serve when a server is available.

In a standard multiclass queueing network, the arrival pattern of customers is not changing, although certain fluctuation is allowed. Essentially the arrival rate is assumed to be a constant throughout the whole time horizon. However, Web traffic characterization studies such as [1, 2] have shown that there can be some sustainable periods during which the traffic volume is significantly larger than other periods. The standard queueing network model fails to capture the time varying characteristics of such communication networks. Our goal therefore is to build a mathematical model for time varying networks and develop a framework to search for an optimal or a nearly optimal scheduling policy for such networks.

More specifically, we consider a multiclass queueing network that operates in a slowly changing environment. The changing environment is modelled as a stochastic process which takes discrete values and each value is referred to as an environment state. Each environment state corresponds to an operating state of the network, and the operating state of the network is described by a set of parameters, such as the arrival rates, the service rates, and the routing matrix. In other words, each environment state corresponds to one set of parameters that describe the dynamics of the network. At each specific environment state, the network operates the same as the standard queuing networks with the associated parameters. If the state of the environment takes only one value, i.e the environment is not changing, the queueing network under our consideration reduces to the standard multiclass queueing network. At some environment states, the network might be overloaded, i.e the traffic intensity of some service stations (i.e servers) might be bigger than one.

We assume that the environment changes very slowly relative to the network dynamics, i.e the customers arrive and depart the network much more frequently than the environment changes states. For example, within one minute, there are hundreds or thousands of customers trying to gain the access of Internet, while the peak time will last tens of minutes or even hours until the network observes off peak time. In other words, there exist different behaviors on different time scales in the computer communication networks. As pointed out in [14], "the relevant time scale for users may be seconds, while the relevant time scale for system transactions may be milliseconds or microseconds".

The existence of different time scales in the Internet results from its large scale and high processing speed, which is also referred to as the network speed. We focus at the time scale of users. We will show that with proper scaling the stochastic network under our consideration can be approximated by a stochastic fluid model when the network speed increases. Given a solution of the stochastic fluid model, we provide a general method to construct an implementable policy. We refer to this methodology as a translation method of fluid model policies. The policy produced by this method is a discrete review type policy, similar to the ones in the literature, e.g [3, 20, 27]. When implementing a discrete review policy, the network reviews its status and makes the scheduling decision at discrete instances of time. However, in this study, the implementation of the discrete review policy will be interrupted by each environment state transition. We prove that the queueing network controlled by this derived policy will converge to the given stochastic fluid model solution. In this way, we say that the translation method we provide is a successful or valid translation method.

Showing that the mutliclass open queueing networks in a slowly changing environment can be approximated by a stochastic fluid model and providing a general method to successfully translate the fluid model policy, we have established a hierarchical framework to search for suitable scheduling policies for such queueing networks by studying their corresponding stochastic fluid models. It is important to note that although our study here is motivated by the Internet setting, the model and results also apply to more general settings, including large manufacturing firms in which events occur at different time scales. Readers can find such discussions in Sethi and Zhang [33].

Similar results of approximating time varying queueing networks can also be found in Choudhury, Mandelbaum, Reiman, and Whitt [14] and Massey [28] and the references therein. But Choudhury et al [14] considers only single class queueing systems, and our results generalize the result of [14] to not only multiclass queueing systems but also to the network setting. And Choudhury et al [14] did not rigorously build the mathematical model of queueing systems in a random environment although some suggestions about the model were made. The approximation results of Massey [28] apply only to the case that the environment is deterministically changing, while we allow the environment to change stochastically. To the best of our knowledge, there is not a general and successful translation method for time varying queuing networks. For queueing networks without a *changing* environment, Maglaras [26, 27] provide a general translation method, although some of the proofs provided in [26, 27] lack mathematical rigor. In this chapter, we provide a rigorous mathematical proof of our result which is similar to [26, 27]; moreover, our result provides a successful translation method in a more general setting, i.e queueing networks in a stochastically *changing* environment. At the same time, we relax some assumptions required in [26, 27]. But we want to acknowledge that results claimed in [26, 27] inspire us to investigate the translating method for queueing networks in a slowly changing environment; and we adopt the uniform acceleration scaling method developed by William Massey.

Bäuerle [5] also provides a general scheme to track the fluid model solutions, but the results therein rely on the piecewise constant structure of the fluid model solutions and the exponential distribution of the inter-arrival times and service times. More discussions about the relation between the fluid model policy and that of queueing networks can also be found in Bäuerle [5], Meyn [30], and the references therein. Note that the research of [5, 26, 27, 30] and the references therein concentrate on standard queueing networks, i.e the queueing networks operating at a single environment state.

The rest of the chapter is organized as follows. In Section 3.2, we build the mathematical model for multiclass open queueing networks in a changing environment. We show in Section 3.2 that such a queueing network can be approximated by a stochastic fluid model when the network speed increases. In Section 3.4, we describe a general method to derive a scheduling policy for queueing networks by modifying the fluid model solutions. We provide the proof in Section 3.5 that the provided method is successful, i.e the queueing network controlled by the derived policy will converge to the given stochastic fluid model solution under the fluid scaling method. The fluid scale asymptotic optimality is introduced in Section 3.6. We provide the proof in Section 3.7 for all the lemmas that appear in Section 3.5. Finally, we give a brief summary of our results in Section 3.8.

3.2 Queueing network model

In this section, we will describe the mathematical model of the queueing network in a changing environment. We first present the primitive data, then we describe the network dynamics.

3.2.1 Primitive data

We consider a queueing network that has S service stations, indexed by $s = 1, \ldots, S$, serving K classes of jobs (or customers), indexed by $k = 1, \ldots, K$. Jobs of class k are served exclusively at station $s = \sigma(k)$, where $\sigma(\cdot)$ is a many-to-one mapping from class to stations. We denote by $C_s = \{k : \sigma(k) = s\}$ the set of classes that are served at station s, and by $C = (c_{sk})$ an $S \times K$ matrix with $c_{sk} = 1$ if $\sigma(k) = s$ and $c_{sk} = 0$ otherwise. C is referred to as the constituency matrix later on. Without loss of generality, we assume that C_s is non-empty for all $s = 1, \ldots, S$.

All the random variables throughout what follows are defined on the same probability space $(\Omega, \mathcal{F}, \mathbf{P})$. Let $X = \{X(t), t \geq 0\}$ denote a stochastic process which represents a changing environment. For each $t \geq 0$, the random variable $X(t)$ takes values in $\mathcal{I}$, where $\mathcal{I}$ is a fixed finite subset of $\mathbb{R}$. For each $i \in \mathcal{I}$, we refer to i as a state of the environment. For each $i \in \mathcal{I}$, there exist $2K$ sequences of independent and identically distributed (i.i.d.) nonnegative random variables $\xi_k(i) = \{\xi_k(i,n), n \geq 1\}$ and $\eta_k(i) = \{\eta_k(i,n), n \geq 1\}$, and K sequences of i.i.d. K-dimensional random vectors $\phi^k(i) = \{\phi^k(i,n), n \geq 1\}$ $(k = 1, \ldots, K)$. We assume that $\xi_k(i,n)$ and $\eta_k(i,n)$ are strictly positive with probability one. The random vector $\phi^k(i,n)$ takes values in the set $\{0, e_1, \ldots, e_K\}$ (where e_k is the kth unit vector in K-dimensional metric space $\mathbb{R}^K$), and $\mathbf{P}\{\phi^k(i,n) = e_l\} = p_{kl}(i)$ and $\mathbf{P}\{\phi^k(i,n) = 0\} = 1 - \sum_{l=1}^{K} p_{kl}(i)$. We use $P(i)$ to denote the $K \times K$ routing matrix associated with i, and its (k,l)th element is $p_{kl}(i)$ for $1 \leq k, l \leq K$. We assume that $\xi_k(i)$, $\eta_k(i)$ and $\phi^k(i)$, $k = 1, \ldots, K$, $i \in \mathcal{I}$, are mutually independent.

For each $i \in \mathcal{I}$, there exists a nonnegative K-dimensional vector $\alpha(i) = (\alpha_k(i), 1 \leq k \leq K)$ and a strictly positive K-dimensional vector $\mu(i) = (\mu_k(i), 1 \leq k \leq K)$. When $\alpha_k(i) > 0$, we refer to $\xi^k(i,n)$ as the time between the $(n-1)$st and nth exogenous arrival of a class k job at state i of the environment and we take $\mathbf{E}[\xi_k(i,n)] = 1/\alpha_k(i)$; when $\alpha_k(i) = 0$, there are no exogenous arrivals of class k jobs at state i of the environment, and we take $\xi_k(i,n) = \infty$ for all $n \geq 1$. We refer to $\alpha_k(i)$ as the exogenous arrival rate to class k at state i of the environment. The random variable $\eta_k(i,n)$ is the required service time for the nth class k job that is served at state i of the environment and satisfies $\mathbf{E}[\eta_k(i,n)] = 1/\mu_k(i)$, where $\mu_k(i)$ is referred to as the service rate for class k at sate i of the environment. The random variable ϕ^k describes the routing mechanism for class k jobs: the nth class k job after service completion turns into a class l job if $\phi^k(i,n) = e_l$, and leaves the network if $\phi^k(i,n) = 0$.

We are also going to use the following notations. For each environment $i \in \mathcal{I}$, and each class $k = 1, \ldots, K$,

$$E_k(i,t) = \sup\{n \geq 0 : \sum_{m=1}^{n} \xi_k(i,m) \leq t\}, \tag{77}$$

$$S_k(i,t) = \sup\{n \geq 0 : \sum_{m=1}^{n} \eta_k(i,m) \leq t\}, \tag{78}$$

$$\Phi^k(i,n) = \sum_{m=1}^{n} \phi^k(i,m). \tag{79}$$

We refer to $\{E_k(i,t), t \geq 0\}$, $\{S_k(i,t), t \geq 0\}$, and $\{\Phi^k(i,n),\, n \geq 0\}$ respectively as the exogenous arrival process, the service process, and the routing process of class k at state i of the environment. Note that $E_k(i,t)$ indicates the number of class k jobs that arrive exogenously at state i of the environment if the network has stayed at this state for t units of time, and $S_k(i,t)$ indicates the number of class k jobs completed at state i of the environment under the head-of-line policy if the station $\sigma(k)$ spends t units of service time on class k at this state. We assume that all policies considered throughout this paper are of head-of-line type, and idling type policies are allowed.

3.2.2 Network dynamics

The performance measure of interest is the K-dimensional queue length process $Z = (Z_1, \ldots, Z_K)'$, where $Z_k = \{Z_k(t), t \geq 0\}$, $1 \leq k \leq K$. Each component of process Z is nonnegative and integer-valued with $Z_k(t)$ indicating the number of class k jobs in the network at time t. We assume that the queueing network operates in the changing environment described by X and that X satisfies the regularity condition, i.e. averagely X has only finite number of state transitions within any finite time. If we denote $N(t)$ the number of state transitions of the environment before time t, then $\mathbf{E}[N(t)] < \infty$ for any finite $t > 0$.

For each $i \in \mathcal{I}$, we use $I(i,t)$ to denote the total time the network has stayed at state i of the environment in the interval $[0,t]$, and $T_k(i,t)$ to indicate the cumulative time that server $\sigma(k)$ has spent on serving class k customers at state i of the environment in the interval $[0,t]$. We also introduce $D_k(i,t)$ to indicate the total number of class k service completions at state i of the environment in the interval $[0,t]$, and $A_k(i,t)$ to indicate the total number of class k arrivals at state i of the environment in the interval $[0,t]$. Recall that $E_k(i) = \{E_k(i,t),\, t \geq 0\}$, $S_k(i) = \{S_k(i,t),\, t \geq 0\}$, and $\Phi^k(i) = \{\Phi^k(i,n),\, n \geq 0\}$ ($i \in \mathcal{I}$, $1 \leq k \leq K$) respectively describe the exogenous arrival process, the service process and the routing process of class k at state i and they are defined as in Section 3.2.1. Note that $\Phi^k_l(i,n)$ is the lth element of the random vector $\Phi^k(i,n)$, and it denotes the total number of class k customers that are routed to class l among the first n customers that departed class k in environment state i. Then we have

$$\begin{aligned} A_k(i,t) &= E_k(i, I(i,t)) + \sum_{l=1}^{K} \Phi^l_k(i, D_l(i,t)), \\ D_k(i,t) &= S_k(i, T_k(i,t)), \\ Z_k(t) &= Z_k(0) + \sum_{i \in \mathcal{I}} A_k(i,t) - \sum_{i \in \mathcal{I}} D_k(i,t) \end{aligned}$$

and

- $I(i,t)$ is nondecreasing in t for all $i \in \mathcal{I}$,
- $T_k(i,t)$ is nondecreasing in t for each $i \in \mathcal{I}$ and $1 \leq k \leq K$, and $T_k(i,t) \leq I(i,t)$,
- $\sum_{k \in C_s} \sum_{i \in \mathcal{I}} (T_k(i,t_2) - T_k(i,t_1)) \leq t_2 - t_1$ for all $0 \leq t_1 \leq t_2$ and $1 \leq s \leq S$.

3.3 The stochastic fluid model approximation

In this section, we are going to present our first result for the queueing network described in Section 3.2. The result shows that the queueing network in a slowly changing environment can be approximated by a stochastic fluid model under appropriate assumptions. We first describe the assumptions on the network data, and then we present our first theorem. We also provide a sequence of lemmas in order to prove this theorem.

3.3.1 Assumptions on the network data

As is traditionally done in fluid limit theorems for open queueing networks, we consider a sequence of queueing networks as described in the previous section, indexed by r, where $r \in \mathbb{R}_+$. For $r \in \mathbb{R}_+$, let the stochastic process $X^r = \{X^r(t),\, t \geq 0\}$ denote the changing environment of the rth network, where $X^r(t)$ takes values in $\mathcal{I}$ for each t; let $r^{-1}\xi_k(i) = \{r^{-1}\xi_k(i,n),\, n \geq 1\}$ and $r^{-1}\eta_k(i) = \{r^{-1}\eta_k(i,n),\, n \geq 1\}$ respectively be the exogenous inter-arrival time sequence and the service time sequence of class k at state i of the environment for rth queueing network. For rth queueing network, the exogenous arrival process of class k at state i of the environment is denoted by $\{E_k^r(i,t), t \geq 0\}$, which is defined in the same way as (77); and the service process of class k at state i of the environment is denoted as $\{S_k^r(i,t), t \geq 0\}$, which is defined in the same way as (78). We assume that the routing processes do not vary with r. For rth queueing network, we use $A_k^r(i,t)$ to denote the total number of class k customers that arrive at state i of the environment until t, $D_k^r(i,t)$ to denote the total number of class k customers that depart at state i of the environment until time t, $T_k^r(i,t)$ to denote the total time spent on serving class k customers at state i of the environment until time t, $1 \leq k \leq K$, $i \in \mathcal{I}$. We also use $I^r(i,t)$ to denote the total time that the queueing network has stayed at state i of the environment until time t for rth queueing system, $i \in \mathcal{I}$. The dynamics of the queueing network satisfy the following set of equations:

$$A_k^r(i,t) = E_k^r(i, I^r(i,t)) + \sum_{1 \leq l \leq K} \Phi_k^l(i, D_l^r(i,t)),\ k = 1, \ldots, K, \tag{80}$$

$$I^r(i,t) = \int_0^t \chi(X^r(s) = i)ds, \tag{81}$$

$$D_k^r(i,t) = S_k^r(i, T_k^r(i,t)),\ k = 1, \ldots, K, \tag{82}$$

$$T_k^r(i,t) \quad \text{is nondecreasing in } t,\ k = 1, \ldots, K,\ i \in \mathcal{I} \tag{83}$$

$$t_2 - t_1 \geq \sum_{k \in C_s} \sum_{i \in \mathcal{I}} (T_k^r(i,t_2) - T_k^r(i,t_1)), \ \text{ for any } t_2 \geq t_1 \geq 0,\ s = 1, \ldots, S, \tag{84}$$

where $\chi(A)$ is a indicator random variable for any $A \in \mathcal{F}$, i.e.

$$\chi(\omega, A) = \begin{cases} 1 & \text{if } \omega \in A, \\ 0 & \text{otherwise.} \end{cases}$$

We use vector $Z^r(t)$ to denote the number of customers in the system at time t, where the kth component corresponds to the number of class k customers. Then we have

$$Z^r(t) = Z^r(0) + \sum_{i \in \mathcal{I}} (A^r(i,t) - D^r(i,t)), \tag{85}$$

where $A^r(i,t) = (A^r_k(i,t), 1 \le k \le K)'$ and $D^r(i,t) = (D^r_k(i,t), 1 \le k \le K)'$.

We assume that the sequence of stochastic processes which describe the changing environments converges almost surely to a stochastic process such that

$$\text{w.p.1} \quad X^r(\cdot) \to X(\cdot) \text{ in } D_{\mathbb{R}}[0,\infty) \text{ as } r \to \infty, \tag{86}$$

where $D_Y[0,\infty)$ is the space of functions defined on $[0,\infty)$ and taking values in a metric space Y, and each function is right continuous on $[0,\infty)$ and have left limits on $(0,\infty)$. $D_Y[0,\infty)$ is endowed with the Skorohod J-1 topology (see Ethier and Kurtz [16]). We assume that the stochastic process X satisfies the regularity condition, i.e. almost surely X has only finite transitions within any finite time. Later on, X is referred to as the limiting environment process. By the strong law of large numbers, we know that for each state $i \in \mathcal{I}$ of the environment and each class k and l,

$$\text{w.p.1} \quad E^r_k(i,t)/r \to \alpha_k(i)t \text{ in } \mathbb{R} \text{ as } r \to \infty, \tag{87}$$

$$\text{w.p.1} \quad S^r_k(i,t)/r \to \mu_k(i)t \text{ in } \mathbb{R} \text{ as } r \to \infty, \tag{88}$$

$$\text{w.p.1} \quad \Phi^k_l(i,n)/n \to p_{kl}(i) \text{ in } \mathbb{R} \text{ as } n \to \infty. \tag{89}$$

3.3.2 Stochastic fluid model approximation

Now we are ready to see our first result, i.e. the limiting points of the scaled queue length processes $\{r^{-1}Z^r(t), t \ge 0\}_{r \ge 0}$ satisfy a stochastic fluid model as the network speed increases. We use $|\mathcal{I}|$ to denote the cardinality of $\mathcal{I}$.

Theorem 14. *If assumption (86) holds, then for almost all $\omega \in \Omega$ (for notational convenience, ω is not specified explicitly in what follows) and for each sequence of $\{r_n, n \geq 1\} \subset \{r, r \in \mathbb{R}_+\}$, there exists a subsequence $\{r_{n_m}, m \geq 1\}$ such that as $m \to \infty$, $r_{n_m} \to \infty$ and*

$$\begin{aligned} & (I^{r_{n_m}}(i,t),\, T^{r_{n_m}}(i,t),\, E^{r_{n_m}}(i,t)/r_{n_m},\, D^{r_{n_m}}(i,t)/r_{n_m},\, A^{r_{n_m}}(i,t)/r_{n_m},\, i \in \mathcal{I}) \\ \to \quad & (I(i,t),\, T(i,t),\, E(i,t),\, D(i,t),\, A(i,t),\, i \in \mathcal{I}) \text{ in } D_{\mathbb{R}^{4K+1+|\mathcal{I}|}}[0,\infty), \end{aligned} \tag{90}$$

where for each $i \in \mathcal{I}$, $(I(i,t), T(i,t), E(i,t), D(i,t), A(i,t))$ satisfies

$$I(i,t) = \int_0^t \chi(X(s) = i)ds, \tag{91}$$

$$A(i,t) = E(i, I(i,t)) + P(i)'D(i,t), \tag{92}$$

$$E(i,t) = \alpha(i)t, \tag{93}$$

$$D(i,t) = (M(i))^{-1}T(i,t), \tag{94}$$

$$T(i,t) \quad \text{is a vector of nondecreasing functions in } t, \tag{95}$$

and $M(i) = diag(\mu_1(i)^{-1}, \ldots, \mu_K(i)^{-1})$. We also have that

$$\sum_{k \in C_s} \sum_{i \in \mathcal{I}} (T_k(i,t_2) - T_k(i,t_1)) \leq t_2 - t_1, \quad \text{for any } 0 \leq t_1 \leq t_2,\ s = 1, \ldots, S. \tag{96}$$

Moreover, if $Z^{r_{n_m}}(0)/r_{n_m} \to Z(0)$ in $\mathbb{R}$, then $Z^{r_{n_m}}(t)/r_{n_m} \to Z(t)$ in $D_{\mathbb{R}_+}[0,\infty)$ and it satisfies that

$$Z(t) = Z(0) + \sum_{i \in \mathcal{I}} A(i,t) - \sum_{i \in \mathcal{I}} D(i,t). \tag{97}$$

Throughout what follows, we use τ_n and τ_n^r to denote the nth transition time of $X(\cdot)$ and $X^r(\cdot)$ respectively, $n \geq 0$, with $\tau_0 = \tau_0^r = 0$ and $\tau_n = \infty$ ($\tau_n^r = \infty$) if X (X^r) has fewer than n transitions.

Before giving the proof of this theorem, we first present the following lemma which will be needed in our analysis and uncovers a condition equivalent to assumption (86).

Lemma 15. *Assumption (86) holds if and only if for every $m \geq 0$, both of the following conditions hold almost surely.*

$$(i)\ \{(\tau_n^r, X^r(\tau_n^r)) : 0 \leq n \leq m\} \to \{(\tau_n, X(\tau_n)) : 0 \leq n \leq m\} \text{ in } \mathbb{R}^{2m+2} \text{ if } \tau_m < \infty, \tag{98}$$

$$(ii)\ \tau_m^r \to \infty \text{ in } \mathbb{R} \text{ if } \tau_m = \infty. \tag{99}$$

We now define an alternative convergence in the space of real valued functions defined on $[0, \infty)$.

Definition 16. Let $f_n(\cdot)$ and $f(\cdot)$ be non-negative real valued functions defined on $[0, \infty)$, $n \geq 1$, then we say $f_n(\cdot) \to f(\cdot)$ uniformly on compact sets (u.o.c) if for any $t \geq 0$,

$$\sup_{u \in [0,t]} |f_n(u) - f(u)| \to 0 \quad \text{as } n \to \infty. \tag{100}$$

The following convergence together theorem is also needed in our proof for Theorem 14 and its proof is given in Billingsley [8]. This result gives a sufficient condition for the convergence of a compound sequence.

Lemma 17. *(convergence together theorem) Assume that $f_n(\cdot)$, $g_n(\cdot)$, $f(\cdot)$, and $g(\cdot)$ are non-negative real valued functions defined on $[0, \infty)$, $n \geq 1$. If as $n \to \infty$, $f_n(\cdot) \to f(\cdot)$ u.o.c, $g_n(\cdot) \to g(\cdot)$ u.o.c, $f(\cdot)$ and $g(\cdot)$ are both continuous, then*

$$f_n(g_n(\cdot)) \to f(g(\cdot)) \quad u.o.c.$$

The following lemma is useful in our analysis and has been proven in Ethier and Kurtz [16].

Lemma 18. *Assume that $f_n \in D_{\mathbb{R}}[0, \infty)$ for each $n \geq 1$ and $f(\cdot)$ is a real valued function which is continuous in $(0, \infty)$ and right continuous at 0. Then*

$$f_n(\cdot) \to f(\cdot) \ \text{ in } D_{\mathbb{R}}[0, \infty) \ \text{ as } n \to \infty$$

if and only if $f_n(\cdot) \to f(\cdot)$ u.o.c.

In our analysis, we also need the following lemma which gives a sufficient condition for (100) to hold and has been proven in Dai [15].

Lemma 19. *Let $\{f_n\}$ be a sequence of nondecreasing real valued functions defined on $\mathbb{R}_+$ and f be a real valued continuous function defined on $\mathbb{R}_+$. Assume that $f_n(t) \to f(t)$ for all rational $t \geq 0$, then $f_n \to f$ u.o.c.*

Before giving the famous Ascoli-Arzela theorem in Lemma 21, we define the equicontinuity first.

Definition 20. A family of real valued functions $f_n : [0, \infty) \to \mathbb{R}$, $n \geq 0$, are equicontinuous if and only if for any $t \geq 0$ and any $\epsilon > 0$, there exists $\delta(t, \epsilon) > 0$ such that for any $t' \geq 0$ and $|t' - t| < \delta(t, \epsilon)$, it satisfies $\sup_{n \geq 0} |f_n(t') - f_n(t)| < \epsilon$.

A particular family of equicontinuous functions satisfy the following Ascoli-Arzela theorem which will be needed in our proof for Theorem 14.

Lemma 21. *(Ascoli-Arzela) Assume that the sequence $\{f_n, n \geq 1\}$ of functions $f_n : [0, \infty) \to \mathbb{R}$ is equicontinuous and the sets $\cup_{n \geq 1} f_n(u)$ are bounded in $\mathbb{R}$ for every $u \in [0, \infty)$. Then there exists a function $f : [0, \infty) \to \mathbb{R}$ which is continuous on $(0, \infty)$ and right continuous at 0 and a subsequence of $\{f_n\}$ which converges to f uniformly on compact sets.*

Now we provide the proof of Lemma 15.

Proof of Lemma 15. Let Λ be a set of strictly increasing, continuous functions $x : \mathbb{R}_+ \to \mathbb{R}_+$ such that $x(0) = 0$ and $\lim_{t \to \infty} x(t) = \infty$. From the definition of Skorohod J-1 convergence, (86) is equivalent to that for every $t > 0$, there exists $\{\gamma^r, r \geq 0\} \subset \Lambda$ such that

$$\lim_{r \to \infty} \sup_{0 \leq u \leq t} |\gamma^r(u) - u| = 0, \tag{101}$$

$$\lim_{r \to \infty} \sup_{0 \leq u \leq t} |X^r(\gamma^r(u)) - X(u)| = 0. \tag{102}$$

First, we show that if (98) and (99) hold for all $m > 0$ then (101) and (102) hold, hence (86) holds. Consider a sample path such that (98) and (99) hold, we show that (101) and (102) also hold for this sample path. We do not specify explicitly this sample path throughout the rest of the proof to avoid lengthy notations. For each $t > 0$, from the regularity condition of X, we know that there exists a finite $m > 0$ such that $\tau_{m-1} \leq t < \tau_m$. Since (98) and (99) are true, then there exists r_0 such that if $r \geq r_0$, then $t < \tau_m^r$ and $\tau_{m-1}^r < t + 1$. We construct a continuous and strictly increasing function $\gamma^r(\cdot) \in \Lambda$ such that it maps τ_n to τ_n^r and $\gamma^r(u) \in [\tau_{n-1}^r, \tau_n^r]$ if $u \in [\tau_{n-1}, \tau_n]$ and $\tau_n < \infty$, $1 \leq n \leq m$. In particular, we construct such a $\gamma^r(\cdot)$ which increases piecewise linearly so that (101) is satisfied. With this $\gamma^r(\cdot)$ and (98), we then show that (102) is also satisfied. We first consider the case of $\tau_m < \infty$ and then the case of $\tau_m = \infty$.

Case 1: If $\tau_m < \infty$, we define $\gamma^r(\cdot)$ such that it maps τ_n to τ_n^r and is linear in the interval $[\tau_{n-1}, \tau_n]$ for all $1 \le n \le m$. More specifically,

$$\gamma^r(u) = \begin{cases} 0 & u = 0, \\ \tau_{n-1}^r + (\tau_n - \tau_{n-1})^{-1}(\tau_n^r - \tau_{n-1}^r)(u - \tau_{n-1}), & u \in (\tau_{n-1}, \tau_n], \ 1 \le n \le m, \\ \tau_m^r + (u - \tau_m), & u > \tau_m. \end{cases}$$

Clearly, $\gamma^r(\cdot) \in \Lambda$. Note that $\tau_0 = 0$ and $t < \tau_m < \infty$. For n such that $1 \le n \le m$, note that the function $\gamma^r(u) - u$ is a linear function of u in the closed interval $[\tau_{n-1}, \tau_n]$, therefore its extreme values (maximum and minimum) will be reached at one of the end points of this interval, i.e. at τ_{n-1} or τ_n. Hence $\max_{u\in[\tau_{n-1},\tau_n]} |\gamma^r(u) - u| = \max\{|\gamma^r(\tau_{n-1}) - \tau_{n-1}|, |\gamma^r(\tau_n) - \tau_n|\}$, $1 \le n \le m$. Since $\gamma^r(\tau_n) = \tau_n^r$ for $n = 0, 1, \ldots, m$, then

$$\sup_{0\le u\le t} |\gamma^r(u) - u| \le \max_{1\le n\le m} \sup_{u\in[\tau_{n-1},\tau_n]} |\gamma^r(u) - u| = \max_{0\le n\le m} |\tau_n^r - \tau_n|.$$

With $\tau_n < \infty$ for all $0 \le n \le m$, (98) and the above inequality imply that (101) holds.

We know that X^r and X are piecewise constant and are right continuous, i.e.

$$\begin{aligned} X^r(u) &= X^r(\tau_n^r), \text{ for } u \in [\tau_n^r, \tau_{n+1}^r), \ \tau_n^r < \infty, \ n \ge 0, \\ X(u) &= X(\tau_n), \text{ for } u \in [\tau_n, \tau_{n+1}), \ \tau_n < \infty, \ n \ge 0. \end{aligned}$$

Also note that $\gamma^r(u) \in [\tau_{n-1}^r, \tau_n^r)$ for any $u \in [\tau_{n-1}, \tau_n)$, $1 \le n \le m$. Then

$$\sup_{0\le u\le t} |X^r(\gamma^r(u)) - X(u)| \le \max_{1\le n\le m} \sup_{u\in[\tau_{n-1},\tau_n]} |X^r(\gamma^r(u)) - X(u)| = \sup_{0\le n\le m} |X^r(\tau_n^r) - X(\tau_n)|.$$

This inequality and (98) imply that (102) holds.

Case 2: If $\tau_m = \infty$, we define $\gamma^r(\cdot)$ similar to case 1, such that it maps τ_n to τ_n^r, is linear in the interval $[\tau_{n-1}, \tau_n]$ for all $1 \le n \le m - 1$, and increases with rate 1 after τ_{m-1}. More specifically,

$$\gamma^r(u) = \begin{cases} 0 & u = 0, \\ \tau_{n-1}^r + (\tau_n - \tau_{n-1})^{-1}(\tau_n^r - \tau_{n-1}^r)(u - \tau_{n-1}), & u \in (\tau_{n-1}, \tau_n], \ 1 \le n \le m-1, \\ \tau_{m-1}^r + (u - \tau_{m-1}), & u > \tau_{m-1}. \end{cases}$$

Clearly $\gamma^r(\cdot) \in \Lambda$. Recall that $\tau_{m-1} \le t$, we have

$$\begin{aligned} \sup_{0\le u\le t} |\gamma^r(u) - u| &= (\max_{1\le n\le m-1} \sup_{u\in[\tau_{n-1},\tau_n]} |\gamma^r(u) - u|) \vee (\sup_{u\in[\tau_{m-1},t]} |\gamma^r(u) - u|) \\ &= \max_{0\le n\le m-1} |\tau_n^r - \tau_n|. \end{aligned}$$

where $a \vee b = \max\{a, b\}$. Note that $\tau_n \le t$ for all $0 \le n \le m-1$, then the above equation and (98) imply that (101) holds.

As in case 1, $\gamma^r(u) \in [\tau^r_{n-1}, \tau^r_n]$ for any $u \in [\tau_{n-1}, \tau_n]$ and $1 \le n \le m-1$. Recall that $X^r(\cdot)$ is right continuous, we have

$$\begin{aligned}
& \sup_{0\le u\le t} |X^r(\gamma^r(u)) - X(u)| \\
= & \ (\max_{1\le n\le m-1} \sup_{u\in[\tau_{n-1},\tau_n]} |X^r(\gamma^r(u)) - X(u)|) \vee (\sup_{u\in[\tau_{m-1},t]} |X^r(\gamma^r(u)) - X(u)|) \\
= & \ (\max_{0\le n\le m-1} |X^r(\tau^r_n) - X(\tau_n)|) \vee (\sup_{u\in[\tau_{m-1},t]} |X^r(\gamma^r(u)) - X(\tau_{m-1})|). \qquad (103)
\end{aligned}$$

Note that $\tau_{m-1} < \infty$, $\tau^r_{m-1} \to \tau_{m-1}$, and $\tau^r_m \to \infty$ in $\mathbb{R}$, hence we can choose $r_0 > 0$ such that for $r > r_0$, $|\tau^r_{m-1} - \tau_{m-1}| < 0.5$ and $\tau^r_m > t + 2$. Hence if $r > r_0$, for all $u \in [\tau_{m-1}, t]$, we have $\tau^r_{m-1} \le \gamma^r(u) < \tau^r_m$ and $X^r(u) = X^r(\tau^r_{m-1})$. Now it is clear that if $r > r_0$,

$$\sup_{u\in[\tau_{m-1},t]} |X^r(\gamma^r(u)) - X(u)| = |X^r(\tau^r_{m-1}) - X(\tau_{m-1})|.$$

Combining this equality and (103), we have

$$\sup_{0\le u\le t} |X^r(\gamma^r(u)) - X(u)| = \sup_{0\le n\le m-1} |X^r(\tau^r_n) - X(\tau_n)|.$$

This equality and (98) imply that (102) holds.

The results of case 1 and case 2 show that (98) and (99) are sufficient conditions for (86) to hold.

We next show that (98) and (99) hold for all $m \ge 1$ are also necessary conditions for (86) to hold.

If (86) holds, then there exist $\{\gamma^r(\cdot), r \ge 0\} \subset \Lambda$ such that (101) and (102) are satisfied. For any $m \ge 1$, we first consider the case of $\tau_m < \infty$ and then the case of $\tau_m = \infty$.

Case 1: We assume that $\tau_m < \infty$ in this case. For any $t < \infty$, there exists $r_0(t) \in \mathbb{R}_+$ such that as $r > r_0(t)$,

$$\sup_{u\in[0,t]} |X^r(\gamma^r(u) - X(u)| < c_0.$$

In particular, we pick t such that $\tau_m < t < \tau_{m+1}$. From the assumption that elements of $\mathcal{I}$ are distinguishable, i.e. the difference between any two distinct values of $\mathcal{I}$ is no smaller than c_0, we know that if $r > r_0$,

$$X^r(\gamma^r(u)) = X(u) = X(\tau_{n-1}) \quad \text{for } u \in [\tau_{n-1}, \tau_n),\, 1 \le n \le m,$$
$$X^r(\gamma^r(u)) = X(u) = X(\tau_m) \quad \text{for } u \in [\tau_m, t].$$

Note that since $\gamma^r(\cdot)$ is strictly increasing, the above equalities are equivalent to

$$X^r(u) = X(\tau_{n-1}) \qquad \text{for } u \in [\gamma^r(\tau_{n-1}), \gamma^r(\tau_n)),\, 1 \le n \le m, \tag{104}$$
$$X^r(u) = X(\tau_m) \qquad \text{for } u \in [\gamma^r(\tau_m), \gamma^r(t)]. \tag{105}$$

From these equalities, we see that the first $m+1$ environment transition times of X^r are $\gamma^r(\tau_n)$, $0 \le n \le m$ if $r > r_0$. That is $\tau_n^r = \gamma^r(\tau_n)$, $0 \le n \le m$. Hence, from (104) and (105), we see that if $r > r_0$,

$$\{X^r(\tau_n^r),\, 0 \le n \le m\} \;\; = \;\; \{X(\tau_n),\, 0 \le n \le m\}. \tag{106}$$

It is clear from (101) that

$$\{\gamma^r(\tau_n),\, 0 \le n \le m\} \to \{\tau_n,\, 0 \le n \le m\},$$

that is

$$\{\tau_n^r,\, 0 \le n \le m\} \to \{\tau_n,\, 0 \le n \le m\}. \tag{107}$$

From (106) and (107), we conclude that (98) holds.

Case 2: We assume that $\tau_m = \infty$ in this case.

Without loss of generality, we assume that m is the smallest integer such that $\tau_m = \infty$, hence $\tau_{m-1} < \infty$. Applying the same technique as in the proof of case 1, and choosing t arbitrarily large such that $\tau_{m-1} < t < \infty$, we still have (104) and (105) by replacing m by $m-1$. From this, we know that the first m environment transition times of X^r are $\tau_n^r = \gamma^r(\tau_n)$, $0 \le n \le m-1$ for $r > r_0(t)$, and the $(m+1)$th environment transition time $\tau_m^r > \gamma^r(t)$. Note that γ^r satisfies (101), then we can choose $r_0(t)$ large enough such that as $r > r_0(t)$, $\gamma^r(t) > t-1$, hence $\tau_m^r > t-1$. Since t is arbitrarily large, we know that $\tau_m^r \to \infty$ in $\mathbb{R}$ as $r \to \infty$.

The results in case 1 and case 2 show that (98) and (99) are necessary conditions for (86) to be true.

□

Now we are ready to prove Theorem 14.

Proof of Theorem 14. We consider any sample path that satisfies (86). For notational convenience, this sample path is not specified explicitly in the corresponding notations that follow. We first show that for this sample path, $I^r(i,t) \to I(i,t)$ in $\mathbb{R}$ as $r \to \infty$. For any $t > 0$, from the regularity condition of X, we know that there exists a finite integer $m \geq 1$ such that $\tau_{m-2} < \tau_{m-1} \leq t < \tau_m$. From Lemma 15, we know there exists $r_0(t) > 0$, such that if $r > r_0(t)$, $\tau^r_{m-2} < t < \tau^r_m$. We define $\tau_{m-2} = \tau^r_{m-2} = 0$ if $m < 2$. Also note that X^r is piecewise constant and right continuous, hence $I^r(i,t)$ can be rewritten as

$$\begin{aligned}
I^r(i,t) &= \int_0^t \chi(X^r(s) = i)ds \\
&= \sum_{n=1}^{m-2} \int_{\tau^r_{n-1}}^{\tau^r_n} \chi(X^r(s) = i)ds + \int_{\tau^r_{m-2}}^{t\wedge\tau^r_{m-1}} \chi(X^r(s) = i)ds + \int_{t\wedge\tau^r_{m-1}}^{t} \chi(X^r(s) = i)ds \\
&= \sum_{n=1}^{m-2} (\tau^r_n - \tau^r_{n-1})\chi(X^r(\tau^r_{n-1}) = i) + (t\wedge\tau^r_{m-1} - \tau^r_{m-2})\chi(X^r(\tau^r_{m-2}) = i) \\
&\quad + (t - t\wedge\tau^r_{m-1})\chi(X^r(\tau^r_{m-1}) = i).
\end{aligned}$$

From (106) in the the proof for Lemma 15, we can choose $r_0(t)$ large enough such that as $r > r_0(t)$, we also have

$$X^r(\tau^r_n) \quad = \quad X(\tau_n), \quad 0 \leq n \leq m-1.$$

Hence, as $r > r_0(t)$,

$$\chi(X^r(\tau^r_n) = i) \quad = \quad \chi(X(\tau_n) = i), \quad i \in \mathcal{I}, \quad 0 \leq n \leq m-1,$$

and

$$\begin{aligned}
I^r(i,t) &= \sum_{n=1}^{m-2} (\tau^r_n - \tau^r_{n-1})\chi(X(\tau_{n-1}) = i) + (t\wedge\tau^r_{m-1} - \tau^r_{m-2})\chi(X(\tau_{m-2}) = i) \\
&\quad + (t - t\wedge\tau^r_{m-1})\chi(X(\tau_{m-1}) = i).
\end{aligned}$$

From Lemma 15, we also know that

$$\{\tau_n^r,\, 0 \le n \le m\} \to \{\tau_n,\, 0 \le n \le m\} \text{ in } \mathbb{R}^{m+1}.$$

Now it is clear that as $r \to \infty$,

$$\begin{aligned} I^r(i,t) &\to \sum_{n=1}^{m-2} (\tau_n - \tau_{n-1})\chi(X(\tau_{n-1}) = i) + (t \wedge \tau_{m-1} - \tau_{m-2})\chi(X(\tau_{m-2}) = i) \\ &\quad +(t - t \wedge \tau_{m-1})\chi(X(\tau_{m-1}) = i) \\ &= \int_0^t \chi(X(s) = i)ds, \end{aligned}$$

where the equality comes from the fact that X is piecewise constant and right continuous at its transition times τ_n, $n \ge 0$. Throughout the rest of the paper, we let $I(i,t) \equiv \int_0^t \chi(X(s) = i)ds$, then $I(i,t)$ is continuous in t and $I^r(i,t)$ is nondecreasing in t, by Lemma 19, we see that

$$I^r(i,t) \to I(i,t) \quad \text{u.o.c} \quad \text{as } r \to \infty. \tag{108}$$

Next, we consider the convergence of $\{T_k^r(i,t),\, r \in \mathbb{R}_+\}$. It is easy to check that $\{T_k^r(i,t),\, r \in \mathbb{R}_+\}$ satisfy the conditions of Ascoli-Arzela lemma, hence for any sequence $\{r_n\}_{n=1}^{\infty} \subset \{r,\, r \in \mathbb{R}_+\}$, there exists a subsequence $\{r_{n'}\}_{n'=1}^{\infty}$ and a continuous function $T(i,t)$, such that as $n' \to \infty$, it satisfies $r_{n'} \to \infty$ and

$$T^{r_{n'}}(i,t) \to T(i,t) \quad \text{u.o.c} \quad \text{as } n' \to \infty. \tag{109}$$

From (83), we know $T(i,t)$ satisfies (95).

By Lemma 19 and (87),(88), (89), we have

$$E_k^r(i,t)/r \to \alpha_k(i)t \quad \text{u.o.c} \quad \text{as } r \to \infty, \tag{110}$$

$$S_k^r(i,t)/r \to \mu_k(i)t \quad \text{u.o.c} \quad \text{as } r \to \infty, \tag{111}$$

$$\Phi_l^k(i,n)/n \to p_{kl}(i) \quad \text{u.o.c} \quad \text{as } n \to \infty. \tag{112}$$

Now, applying the convergence together theorem (Lemma 17) to (110) and (108), we see that

$$E_k^r(i, I^r(i,t))/r \to \alpha_k(i)I(i,t) \quad \text{u.o.c} \quad \text{as } r \to \infty. \tag{113}$$

Again, applying Lemma 17 to (109) and (111), we see that

$$S_k^{r_{n'}}(i, T^{r_{n'}}(i,t))/r_{n'} \to \mu_k(i)T(i,t) \quad \text{u.o.c} \quad \text{as } n' \to \infty. \tag{114}$$

Throughout the rest of the paper we let $D_k(i,t) \equiv \mu_k(i)T_k(i,t)$, recall that $D_k^r(i,t) = S_k^r(i, T_k^r(i,t))$, we see (114) is equivalent to

$$D_k^{r_{n'}}(i,t)/r_{n'} \to D_k(i,t) \quad \text{u.o.c} \quad \text{as } n' \to \infty. \tag{115}$$

If $D_k(i,t) = 0$, note that $0 \le \Phi_l^k(D_k^{r_{n'}}(i,t))/r_{n'} \le D_k^{r_{n'}}(i,t)/r_{n'}$, from the Sandwich theorem and (115), we see $\Phi_l^k(D_k^{r_{n'}}(i,t))/r_{n'} \to 0$ u.o.c. If $D_k(i,t) > 0$, then (115) tells us that $D^{r_{n'}}(i,t) \to \infty$ as $n' \to \infty$. From (112) and (115), we have

$$\begin{aligned} \frac{\Phi_l^k(D_k^{r_{n'}}(i,t))}{r_{n'}} &= \left(\frac{\Phi_l^k(D_k^{r_{n'}}(i,t))}{D_k^{r_{n'}}(i,t)}\right)\left(\frac{D_k^{r_{n'}}(i,t)}{r_{n'}}\right) \\ &\to p_{kl}(i)D_k(i,t) \quad \text{u.o.c.} \end{aligned}$$

Combining the results for both cases of $D_k(i,t) = 0$ and $D_k(i,t) > 0$, we have

$$\Phi_l^k(D_k^{r_{n'}}(i,t))/r_{n'} \to p_{kl}(i)D_k(i,t) \quad \text{u.o.c} \quad \text{as } n' \to \infty. \tag{116}$$

Since $\{r_{n'}\}_{n'=1}^{\infty}$ is a subsequence of $\{r, r \in \mathbb{R}_+\}$, (113) also holds if we replace r by $r_{n'}$ and let $n' \to \infty$. Without loss of generality, we can choose $\{r_{n'}\}_{n'=1}^{\infty}$ such that (113)-(116) hold for all $k, l = 1, \ldots, K$.

From equation (80) and the results that (113) and (116) hold for all $k = 1, \ldots K$, we have

$$A^{r_{n'}}(i,t)/r_{n'} \to A(i,t) \quad \text{u.o.c} \quad \text{as } n' \to \infty, \tag{117}$$

and $A(i,t)$ satisfies (92).

We can choose a common subsequence of $\{r, r \in \mathbb{R}_+\}$, still denoted as $\{r_{n'}\}_{n'=1}^{\infty}$ for notational convenience, such that (108), (109), and (113)-(117) hold for all $i \in I$ and $k = 1, \ldots, K$. Note that when we take sum over $i \in \mathcal{I}$, the summands are all nonnegative real numbers, so we can exchange the order between the limit operation and the summation. Now from the hypothesis that $Z^r(0)/r \to Z(0)$, we can see that $Z^{r_{n'}}(t)/r_{n'} \to Z(t)$ u.o.c and $Z(t)$ satisfies (97) from (80) -(82), (85), and the above results. By Lemma 18, all the above convergence results hold in the Skorohod J-1 topology. It is also clear that (96) follows from (84) and (109).

$\square$

From the proof for Theorem 14, we have the following corollary which will be used for our future analysis.

Corollary 22. *Let $\{f^r(\cdot), r \geq 0\}$ be real valued functions defined on $[0, \infty)$. Let $\mathbf{1}(\cdot)$ be the identity function defined on $[0, \infty)$, i.e. $\mathbf{1}(u) = u$ for all $u \geq 0$. If the assumption (86) holds, then as $r \to \infty$, $f^r(\cdot) \to \mathbf{1}(\cdot)$ u.o.c implies that w.p.1,*

$$r^{-1}(Z^r(\cdot) - Z^r(f^r(\cdot))) \to 0 \quad u.o.c. \tag{118}$$

Proof of Corollary 22. We prove it by contradiction. If the result is not true, then there exists a subset of Ω, say Ω_0, such that $P(\Omega_0) > 0$ and the above result does not hold for every $\omega \in \Omega_0$. Hence, for every $\omega \in \Omega_0$ (where ω is not specified in what follows for notational convenience), there exists an $\epsilon_0 > 0$, a $t_0 \geq 0$ and a subsequence $\{r_n, n \geq 1\}$ such that $r_n \to \infty$ as $n \to \infty$, and for all $n \geq 1$,

$$\sup_{0 \leq s \leq t_0} r_n^{-1} |Z^{r_n}(s) - Z^{r_n}(f^{r_n}(s))| > \epsilon_0.$$

Let $\check{Z}^r(\cdot) = Z^r(\cdot) - Z^r(0)$, then for all $n \geq 1$, we know

$$\sup_{0 \leq s \leq t_0} r_n^{-1} |\check{Z}^{r_n}(s) - \check{Z}^{r_n}(f^{r_n}(s))| > \epsilon_0. \tag{119}$$

Since $\check{Z}^r(0) = 0$, then from Theorem 14, we know that there exists a subsequence of $\{r_n, n \geq 1\}$, still denoted as $\{r_n, n \geq 1\}$ for notational convenience, and a continuous function $\check{Z}(\cdot)$ such that as $n \to \infty$,

$$r_n^{-1} \check{Z}^{r_n}(\cdot) \to \check{Z}(\cdot) \quad \text{u.o.c} \tag{120}$$

Since $f^r(\cdot) \to \mathbf{1}(\cdot)$ u.o.c as $r \to \infty$, then from the convergence together theorem (Lemma 17), we know that

$$r_n^{-1} \check{Z}^{r_n}(f^{r_n}(\cdot)) \to \check{Z}(\mathbf{1}(\cdot)) = \check{Z}(\cdot) \quad \text{u.o.c} \tag{121}$$

From (120) and (121), we see that for any finite $t \geq 0$, as $n \to \infty$,

$$\sup_{0 \leq u \leq t} r_n^{-1} |\check{Z}^{r_n}(u) - \check{Z}^{r_n}(f^{r_n}(u)))| \to 0,$$

which contradicts to (119). This concludes the proof. $\square$

3.4 Fluid tracking policy for queueing networks in a slowly changing environment

In this section, we provide a method to construct an implementable scheduling policy for the queueing network in a changing environment if a stochastic fluid model solution is given.

We assume that the stochastic fluid model solution is given as $\Psi = \{(\Psi(t; z, i), t \geq 0, z \geq 0, i \in \mathcal{I}\}$ or $T^{\Psi} = \{T^{\Psi}(t; z, i)), t \geq 0, z \geq 0, i \in \mathcal{I}\}$, where $\Psi(t; z, i)$ and $T^{\Psi}(t; z, i)$ are K dimensional vectors of real numbers. For any $1 \leq k \leq K$, the real number $\Psi_k(t; z, i)$ denotes the fluid level of class k at time t and $T_k^{\Psi}(t; z, i)$ denotes the total time spent on serving class k during $[0, t)$ if the initial fluid level vector is z and the environment state is i in the interval $[0, t)$. Note that

$$\Psi(t; z, i) = z + \alpha(i)t - (I - P(i)')(M(i))^{-1} T^{\Psi}(t; z, i), \tag{122}$$

where I denotes the $K \times K$ identical matrix Since $(I - P(i)')$ is invertible, then form either of Ψ and T^{Ψ}, we know the other one. Note that with slight adaptation of notations, the functions in set T^{Ψ} satisfy (95) and (96).

Now, we construct a scheduling policy for the queueing network so that we review the queueing network periodically. At the beginning of each review period, we check the queue length and the state of the environment, then we implement the policy as described below.

We initially set the planned review period length as $l(r)$ such that for any $n > 0$ and any δ such that $0 < \delta < 1$

$$l(r) \to 0 \text{ and } \frac{r(l(r))^n}{r^{\delta}} \to \infty \text{ as } r \to \infty. \tag{123}$$

This assumption in fact can be relaxed (see Remark 32). We also set a safety stock level $\theta^r(i)$ for each state i of the environment such that $\theta^r(i) = \beta(i) r l(r)$, where $\beta(i)$ is a K-dimensional vector of real numbers and $\beta(i) > \mu(i)$. We use j to denote the index of the review period, $t^r(j)$ to denote the beginning time of the jth review period, and $q^r(j)$ to denote the queue length vector at $t^r(j)$. Set $j = 0$ and $t^r(0) = 0$ initially, then we implement the following policy. For notational convenience, we denote $X^r(t^r(j))$ by X^r_j.

- At the beginning of the review period $t^r(j)$, observe the environment state, say it is i, i.e. $X^r_j = i$.

- If $q^r(j) \geq \theta^r(i)$, then we plan a policy for the next $l(r)$ time units , referred to as the *fluid policy*, according to the fluid model solution. The exact procedure is given as follows.

  First, let

$$\begin{aligned}
\bar{q}^r(j) &= (q^r(j) - \theta^r(i))^+/r, && (124)\\
x^r_k(j) &= T^\Psi_k(l(r); \bar{q}^r(j), i),\ 1 \leq k \leq K, && (125)\\
p^r_k(j) &= \lfloor r\mu_k(i)x^r_k(j)\rfloor,\ 1 \leq k \leq K, && (126)\\
u^r_s(j) &= l(r) - \sum_{k\in C_s} x^r_k(j),\ 1 \leq s \leq S, && (127)
\end{aligned}$$

  where $a^+ = \max\{a, 0\}$ and $\lfloor a\rfloor$ is the maximum integer that is smaller than or equal to a. We use $x^r_k(j)$ to denote the planned time to spend on serving class k and $u^r_s(j)$ the planned idle time for server s during the jth review period, which are estimated through the stochastic fluid model solution. Based on the average service rate during environment state i, we schedule $p^r_k(j)$ amount of class k jobs to be processed during jth review period for each $1 \leq k \leq K$.

  Server s processes $p^r_k(j)$ jobs of class k for each $k \in C_s$, $1 \leq s \leq S$. Let

$$b^r_s(j) = \sum_{k\in C_s} \sum_{n=\iota^r_k(i,j)+2}^{\iota^r_k(i,j)+p^r_k(j)} \eta^r_k(i, n) + \tilde{\eta}^r_k(i, \iota^r_k(i, j) + 1), \qquad (128)$$

  where $\iota^r_k(i, j)$ denotes the number of class k jobs that has departed at environment state i until $t^r(j)$ and $\tilde{\eta}^r_k(i, \iota^r_k(i, j) + 1)$ denotes the remaining service time of the $\iota^r_k(i, j) + 1$st class k customer served at environment state i. Then $b^r_s(j)$ denotes the total service time of the planned jobs for server s. After finishing the scheduled jobs, server s idles $\min\{u_s^r(j), (l(r) - b_s^r(j))^+\}$ time units and then sends the *finish* signal. As soon as every server sends the *finish* signal or the environment changes to another state, we start a new review period. When the environment changes its state, we say that an environment transition happens.

Now we give the expression of the starting time of the next review period. Note that the initial queue length is no less than the safety stock level, i.e. $q^r(j) \geq \theta^r(i)$, which guarantees that $q^r(j) \geq p^r(j)$, hence each server can continuously serve the scheduled jobs without having to wait for additional arrivals. Let

$$
\begin{aligned}
e_s^{r,F}(j) &= b_s^r(j) + \min\{u_s^r(j), (l(r) - b_s^r(j))^+\} \\
&= \begin{cases} b_s^r(j) + u_s^r(j) & \text{if } b_s^r(j) + u_s^r(j) \leq l(r), \\ l(r) & \text{if } b_s^r(j) \leq l(r) < b_s^r(j) + u_s^r(j), \\ b_s^r(j) & \text{if } l(r) < b_s^r(j), \end{cases}
\end{aligned} \tag{129}
$$

which denotes the time elapsed until server s sends the *finish* signal if there is no environment transition to interrupt the policy implementation. Let $\tau^r(j)$ denote the first environment transition time after $t^r(j)$, i.e. $\tau^r(j) = \min\{\tau_n^r : \tau_n^r > t^r(j),\, n \geq 1\}$, where τ_n^r is the nth transition time of $\{X^r(t), t \geq 0\}$. Then the $(j+1)$th review period starts at

$$
t^r(j+1) = \min\{t^r(j) + \max_{1 \leq s \leq S} e_s^{r,F}(j),\, \tau^r(j)\}, \tag{130}
$$

- If $q^r(j) \ngeq \theta^r(i)$, then we implement a policy, referred to as the *target idle* policy, such that the queue length of each active class is above or equal to the safety stock level at the end of the review period if there is no environment transition during this review period. The meaning of an active class will be clear in the next paragraph. Though various policies can achieve this goal, we particularly choose to implement the following one with the purpose of having a clear description of the policy and a rigorous proof for our next result. The procedure is given as follows.

  Recall that $(I - P(i))$ is invertible, hence there is a unique solution of $\lambda(i)$ such that $\lambda(i) = \alpha(i) + P(i)'\lambda(i)$, which denotes the nominal arrival rate vector in environment state i. For each class k such that $\lambda_k(i) > 0$, i.e class k has positive nominal arrival rate at the state i of the environment, class k is called an active class.

First, at each state of the environment, say state i , we want to specify a path for each active class k through which an exogenous job arrives at buffer k. Let o denote a dummy node that represents the outside of the network from where all the exogenous arrivals come. Then for each $k \in \mathcal{A}(i)$, either $\alpha_k(i) > 0$ or there exists $1 \le m \le K$ and $\{k_1, \ldots, k_m\} \subset \mathcal{A}(i)$ such that $\alpha_{k_1}(i) > 0$ and $p_{k_1 k_2}(i) \cdots p_{k_m k}(i) > 0$. For the first case, we say k is connected to the exogenous source o through a path (o, k) at the state i of the environment; for the second case, we say k is connected to the exogenous source o through a path $(o, k_1, \ldots, k_m, k)$ at the state i of the environment. Note that an active class may be connected to o through multiple paths. For each k, we select *only one* path for it. Then the dummy node o, the active classes and the selected paths compose a tree. If the selected path for k is of the form (o, k), then we say k is a root class. We use $\mathcal{R}(i)$ to denote the set of root classes at state i of the environment. If at the state i of the environment, the selected path for k is of the form $(o, k_1, \ldots, k_m, k)$, $m \ge 1$, then we say $k \succ k_n$ for $n = 1, \ldots, m$; and we say k is a *immediate* child class of k_m and k_m is the *immediate* parent class of k at the state i of the environment. We use $\mathcal{C}_k(i)$ to denote the set of all immediate child classes of class k at the state i of the environment. For each class k, we also let $\mathcal{T}_k(i)$ denote the set of classes that succeed k in the preselected paths and class k itself, i.e $\mathcal{T}_k(i) = \{l : l \succ k\} \cup \{k\}$, $1 \le k \le K$. Then $\mathcal{T}_k(i)$ is a subtree of the preselected tree with root node being k. For a class k, if among all the preselected paths, there is *not* a class l such that $l \succ k$, we say k is a leaf of the preselected tree.

We implement the following policy with the preselected tree.

- All classes have two status: *working* or *finish*.
- Jobs of a class in *working* status are processed once the server for this class is available. Jobs of a class in *finish* status are not processed.
- At the beginning of the review period, all active classes having child classes set

their status as *working*, and all the other classes set their status as *finish*.

– A class switches status from *working* to *finish* at the time when the status of its child classes are all *finish* and the queue length of its child classes are all above or equal to their respective safety stock levels.

– A new review period starts at the time when the status of all root classes are *finish* and the queue length of all root classes are above or equal to their respective safety stock levels or at the time when an environment transition happens, whichever happens first.

Without loss of generality and for notational convenience, we assume that for every state $i \in \mathcal{I}$, every class k is an active class, where $1 \leq k \leq K$.

Now we estimate the upper bound of the duration of the jth review period during which the *target idle* policy is implemented. Let $\tilde{p}_k^r(j)$ denote the number of class k service completions during the jth review period during which the target-idle policy is implemented. Each root class k needs to wait for at most $(\lceil \theta_k^r(i) \rceil + \tilde{p}_k^r(j) - q_k^r(j))^+$ exogenous class k arrivals in order to fulfill the *target idle* policy. For $k \in \mathcal{R}(i)$, let $e_k^r(j) = (\lceil \theta_k^r(i) \rceil + \tilde{p}_k^r(j) - q_k^r(j))^+$, then the total time spent on waiting for exogenous arrivals is at most

$$\max_{k \in \mathcal{R}(i)} \sum_{n=\kappa_k^r(i,j)+2}^{\kappa_k^r(i,j)+e_k^r(j)} \xi_k^r(i,n) + \tilde{\xi}_k^r(\kappa_k^r(i,j)+1),$$

where $\kappa_k^r(i,j)$ is the total number of class k jobs that has arrived exogenously until $t^r(j)$ at environment state i, and $\tilde{\xi}^r(i, \kappa^r(i,j)+1)$ is the remaining time to wait for the $\kappa_k^r(i,j)+1$st job that arrives at the environment state i after time $t^r(j)$. The total processing time during the jth review period is at most

$$\sum_{k=1}^{K} \sum_{n=\iota_k^r(i,j)+2}^{\iota_k^r(i,j)+\tilde{p}_k^r(j)} \eta_k^r(i,n) + \tilde{\eta}_k^r(i, \iota_k^r(i,j)+1),$$

which is the total processing time if there is only one server for all these classes. Let $e^{r,I}(j)$ denote the duration of the jth review period during which the *target idle* policy is implemented, then

$$\begin{aligned} e^{r,I}(j) &\le \max_{k\in\mathcal{R}(i)} \sum_{n=\kappa_k^r(i,j)+2}^{\kappa_k^r(i,j)+e_k^r(j)} \xi_k^r(i,n) + \tilde{\xi}_k^r(\kappa_k^r(i,j)+1) \\ &\quad + \sum_{k=1}^{K} \sum_{n=\iota_k^r(i,j)+2}^{\iota_k^r(i,j)+\tilde{p}_k^r(j)} \eta_k^r(i,n) + \tilde{\eta}_k^r(i,\iota_k^r(i,j)+1). \end{aligned} \tag{131}$$

Then the ending time of the jth review period (i.e the beginning of the $j+1$st review period) satisfies

$$t^r(j+1) \le \min\{t^r(j) + e^{r,I}(j),\, \tau^r(j)\}, \tag{132}$$

where $\tau^r(j)$ denotes the earliest environment transition time after $t^r(j)$.

If a review period is ended due to an environment transition, then we say this review period is an interrupted period, otherwise we say it is uninterrupted period. We refer to an uninterrupted review period during which the *fluid policy* is implemented as a normal review period.

3.5 Main result of the stochastic fluid tracking method

In this section, we will show that the method provided in Section 3.4 successfully translates a very general stochastic fluid model solution and constructs a scheduling policy for the original queueing network. That is, under some mild conditions the queue length of the network in the changing environment under the constructed policy will converge to the given stochastic fluid model solution as the network speed increases. We begin with the notation conventions in this study.

Throughout the rest of the manuscript, we adopt the following notations. For each class k ($1 \le k \le K$), and each state i of the environment ($i \in \mathcal{I}$), we let

$$\begin{aligned} \bar{\mu}_k &= \sup\{\mu_k(i) : \mu_k(i) > 0,\, i \in \mathcal{I}\}, & \underline{\mu}_k &= \inf\{\mu_k(i) : \mu_k(i) > 0,\, i \in \mathcal{I}\}, \\ \bar{\alpha}_k &= \sup\{\alpha_k(i) : \alpha_k(i) > 0,\, i \in \mathcal{I}\}, & \underline{\alpha}_k &= \inf\{\alpha_k(i) : \alpha_k(i) > 0,\, i \in \mathcal{I}\}, \\ \bar{\beta}_k &= \sup\{\beta_k(i) : \beta_k(i) > 0,\, i \in \mathcal{I}\}, & \underline{\beta}_k &= \inf\{\beta_k(i) : \beta_k(i) > 0,\, i \in \mathcal{I}\}, \\ g_k(i,x) &= \mathbf{E}[(\xi_k(i,1))^2 \chi(\xi_k(i,1) > x)], & h_k(i,x) &= \mathbf{E}[(\eta_k(i,1))^2 \chi(\eta_k(i,1) > x)]. \end{aligned}$$

We also let

$$\bar{\mu} = \max\{\bar{\mu}_k,\, 1 \le k \le K\}, \qquad \underline{\mu} = \min\{\bar{\mu}_k, 1 \le k \le K\},$$

$$\bar{\alpha} = \max\{\bar{\alpha}_k,\, 1 \le k \le K\}, \qquad \underline{\alpha} = \min\{\bar{\alpha}_k,\, 1 \le k \le K\}$$

$$\bar{\beta} = \max\{\bar{\beta}_k,\, 1 \le k \le K\}, \qquad \underline{\beta} = \min\{\bar{\beta}_k, 1 \le k \le K\}.$$

We also define the notation $o(\cdot)$ as follows.

Definition 23. For two real valued functions $f_1(x)$ and $f_2(x)$, if $\lim_{x\to\infty} f_2(x)/f_1(x) = 0$, then, we say $f_2(x) = o(f_1(x))$.

Remark 24. Note that if $f_2(x) = o(f_1(x))$ and $f_3(x) = o(f_1(x))$, then $f_2(x) + f_3(x) = o(f_1(x))$.

Throughout this section, we make the following assumptions.

- For each class k, $1 \le k \le K$, there exists $\gamma > 0$ such that

$$\sup_{i\in\mathcal{I}} g_k(i,x) = o(x^{-\gamma}), \quad \sup_{i\in\mathcal{I}} h_k(i,x) = o(x^{-\gamma}). \tag{133}$$

  Note that since $\mathcal{I}$ is finite, the assumption (133) is satisfied if $\mathbf{E}[(\xi_k(i,1))^{2+\gamma}] < \infty$ and $\mathbf{E}[(\eta_k(i,1))^{2+\gamma}] < \infty$ for each $i \in \mathcal{I}$ and some $\gamma > 0$. Without loss of generality, we assume that $\gamma \le 1$.

  Remark 25. Assumption (133) implies that there exists c_1 and c_2 such that

$$\sup_{i\in\mathcal{I}} g_k(i,x) \le \frac{c_1}{x^\gamma} \quad \text{and} \quad \sup_{i\in\mathcal{I}} h_k(i,x) \le \frac{c_2}{x^\gamma}.$$

- There exist real valued functions $\hat{c}_1(\cdot) > 0$ and $\hat{c}_2(\cdot) > 0$ such that

$$\begin{aligned} \sup_{1\le k\le K} \sup_{i\in\mathcal{I}} \sup_{0\le u\le t} \mathbf{E}[\xi_k(i,n)|\xi_k(i,n) > u] - u &< \hat{c}_1(t), \\ \sup_{1\le k\le K} \sup_{i\in\mathcal{I}} \sup_{0\le u\le t} \mathbf{E}[\eta_k(i,n)|\eta_k(i,n) > u] - u &< \hat{c}_2(t). \end{aligned}$$

  If the inter-arrival times $\xi_k(i,n)$, $i \in \mathcal{I}$, $k = 1, \ldots, K$, are all *new better than used in expectation* (NBUE) ([34], page 68), then $\hat{c}_1(t)$ can be a constant, e.g $\hat{c}_1(t) = \check{c}_1 \sup_{1\le k\le K} \sup_{i\in\mathcal{I}} \mathbf{E}[\xi_k(i,n)]$ for any constant $\check{c}_1 > 1$. The exponential random variable is a special case of NBUE type of random variables. This conclusion also applies to the service times $\eta_k(i,n)$, $i \in \mathcal{I}$, $k = 1, \ldots, K$.

Remark 26. Since $\xi_k^r(i,n) = r^{-1}\xi_k(i,n)$ and $\eta_k^r(i,n) = r^{-1}\eta_k(i,n)$, the above assumption implies that for any $n \geq 1$

$$\sup_{1\leq k\leq K}\sup_{i\in\mathcal{I}}\sup_{0\leq u\leq t} \mathbf{E}[\xi_k^r(i,n)|\xi_k^r(i,n) > u] - u \quad < \quad r^{-1}\hat{c}_1(rt), \tag{134}$$

$$\sup_{1\leq k\leq K}\sup_{i\in\mathcal{I}}\sup_{0\leq u\leq t} \mathbf{E}[\eta_k^r(i,n)|\eta_k^r(i,n) > u] - u \quad < \quad r^{-1}\hat{c}_2(rt). \tag{135}$$

- We assume that

$$r^{-1}\hat{c}_n(rt) = o(l(r)) \tag{136}$$

for any fixed $t > 0$, $n = 1, 2$.

Note that if the inter-arrival times $\xi_k(i,n)$ and service times $\eta_k(i,n)$ are all NUBE type of random variables (e.g exponential random variables), then $\hat{c}_n(rt)$, $n = 1, 2$ are constants and the assumption (136) holds since $l(r)$ satisfies the assumption (123).

Remark 27. This assumption can be relaxed such as there exist real valued functions $\check{c}_n(t)$, such that $r^{-1}\hat{c}_n(rt) \leq \check{c}_n(t)$, $n = 1, 2$.

- There exists $n(\cdot)$ such that for any $t \geq 0$,

$$\mathbf{E}N^r(t) \leq n(t), \tag{137}$$

where $N^r(t)$ denotes the number of environment transitions until time t for rth system.

Note that if the transition of environment follows a poisson process with rate $\bar{n}^r$, then $\mathbf{E}N^r(t) = \bar{n}^r t$. For this case, the assumption (137) is satisfied if $\{\bar{n}^r : r \geq 0\}$ is bounded.

We consider a given set of functions Ψ, then it satisfies

- Initial condition:

$$\Psi(0; z, i) = z. \tag{138}$$

- Continuity in t:

$$\lim_{s\to t}\Psi(s; z, i) = \Psi(t; z, i). \tag{139}$$

Furthermore, we assume Ψ satisfies

- Consistency:

$$\Psi(t+s;z,i)=\Psi(t;\Psi(s;z,i),i); \tag{140}$$

- Equi-continuity in z:

$$|\Psi(s;z_1,i)-\Psi(s;z_2,i)|\le L(s)|z_1-z_2|, \tag{141}$$

  where $\sup_{0\le s\le t} L(s)<\infty$ for any $t<\infty$.

With the assumptions above, we have the following result.

Theorem 28. *Let $Z^r(t;\Psi)$ denote the queue length vector at time t of the rth system operating under the policy constructed from the stochastic fluid model solution (Ψ, T^{Ψ}) which has the properties listed by (138)-(141). If the assumptions (86), (133), (136) and (137) hold, and $Z^r(0)/r\to\bar{Z}(0)$ w.p.1, then w.p.1,*

$$Z^r(t;\Psi)/r\to\Psi(t;\bar{Z}(0),X)\ \ \textit{in}\ D_{\mathbb{R}_+^K}[0,\ \infty)\ \ \textit{as}\ r\to\infty, \tag{142}$$

where $\Psi(t;\bar{Z}(0),X)$ is defined recursively in the following way:

$$\begin{aligned}
\Psi(0;z,X) &= z, &(143)\\
\Psi(t;z,X) &= \Psi(t-\tau_{N(t)};\Psi(\tau_{N(t)};z,X),X(\tau_{N(t)})), &(144)\\
\Psi(\tau_{n+1};z,X) &= \Psi\left(\tau_{n+1}-\tau_n,\Psi(\tau_{n-1};z,X),X(\tau_n)\right), &(145)
\end{aligned}$$

and $\tau_{N(t)}$ is the last transition time of $X(\cdot)$ before t.

Remark 29. Note that since (Ψ, T^{Ψ}) satisfies (122) and T^{Ψ} satisfies (95) and (96), then we know $(X,\Psi(;\bar{Z}(0),X),T^{\Psi})$ satisfies the stochastic fluid model equations (91)-(97) with replacing $Z(0)$ by $\bar{Z}(0)$ and $Z(t)$ by $\Psi(t;\bar{Z}(0),X)$.

In order to prove Theorem 28, we first develop some lemmas as follow and the proof of them is provided in Section 3.7.

We first develop the following lemma to provide a probabilistic deviation bound between random variable and its mean value. The proof is similar to that of Lemma 4.3 in Bramson [9]. But for the sake of completeness, we also provide the proof in Section 3.7.

Lemma 30. *For any $\epsilon > 0$ any $i \in \mathcal{I}$, and any $x > 0$, we have*

$$P(|\sum_{n=1}^{m} \eta_k(i,n) - \frac{m}{\mu_k(i)}| \geq \epsilon x) \quad \leq \quad \hat{h}_k(\epsilon, x) \quad \textit{for all integer } 0 \leq m \leq x, \tag{146}$$

$$P(|\sum_{n=1}^{m} \xi_k(i,n) - \frac{m}{\alpha_k(i)}| \geq \epsilon x) \quad \leq \quad \hat{g}_k(\epsilon, x) \quad \textit{for all integer } 0 \leq m \leq x, \tag{147}$$

where

$$\hat{h}_k(\epsilon, x) = x^{-1}(2^8\epsilon^{-4}x^{-1/2} + 4\epsilon^{-2}h_k(x^{-1/8}))$$

$$\hat{g}_k(\epsilon, x) = x^{-1}(2^8\epsilon^{-4}x^{-1/2} + 4\epsilon^{-2}g_k(x^{-1/8})).$$

Remark 31. From (133), we can see that for any fixed $\epsilon > 0$,

$$\hat{h}_k(\epsilon, x) = o(x^{-(1+\gamma/8)}), \qquad \hat{g}_k(\epsilon, x) = o(x^{-(1+\gamma/8)}).$$

Remark 32. We can relax (123) such as $l(r) \to 0$ if $r \to \infty$ and there exists $0 < \epsilon < \gamma/8$ such that $r^{(\gamma/8-\epsilon)}(l(r))^n \to \infty$ for $1 \leq n \leq 2$ if $r \to \infty$. In particular, we can set $\epsilon = \gamma/8 - \gamma/9$.

We also provide the following lemma for later reference.

Lemma 33. *For any $\epsilon > 0$ any $x > 0$, any integer m such that $0 \leq m \leq x$, any $1 \leq l \leq K$, $1 \leq k \leq K$, and any $i \in \mathcal{I}$,*

$$P(|\Phi_k^l(i,m) - mp_{lk}(i)| > \epsilon x) \leq \frac{1}{\epsilon^4 x^2}.$$

For notational convenience, we denote $X^r(t^r(j))$ by X_j^r, then $X_j^r = X_{j+1}^r$ means that the jth review period is not interrupted by an environment transition. We first define a filtration. For $n \geq 0$, let $\mathcal{F}_n$ be the σ-field generated from random variables $\{(t^r(j), q^r(j), X_j^r), j = 0, 1, \ldots, n\}$, i.e

$$\mathcal{F}_n^r = \sigma\{t^r(0), q^r(0), X_0^r, t^r(1), q^r(1), X_1^r, \ldots, t^r(n), q^r(n), X_n^r\}. \tag{148}$$

Recall that $\iota_k^r(i,j)$ denotes the number of class k jobs that have departed at environment state i until $t^r(j)$. Let $\tilde{\eta}_k^r(i,n)$ denote the remaining service time of the nth class k customer that is served at environment state i for every $i \in \mathcal{I}$, $1 \leq k \leq K$ and $n > 0$. For every $i \in \mathcal{I}$ and $j \geq 0$, we let $\Upsilon^r(i,j,\epsilon)$ denote the event that the remaining service time of the

$(\iota_k^r(i,j)+1)$st class k customer served at environment state i is less than $(2K)^{-1}\epsilon l(r)$ for all k, i.e

$$\Upsilon^r(i,j,\epsilon) = \{\tilde{\eta}_k^r(i, \iota_k^r(i,j)+1) < (2K)^{-1}\epsilon l(r), 1 \le k \le K\}. \tag{149}$$

The following lemma provides a probabilistic estimation between the difference of the actual review period length and the planned review period length if the fluid policy is implemented during the review period, the review period is not interrupted by the environment transition and the remaining service times are sufficiently short (i.e $\Upsilon^r(X_j^r, j, \epsilon)$ happens).

For notational convenience, we let

$$\begin{aligned} \Gamma^r(i,j) &= \{X_j^r = i = X_{j+1}^r,\, q^r(j) \ge \theta^r(i)\} \\ \Gamma^r(j) &= \{X_j^r = X_{j+1}^r,\, q^r(j) \ge \theta^r(X_j^r)\}. \end{aligned}$$

Then $\Gamma^r(j)$ denotes that the jth review period is a normal review period and $\Gamma^r(i,j)$ denotes that the jth review period is a normal review period and the environment is at state i during that period.

Lemma 34. *For any $\epsilon > 0$, there exists $r(\epsilon) > 0$, such that if $r > r(\epsilon)$,*

$$\begin{aligned} &\mathbf{E}\Big[\chi(|t^r(j+1) - t^r(j) - l(r)| > \epsilon l(r))\, \chi(\Upsilon^r(X_j^r, j, \epsilon))\, \chi(\Gamma^r(j)) \Big| \mathcal{F}_j^r\Big] \\ \le &\sum_{k=1}^K \hat{h}_k(\epsilon(4K\bar{\mu}_k)^{-1}, \lfloor \underline{\mu}_k r l(r)\rfloor)\chi(q^r(j) \ge \theta^r(X_j^r)), \end{aligned}$$

where $\hat{h}_k(\cdot,\cdot)$, $1 \le k \le K$ is defined in Lemma 30.

The next lemma is a special case of Lemma 34, but it does not require that the remaining service times are small. In this lemma, we provide a probabilistic lower bound for the duration of a normal review period.

Lemma 35. *For any $\epsilon > 0$, there exists $r(\epsilon) > 0$, such that if $r > r(\epsilon)$,*

$$\begin{aligned} &\mathbf{E}\Big[\chi(t^r(j+1) - t^r(j) < (1-\epsilon) l(r))\, \chi(\Gamma^r(j)) \Big| \mathcal{F}_j^r\Big] \\ \le &\sum_{k=1}^K \hat{h}_k(\epsilon(2K\bar{\mu}_k)^{-1}, \lfloor \underline{\mu}_k r l(r)\rfloor)\chi(q^r(j) \ge \theta^r(X_j^r)), \end{aligned}$$

where $\hat{h}_k(\cdot,\cdot)$, $1 \le k \le K$ is defined in Lemma 30.

The next lemma is also a special case of Lemma 34, where we only provided a probabilistic upper bound for the duration of a review period when the fluid policy is implemented, but this review period might be interrupted by an environment transition.

Lemma 36. *For any $\epsilon > 0$, there exists $r(\epsilon) > 0$, such that if $r > r(\epsilon)$,*

$$\begin{aligned} & \mathbf{E}\Big[\chi(t^r(j+1) - t^r(j) > (1+\epsilon)l(r))\chi(\Upsilon^r(X_j^r, j, \epsilon)\,\chi(q^r(j) \geq \theta^r(X_j^r))\Big|\mathcal{F}_j^r\Big] \\ \leq & \sum_{k=1}^{K} \hat{h}_k(\epsilon(4K\bar{\mu}_k)^{-1}, \lfloor \underline{\mu}_k rl(r)\rfloor)\chi(q^r(j) \geq \theta^r(X_j^r)), \end{aligned}$$

where $\hat{h}_k(\cdot,\cdot)$, $1 \leq k \leq K$ is defined in Lemma 30.

For the rest of the manuscript, we adopt the following conventions. For a real number x, $|x| = \max\{x, -x\}$, $\lceil x \rceil$ denotes the smallest integer that is bigger than or equal to x, and $\lfloor x \rfloor$ denotes the largest integer that is smaller than or equal to x.

The following lemma provides a probabilistic bound between the difference of the actual exogenous arrivals and average number of exogenous arrivals.

Lemma 37. *For any $x > 0$, any $t \geq 0$, any state $i \in \mathcal{I}$, and any class k such that $\alpha_k(i) > 0$,*

$$\begin{aligned} (i) \qquad & \mathbf{P}(E_k^r(i,t) - r\alpha_k(i)t > x) \leq \hat{g}_k(\frac{x}{\alpha_k(i)\lceil r\alpha_k(i)t + x\rceil}, \lceil r\alpha_k(i)t + x\rceil), \\ (ii) \qquad & \mathbf{P}(E_k^r(i,t) - r\alpha_k(i)t < -x) \leq \hat{g}_k(\frac{x}{\alpha_k(i)\lfloor r\alpha_k(i)t - x\rfloor}, \lfloor r\alpha_k(i)t - x\rfloor), \\ (iii) \qquad & P(|E_k^r(i,t) - r\alpha_k(i)t| > x) \leq 2\hat{g}_k(\frac{x}{\alpha_k(i)(\lceil \alpha_k(i)rt + x\rceil)}, \lfloor \alpha_k(i)rt + x\rfloor), \end{aligned}$$

where $\hat{g}_k(x,y)$ is defined in Lemma 30.

The following lemma provides the probabilistic bound between the difference of the actual service completions and the mean service completions.

Lemma 38. *For any $\epsilon > 0$, any $t \geq 0$, any state $i \in \mathcal{I}$, and any class k such that $\mu_k(i) > 0$,*

$$P(|S_k^r(i,t) - r\mu_k(i)t| > x) \quad \leq \quad 2\hat{h}_k(\frac{x}{\bar{\mu}_k(\lceil \bar{\mu}_k rt + x\rceil)}, \lfloor \underline{\mu}_k rt + x\rfloor).$$

Let $E_k^r(t) = \sum_{i\in\mathcal{I}} E_k^r(i, I^r(i,t))$, then $E_k^r(t)$ denotes the total number of exogenous arrivals to class k until time t. The following lemma provides us a probabilistic bound on the total number of external arrivals within a finite time period.

Lemma 39. *For each class k with $\bar{\alpha}_k > 0$, for any $t > 0$, and any $\epsilon > 0$, there exists $r_k(t,\epsilon) > 0$, for all $r > r_k(t,\epsilon)$, such that*

$$\mathbf{P}(E^r_k(t) > r(\bar{\alpha}_k + \epsilon)t) \quad \leq \quad \hat{g}_k(\frac{\epsilon}{2(|\mathcal{I}|(\bar{\alpha}_k)^2 + \epsilon)}, \epsilon rt/|\mathcal{I}|).$$

Let $S^r_k(t) = \sum_{i\in\mathcal{I}} S^r_k(i, T^r_k(i,t))$, then $S^r_k(t)$ denotes the total number of service completions of class k until time t. The following lemma provides a probabilistic bound on the total number of service completions of class k within a finite time period.

Lemma 40. *For each class k with $\bar{\mu}_k > 0$, any $t > 0$, and any $\epsilon > 0$, there exists $r_k(t,\epsilon) > 0$, for all $r > r_k(t,\epsilon)$, such that*

$$\mathbf{P}(S^r_k(t) > r(\bar{\mu}_k + \epsilon)t) \quad \leq \quad \hat{h}_k(\frac{\epsilon}{2(|\mathcal{I}|(\bar{\mu}_k)^2 + \epsilon)}, \epsilon rt/|\mathcal{I}|).$$

The following lemma provides a probabilistic bound on the maximum inter-arrival time between consecutive customers who have come before time t.

Lemma 41. *For each class k such that $\bar{\alpha}_k > 0$, any $t > 0$ and any $\epsilon > 0$, there exists $r(t,\epsilon)$ such that if $r > r(t,\epsilon)$, then*

$$\mathbf{P}\Big(\max_{i\in\mathcal{I}} \max_{1\leq n\leq E^r_k(i,I^r(i,t))} \xi^r_k(i,n) > \epsilon l(r)\Big) \quad \leq \quad \hat{f}_1(\epsilon,t,r),$$

where

$$\hat{f}_1(\epsilon,t,r) = \frac{c_1(\bar{\alpha}_k+\epsilon)t}{r^{1+\gamma}(l(r))^2} + \hat{g}_k\Big(\frac{\epsilon}{2(|\mathcal{I}|(\bar{\alpha}_k)^2 + \epsilon)}, \epsilon rt/|\mathcal{I}|\Big).$$

Remark 42. From Remark 31, we know that $\hat{f}_1(\epsilon,t,r) = o(r^{-(1+\gamma/8)})$ for any fixed $\epsilon > 0$ and $t > 0$.

The next lemma provides a probabilistic upper bound on the service time of customers who have been served before time t.

Lemma 43. *For each class k such that $\bar{\mu}_k > 0$, any $t > 0$ and any $\epsilon > 0$, there exists $r(t,\epsilon)$ such that if $r > r(t,\epsilon)$, then*

$$\mathbf{P}\Big(\max_{i\in\mathcal{I}} \max_{1\leq n\leq S^r_k(i,T^r(i,t))} \eta^r_k(i,n) > \epsilon l(r)\Big) \quad \leq \quad \hat{f}_2(\epsilon,t,r),$$

where

$$\hat{f}_2(\epsilon,t,r) = \frac{c_2(\bar{\mu}_k+\epsilon)t}{r^{1+\gamma}(l(r))^2} + \hat{h}_k\Big(\frac{\epsilon}{2(|\mathcal{I}|(\bar{\mu}_k)^2 + \epsilon)}, \epsilon rt/|\mathcal{I}|\Big).$$

Remark 44. From Remark 31 and $r^{7\gamma/8}(l(r))^2 \to 0$, we know that $\hat{f}_2(\epsilon, t, r) = o(r^{-(1+\gamma/8)})$ for any fixed $\epsilon > 0$ and $t > 0$.

Recall that $\kappa_k^r(i,j)$ denotes the number of class k customers that arrive at environment state i before time $t^r(j)$, and $\tilde{\xi}_k^r(i, \kappa_k^r(i,j)+1)$ denotes the remaining time until the first class k customer arrives at environment state i after $t^r(j)$. For every $i \in \mathcal{I}$ and $j \geq 0$, we let $\Lambda^r(i,j,\epsilon)$ denote the event that the remaining time of the $(\kappa_k^r(i,j)+1)$st class k customer to arrive at environment state i after $t^r(j)$ is less than $(2K)^{-1}\epsilon l(r)$ for all k, i.e

$$\Lambda^r(i,j,\epsilon) = \{\tilde{\xi}_k^r(i, \kappa_k^r(i,j)+1) < (2K)^{-1}\epsilon l(r), 1 \leq k \leq K\}. \tag{150}$$

At the beginning of the jth review period, if the initial queue length is above the chosen safety stock level, we will implement the fluid policy as is mentioned in Section 3.4. Let the planned queue length at the end of the jth review period be denoted by $z^r(j+1)$, then it satisfies

$$z^r(j+1) \quad = \quad \Psi(l(r); \bar{q}^r(j), X_j^r).$$

The following lemma provides an estimate for the difference between the planned queue length and the actual queue length at the end of the jth review period if the jth review period is a normal review period (i.e the jth review period is not interrupted by a change of the environment state and the fluid policy is implemented during this period) and the remaining inter-arrival times for each class are sufficiently short (i.e $\Lambda^r(i,j,\epsilon)$ happens).

Lemma 45. *For each class k, any $\epsilon > 0$, there exists $r_k(\epsilon)$, such that if $r > r_k(\epsilon)$, we have*

$$\begin{aligned}
&\mathbf{P}(\{|\bar{q}_k^r(j+1) - z_k^r(j+1)| \geq \epsilon r^{-1}\theta_k^r(X_j^r)\} \cap \Gamma^r(j) \cap \Upsilon(X_j^r, j, \frac{\epsilon \underline{\beta}_k}{16\bar{\alpha}_k}) \cap \Lambda(X_j^r, j, \frac{\epsilon \underline{\beta}_k}{16\bar{\alpha}_k})|\mathcal{F}_j^r) \\
\leq \quad & \hat{f}_3(\epsilon, r)\chi(q^r(j) \geq \theta^r(X_j^r)),
\end{aligned}$$

where

$$\begin{aligned}
\hat{f}_3(\epsilon, r) \quad = \quad & 2\hat{g}_k((\frac{\epsilon \underline{\beta}_k}{32\bar{\alpha}_k})(\bar{\alpha}_k + \frac{\epsilon \bar{\beta}_k}{4})^{-1}, \underline{\alpha}_k r l(r)) \\
& + \sum_{k=1}^{K} \hat{h}_k(\frac{\epsilon \underline{\beta}_k}{64K\bar{\alpha}_k\bar{\mu}_k}, \lfloor \underline{\mu}_k r l(r) \rfloor) + \frac{(4K)^4}{\epsilon^4}\frac{1}{(rl(r)\underline{\beta}_k)^2}
\end{aligned}$$

Remark 46. From (123) and Remark 31, we know that $\hat{f}_3(\epsilon, r) = o(r^{-(1+\gamma/9)})$ for any fixed $\epsilon > 0$.

The next lemma provides an estimate for the queue length at the end of a normal review period, i.e. during the review period the fluid policy is implemented without interruption. It shows that if at the end of a normal review period the actual queue length is close to the planned queue length, then the actual queue length will be either above the safety stock level, or very close to the safety stock level when it is not above the safety stock level.

Lemma 47. *For any $\epsilon > 0$, there exists $r(\epsilon) > 0$ such that if $r > r(\epsilon)$,*

$$\begin{aligned} & \mathbf{P}(\{q^r(j+1) \not\geq (1-\epsilon)\theta^r(X^r_{j+1})\} \cap \Gamma^r(j) \cap \Upsilon(X^r_j, j, \frac{\epsilon \underline{\beta}}{16\bar{\alpha}}) \cap \Lambda(X^r_j, j, \frac{\epsilon \underline{\beta}}{16\bar{\alpha}}) | \mathcal{F}^r_j) \\ \leq \quad & \hat{f}_4(\epsilon, r)\chi(q^r(j) \geq \theta^r(X^r_j)), \end{aligned}$$

where

$$\begin{aligned} \hat{f}_4(\epsilon, r) \quad = \quad & \sum_{k=1}^{K} \Big(2\hat{g}_k((\frac{\epsilon \underline{\beta}_k}{32\bar{\alpha}_k})(\bar{\alpha}_k + \frac{\epsilon \bar{\beta}_k}{4})^{-1}, \underline{\alpha}_k rl(r)) \\ & + \hat{h}_k(\frac{\epsilon \underline{\beta}_k}{64K\bar{\alpha}_k\bar{\mu}_k}, \lfloor \underline{\mu}_k rl(r) \rfloor) + \frac{(4K)^4}{\epsilon^4}\frac{1}{(rl(r)\underline{\beta}_k)^2} \Big). \end{aligned}$$

Remark 48. From (123) and Remark 31, we know that $\hat{f}_4(\epsilon, r) = o(r^{-(1+\gamma/9)})$ for any fixed $\epsilon > 0$.

At the beginning of a review period, if the queue length of some class is not above its safety stock level, we implement the target-idle policy so that each class will be above its safety stock level. The following lemma provides us with an estimate on the duration of such a review period during which the target-idle policy is implemented. For notational convenience, we let

$$\tilde{\Gamma}^r(j, \epsilon) = \{q^r(j) \not\geq \theta^r(X^r_j),\, q^r(j) \geq (1-\epsilon)\theta^r(X^r_j)\} \tag{151}$$

Lemma 49. *For any any $\epsilon > 0$, we have*

$$\begin{aligned} & \mathbf{P}\Big(\{t^r(j+1) - t^r(j) > 2K\tilde{C}_1\epsilon l(r)\} \cap \Lambda^r(X^r_j, j, \epsilon) \cap \Upsilon^r(X^r_j, j, \epsilon) \cap \tilde{\Gamma}^r(j, \epsilon) | \mathcal{F}^r_j\Big) \\ \leq \quad & \hat{f}_5(\epsilon, r)\chi(\tilde{\Gamma}^r(j, \epsilon)), \end{aligned}$$

where $\tilde{C}_1$ is a constant which is defined in (223),

$$\hat{f}_5(\epsilon, r) = \sum_{1 \le k \le K} \hat{g}_k(\frac{\epsilon}{\bar{\alpha}}, \underline{\mu} rl(r)) + \sum_{1 \le k \le K} \hat{h}_k(\frac{\epsilon}{\bar{\mu}}, 2\underline{\mu} rl(r)) + K^3(\frac{\epsilon}{2\tilde{c}_{\max}})^{-4}(\underline{\beta} rl(r))^{-2},$$

and $\tilde{c}_{\max}$ is a constant which is defined in (222).

Remark 50. From (123) and Remark 31, we know that $\hat{f}_5(\epsilon, r) = o(r^{-(1+\gamma/9)})$ for any fixed $\epsilon > 0$.

Next, we provide a probabilistic upper bound on the difference between the queue length at the end of a target-idle review period and the queue length at the beginning of this review period.

Lemma 51. *For any $\epsilon > 0$, there exists $r(\epsilon) > 0$ such that if $r > r(\epsilon)$, then*

$$\mathbf{P}(\{|q^r(j+1) - q^r(j)| > \tilde{C}_2 \epsilon rl(r)\mathbf{e}\} \cap \Lambda^r(X_j^r, j, \epsilon)\Upsilon^r(X_j^r, j, \epsilon)\tilde{\Gamma}^r(j, \epsilon)|\mathcal{F}_j^r) \le \hat{f}_6(\epsilon, r),$$

where $\tilde{C}_2 = 4K\tilde{C}_1\bar{\alpha} + 2(K+1)\tilde{c}_{\max}\bar{\beta}$. The notation $\mathbf{e}$ is a K dimensional vector with all elements equal to 1 and

$$\hat{f}_6(\epsilon, r) = \sum_{1 \le k \le K} \Big(\hat{g}_k(\frac{1}{4\bar{\alpha}}, 4K\tilde{C}_1\bar{\alpha}\epsilon rl(r)) + K^2(\frac{\epsilon}{2\tilde{c}_{\max}})^{-4}(\underline{\beta}\epsilon rl(r))^{-2} + \hat{f}_5(\epsilon, r)\Big).$$

Remark 52. From (123), Remark 31 and Remark 50, we know that $\hat{f}_6(\epsilon, r) = o(r^{-(1+\gamma/9)})$ for any fixed $\epsilon > 0$.

The next lemma generalizes the result of Lemma 45 and provides a probabilistic upper bound on the difference between the queue lengths at the end and at the beginning of a review period when the fluid policy is implemented. This result applies to both interrupted and uninterrupted review periods.

Lemma 53. *For any $\epsilon > 0$, there exists $r(\epsilon) > 0$ such that if $r > r(\epsilon)$, then*

$$\begin{aligned} & \mathbf{P}(\{|q^r(j+1) - q^r(j)| > \hat{C}_0 rl(r)\mathbf{e}\} \cap \Upsilon^r(X_j^r, j, \epsilon) \cap \chi(q^r(j) \ge \theta^r(X_j^r))|\mathcal{F}_j^r) \\ \le \quad & \hat{f}_7(\epsilon, r)\chi(q^r(j) \ge \theta^r(X_j^r)), \end{aligned}$$

where $\hat{C}_0 = 2\bar{\alpha}(1+\epsilon) + (K+1)\bar{\mu}$ and

$$\hat{f}_7(\epsilon, r) = \sum_{k=1}^{K} \hat{g}_k(\frac{1}{4\bar{\alpha}}, 2\bar{\alpha}(1+\epsilon)rl(r)) + \sum_{k=1}^{K} \hat{h}_k(\epsilon(4K\bar{\mu}_k)^{-1}, \lfloor \mu_k rl(r) \rfloor).$$

Remark 54. From (123), Remark 31 and Remark 50, we know that $\hat{f}_7(\epsilon, r) = o(r^{-(1+\gamma/9)})$ for any fixed $\epsilon > 0$.

Let $j^r(t)$ denote the index of the first review period that starts after t, i.e.

$$j^r(t) \quad = \quad \min\{j \geq 0 : t^r(j) \geq t\}. \tag{152}$$

The next lemma gives us an upper bound on the expected duration of a review period before t.

Lemma 55. *There exists a real valued function $f(\cdot)$ defined on $[0, \infty)$ such that for every $0 \leq j \leq j^r(t) - 1$,*

$$E[t^r(j+1) - t^r(j)] < f(t, r),$$

where

$$\begin{aligned} f(t,r) \;=\; & \max\Big\{(K+1)l(r) + r^{-1}\hat{c}_2(rt), \\ & \Big(2\tilde{c}_{\max}K^2 c_3 \beta_{\max} \sum_{k=1}^{K}(\bar{\alpha}_k + \bar{\mu}_k) + \beta_{\max}\sum_{k=1}^{K}\bar{\alpha}_k\Big)l(r) + 4Kr^{-1} + r^{-1}\hat{c}_1(rt) + r^{-1}\hat{c}_2(rt)\Big\}. \end{aligned}$$

Remark 56. From the assumptions on $\hat{c}_n(\cdot)$, $n = 1, 2$, we know that there exists a constant $c_4 > 0$ such that $f(r) \leq c_4 l(r)$.

We have the following lemma which provides us an upper bound on the expected number of review periods until t.

Lemma 57. *For any ϵ such that $0 < \epsilon < 1$, there exists $r(\epsilon)$, such that if $r > r(\epsilon)$, we have*

$$E[j^r(t)] \quad \leq \quad \hat{f}_7(\epsilon, t, r),$$

where

$$\hat{f}_7(\epsilon, t, r) = \frac{2(t + f(t,r)) + 3(1-\epsilon)l(r)E[N^r(t)]}{(1-\epsilon)^2 l(r)},$$

and $f(t, r)$ is defined in Lemma 55.

Remark 58. From (56) and the assumption that $\mathbf{E}[N^r(t)] < n(t)$, we know that there exists a real valued function $\tilde{f}_7(\epsilon, t)$ such that $\tilde{f}_7(\epsilon, t) > 0$ and $\hat{f}_7(\epsilon, t, r) < \tilde{f}_7(\epsilon, t)$ for all $r > 0$.

For the rest of the manuscript, we adopt the following notation:

$$\begin{aligned}
\Lambda^r(t,\epsilon) &= \{\max_{i\in\mathcal{I}} \max_{1\le n\le E_k^r(i,I^r(i,t))} \xi_k^r(i,n) > \epsilon l(r)\}^c, \\
\Upsilon^r(t,\epsilon) &= \{\max_{i\in\mathcal{I}} \max_{1\le n\le S_k^r(i,T^r(i,t))} \eta_k^r(i,n) > \epsilon l(r)\}^c.
\end{aligned}$$

It satisfies that

$$\Lambda^r(t,\epsilon) \subset \bigcap_{1\le j\le j^r(t)} \Lambda(X_{j-1}^r, j-1, \epsilon), \quad \text{and} \quad \Upsilon^r(t,\epsilon) \subset \bigcap_{1\le j\le j^r(t)} \Upsilon(X_{j-1}^r, j-1, \epsilon). \tag{153}$$

The following lemma reveals that the queue length at the end of an uninterrupted review period will be above or will be close to the safety stock level if the network speed is large enough. That is, the queue length at the end of a review period will not be far below the safety stock level.

Lemma 59. *For every $t > 0$, $\epsilon > 0$, and almost every ω, there exists $r(\omega, t, \epsilon)$, such that if $r > r(\omega, t, \epsilon)$, then for all $1 \le j \le j^r(t)$,*

$$\chi(X_j^r = X_{j-1}^r) q^r(j) \quad \ge \quad \chi(X_j^r = X_{j-1}^r)(1-\epsilon)\theta^r(X_j^r).$$

The following lemma estimates the duration of an uninterrupted review period during which the *fluid policy* is implemented. The lemma tells us that the duration will be close to the initially planned review period length almost surely if the network speed is large enough.

Lemma 60. *For every $t > 0$, $\epsilon > 0$, and almost every ω, there exists $r(\omega, t, \epsilon)$, such that if $r > r(\omega, t, \epsilon)$, then for all $1 \le j \le j^r(t)$,*

$$\chi\Big(X_j^r = X_{j-1}^r,\, q^r(j-1) \ge \theta^r(X_{j-1}^r)\Big)|(t^r(j) - t^r(j-1)) - l(r)| \quad < \quad \epsilon l(r).$$

The next lemma provides an upper bound on the duration of a review period when the fluid policy is implemented. This bound holds almost surely compared to the probabilistic one provided in Lemma 36.

Lemma 61. *For every $t > 0$, $\epsilon > 0$, and almost every ω, there exists $r(\omega, t, \epsilon)$, such that if $r > r(\omega, t, \epsilon)$, then for all $1 \le j \le j^r(t)$,*

$$\chi\Big(q^r(j-1) \ge \theta^r(X_{j-1}^r)\Big)(t^r(j) - t^r(j-1)) \quad < \quad (1+\epsilon)l(r).$$

The following lemma estimates the duration of an uninterrupted review period when the *target idle* policy is implemented. Almost surely, the duration will be at most in the same order of $l(r)$ compared to the probabilistic bound provided in Lemma 49 .

Lemma 62. *For every $t > 0$, $\epsilon > 0$, and almost every ω, there exists $r(\omega, t, \epsilon)$ and if $r > r(\omega, t, \epsilon)$, then for all $1 \le j \le j^r(t)$,*

$$\chi\Big(q^r(j-1) \ngeq \theta^r(X^r_{j-1}),\, q^r(j-1) \ge (1-\epsilon)\theta^r(X^r_{j-1})\Big)(t^r(j) - t^r(j-1)) < 2K\tilde{C}_1\epsilon l(r).$$

The following lemma follows the results of Lemma 61 and Lemma 62 and provides an upper bound on the duration of a review period, including those when the fluid policy is implemented and those when the target idle policy is implemented.

Lemma 63. *For every $t > 0$, and almost every ω, there exists $r(\omega, t)$, such that if $r > r(\omega, t)$, then for all $1 \le j \le j^r(t)$,*

$$t^r(j) - t^r(j-1) \quad < \quad 2K\tilde{C}_1 l(r).$$

The following lemma provides an upper bound on the difference between the actual queue length and the planned queue length at the end of a normal review period, i.e an uninterrupted review period when the *fluid policy* is implemented. In particular, it shows that this upper bound holds almost surely compared to the probabilistic bound provided by Lemma 45. This lemma reveals that the actual queue length will be almost surely close to the planned level obtained from the stochastic fluid model solution when the network processing speed is fast enough.

Lemma 64. *For every t, $\epsilon > 0$, and almost every ω, there exists $r(\omega, t, \epsilon)$, such that if $r > r(\omega, t, \epsilon)$, then for all $1 \le j \le j^r(t)$,*

$$\chi\Big(X^r_j = X^r_{j-1},\, q^r(j-1) \ge \theta^r(X^r_{j-1})\Big)|z^r(j) - \bar{q}^r(j)| \quad \le \quad \epsilon r^{-1}\theta^r(X^r_{j-1}),$$

where $z^r(j) = \Psi(l(r); \bar{q}^r(j-1), X^r_{j-1})$.

The following lemma provides an estimation on the difference between the queue lengths at the beginning and at the end of a target idle review period, i.e a review period when the target idle policy is implemented. In particular, it shows that this difference will be at most of the same order as $rl(r)$ almost surely, while we have provided a probabilistic upper bound on this difference in Lemma 51.

Lemma 65. *For every $t > 0$, $\epsilon > 0$, and almost every ω, there exists $r(\omega, t, \epsilon)$, such that if $r > r(\omega, t, \epsilon)$, then for all $1 \le j \le j^r(t)$,*

$$\chi\Big(q^r(j-1) \ngeq \theta^r(X^r_{j-1}),\, q^r(j-1) \ge (1-\epsilon)\theta^r(X^r_{j-1})\Big)|q^r(j) - q^r(j-1)| \quad \le \quad \tilde{C}_2 \epsilon r l(r)\mathbf{e}.$$

Remark 66. Recall that $\bar{q}^r(j) = r^{-1}(q^r(j) - \theta^r(X^r_j))$ for all $j \ge 0$. So if $X^r_j = X^r_{j+1}$, then $\theta^r(X^r_j) = \theta^r(X^r_{j+1})$ and $\bar{q}^r(j+1) - \bar{q}^r(j) = r^{-1}(q^r(j+1) - q^r(j))$. Therefore, for every $t > 0$, $\epsilon > 0$, and almost every ω, there exists $r(\omega, t, \epsilon)$, such that if $r > r(\omega, t, \epsilon)$, for all $1 \le j \le j^r(t)$,

$$\begin{aligned} &\chi\Big(q^r(j-1) \ngeq \theta^r(X^r_{j-1}),\, q^r(j-1) \ge (1-\epsilon)\theta^r(X^r_{j-1}), X^r_j = X^r_{j-1}\Big)|\bar{q}^r(j) - \bar{q}^r(j-1)| \\ \le \quad & \tilde{C}_2 \epsilon l(r)\mathbf{e}. \end{aligned}$$

The following lemma provides an estimate for the difference between the queue lengths at the end and at the beginning of a fluid review period, i.e a review period when the fluid policy is implemented. This result applies to both interrupted and uninterrupted fluid review periods. In particular, it shows that this difference will be at most of the same order as $rl(r)$ almost surely, while we have provided a probabilistic upper bound on this difference in Lemma 53.

Lemma 67. *For every $t > 0$, $\epsilon > 0$, and almost every ω, there exists $r(\omega, t, \epsilon)$, such that if $r > r(\omega, t, \epsilon)$, for all $1 \le j \le j^r(t)$,*

$$\chi\Big(q^r(j-1) \ge \theta^r(X^r_{j-1})\Big)|q^r(j) - q^r(j-1)| \quad \le \quad \hat{C}_0 r l(r)\mathbf{e}.$$

Following the results of Lemma 65 and Lemma 67, we provide an upper bound on the difference between the queue lengths at the end and at the beginning of a review period. This upper bound holds almost surely for each review period.

Lemma 68. *For every $t > 0$ and almost every ω, there exists $r(\omega, t)$, such that as $r > r(\omega, t)$, for all $1 \leq j \leq j^r(t)$,*

$$|q^r(j) - q^r(j-1)| \quad \leq \quad \tilde{C}_3 r l(r) \mathbf{e},$$

where $\tilde{C}_3$ is a constant and $\tilde{C}_3 = \max\{\tilde{C}_2, 4\bar{\alpha} + (K+1)\bar{\mu}\}$.

Remark 69. Following Lemma 68, and noting that

$$|\bar{q}^r)(j) - \bar{q}^r(j-1)| \leq r^{-1}|q^r(j) - q^r(j-1)| + r^{-1}(\theta^r(X_j^r) + \theta^r(X_{j-1}^r)),$$

we have the following result.

For every $t > 0$ and almost every ω, there exists $r(\omega, t)$, such that if $r > r(\omega, t)$, then for all $1 \leq j \leq j^r(t)$,

$$|\bar{q}^r(j) - \bar{q}^r(j-1)| \leq \tilde{C}_4 l(r) \mathbf{e},$$

where $\tilde{C}_4 = \tilde{C}_3 + 2\bar{\beta}$.

Next lemma reveals that the first review period after time s will start almost at time s if the network speed is large enough. The proof follows the result of Lemma 63.

Lemma 70. *For any $t > 0$, almost every ω, and every $\epsilon > 0$, there exists $r(\omega, t, \epsilon)$, such that if $r > r(\omega, t, \epsilon)$, then*

$$\sup_{0 \leq s \leq t} |t^r(j^r(s)) - s| \leq \epsilon.$$

That is,

$$t^r(j^r(\cdot)) \to \mathbf{1}(\cdot) \quad u.o.c \quad as\ r \to \infty.$$

The following lemma says that the *fluid policy* is implemented in an uninterrupted fashion most of the time.

Lemma 71. *Let $n^r(j,t) = \sum_{n=j}^{j^r(t)-1} \chi(q^r(n) \geq \theta^r(X_n^r),\, X_{n+1}^r = X_n^r)$. With probability one,*

$$n^r(0, \cdot) l(r) \to \mathbf{1}(\cdot) \quad u.o.c \quad as\ r \to \infty.$$

Remark 72. From the definition, $n^r(j,t)$ denotes the total number of normal review periods (uninterrupted review periods when the *fluid policy* is implemented) up to time t.

To make the proof of Theorem 28 more compact, we also develop the following lemma.

Lemma 73. *Let $j^r(s)$ denote the index of the first review period after s for $s \geq 0$. For every $t > 0$, $\epsilon > 0$ and almost all ω, there exists $r(\omega, t, \epsilon)$, such that if $r > r(\omega, t, \epsilon)$,*

$$\sup_{0 \leq s \leq t} |r^{-1} Z^r(s) - \bar{q}^r(j^r(s))| \leq \epsilon,$$

where

$$\bar{q}^r(j) = r^{-1}(q^r(j) - \theta^r(X_j^r))^+, \quad \textit{for all } j \geq 0.$$

In other words, with probability one,

$$\lim_{r \to \infty} \sup_{0 \leq s \leq t} |r^{-1} Z^r(s) - \bar{q}^r(j^r(s))| = 0.$$

The proof of Theorem 28 is as follows.

Proof of Theorem 28. For each fixed time $t_0 \geq 0$, we consider a sample path ω such that for any $\epsilon > 0$, there exists $r(\omega, t_0, \epsilon) > 0$ and if $r > r(\omega, t_0, \epsilon)$, the results in Lemma 59 to Lemma 71 hold. Note that such sample paths exist almost surely. Throughout the rest of the proof, we consider this sample path ω though it is not spelled out explicitly for notational convenience.

For every fixed $t_0 \geq 0$, we let $\gamma^r(\cdot)$ be defined as in the proof of Lemma 15 such that (101) is satisfied, then we want to show that

$$\lim_{r \to \infty} \sup_{0 \leq s \leq t_0} |r^{-1} Z^r(\gamma^r(s)) - \Psi(s; \bar{Z}(0), X)| \quad = \quad 0, \tag{154}$$

where $\Psi(\cdot; \bar{Z}(0), X)$ is defined through (143)-(145). We will prove it through induction.

Since $X(\cdot)$ satisfies the regularity condition, we know that for fixed $t_0 \geq 0$, there exists a finite $m \geq 0$, such that

$$\tau_m \leq t_0 < \tau_{m+1}. \tag{155}$$

We use induction. First, we want to show that for any finite $t \in [0, \tau_1]$,

$$\lim_{r\to\infty} \sup_{0\le s\le t} |r^{-1}Z^r(\gamma^r(s)) - \Psi(s; \bar{Z}(0), X)| \quad = \quad 0. \tag{156}$$

Then, assuming that for any finite $t \in [0, \tau_n]$,

$$\lim_{r\to\infty} \sup_{0\le s\le t} |r^{-1}Z^r(\gamma^r(s)) - \Psi(s; \bar{Z}(0), X)| \quad = \quad 0, \tag{157}$$

we will show that for any finite $t \in [0, \tau_{n+1}]$,

$$\lim_{r\to\infty} \sup_{0\le s\le t} |r^{-1}Z^r(\gamma^r(s)) - \Psi(s; \bar{Z}(0), X)| \quad = \quad 0. \tag{158}$$

This will conclude the proof.

We first show that (156) holds. Consider any finite t such that $t \in [0, \tau_1]$. First, following the convergence together theorem (Lemma 17), Lemma 73 and (101) imply that

$$\lim_{r\to\infty} \sup_{0\le s\le t} |r^{-1}Z^r(\gamma^r(s)) - \bar{q}^r(j^r(\gamma^r(s)))| = 0. \tag{159}$$

Next, we compare the difference between $\bar{q}^r(j^r(\gamma^r(s)))$ and $\Psi(n^r(0,\gamma^r(s))l(r); \bar{q}^r(0), X)$ for all s such that $0 \le s \le t$. Note that $\bar{q}^r(j^r(\gamma^r(s)))$ is the scaled actual queue length at the end of the last review period by $\gamma^r(s)$, $\Psi(n^r(0,\gamma^r(s))l(r); \bar{q}^r(0), X)$ is from the given stochastic fluid model solution with the initial fluid level being $\bar{q}^r(0)$; and $n^r(0,\cdot)$ is defined in Lemma 71. We will find out how far the actual queue length is from the planned level by analyzing their difference. We first illustrate the idea of characterizing the difference inductively. Then we give a complete representation of the difference. Recall the definition of $z^r(j)$ given in Lemma 64. Generally, $z^r(j)$ denotes the planned fluid level at the end of the $(j-1)$th review period if the fluid level at the beginning of the $(j-1)$th review period is $\bar{q}^r(j-1)$, i.e

$$z^r(j) \quad = \quad \Psi(l(r); \bar{q}^r(j-1), X^r_{j-1}).$$

To illustrate the idea, without loss of generality, we assume $X^r_{j-1} = X^r_{j-2} = i$. Note that Ψ satisfies the initial condition. Thus,

$$\begin{aligned}
\bar{q}^r(j) \quad &= \quad \Psi(0; \bar{q}^r(j), i) \\
&= \quad \Psi(0; \bar{q}^r(j), i) - \Psi(0; z^r(j), i) + \Psi(0; z^r(j), i) \\
&= \quad \Psi(0; \bar{q}^r(j), i) - \Psi(0; z^r(j), i) + z^r(j).
\end{aligned} \tag{160}$$

From the definition of $z^r(j)$ and the assumption that $X^r_{j-1} = i$, we have

$$
\begin{aligned}
z^r(j) &= \Psi(l(r); \bar{q}^r(j-1), i) \\
&= (\Psi(l(r); \bar{q}^r(j-1), i) - \Psi(l(r); z^r(j-1), i)) + \Psi(l(r); z^r(j-1), i). \quad (161)
\end{aligned}
$$

From the assumption that $X^r_{j-2} = i$ and the definition that of $z^r(\cdot)$, we know $z^r(j-1) = \Psi(l(r); \bar{q}^r(j-2), i)$. Recall that the stochastic fluid model solution Ψ also satisfies the *consistency* condition. That is

$$
\Psi(l(r); \bar{z}^r(j-1), i) = \Psi(2l(r); \bar{q}^r(j-2), i). \quad (162)
$$

Now, from (160), (161), and (162), we get

$$
\begin{aligned}
&\bar{q}^r(j) - \Psi(2l(r); \bar{q}^r(j-2), i) \\
&= \Big(\Psi(0; \bar{q}^r(j), i) - \Psi(0; z^r(j), i)\Big) + \Big(\Psi(l(r); \bar{q}^r(j-1), i) - \Psi(l(r); z^r(j-1), i)\Big). (163)
\end{aligned}
$$

With this idea, we can characterize the difference between the scaled actual queue length and the fluid trajectory inductively. We characterize this difference as follows.

For any s such that $0 \le s \le t$ ($t \le \tau_1$), we characterize the difference between $\bar{q}^r(j^r(\gamma^r(s)))$ and $\Psi(n^r(0, \gamma^r(s))l(r); \bar{q}^r(0), i_0)$. Note that $n^r(0, \gamma^r(s))$ counts the total number of normal review periods until time $\gamma^r(s)$ and $l(r)$ is the planned review period length. Thus, $n^r(0, \gamma^r(s))l(r)$ estimates the total time during which the fluid policy is implemented. (At the beginning of a review period, if the queue length is above the safety stock level, then the fluid policy is implemented during this review period. If this review period is not interrupted by the environment transition, then we call such a review period a normal review period.)

From the proof of Lemma 15, we know that there exists $r_1(\omega, t) > 0$, such that if $r > r_1(\omega, t)$, for all $s \in [\tau_n, \tau_{n+1})$ and $0 \le n \le m$

$$
\gamma^r(s) \in [\tau^r_n, \tau^r_{n+1}), \quad \gamma^r(\tau_n) = \tau^r_n, \quad X^r(\gamma^r(s)) = X(s) = i_n. \quad (164)
$$

Recall that $j^r(s)$ denotes the index of the first review period that starts after or at time s and $t^r(j)$ denotes the beginning time of jth review period. Thus, $t^r(j^r(\gamma^r(s)) - 1) < \gamma^r(s)$.

Hence for any $s \in [0,t]$ $(t \leq \tau_1)$, and any j such that $0 \leq j \leq j^r(\gamma^r(s)) - 1$, $X_j^r = i_0$, i.e the state of the environment is the same as the initial state i_0 until $\gamma^r(s)$. Therefore, for any $s \in [0,t]$ $(t \leq \tau_1)$, from the *consistency* assumption on Ψ, with the same idea as (163), we have

$$
\begin{aligned}
& \bar{q}^r(j^r(\gamma^r(s))) - \Psi(n^r(0,\gamma^r(s))l(r); \bar{q}^r(0), i_0) && (165)\\
= \;& \bar{q}^r(j^r(\gamma^r(s))) - \bar{q}^r(j^r(\gamma^r(s)) - 1) \\
& + \sum_{j=1}^{j^r(\gamma^r(s))-1} \chi(q^r(j-1) \geq \theta^r(i_0)) \mathbb{I}_1^r(i_0, j, s) && (166)\\
& + \sum_{j=1}^{j^r(\gamma^r(s))-1} \chi(q^r(j-1) \geq \theta^r(i_0)) \mathbb{I}_2^r(i_0, j, s),
\end{aligned}
$$

where

$$
\begin{aligned}
\mathbb{I}_1^r(i_0, j, s) &= \Big(\Psi(n^r(j, \gamma^r(s))l(r); \bar{q}^r(j), i_0) - \Psi(n^r(j, \gamma^r(s))l(r); z^r(j), i_0) \Big) \\
\mathbb{I}_2^r(i_0, j, s) &= \Big(\Psi(n^r(j, \gamma^r(s))l(r); \bar{q}^r(j), i_0) - \Psi(n^r(j, \gamma^r(s))l(r); \bar{q}^r(j-1), i_0) \Big).
\end{aligned}
$$

From the assumption that Ψ satisfies the equi-continuity condition, we have

$$
\begin{aligned}
|\mathbb{I}_1^r(i_0, j, s)| &\leq L(i, n^r(j, \gamma^r(s))l(r)) |\bar{q}^r(j) - z^r(j)|, \\
|\mathbb{I}_2^r(i_0, j, s)| &\leq L(i, n^r(j, \gamma^r(s))l(r)) |\bar{q}^r(j) - \bar{q}^r(j-1)|.
\end{aligned}
$$

Therefore,

$$
\begin{aligned}
& |\bar{q}^r(j^r(\gamma^r(s))) - \Psi(n^r(0,\gamma^r(s))l(r); \bar{q}^r(0), i_0)| \\
\leq \;& |\bar{q}^r(j^r(\gamma^r(s))) - \bar{q}^r(j^r(\gamma^r(s)) - 1)| \\
& + \sum_{j=1}^{j^r(\gamma^r(s))-1} \chi(q^r(j-1) \geq \theta^r(i_0)) L(i_0, n^r(j, \gamma^r(s))l(r)) |\bar{q}^r(j) - z^r(j)| && (167)\\
& + \sum_{j=1}^{j^r(\gamma^r(s))-1} \chi(q^r(j-1) \ngeq \theta^r(i_0)) L(i_0, n^r(j, \gamma^r(s))l(r)) |\bar{q}^r(j) - \bar{q}^r(j-1)|.
\end{aligned}
$$

From Lemma 71, we know that

$$
\lim_{r \to \infty} \sup_{0 \leq s \leq t} |n^r(0, \gamma^r(s))l(r) - s| = 0. \tag{168}
$$

Therefore, there exists $r_2(\omega, t, \epsilon) > 0$ such that if $r > r_2(\omega, t, \epsilon)$,

$$0 \leq \sup_{0 \leq s \leq t} n^r(0, \gamma^r(s)) l(r) \leq t + \epsilon.$$

The definition of $n^r(j, s)$ (given in Lemma 71) implies that it is decreasing in j, therefore $n^r(j, \gamma^r(s)) \leq n^r(0, \gamma^r(s))$ for all $s \geq 0$. Hence, for all $0 \leq j \leq j^r(\gamma^r(s))$,

$$0 \leq \sup_{0 \leq s \leq t} n^r(j, \gamma^r(s)) l(r) \leq \sup_{0 \leq s \leq t} n^r(0, \gamma^r(s)) l(r) \leq t + \epsilon, \tag{169}$$

and

$$L(i, n^r(j, \gamma^r(s)) l(r)) \leq \sup_{0 \leq u \leq t+\epsilon} L(i, u). \tag{170}$$

Inequalities (167), (170) imply that if $r > \max\{r(\omega, t, \epsilon), r_1(\omega, t), r_2(\omega, t, \epsilon)\}$, for any $s \in [0, t]$ ($t \leq \tau_1$),

$$\begin{aligned} & |\bar{q}^r(j^r(\gamma^r(s))) - \Psi(n^r(0, \gamma^r(s)) l(r); \bar{q}^r(0), i_0)| \\ \leq \; & |\bar{q}^r(j^r(\gamma^r(s))) - \bar{q}^r(j^r(\gamma^r(s)) - 1)| \\ & + \sup_{0 \leq u \leq t+\epsilon} L(i, u) \Big(\sum_{j=1}^{j^r(\gamma^r(s))-1} \chi(q^r(j-1) \geq \theta^r(i_0)) |\bar{q}^r(j) - z^r(j)| \\ & \qquad + \sum_{j=1}^{j^r(\gamma^r(s))-1} \chi(q^r(j-1) \ngeq \theta^r(i_0)) |\bar{q}^r(j) - \bar{q}^r(j-1)| \Big). \end{aligned} \tag{171}$$

Note that for all $s \in [0, t]$ ($t \leq \tau_1$) and for all $1 \leq j \leq j^r(\gamma^r(s)) - 1$, $X^r_j = X^r_{j-1} = i_0$, therefore

$$\begin{aligned} & \sum_{j=1}^{j^r(\gamma^r(s))-1} \chi(q^r(j-1) \geq \theta^r(i_0)) |\bar{q}^r(j) - z^r(j)| \\ = \; & \sum_{j=1}^{j^r(\gamma^r(s))-1} \chi(q^r(j-1) \geq \theta^r(i_0),\, X^r_j = X^r_{j-1} = i_0) |\bar{q}^r(j) - z^r(j)|, \end{aligned}$$

and

$$\begin{aligned} & \sum_{j=2}^{j^r(\gamma^r(s))-1} \chi(q^r(j-1) \ngeq \theta^r(i_0)) |\bar{q}^r(j) - \bar{q}^r(j-1)| \\ = \; & \sum_{j=2}^{j^r(\gamma^r(s))-1} \chi(q^r(j-1) \ngeq \theta^r(i_0),\, X^r_{j-2} = X^r_{j-1} = i_0) |\bar{q}^r(j) - \bar{q}^r(j-1)|. \end{aligned}$$

First,

$$
\begin{aligned}
& \sum_{j=1}^{j^r(\gamma^r(s))-1} \chi(q^r(j-1) \geq \theta^r(i_0),\, X_j^r = X_{j-1}^r = i_0)|\bar{q}^r(j) - z^r(j)| \\
\leq & \sum_{j=1}^{j^r(\gamma^r(s))-1} \chi(q^r(j-1) \geq \theta^r(i_0),\, X_j^r = X_{j-1}^r = i_0)\epsilon r^{-1}\theta^r(i_0) \\
= & \ \epsilon\beta(i_0)l(r) \sum_{j=1}^{j^r(\gamma^r(s))-1} \chi(q^r(j-1) \geq \theta^r(i_0),\, X_j^r = X_{j-1}^r = i_0) \\
\leq & \ \epsilon\beta(i_0)l(r)n^r(0,\gamma^r(s)), \qquad (172)
\end{aligned}
$$

where the first inequality is from Lemma 64, the equality is from the definition of $\theta^r(i)$, $i \in \mathcal{I}$, and the last equality is from the definition of $n^r(0,\cdot)$ given in Lemma 71.

From Lemma 59, we know that for $2 \leq j \leq j^r(\gamma^r(s)) - 1$,

$$
\chi(X_{j-2}^r = X_{j-1}^r = i_0) = \chi\Big(X_{j-2}^r = X_{j-1}^r = i_0, q^r(j-1) \geq (1-\epsilon)\theta^r(i_0)\Big).
$$

Therefore,

$$
\begin{aligned}
& \sum_{j=1}^{j^r(\gamma^r(s))-1} \chi(q^r(j-1) \ngeq \theta^r(i_0),\, X_{j-2}^r = X_{j-1}^r = i_0)|\bar{q}^r(j) - \bar{q}^r(j-1)| \\
= & \sum_{j=2}^{j^r(\gamma^r(s))-1} \chi(\tilde{\Gamma}^r(j-1,\epsilon))\chi(X_{j-2}^r = X_{j-1}^r = i_0)|\bar{q}^r(j) - \bar{q}^r(j-1)| \\
& + \chi(q^r(0) \ngeq \theta^r(i_0),\, X_1^r = X_0^r = i_0)|\bar{q}^r(1) - \bar{q}^r(0)| \\
\leq & \sum_{j=2}^{j^r(\gamma^r(s))-1} \chi(q^r(j-1) \ngeq \theta^r(i_0),\, X_{j-2}^r = X_{j-1}^r = i_0)\tilde{C}_2\epsilon l(r)\mathbf{e} + |\bar{q}^r(1) - \bar{q}^r(0)| \\
\leq & \sum_{j=2}^{j^r(\gamma^r(s))-1} \chi(q^r(j-2) \geq \theta^r(i_0),\, X_{j-2}^r = X_{j-1}^r = i_0)\tilde{C}_2\epsilon l(r)\mathbf{e} + |\bar{q}^r(1) - \bar{q}^r(0)| \\
\leq & \ n^r(0,\gamma^r(s))\tilde{C}_2\epsilon l(r)\mathbf{e} + |\bar{q}^r(1) - \bar{q}^r(0)| \qquad (173)
\end{aligned}
$$

where the first inequality is from Remak 66, and the third inequality is from the definition of $n^r(0,\cdot)$. Note that from the designed policy, the queue length at the end of an uninterrupted review period is below safety stock level implies that the fluid policy was implemented during this review period, i.e the queue length at the beginning of this review period is above the safety stock. This implies the second inequality above.

Combining (171), (172), and (173), we have

$$
\begin{aligned}
& |\bar{q}^r(j^r(\gamma^r(s))) - \Psi(n^r(0,\gamma^r(s))l(r);\bar{q}^r(0),i_0)| \\
\leq\ & |\bar{q}^r(j^r(\gamma^r(s))) - \bar{q}^r(j^r(\gamma^r(s))-1)| + \sup_{0\leq u\leq t+\epsilon} L(i,u)\Big(|\bar{q}^r(1)-\bar{q}^r(0)| \\
& +\epsilon\beta(i_0)l(r)n^r(0,\gamma^r(s)) + n^r(0,\gamma^r(s))\tilde{C}_2\epsilon l(r)\mathbf{e}\Big). \qquad (174)
\end{aligned}
$$

From Remark 69, we have

$$
|\bar{q}^r(j^r(\gamma^r(s))) - \bar{q}^r(j^r(\gamma^r(s))-1)| \leq \tilde{C}_4 l(r)\mathbf{e}, \qquad |\bar{q}^r(1)-\bar{q}^r(0)| \leq \tilde{C}_4 l(r)\mathbf{e}.
$$

Note that $\tilde{C}_2 > \bar{\beta} \geq \beta(i_0)$, therefore we have,

$$
\begin{aligned}
& |\bar{q}^r(j^r(\gamma^r(s))) - \Psi(n^r(0,\gamma^r(s))l(r);\bar{q}^r(0),i_0)| \\
\leq\ & \max\{1, \sup_{0\leq u\leq t+\epsilon} L(i_0,u)\}\Big(2\tilde{C}_4 l(r)\mathbf{e} + 2\tilde{C}_2\epsilon l(r)n^r(0,\gamma^r(s)\mathbf{e})\Big) \\
\leq\ & \max\{1, \sup_{0\leq u\leq t+\epsilon} L(i_0,u)\}\Big(2\tilde{C}_4 l(r) + 2\tilde{C}_2\epsilon(t+\epsilon)\Big)\mathbf{e}, \qquad (175)
\end{aligned}
$$

where the second inequality is from (169). We choose $r(\omega,t,\epsilon)$ large enough, such that if $r > r(\omega,t,\epsilon)$, then $l(r) < \epsilon$, therefore (175) reduces to

$$
\begin{aligned}
& |\bar{q}^r(j^r(\gamma^r(s))) - \Psi(n^r(0,\gamma^r(s))l(r);\bar{q}^r(0),i_0)| \\
\leq\ & \epsilon\Big(\max\{1, \sup_{0\leq u\leq t+\epsilon} L(i_0,u)\}(2\tilde{C}_4 + 2\tilde{C}_2(t+\epsilon))\Big)\mathbf{e}. \qquad (176)
\end{aligned}
$$

Since ϵ is chosen arbitrarily, $\sup_{0\leq u\leq t+\epsilon} L(i_0,u)\} < \infty$ and the above inequality holds for all $s\in[0,t]$ ($t\leq\tau_1$), we see that

$$
\lim_{r\to\infty}\sup_{0\leq s\leq t} |\bar{q}^r(j^r(\gamma^r(s))) - \Psi(n^r(0,\gamma^r(s))l(r);\bar{q}^r(0),i_0)| = 0. \qquad (177)
$$

For any $s\in[0,t]$ ($t\leq\tau_1$), from the assumption that Ψ satisfies equi-continuity,

$$
\begin{aligned}
& |\Psi(n^r(0,\gamma^r(s))l(r);\bar{q}^r(0),i_0) - \Psi(n^r(0,\gamma^r(s))l(r);\bar{Z}(0),i_0)| \\
\leq\ & L(i_0, n^r(0,\gamma^r(s))l(r))|\bar{q}^r(0)-\bar{Z}(0)| \\
\leq\ & \sup_{0\leq s\leq t+\epsilon} L(i_0,s)|\bar{q}^r(0)-\bar{Z}(0)|,
\end{aligned}
$$

where the second inequality is from (169). From the hypothesis that $|\bar{q}^r(0)-\bar{Z}(0)|\to 0$ as $r\to\infty$ and $\sup_{0\leq s\leq t+\epsilon} L(i_0,s) < \infty$, we have

$$
\lim_{r\to\infty}\sup_{0\leq s\leq t} |\Psi(n^r(0,\gamma^r(s))l(r);\bar{q}^r(0),i_0) - \Psi(n^r(0,\gamma^r(s))l(r);\bar{Z}(0),i_0)| = 0. \qquad (178)
$$

The result of (168), the assumption that $\Phi(\cdot; z, i)$ is continuous (hence $\Phi(\cdot; z, i)$ is uniformly continuous on compact sets) for any fixed z, i and the convergence together theorem (Lemma 17) imply that

$$\lim_{r\to\infty} \sup_{0\le s\le t} |\Psi(n^r(0, \gamma^r(s))l(r); \bar{Z}(0), i_0) - \Psi(s; \bar{Z}(0), i_0)| = 0. \tag{179}$$

The triangular inequality, (159), (177), (178), and (179) imply (156).

Now assuming (157) holds and we show (158) holds. We consider any finite $t \in [\tau_n, \tau_{n+1}]$. We go through the same procedure as the proof of (156). We first compare the difference between the scaled queue length and the estimated fluid level, i.e the difference between $\bar{q}^r(j^r(\gamma^r(s)))$ and $\Psi(n^r(j^r(\gamma^r(\tau_n)), \gamma^r(s))l(r); \bar{q}^r(j^r(\tau_n)), i_n)$, for all $s \in [\tau_n, t]$.

From (164), for all $s \in [\tau_n, t]$, and any j such that $j^r(\gamma^r(\tau_n)) \le j \le j^r(\gamma^r(s)) - 1$, $X_j^r = i_n$. Similar to (166),

$$\begin{aligned}
& \bar{q}^r(j^r(\gamma^r(s))) - \Psi(n^r(j^r(\gamma^r(\tau_n)), \gamma^r(s))l(r); \bar{q}^r(j^r(\gamma^r(\tau_n))), i_n) \\
= \;& \bar{q}^r(j^r(\gamma^r(s))) - \bar{q}^r(j^r(\gamma^r(s)) - 1) \\
& + \sum_{j=j^r(\gamma^r(\tau_n))}^{j^r(\gamma^r(s))-1} \chi(q^r(j-1) \ge \theta^r(i_n)) \mathbb{I}_1^r(i_n, j, s) \\
& + \sum_{j=j^r(\gamma^r(\tau_n))}^{j^r(\gamma^r(s))-1} \chi(q^r(j-1) \not\ge \theta^r(i_n)) \mathbb{I}_2^r(i_n, j, s),
\end{aligned}$$

where

$$\begin{aligned}
\mathbb{I}_1^r(i_n, j, s) &= \Psi(n^r(j, \gamma^r(s))l(r); \bar{q}^r(j), i_n) - \Psi(n^r(j, \gamma^r(s))l(r); z^r(j), i_n) \\
\mathbb{I}_2^r(i_n, j, s) &= \Psi(n^r(j, \gamma^r(s))l(r); \bar{q}^r(j), i_n) - \Psi(n^r(j, \gamma^r(s))l(r); \bar{q}^r(j-1), i_n)
\end{aligned}$$

Going through the same procedure as in the proof of (171), (172), (173), (174), (175), and (176) we have $r_n > 0$ such that if $r > r_n$, then for any $s \in [\tau_n, t]$,

$$\begin{aligned}
& |\bar{q}^r(j^r(\gamma^r(s))) - \Psi(n^r(j^r(\tau_n^r), \gamma^r(s))l(r); \bar{q}^r(j^r(\tau_n^r), i_n)| \\
\le \;& \epsilon\Big(\max\{1, \sup_{0\le u\le t+\epsilon} L(i_0, u)\}(2\tilde{C}_4 + 2\tilde{C}_2(t+\epsilon))\Big)\mathbf{e}.
\end{aligned}$$

Since ϵ is arbitrarily chosen, we have

$$\lim_{r\to\infty} \sup_{\tau_n\le s\le t} |\bar{q}^r(j^r(\gamma^r(s))) - \Psi(n^r(j^r(\tau_n^r), \gamma^r(s))l(r); \bar{q}^r(j^r(\tau_n^r), i_n)| = 0. \tag{180}$$

From (159) and the induction hypothesis (157), we have

$$\lim_{r\to\infty} \bar{q}^r(j^r(\gamma^r(\tau_n))) = \lim_{r\to\infty} r^{-1} Z^r(\gamma^r(\tau_n)) = \Psi(\tau_n; \bar{Z}(0), X).$$

Note that $\tau_n^r = \gamma^r(\tau_n)$. Hence,

$$\lim_{r\to\infty} \bar{q}^r(j^r(\tau_n^r)) = \Psi(\tau_n; \bar{Z}(0), X). \tag{181}$$

From the equi-continuity assumption on Ψ, for all $s \in [0, t]$, we have

$$\begin{aligned}
& |\Psi(n^r(j^r(\tau_n^r), \gamma^r(s))l(r); \bar{q}^r(j^r(\tau_n^r), i_n) - \Psi(n^r(j^r(\tau_n^r), \gamma^r(s))l(r); \Psi(\tau_n; \bar{Z}(0), X), i_n)| \\
\le\ & L(i_n, n^r(j^r(\tau_n^r), \gamma^r(s))l(r))|\bar{q}^r(j^r(\tau_n^r)) - \Psi(\tau_n; \bar{Z}(0), X)| \\
\le\ & \sup_{0\le u\le t+\epsilon} L(i_n, u)|\bar{q}^r(j^r(\tau_n^r)) - \Psi(\tau_n; \bar{Z}(0), X)|,
\end{aligned}$$

where the second inequality follows from the fact that $n^r(j^r(\tau_n^r), \gamma^r(s)) \le n^r(0, \gamma^r(s))$ and inequality (169). From (181), $\sup_{0\le u\le t+\epsilon} L(i_n, u) < \infty$, and the fact that the above inequality holds for all $s \in [\tau_n, t]$ ($t \le \tau_{n+1}$), if $r \to \infty$, then we have

$$\begin{aligned}
& \sup_{\tau_n\le s\le t} |\Psi(n^r(j^r(\tau_n^r), \gamma^r(s))l(r); \bar{q}^r(j^r(\tau_n^r), i_n) - \Psi(n^r(j^r(\tau_n^r), \gamma^r(s))l(r); \Psi(\tau_n; \bar{Z}(0), X), i_n)| \\
& \to 0.
\end{aligned} \tag{182}$$

From the definition of $n^r(j, s)$ given in Lemma 71, for any s such that $\gamma^r(s) \ge \tau_n^r$,

$$\begin{aligned}
n^r(j^r(\tau_n^r), \gamma^r(s)) &= \sum_{j=j^r(\tau_n^r)}^{j^r(\gamma^r(s))-1} \chi(q_j^r \ge \theta^r,\ X_{j+1}^r = X_j^r) \\
&= \sum_{j=0}^{j^r(\gamma^r(s))-1} \chi(q_j^r \ge \theta^r,\ X_{j+1}^r = X_j^r) - \sum_{j=0}^{j^r(\tau_n^r)-1} \chi(q_j^r \ge \theta^r,\ X_{j+1}^r = X_j^r) \\
&= n^r(0, \gamma^r(s)) - n^r(0, \tau_n^r).
\end{aligned}$$

From (164), we know that $\tau_n^r = \gamma^r(\tau_n)$ and that $\gamma^r(s) \ge \tau_n^r$ is equivalent to $s \ge \tau_n$. Hence for any $s \ge \tau_n$, we have $n^r(j^r(\tau_n^r), \gamma^r(s)) = n^r(0, \gamma^r(s)) - n^r(0, \gamma^r(\tau_n))$. With the result of (168), we have

$$\lim_{r\to\infty} \sup_{\tau_n\le s\le t} |(n^r(0, \gamma^r(s)) - n^r(0, \gamma^r(\tau_n)))l(r) - (s - \tau_n)| = 0.$$

That is

$$\lim_{r\to\infty} \sup_{\tau_n\le s\le t} |n^r(j^r(\tau_n^r), \gamma^r(s))l(r) - (s - \tau_n)| = 0. \tag{183}$$

As in the the proof of (179), the continuity of $\Psi(\cdot; z, i)$ (thus uniformly continuous on compact set) for each fixed z and i, (183) and the convergence together theorem (Lemma 17) imply that if $r \to \infty$ then

$$\sup_{\tau_n \le s \le t} |\Psi(n^r(j^r(\tau_n^r), \gamma^r(s))l(r); \Psi(\tau_n; \bar{Z}(0), X), i_n) - \Psi(s - \tau_n; \Psi(\tau_n; \bar{Z}(0), X), i_n)| \to 0. \quad (184)$$

Note that the stochastic fluid model solution Ψ satisfies that $\Psi(s - \tau_n; \Psi(\tau_n; \bar{Z}(0), X), i_n) = \Psi(s; \bar{Z}(0), X)$ for all $s \in [\tau_n, t]$ ($t \in [t_n, \tau_{n+1}]$). The result of (184) is the same as

$$\lim_{r\to\infty} \sup_{\tau_n \le s \le t} |\Psi(n^r(j^r(\tau_n^r), \gamma^r(s))l(r); \Psi(\tau_n; \bar{Z}(0), X), i_n) - \Psi(s; \bar{Z}(0), X)| = 0. \quad (185)$$

Using the triangular inequality, (180), (182), and (185), we have

$$\lim_{r\to\infty} \sup_{\tau_n \le s \le t} |\bar{q}^r(j^r(\gamma^r(s))) - \Psi(s; \bar{Z}(0), X)| = 0.$$

With this result and the induction hypothesis (157), we obtain (158). □

3.6 Fluid scale asymptotic optimality of the tracking policy

In this section, we show that the tracking method provided in Section 3.4 produces fluid scale asymptotically optimal scheduling policies for queueing networks in a slowly changing environment whenever the given stochastic fluid model solution is optimal.

Let $Z(t)$ denote the K dimensional queue length vector of the queueing network in a slowly changing environment. For a given cost function $g(x) \ge 0$ for any $x \ge 0$, a natural objective is to find a non-anticipating scheduling policy that minimizes the average total cost

$$\mathbf{E} \int_0^{T_0} g(Z(t))dt, \quad (186)$$

where $T_0 > 0$ is a constant. Note that since we allow the network to be overloaded at some environment states, we consider only a finite time horizon problem. Moreover, we assume that $g(\cdot)$ is continuous.

In this study, we restrict our attention to non-anticipating head-of-line policies, and plan to show that the tracking method provided in Section 3.4 produces a fluid scale asymptotically optimal policy. We focus on the objective (186) and define the asymptotic optimality with respect to this criteria.

To define the asymptotic optimality, we consider a sequence of speeded networks as we have done in the earlier sections. Let π^r denote a scheduling policy for the rth network and $Z^r(t;\pi^r)$ denote the K dimensional queue length vector of the rth network at time t under the π^r policy. Let $\bar{Z}^r(t;\pi^r) = r^{-1}Z^r(t;\pi^r)$. We define the fluid scale asymptotic optimality as follows.

Definition 74 (Fluid scale asymptotic optimality). For a given cost function $g(\cdot)$, a sequence of scheduling policies $\{\pi_*^r, r > 0\}$ possesses the fluid scale asymptotic optimality if

$$\limsup_{r\to\infty} \mathbf{E} \int_0^{T_0} g(\bar{Z}^r(t;\pi_*^r)) \le \liminf_{r\to\infty} \mathbf{E} \int_0^{T_0} g(\bar{Z}^r(t;\pi^r)) \tag{187}$$

for any other sequence of non-anticipating head-of-line scheduling policies $\{\pi^r, r > 0\}$.

To produce a sequence of fluid scale asymptotically optimal policy, we consider the optimization problem of the stochastic fluid model, i.e

$$\min_{\{\bar{T}(i,t),\, t\ge 0,\, i\in\mathcal{I}\}} \qquad \mathbf{E} \int_0^{T_0} g(\bar{Z}(t))dt, \tag{188}$$

where $\{\bar{T}(i,t), \bar{Z}(t), t \ge 0, i \in \mathcal{I}\}$ satisfies (90)-(95) with slight adaptation of notations. Assume that the optimal solution is given in the form of Ψ_* or T^{Ψ_*}, as discussed in Section 3.4. Then the optimal fluid level at time t is $\bar{Z}_*(t)$ and $\bar{Z}_*(t) = \Psi_*(t; \bar{Z}(0), X)$ as defined through (143)-(145).

Remark 75. Note that we assume $\mathcal{I}$ is finite, and since the fluid level of the stochastic fluid model changes continuously and T_0 is a finite constant, we know that $\mathbf{E}\int_0^{T_0} g(\bar{Z}(t))dt$ is finite for any fluid trajectory $\{\bar{Z}(t), t \ge 0\}$ if $\mathbf{E}[\bar{Z}(0)]$ is finite.

For any continuous cost function $g(\cdot)$, let $\{\pi_*^r, r > 0\}$ denote the discrete review policies produced by applying the tracking method provided in Section 3.4 to the optimal stochastic fluid model solution Ψ_*, then we have the following theorem.

Theorem 76. *Assume the conditions of Theorem 28, if Ψ_* satisfies (140) and (141), then $\{\pi_*^r, r > 0\}$ possesses the fluid scale asymptotic optimality, i.e (187) is satisfied.*

Proof of Theorem 76. The proof follows from Theorem 28, Fatou's Lemma, and the continuity of $g(\cdot)$. In particular, since Ψ_* and T^{Ψ_*} satisfy (122), the constraints of (95) and

(96) are satisfied by T^{Ψ_*}, then all the functions in Ψ_* satisfy (138) and (139), therefore Theorem 28 holds, i.e $\bar{Z}^r(t;\pi^r_*) \to \bar{Z}_*(t)$.

Without loss of generality, we assume that $\mathbf{E}[\bar{Z}(0)] < \infty$, then

$$\limsup_{r\to\infty} \mathbf{E}\int_0^{T_0} g(\bar{Z}^r(t;\pi^r_*)),\, dt \le \mathbf{E}\int_0^{T_0} \limsup_{r\to\infty} g(\bar{Z}^r(t;\pi^r_*))\, dt = \mathbf{E}\int_0^{T_0} g(\bar{Z}_*(t))\, dt,$$

where the first inequality is from Remark 75 and Fatou's lemma. From Theorem 28, we know that $\bar{Z}^r(t;\pi^r_*) \to \bar{Z}_*(t)$ with probability one, and the continuity of $g(\cdot)$ implies the above equality.

For any sequence of non-anticipating head-of-line policies $\{\pi^r, r > 0\}$, from Theorem 14, we know that $\liminf_{r\to\infty} \bar{Z}^r(t;\pi^r) = \bar{Z}(t)$ for some $\bar{Z}(t)$ such that $\bar{Z}(t)$ satisfies (90)-(95). From Fatou's Lemma, since $g(\cdot)$ is nonnegative,

$$\liminf_{r\to\infty} \mathbf{E}\int_0^{T_0} g(\bar{Z}^r(t;\pi^r))\, dt \ge \mathbf{E}\int_0^{T_0} \liminf_{r\to\infty} g(\bar{Z}^r(t;\pi^r))\, dt = \mathbf{E}\int_0^{T_0} g(\bar{Z}(t))\, dt,$$

where the equality is from the continuity of $g(\cdot)$.

The desired result follows from the fact that $\bar{Z}_*$ is the optimal fluid trajectory. $\square$

3.7 Proof of the lemmas

In this section, we provide the proof of the lemmas that appear in Section 3.5.

Proof of Lemma 30. We only provide the proof of (146). The proof of (147) will be similar.

Let $\hat{\eta}_k(i,n) = \eta_k(i,n)\chi(\eta_k(i,n) \le y)$ and $\tilde{\eta}_k(i,n) = \eta_k(i,n)\chi(\eta_k(i,n) < y)$, then $\eta_k(i,n) = \hat{\eta}_k(i,n) + \tilde{\eta}_k(i,n)$. Applying Chebyshev's inequality, we have

$$\begin{aligned} \mathbf{P}(|\sum_{n=1}^{m} \hat{\eta}_k(i,n) - \mathbf{E}[\hat{\eta}_k(i,n)]| \ge \frac{\epsilon x}{2}) &\le (\frac{\epsilon x}{2})^{-4}\mathbf{E}[(\sum_{n=1}^{m} \hat{\eta}_k(i,n) - \mathbf{E}[\hat{\eta}_k(i,n)])^4] \\ &\le (\frac{\epsilon x}{2})^{-4} m^2 (2y)^4 \quad \le \quad 2^8 \epsilon^{-4} x^{-2} y^4 \end{aligned} \tag{189}$$

where the third inequality is from $m \le x$. Since $\{\hat{\eta}_k(i,n) - \mathbf{E}[\hat{\eta}_k(i,n)],\, n \ge 1\}$ is a sequence of independent and identically distributed random variables with mean value being 0 and the summands are bounded by y, we have the second inequality in the above result.

On the other hand, $\mathbf{Var}(\tilde{\eta}_k(i,n)) \leq \mathbf{E}[(\tilde{\eta}_k(i,n))^2] = h_k(i,y)$. Applying Chebyshev's inequality, we have

$$\mathbf{P}(|\sum_{n=1}^{m} \tilde{\eta}_k(i,n) - \mathbf{E}[\tilde{\eta}_k(i,n)]| \geq \frac{\epsilon x}{2}) \quad \leq \quad (\frac{\epsilon x}{2})^{-2}\mathbf{Var}(\tilde{\eta}_k(i,n)) \quad \leq \quad 4\epsilon^{-2}x^{-2}h_k(i,y) \tag{190}$$

Setting $y = x^{1/8}$ and combining (189) and (190), we have the desired result. $\square$

Proof of Lemma 33. Note that $\Phi_k^l(i,m) = \sum_{u=1}^{m} \phi_k^l(i,u)$ and $\{\phi_k^l(i,u), u \geq 1\}$ is a sequence of independent and identical Bernoulli random variables with mean value being $p_{lk}(i)$. From Chebyshev's inequality,

$$\begin{aligned}
& \mathbf{P}(|\Phi_k^l(i,m) - mp_{lk}(i)| > \epsilon n) \\
\leq \; & \mathbf{E}[\Big(\sum_{u=1}^{m}(\phi_k^l(i,u) - p_{lk}(i))\Big)^4]/(\epsilon^4 x^4) \\
= \; & 2 \sum_{1 \leq u < u' \leq m} \Big(\mathbf{E}(\phi_k^l(i,u) - p_{lk}(i))^2(\phi_k^l(i,u') - p_{lk}(i))^2\Big)/(\epsilon^4 x^4) \\
& + \sum_{u=1}^{m} \mathbf{E}(\phi_k^l(i,u) - p_{lk}(i))^4/(\epsilon^4 x^4)
\end{aligned}$$

Note that for any u,

$$\mathbf{E}((\phi_k^l(i,u) - p_{lk}(i))^2) = p_{lk}(i)(1 - p_{lk}(i)) \leq 1,$$

and

$$\begin{aligned}
\mathbf{E}((\phi_k^l(i,u) - p_{lk}(i))^4) \quad &= \quad (1 - p_{lk}(i))^4 p_{lk}(i) + p_{lk}(i)^4(1 - p_{lk}(i)) \\
&\leq \quad p_{lk}(i)(1 - p_{lk}(i))((1 - p_{lk}(i))^3 + p_{lk}(i)^3) \leq 1.
\end{aligned}$$

Thus, for any $m \leq x$,

$$\mathbf{P}(|\Phi_k^l(i,m) - mp_{lk}(i)| > \epsilon n) \quad \leq \quad m^2/(\epsilon^4 x^4) \leq 1/(\epsilon^4 x^2).$$

This concludes the proof. $\square$

Proof of Lemma 34. In this lemma, we estimate the duration of a review period during which the fluid policy is implemented. Without loss of generality, we consider jth review period and assume that the state of the environment is i. If the queue length at the beginning of the jth review period is above the safety stock level, i.e $q^r(j) \geq \theta^r(i)$, then the fluid policy is implemented. According to the fluid policy, we schedule a number of jobs for each class to process, where the number is calculated through the given stochastic fluid model solution (124)-(127). These jobs are intended to be processed within $l(r)$ amount of time. If the state of the environment is still i after the fluid policy is completed during this review period, then we refer to this review period as a normal review period. For the rest of the proof, we assume that the jth review period is a normal review period. We will compare the actual duration of a normal review period, $t^r(j+1) - t^r(j)$, with the planned duration, $l(r)$. Therefore, we consider sample paths that satisfy

$$X_j^r = i = X_{j+1}^r, \quad q^r(j) \geq \theta^r(i). \tag{191}$$

For these sample paths, from (130), we know that the actual duration of the jth review period is

$$t^r(j+1) - t^r(j) \quad = \quad \max_{1 \leq s \leq S} e_s^{r,F}(j), \tag{192}$$

where $e_s^{r,F}(j)$ is the fluid policy implementation time of server s and it is defined in (129). Recall that $b_s^r(j)$ (defined in (128)) denotes the actual busy time of server s during jth review period and $u_s^r(j)$ (defined in (127)) denotes the planned idle time for server s during the jth review period. From (129), we know that

$$\begin{aligned}
\{|e_s^{r,F}(j) - l(r)| > \epsilon l(r)\} \quad &= \quad \{|b_s^r(j) + u_s^r(j) - l(r)| > \epsilon l(r),\, b_s^r(j) + u_s^r(j) \leq l(r)\} \\
&\quad \cup\{|l(r) - l(r)| > \epsilon l(r),\, b_s^r(j) \leq l(r) < b_s^r(j) + u_s^r(j)\} \\
&\quad \cup\{|b_s^r(j) - l(r)| > \epsilon l(r),\, l(r) < b_s^r(j)\} \\
&= \quad \{l(r) - (b_s^r(j) + u_s^r(j)) > \epsilon l(r),\, b_s^r(j) + u_s^r(j) \leq l(r)\} \\
&\quad \cup\{b_s^r(j) - l(r) > \epsilon l(r),\, l(r) < b_s^r(j)\} \\
&\subset \quad \{l(r) - (b_s^r(j) + u_s^r(j)) > \epsilon l(r),\, b_s^r(j) + u_s^r(j) \leq l(r)\} \\
&\quad \cup\{b_s^r(j) + u_s^r(j) - l(r) > \epsilon l(r),\, l(r) < b_s^r(j)\} \\
&\subset \quad \{l(r) - (b_s^r(j) + u_s^r(j)) > \epsilon l(r),\, b_s^r(j) + u_s^r(j) \leq l(r)\}
\end{aligned}$$

$$\cup\{b_s^r(j)+u_s^r(j)-l(r)>\epsilon l(r),\, l(r)<b_s^r(j)+u_s^r(j)\}$$
$$= \{|b_s^r(j)+u_s^r(j)-l(r)|>\epsilon l(r)\}. \tag{193}$$

From (127) and (128), we have

$$\begin{aligned}
&\{|b_s^r(j)+u_s^r(j)-l(r)|>\epsilon l(r)\}\\
=\ &\{|b_s^r(j)-\sum_{k\in C_s}x_k^r(j)|>\epsilon l(r)\}\\
=\ &\{|\sum_{k\in C_s}\Big(\sum_{n=\iota_k^r(i,j)+2}^{\iota_k^r(i,j)+p_k^r(j)}\eta_k^r(i,n)+\tilde\eta_k^r(i,\iota_k^r(i,j)+1)\Big)-\sum_{k\in C_s}x_k^r(j)|>\epsilon l(r)\}\\
=\ &\{|\sum_{k\in C_s}\Big(\sum_{n=\iota_k^r(i,j)+2}^{\iota_k^r(i,j)+p_k^r(j)}\eta_k^r(i,n)+\tilde\eta_k^r(i,\iota_k^r(i,j)+1)-x_k^r(j)\Big)|>\epsilon l(r)\}\\
\subset\ &\bigcup_{k\in C_s}\{\tilde\eta_k^r(i,\iota_k^r(i,j)+1)\geq(2K)^{-1}\epsilon l(r)\}\bigcup\\
&\{|\sum_{k\in C_s}\Big(\sum_{n=\iota_k^r(i,j)+2}^{\iota_k^r(i,j)+p_k^r(j)}(\eta_k^r(i,n)-\frac{1}{r\mu_k(i)})+(\frac{p_k^r(j)-1}{r\mu_k(i)}-x_k^r(j))\Big)|>\frac{\epsilon}{2}l(r)\}\\
\subset\ &\bigcup_{k\in C_s}\Big(\{\tilde\eta_k^r(i,\iota_k^r(i,j)+1)\geq(2K)^{-1}\epsilon l(r)\}\cup\\
&\quad\{|\sum_{n=\iota_k^r(i,j)+2}^{\iota_k^r(i,j)+p_k^r(j)}(\eta_k^r(i,n)-\frac{1}{r\mu_k(i)})+(\frac{p_k^r(j)-1}{r\mu_k(i)}-x_k^r(j))|>\frac{\epsilon}{2K}l(r)\}\Big).
\end{aligned}$$

From the definition of $p_k^r(j)$ given in (126), we know that

$$\Big|\frac{p_k^r(j)-1}{r\mu_k(i)}-x_k^r(j)\Big|\leq\frac{2}{r\mu_k(i)}\leq\frac{2}{\underline{\mu}_k r}$$

Note that $rl(r)\to\infty$ and $\underline{\mu}_k>0$, hence there exists $r(\epsilon)$, such that if $r>r(\epsilon)$ then $2(\underline{\mu}_k r)^{-1}<(4K)^{-1}\epsilon l(r)$ for all $1\leq k\leq K$. For all $r>r(\epsilon)$, since there are no more than K classes at each station, then

$$\sum_{k\in C_s}\Big|\frac{p_k^r(j)-1}{r\mu_k(i)}-x_k^r(j)\Big|\ \leq\ \frac{\epsilon l(r)}{4}.$$

Therefore, if $r>r(\epsilon)$, then

$$\{|b_s^r(j)+u_s^r(j)-l(r)|>\epsilon l(r)\}$$
$$\subset\bigcup_{k\in C_s}\Big(\{\tilde\eta_k^r(i,\iota_k^r(i,j)+1)\geq\frac{\epsilon l(r)}{2K}\}\cup\{|\sum_{n=\iota_k^r(i,j)+2}^{\iota_k^r(i,j)+p_k^r(j)}(\eta_k^r(i,n)-\frac{1}{r\mu_k(i)})|>\frac{\epsilon}{4K}l(r)\}\Big). \tag{194}$$

From (192), we have

$$\begin{aligned}\{|t^r(j+1)-t^r(j)-l(r)|>\epsilon l(r)\} &= \{|\max_{1\le s\le S} e_s^{r,F}(j)-l(r)|>\epsilon l(r)\}\\ &\subset \{\max_{1\le s\le S}|e_s^{r,F}(j)-l(r)|>\epsilon l(r)\}\\ &\subset \bigcup_{1\le s\le S}\{|e_s^{r,F}(j)-l(r)|>\epsilon l(r)\}.\end{aligned}$$

Combining this result with (193) and (194), we have

$$\begin{aligned}&\{|t^r(j+1)-t^r(j)-l(r)|>\epsilon l(r)\}\\ &\subset \bigcup_{1\le s\le S}\bigcup_{k\in C_s}\Big(\{\tilde{\eta}_k^r(i,\iota_k^r(i,j)+1)\ge\frac{\epsilon l(r)}{2K}\}\cup\{|\sum_{n=\iota_k^r(i,j)+2}^{\iota_k^r(i,j)+p_k^r(j)}(\eta_k^r(i,n)-\frac{1}{r\mu_k(i)})|>\frac{\epsilon l(r)}{4K}\}\Big)\\ &= \bigcup_{1\le k\le K}\Big(\{\tilde{\eta}_k^r(i,\iota_k^r(i,j)+1)\ge\frac{\epsilon l(r)}{2K}\}\cup\{|\sum_{n=\iota_k^r(i,j)+2}^{\iota_k^r(i,j)+p_k^r(j)}(\eta_k^r(i,n)-\frac{1}{r\mu_k(i)})|>\frac{\epsilon l(r)}{4K}\}\Big). \qquad (195)\end{aligned}$$

Considering only the sample paths such that all the remaining service times are less than $(2K)^{-1}l(r)$, from (195) we have

$$\begin{aligned}&\{|t^r(j+1)-t^r(j)-l(r)|>\epsilon l(r)\}\cap\Upsilon^r(i,j)\\ &\subset \bigcup_{1\le k\le K}\{|\sum_{n=\iota_k^r(i,j)+2}^{\iota_k^r(i,j)+p_k^r(j)}(\eta_k^r(i,n)-\frac{1}{r\mu_k(i)})|>\frac{\epsilon l(r)}{4K}\},\end{aligned}$$

which is based on the assumption in (191). Using the indicator function and presenting the result in a self-contained form, we have

$$\begin{aligned}&\chi(|t^r(j+1)-t^r(j)-l(r)|>\epsilon l(r))\,\chi(\Upsilon^r(i,j))\,\chi(X_j^r=i=X_{j+1}^r,q^r(j)\ge\theta^r(i))\\ \le\;& \sum_{k=1}^K\chi(\Big|\sum_{n=\iota_k^r(i,j)+2}^{\iota_k^r(i,j)+p_k^r(j)}(\eta_k^r(i,n)-\frac{1}{r\mu_k(i)})\Big|>\frac{\epsilon l(r)}{4K})\,\chi(X_j^r=i=X_{j+1}^r,q^r(j)\ge\theta^r(i)).\end{aligned}$$

Therefore, noting that $\Gamma^r(j)=\{X_j^r=X_{j+1}^r,q^r(j)\ge\theta^r(X_j^r)\}$ and $\Gamma^r(i,j)=\{X_j^r=i=X_{j+1}^r,q^r(j)\ge\theta^r(i)\}$, we have

$$\begin{aligned}&\mathbf{E}\Big[\chi(|t^r(j+1)-t^r(j)-l(r)|>\epsilon l(r))\,\chi(\Upsilon^r(X_j^r,j))\,\chi(\Gamma^r(j))\Big|\mathcal{F}_j^r,X_j^r=i\Big]\\ =\;& \mathbf{E}\Big[\chi(|t^r(j+1)-t^r(j)-l(r)|>\epsilon l(r))\,\chi(\Upsilon^r(i,j))\,\chi(\Gamma^r(i,j))\Big|\mathcal{F}_j^r,X_j^r=i\Big]\\ \le\;& \mathbf{E}\Big[\sum_{k=1}^K\chi(|\sum_{n=\iota_k^r(i,j)+2}^{\iota_k^r(i,j)+p_k^r(j)}(\eta_k^r(i,n)-\frac{1}{r\mu_k(i)})|>\frac{\epsilon l(r)}{4K})\Big|\mathcal{F}_j^r,\,X_j^r=i\Big]\chi(q^r(j)\ge\theta^r(i))\end{aligned}$$

$$
\begin{aligned}
&= \mathbf{E}\Big[\sum_{k=1}^{K} \chi(|\sum_{n=\iota_k^r(i,j)+2}^{\iota_k^r(i,j)+\lfloor r\mu_k(i)x_k^r(j)\rfloor} (\eta_k^r(i,n) - \frac{1}{r\mu_k(i)})| > \frac{\epsilon l(r)}{4K})\Big|\mathcal{F}_j^r,\, X_j^r = i\Big]\chi(q^r(j) \ge \theta^r(i)) \\
&= \sum_{k=1}^{K} \mathbf{E}\Big[\chi(|\sum_{n=\iota_k^r(i,j)+2}^{\iota_k^r(i,j)+\lfloor r\mu_k(i)x_k^r(j)\rfloor} (\eta_k(i,n) - \frac{1}{\mu_k(i)})| > \frac{\epsilon r l(r)}{4K})\Big|\mathcal{F}_j^r,\, X_j^r = i\Big]\chi(q^r(j) \ge \theta^r(i)) \\
&\le \sum_{k=1}^{K} \hat{h}_k((4K\mu_k(i))^{-1}\epsilon, \lfloor \mu_k(i) r l(r)\rfloor)\chi(q^r(j) \ge \theta^r(i)), \qquad (196)
\end{aligned}
$$

where the second equality is from the definition of $p_k^r(j)$ given in (126) and $x_k^r(j)$ is defined in (125), the third equality follows from the fact that $\eta^r(i,n) = r^{-1}\eta(i,n)$, and the last inequality follows from the fact that the service times that happen after $t^r(j)$ are independent from $\mathcal{F}_j^r$ and from Lemma 30 since $x_k^r(j) \le l(r)$. Note that $\hat{h}_k(x,y)$ is decreasing in x and y for all $1 \le k \le K$. Thus,

$$
\begin{aligned}
&\sum_{k=1}^{K} \hat{h}_k((4K\mu_k(i))^{-1}\epsilon, \lfloor \mu_k(i) r l(r)\rfloor)\chi(q^r(j) \ge \theta^r(i)) \\
\le\ &\sum_{k=1}^{K} \hat{h}_k((4K\bar{\mu}_k)^{-1}\epsilon, \lfloor \underline{\mu}_k r l(r)\rfloor)\chi(q^r(j) \ge \theta^r(i)).
\end{aligned}
$$

Combining this result with the last one, and from the properties of conditional expectation, we have

$$
\begin{aligned}
&\mathbf{E}\Big[\chi(|t^r(j+1) - t^r(j) - l(r)| > \epsilon l(r))\,\chi(\Upsilon^r(X_j^r, j))\,\chi(X_j^r = X_{j+1}^r, q^r(j) \ge \theta^r(X_j^r))\Big|\mathcal{F}_j^r\Big] \\
\le\ &\sum_{k=1}^{K} \hat{h}_k(\epsilon(4K\bar{\mu}_k)^{-1}, \lfloor \underline{\mu}_k r l(r)\rfloor)\chi(q^r(j) \ge \theta^r(X_j^r)).
\end{aligned}
$$

This concludes the proof.

□

Proof of Lemma 35. Going through the same procedure as we prove (193) and (194) in the proof of Lemma 34, we can choose $r(\epsilon) > 0$ such that if $r > r(\epsilon)$, then

$$
\sum_{k\in C_s} \Big|\frac{p_k^r(j) - 1}{r\mu_k(i)} - x_k^r(j)\Big| \ \le\ \frac{\epsilon l(r)}{2}.
$$

and

$$
\begin{aligned}
\{e_s^{r,F}(j) - l(r) < -\epsilon l(r)\} &= \{b_s^r(j) + u_s^r(j) - l(r) < -\epsilon l(r)\} \\
&\subset \bigcup_{k\in C_s}\{\sum_{n=\iota_k^r(i,j)+2}^{\iota_k^r(i,j)+p_k^r(j)} (\eta_k^r(i,n) - \frac{1}{r\mu_k(i)}) < -\frac{\epsilon}{2K} l(r)\}.
\end{aligned}
$$

From (192),

$$
\begin{aligned}
& \{t^r(j+1) - t^r(j) < (1-\epsilon)l(r)\} \cap \Gamma^r(j) \\
\subset \quad & \bigcap_{1\le s\le S} \bigcup_{k\in C_s} \{ \sum_{n=\iota_k^r(i,j)+2}^{\iota_k^r(i,j)+p_k^r(j)} (\eta_k^r(i,n) - \frac{1}{r\mu_k(i)}) < -\frac{\epsilon}{2K} l(r)\}.
\end{aligned}
$$

Similar to the proof of (196), we have

$$
\begin{aligned}
& \mathbf{E}[\chi(t^r(j+1) - t^r(j) < (1-\epsilon)l(r))\,\chi(\Gamma^r(j))|\mathcal{F}_j^r, X_j^r = i] \\
\le \quad & \sum_{1\le k\le K} \mathbf{E}[\sum_{n=\iota_k^r(i,j)+2}^{\iota_k^r(i,j)+p_k^r(j)} (\eta_k^r(i,n) - \frac{1}{r\mu_k(i)}) < -\frac{\epsilon}{2K} l(r) |\mathcal{F}_j^r, X_j^r = i] \chi(q^r(j) \ge \theta^r(i)) \\
\le \quad & \sum_{k=1}^{K} \hat{h}_k((2K\mu_k(i))^{-1}\epsilon, \lfloor \mu_k(i) r l(r) \rfloor) \chi(q^r(j) \ge \theta^r(i)),
\end{aligned}
$$

From that $\hat{h}_k(x,y)$ is decreasing in x and y for all $1 \le k \le K$, we have the conclusion of the lemma. □

Proof of Lemma 36. This proof is the same as that of Lemma 34, except that that the actual duration of the jth review period satisfies that

$$
t^r(j+1) - t^r(j) \quad \le \quad \max_{1\le s\le S} e_s^{r,F}(j),
$$

instead of (192) since we do not assume $X_{j+1}^r = X_j^r$. □

Proof of Lemma 37. Since $E_k^r(i,t)$ attains only nonnegative integer values, then

$$
\{E_k^r(i,t) - r\alpha_k(i)t > x\} \quad = \quad \{E_k^r(i,t) \ge \lceil r\alpha_k(i)t + x \rceil\}. \tag{197}
$$

Recall that for each environment state $i \in \mathcal{I}$ and each class k such that $\alpha_k(i) > 0$,

$$
E_k^r(i,t) = \max\{n : \sum_{m=1}^{n} \xi_k^r(i,m) \le t\} = \max\{n : \sum_{m=1}^{n} \xi_k(i,m) \le rt\}.
$$

Therefore,

$$
\begin{aligned}
\{E_k^r(i,t) \ge \lceil r\alpha_k(i)t + x \rceil\} \quad &= \quad \{ \sum_{m=1}^{\lceil r\alpha_k(i)t+x \rceil} \xi_k^r(i,m) \le t\} \\
&= \quad \{ \sum_{m=1}^{\lceil r\alpha_k(i)t+x \rceil} \xi_k(i,m) \le rt\}
\end{aligned}
$$

$$
\begin{aligned}
&= \{\sum_{m=1}^{\lceil r\alpha_k(i)t+x\rceil} \xi_k(i,m) - \frac{\lceil r\alpha_k(i)t+x\rceil}{\alpha_k(i)} \le rt - \frac{\lceil r\alpha_k(i)t+x\rceil}{\alpha_k(i)}\} \\
&\subset \{\sum_{m=1}^{\lceil r\alpha_k(i)t+x\rceil} \xi_k(i,m) - \frac{\lceil r\alpha_k(i)t+x\rceil}{\alpha_k(i)} \le -\frac{x}{\alpha_k(i)}\} \\
&\subset \{\Big|\sum_{m=1}^{\lceil r\alpha_k(i)t+x\rceil} \xi_k(i,m) - \frac{\lceil r\alpha_k(i)t+x\rceil}{\alpha_k(i)}\Big| \ge \frac{x}{\alpha_k(i)}\}. \qquad (198)
\end{aligned}
$$

From Lemma 30, we have

$$
\begin{aligned}
&\mathbf{P}(\Big|\sum_{m=1}^{\lceil r\alpha_k(i)t+x\rceil} \xi_k(i,m) - \frac{\lceil r\alpha_k(i)t+x\rceil}{\alpha_k(i)}\Big| \ge \frac{x}{\alpha_k(i)}) \\
\le\ & \hat{g}(\frac{x}{\alpha_k(i)\lceil r\alpha_k(i)t+x\rceil}, \lceil r\alpha_k(i)t+x\rceil). \qquad (199)
\end{aligned}
$$

Combining the results of (197)-(199), we have

$$
\mathbf{P}(E_k^r(i,t) - r\alpha_k(i)t > x) \quad \le \quad \hat{g}(\frac{x}{\alpha_k(i)\lceil r\alpha_k(i)t+x\rceil}, \lceil r\alpha_k(i)t+x\rceil),
$$

which concludes the proof of (i).

Note that since

$$
\{E_k^r(i,t) - r\alpha_k(i)t < -x\} \quad = \quad \{E_k^r(i,t) \le \lfloor r\alpha_k(i)t - x\rfloor\},
$$

using the procedure above, we get (ii). Combining (i) and (ii), and using the fact that $\hat{g}(x,y)$ is decreasing in both x and y, we have (iii).

□

Proof of Lemma 38. The proof is the same as that of Lemma 37 except that the arrival rate is replaced by the service rate. □

Proof of Lemma 39. Note that $\sum_{i\in\mathcal{I}} I^r(i,t) = t$ and $\bar{\alpha}_k \ge \alpha_k(i)$. For any $i \in \mathcal{I}$, we have

$$
\begin{aligned}
\{E_k^r(t) > r(\bar{\alpha}_k+\epsilon)t\} &= \{\sum_{i\in\mathcal{I}} E_k^r(i, I^r(i,t)) > \sum_{i\in\mathcal{I}} r\bar{\alpha}_k I^r(i,t) + \epsilon r t\} \\
&\subset \bigcup_{i\in\mathcal{I}} \{E_k^r(i, I^r(i,t)) > r\alpha_k(i) I^r(i,t) + \frac{\epsilon r t}{|\mathcal{I}|}\} \\
&= \bigcup_{i\in\mathcal{I}} \{E_k^r(i, I^r(i,t)) - r\alpha_k(i) I^r(i,t) > \frac{\epsilon r t}{|\mathcal{I}|}\}.
\end{aligned}
$$

Applying Lemma 37, we have

$$\begin{aligned}
\mathbf{P}(E_k^r(t) > r(\bar{\alpha}_k + \epsilon)t) &\leq \sum_{i\in\mathcal{I}} \mathbf{P}(E_k^r(i, I^r(i,t)) - r\alpha_k(i)I^r(i,t) > \frac{\epsilon rt}{|\mathcal{I}|}) \\
&\leq \sum_{i\in\mathcal{I}} \hat{g}_k(\frac{\epsilon rt/|\mathcal{I}|}{\alpha_k(i)\lceil r\alpha_k(i)I^r(i,t) + \epsilon rt/|\mathcal{I}|\rceil}, \lceil r\alpha_k(i)I^r(i,t) + \epsilon rt/|\mathcal{I}|\rceil).
\end{aligned}$$

Since $I^r(i,t) \leq t$, we have

$$\frac{\epsilon rt/|\mathcal{I}|}{\alpha_k(i)\lceil r\alpha_k(i)I^r(i,t) + \epsilon rt/|\mathcal{I}|\rceil} \geq \frac{\epsilon rt/|\mathcal{I}|}{\bar{\alpha}_k\lceil r\bar{\alpha}_k t + \epsilon rt/|\mathcal{I}|\rceil}.$$

Choose $r_k(\epsilon, t)$ large enough such that if $r > r_k(\epsilon, t)$, then

$$\frac{\epsilon rt/|\mathcal{I}|}{\bar{\alpha}_k\lceil r\bar{\alpha}_k t + \epsilon rt/|\mathcal{I}|\rceil} \geq \frac{\epsilon}{2(|\mathcal{I}|(\bar{\alpha}_k)^2 + \epsilon)}.$$

Note that $\lceil r\alpha_k(i)I^r(i,t) + \epsilon rt/|\mathcal{I}|\rceil \geq \epsilon rt/|\mathcal{I}|$ and $\hat{g}_k(x,y)$ is decreasing in x and y, then for $r > r_k(\epsilon,t)$, we have

$$\hat{g}_k(\frac{\epsilon rt/|\mathcal{I}|}{\alpha_k(i)\lceil r\alpha_k(i)I^r(i,t) + \epsilon rt/|\mathcal{I}|\rceil}, \lceil r\alpha_k(i)I^r(i,t) + \epsilon rt/|\mathcal{I}|\rceil) \leq \hat{g}_k(\frac{\epsilon}{2(|\mathcal{I}|(\bar{\alpha}_k)^2 + \epsilon)}, \epsilon rt/|\mathcal{I}|).$$

Combining the above inequalities, we get the conclusion of the lemma.

□

Proof of Lemma 40. Using the same analysis as in the proof of Lemma 39 and noting that $T_k^r(i,t) \leq I^r(i,t) \leq t$, we obtain the desired result. □

Proof of Lemma 41. For each $i \in \mathcal{I}$ and any $n \geq 1$, applying techniques similar to those used in the proof of Chebyshev's inequality, we have

$$\begin{aligned}
\mathbf{P}(\xi_k^r(i,n) > \epsilon l(r)) &= \mathbf{P}(\xi_k(i,n) > rl(r)) \\
&\leq \frac{\mathbf{E}((\xi_k(i,n))^2\chi(\xi_k(i,n) > rl(r)))}{(rl(r))^2} \\
&\leq \frac{g_k(i, rl(r))}{(rl(r))^2} \leq \frac{c_1}{r^{2+\gamma}(l(r))^2}, \qquad (200)
\end{aligned}$$

where the last inequality is from Remark 25.

Recall that we let $E_k^r(t)$ denote the total number of external arrivals to class k until time t and it satisfies $E_k^r(t) = \sum_{i\in\mathcal{I}} E^r(i, I^r(i,t))$, then

$$\mathbf{P}(\max_{i\in\mathcal{I}} \max_{1\leq n\leq E_k^r(i,I^r(i,t))} \xi_k^r(i,n) > \epsilon l(r))$$

$$
\begin{aligned}
&\le \mathbf{P}(\max_{i\in\mathcal{I}} \max_{1\le n\le E^r_k(i,I^r(i,t))} \xi^r_k(i,n) > \epsilon l(r), \sum_{i\in\mathcal{I}} E^r_k(i,I^r(i,t)) \le r(\bar{\alpha}_k+\epsilon)t) \\
&\quad +\mathbf{P}\Big(\sum_{i\in\mathcal{I}} E^r_k(i,I^r(i,t)) > r(\bar{\alpha}_k+\epsilon)t\Big) \\
&\le r(\bar{\alpha}_k+\epsilon)t\,\mathbf{P}(\xi^r_k(i,n) > \epsilon l(r)) + \mathbf{P}(E^r_k(t) > r(\bar{\alpha}_k+\epsilon)t) \\
&\le \frac{r(\bar{\alpha}_k+\epsilon)tc_1}{r^{2+\gamma}(l(r))^2} + \hat{g}_k(\frac{\epsilon}{2(|\mathcal{I}|(\bar{\alpha}_k)^2+\epsilon)}, \epsilon rt/|\mathcal{I}|),
\end{aligned}
$$

where the last inequality is from (200) and Lemma 39. Simplifying the last expression, we get the desired result. □

Proof of Lemma 43. The proof is the same as that of Lemma 41 except that the arrival rate is replaced by the service rate. □

Proof of Lemma 45. Without loss of generality, we assume that at the beginning of the jth review period the environment is at state i and the queue length is above the selected safety stock level. That is

$$
X^r_j = i, \quad q^r(j) \ge \theta^r(i). \tag{201}
$$

So we will implement the fluid policy characterized by (124)-(127) in Section 3.4. The targeted fluid level vector at the end of this review period is from the given stochastic fluid model solution, and it satisfies that

$$
z^r(j+1) = \Psi(l(r); \bar{q}^r(j), i) = \bar{q}^r(j) + \alpha(i)l(r) - (I - P'(i))M^{-1}(i)x^r(j).
$$

Thus, for each class k, we have

$$
z^r_k(j+1) \quad = \quad \bar{q}^r_k(j) + \alpha_k(i)l(r) - \mu_k(i)x^r_k(j) + \sum_{l=1}^{K} p_{lk}(i)\mu_l(i)x^r_l(j).
$$

We assume that the fluid policy is completely implemented during this review period, i.e there is no environment transition to interrupt the review period. This assumption and (201) imply (191). We assume (191), then the actual queue length of class k at the end of the jth review period is

$$
\begin{aligned}
q^r_k(j+1) \quad = \quad & q^r_k(j) + (E^r_k(t^r(j+1)) - E^r_k(t^r(j))) - p^r_k(j) \\
& + \sum_{l=1}^{K} (\Phi^l_k(i, \iota^r_l(i,j) + p^r_l(j)) - \Phi^l_k(i, \iota^r_l(i,j))),
\end{aligned}
$$

where $\iota_l^r(i,j)$ denotes the number of class l jobs that have departed at environment state i until $t^r(j)$ and it is defined in Section 3.4. Recall that $\bar{q}_k^r(j) = r^{-1}(q_k^r(j) - \theta_k^r(X_j^r))$ for all $j \geq 0$. From assumption (191), $X_{j+1}^r = X_j^r = i$, then

$$\begin{aligned}
\bar{q}_k^r(j+1) &= \bar{q}_k^r(j) + r^{-1}\Big(E_k^r(t^r(j+1)) - E_k^r(t^r(j)) - p_k^r(j) \\
&\quad + \sum_{l=1}^{K}(\Phi_k^l(i, \iota_l^r(i,j) + p_l^r(j)) - \Phi_k^l(i, \iota_l^r(i,j)))\Big).
\end{aligned}$$

Comparing the difference between $z_k^r(j+1)$ and $\bar{q}_k^r(j+1)$, we have

$$\begin{aligned}
& |\bar{q}_k^r(j+1) - z_k^r(j+1)| \\
\leq\ & r^{-1}\Big(|E_k^r(t^r(j+1)) - E_k^r(t^r(j)) - \alpha_k(i)rl(r)| + |p_k^r(j) - r\mu_k(i)x_k^r(j)| \\
& + \sum_{l=1}^{K} |\Phi_k^l(i, \iota_l^r(i,j) + p_l^r(j)) - \Phi_k^l(i, \iota_l^r(i,j)) - p_{lk}(i)r\mu_l(i)x_l^r(j)|\Big) \\
\leq\ & |E_k^r(t^r(j+1)) - E_k^r(t^r(j)) - \alpha_k(i)rl(r)| + |p_k^r(j) - r\mu_k(i)x_k^r(j)| \\
& + \sum_{l=1}^{K} |\Phi_k^l(i, \iota_l^r(i,j) + p_l^r(j)) - \Phi_k^l(i, \iota_l^r(i,j)) - p_{lk}(i)p_l^r(j)| \\
& + \sum_{l=1}^{K} p_{lk}(i)|p_l^r(j) - r\mu_l(i)x_l^r(j)|.
\end{aligned}$$

From (126), we know that $|p_l^r(j) - r\mu_l(i)x_l^r(j)| < 1$. Note that $\sum_{l=1}^{K} p_{lk}(i) \leq 1$ and

$$\begin{aligned}
|\bar{q}_k^r(j+1) - z_k^r(j+1)| \ &\leq\ r^{-1}\Big(|E_k^r(t^r(j+1)) - E_k^r(t^r(j)) - \alpha_k(i)rl(r)| \\
&\quad + \sum_{l=1}^{K} |\Phi_k^l(i, \iota_l^r(i,j) + p_l^r(j)) - \Phi_k^l(i, \iota_l^r(i,j)) - p_{lk}(i)p_l^r(j)| + 2\Big).
\end{aligned}$$

Since $\theta_k^r(i) = \beta_k(i)rl(r)$, $\underline{\beta}_k = \inf_{i\in\mathcal{I}} \beta_k > 0$, and $rl(r) \to \infty$ as $r \to \infty$, there exists $r_k(\epsilon) > 0$, such that if $r > r_k(\epsilon)$, $2r^{-1} \leq \epsilon\underline{\beta}_k l(r)/2 \leq \epsilon\beta_k(i)l(r)/2 \leq \epsilon r^{-1}\theta_k^r(i)/2$. This implies that if $r > r_k(\epsilon)$, then

$$\begin{aligned}
& \{|\bar{q}_k^r(j+1) - z_k^r(j+1)| \geq \epsilon r^{-1}\theta_k^r(i)\} \\
\subset\ & \{|E_k^r(t^r(j+1)) - E_k^r(t^r(j)) - \alpha_k(i)rl(r)| \geq 4^{-1}\epsilon\theta_k^r(i)\} \\
& \cup\{\sum_{l=1}^{K} |\Phi_k^l(i, \iota_l^r(i,j) + p_l^r(j)) - \Phi_k^l(i, \iota_l^r(i,j)) - p_{lk}(i)p_l^r(j)| \geq 4^{-1}\epsilon\theta_k^r(i)\} \\
\subset\ & \{|E_k^r(t^r(j+1)) - E_k^r(t^r(j)) - \alpha_k(i)rl(r)| \geq 4^{-1}\epsilon\theta_k^r(i)\} \\
& \cup_{l=1}^{K}\{|\Phi_k^l(i, \iota_l^r(i,j) + p_l^r(j)) - \Phi_k^l(i, \iota_l^r(i,j)) - p_{lk}(i)p_l^r(j)| \geq (4K)^{-1}\epsilon\theta_k^r(i)\}.
\end{aligned}$$

Note that the above result holds under the assumption (191). In other words,

$$
\begin{aligned}
& \{|\bar{q}_k^r(j+1) - z_k^r(j+1)| \geq \epsilon r^{-1}\theta_k^r(i)\} \cap \Gamma^r(i,j) \\
\subset \; & \Big(\cup_{l=1}^K \{|\Phi_k^l(i, \iota_l^r(i,j) + p_l^r(j)) - \Phi_k^l(i, \iota_l^r(i,j)) - p_{lk}(i)p_l^r(j)| \geq (4K)^{-1}\epsilon\theta_k^r(i)\} \\
& \cup \{|E_k^r(t^r(j+1)) - E_k^r(t^r(j)) - \alpha_k(i) r l(r)| \geq 4^{-1}\epsilon\theta_k^r(i)\}\Big) \cap \Gamma^r(i,j).
\end{aligned}
$$

Therefore, for any $\tilde{\epsilon} > 0$,

$$
\begin{aligned}
& \{|\bar{q}_k^r(j+1) - z_k^r(j+1)| \geq \epsilon r^{-1}\theta_k^r(i)\} \cap \Gamma^r(i,j) \cap \Upsilon(i,j,\tilde{\epsilon}) \cap \Lambda(i,j,\tilde{\epsilon}) \\
\subset \; & \Big(\cup_{l=1}^K \{|\Phi_k^l(i, \iota_l^r(i,j) + p_l^r(j)) - \Phi_k^l(i, \iota_l^r(i,j)) - p_{lk}(i)p_l^r(j)| \geq (4K)^{-1}\epsilon\theta_k^r(i)\} \\
& \cup \{|E_k^r(t^r(j+1)) - E_k^r(t^r(j)) - \alpha_k(i) r l(r)| \geq 4^{-1}\epsilon\theta_k^r(i)\}\Big) \\
& \cap \Gamma^r(i,j) \cap \Upsilon(i,j,\tilde{\epsilon}) \cap \Lambda(i,j,\tilde{\epsilon}).
\end{aligned} \tag{202}
$$

We first provide a probabilistic bound on the difference between the actual number of customers that are routed from l to k and its expected value. Then we derive a probabilistic bound on the difference between the actual number of external arrivals to class k during the jth review period and its expected value.

From assumption (201), we have $p_l^r(j) \leq \lfloor r l(r)\beta_k(i) \rfloor \leq \theta_k^r(i)$. Note that $\phi_k^l(\iota_l^r(i,j)+n)$ is independent of $\mathcal{F}_j^r$ for $n \geq 1$, then from Lemma 33,

$$
\begin{aligned}
& \mathbf{P}(|\Phi_k^l(i, \iota_l^r(i,j) + p_l^r(j)) - \Phi_k^l(i, \iota_l^r(i,j)) - p_{lk}(i)p_l^r(j)| \geq (4K)^{-1}\epsilon\theta_k^r(i) \,\Big|\, \mathcal{F}_j^r, \, X_j^r = i) \\
= \; & \mathbf{P}(|\sum_{n=\iota_l^r(i,j)+1}^{\iota_l^r(i,j)+p_l^r(j)} (\phi_k^l(i,n) - p_{lk}(i))| \geq (4K)^{-1}\epsilon\theta_k^r(i) \,|\mathcal{F}_j^r, \, X_j^r = i) \\
\leq \; & \frac{(4K)^4}{\epsilon^4} \frac{1}{(r l(r)\beta_k(i))^2} \quad \leq \frac{(4K)^4}{\epsilon^4} \frac{1}{(r l(r)\underline{\beta}_k)^2}.
\end{aligned}
$$

Note that $\{q^r(j) \geq \theta^r(X_j^r)\} \in \mathcal{F}_j^r$, we have

$$
\begin{aligned}
& \mathbf{P}(|\Phi_k^l(i, \iota_l^r(i,j) + p_l^r(j)) - \Phi_k^l(i, \iota_l^r(i,j)) - p_{lk}(i)p_l^r(j)| \geq \frac{\epsilon}{4K}\theta_k^r(i), \, q^r(j) \geq \theta^r(i) \Big| \mathcal{F}_j^r, \, X_j^r = i) \\
\leq \; & \frac{(4K)^4}{\epsilon^4} \frac{1}{(r l(r)\underline{\beta}_k)^2} \chi(q^r(j) \geq \theta^r(i)).
\end{aligned} \tag{203}
$$

Now we estimate the exogenous arrivals during this review period. We consider class k with $\alpha_k(i) > 0$. Recall that $\kappa_k^r(i,j)$ denotes the total number of class k jobs that have arrived until $t^r(j)$ at the environment state i. Let $\tilde{\xi}_k^r(i,\kappa_k^r(i,j)+1)$ denote the remaining inter-arrival time of the first customer that arrives after $t^r(j)$. Then

$$
\begin{aligned}
&\{|E_k^r(t^r(j+1)) - E_k^r(t^r(j)) - \alpha_k(i)rl(r)| \geq 4^{-1}\epsilon\theta_k^r(i)\}\Lambda^r(i,j,\tilde{\epsilon})\Upsilon(i,j,\tilde{\epsilon}) \\
= &\{E_k^r(t^r(j+1)) - E_k^r(t^r(j)) \geq \lceil \alpha_k(i)rl(r) + \epsilon\theta_k^r(i)/4 \rceil\}\Lambda^r(i,j,\tilde{\epsilon})\Upsilon(i,j,\tilde{\epsilon}) \\
&\cup\{E_k^r(t^r(j+1)) - E_k^r(t^r(j)) \leq \lfloor \alpha_k(i)rl(r) - \epsilon\theta_k^r(i)/4 \rfloor\}\Lambda^r(i,j,\tilde{\epsilon})\Upsilon(i,j,\tilde{\epsilon}) \\
= &\{\sum_{n=\kappa_k^r(i,j)+2}^{\kappa_k^r(i,j)+\lceil \alpha_k(i)rl(r)+\epsilon\theta_k^r(i)/4 \rceil} \xi_k^r(i,n) + \tilde{\xi}_k^r(i,\kappa_k^r(i,j)+1) \leq t^r(j+1) - t^r(j)\}\Lambda^r(i,j,\tilde{\epsilon})\Upsilon(i,j,\tilde{\epsilon}) \\
&\cup\{\sum_{n=\kappa_k^r(i,j)+2}^{\kappa_k^r(i,j)+\lfloor \alpha_k(i)rl(r)-\epsilon\theta_k^r(i)/4 \rfloor} \xi_k^r(i,n) + \tilde{\xi}_k^r(i,\kappa_k^r(i,j)+1) > t^r(j+1) - t^r(j)\}\Lambda^r(i,j,\tilde{\epsilon})\Upsilon(i,j,\tilde{\epsilon}).
\end{aligned}
$$

Since $\tilde{\xi}_k^r(i,\kappa_k^r(i,j)+1) \geq 0$, then

$$
\begin{aligned}
&\{\sum_{n=\kappa_k^r(i,j)+2}^{\kappa_k^r(i,j)+\lceil \alpha_k(i)rl(r)+\epsilon\theta_k^r(i)/4 \rceil} \xi_k^r(i,n) + \tilde{\xi}_k^r(i,\kappa_k^r(i,j)+1) \leq t^r(j+1) - t^r(j)\}\Lambda^r(i,j,\tilde{\epsilon})\Upsilon(i,j,\tilde{\epsilon}) \\
\subset\ &\{\sum_{n=\kappa_k^r(i,j)+2}^{\kappa_k^r(i,j)+\lceil \alpha_k(i)rl(r)+\epsilon\theta_k^r(i)/4 \rceil} \xi_k^r(i,n) \leq t^r(j+1) - t^r(j)\}\Upsilon(i,j,\tilde{\epsilon}) \\
\subset\ &\{\sum_{n=\kappa_k^r(i,j)+2}^{\kappa_k^r(i,j)+\lceil \alpha_k(i)rl(r)+\epsilon\theta_k^r(i)/4 \rceil} \xi_k^r(i,n) \leq (1+\tilde{\epsilon})l(r)\} \\
&\cup\{|t^r(j+1) - t^r(j) - l(r)| > \tilde{\epsilon}l(r)\}\Upsilon(i,j,\tilde{\epsilon}).
\end{aligned}
$$

From the definition of $\Upsilon(i,j,\tilde{\epsilon})$, we have

$$
\begin{aligned}
&\{\sum_{n=\kappa_k^r(i,j)+2}^{\kappa_k^r(i,j)+\lfloor \alpha_k(i)rl(r)-\epsilon\theta_k^r(i)/4 \rfloor} \xi_k^r(i,n) + \tilde{\xi}_k^r(i,\kappa_k^r(i,j)+1) > t^r(j+1) - t^r(j)\}\Lambda^r(i,j,\tilde{\epsilon})\Upsilon(i,j,\tilde{\epsilon}) \\
\subset\ &\{\sum_{n=\kappa_k^r(i,j)+2}^{\kappa_k^r(i,j)+\lfloor \alpha_k(i)rl(r)-\epsilon\theta_k^r(i)/4 \rfloor} \xi_k^r(i,n) > t^r(j+1) - t^r(j) - \frac{\tilde{\epsilon}l(r)}{2K}\}\Upsilon(i,j,\tilde{\epsilon}) \\
\subset\ &\Big(\{\sum_{n=\kappa_k^r(i,j)+2}^{\kappa_k^r(i,j)+\lfloor \alpha_k(i)rl(r)-\epsilon\theta_k^r(i)/4 \rfloor} \xi_k^r(i,n) > (1-\tilde{\epsilon})l(r) - \frac{\tilde{\epsilon}l(r)}{2K}\}\Upsilon(i,j,\tilde{\epsilon})\Big) \\
&\cup\Big(\{|t^r(j+1) - t^r(j) - l(r)| \geq \tilde{\epsilon}l(r)\}\Upsilon(i,j,\tilde{\epsilon})\Big) \\
\subset\ &\{\sum_{n=\kappa_k^r(i,j)+2}^{\kappa_k^r(i,j)+\lfloor \alpha_k(i)rl(r)-\epsilon\theta_k^r(i)/4 \rfloor} \xi_k^r(i,n) > (1-2\tilde{\epsilon})l(r)\} \\
&\cup\Big(\{|t^r(j+1) - t^r(j) - l(r)| \geq \tilde{\epsilon}l(r)\}\Upsilon(i,j,\tilde{\epsilon})\Big).
\end{aligned}
$$

Therefore,

$$
\begin{aligned}
& \{|E_k^r(t^r(j+1)) - E_k^r(t^r(j)) - \alpha_k(i)rl(r)| \geq 4^{-1}\epsilon\theta_k^r(i)\}\Lambda^r(i,j,\tilde{\epsilon})\Upsilon(i,j,\tilde{\epsilon}) \\
\subset \ & \{\sum_{n=\kappa_k^r(i,j)+2}^{\kappa_k^r(i,j)+\lceil\alpha_k(i)rl(r)+\epsilon\theta_k^r(i)/4\rceil} \xi_k^r(i,n) \leq (1+\tilde{\epsilon})l(r)\} \\
& \cup\{\sum_{n=\kappa_k^r(i,j)+2}^{\kappa_k^r(i,j)+\lfloor\alpha_k(i)rl(r)-\epsilon\theta_k^r(i)/4\rfloor} \xi_k^r(i,n) > (1-2\tilde{\epsilon})l(r)\} \\
& \cup(\{|t^r(j+1)-t^r(j)-l(r)| > \tilde{\epsilon}l(r)\}\Upsilon(i,j,\tilde{\epsilon})).
\end{aligned}
$$

Recall that we assume (201). Then

$$
\begin{aligned}
& \{|E_k^r(t^r(j+1)) - E_k^r(t^r(j)) - \alpha_k(i)rl(r)| \geq 4^{-1}\epsilon\theta_k^r(i)\} \cap \Lambda^r(i,j,\tilde{\epsilon}) \cap \Upsilon(i,j,\tilde{\epsilon}) \cap \Gamma^r(i,j) \\
\subset \ & \Big(\{\sum_{n=\kappa_k^r(i,j)+2}^{\kappa_k^r(i,j)+\lceil\alpha_k(i)rl(r)+\epsilon\theta_k^r(i)/4\rceil} \xi_k^r(i,n) \leq (1+\tilde{\epsilon})l(r)\} \qquad (204) \\
& \cup\{\sum_{n=\kappa_k^r(i,j)+2}^{\kappa_k^r(i,j)+\lfloor\alpha_k(i)rl(r)-\epsilon\theta_k^r(i)/4\rfloor} \xi_k^r(i,n) > (1-2\tilde{\epsilon})l(r)\} \\
& \cup(\{|t^r(j+1)-t^r(j)-l(r)| > \tilde{\epsilon}l(r)\} \cap \Upsilon(i,j,\tilde{\epsilon}))\Big) \cap \Gamma^r(i,j). \qquad (205)
\end{aligned}
$$

We will provide a probabilistic bound on each component separately. First,

$$
\begin{aligned}
& \mathbf{P}\Big(\sum_{n=2}^{\lceil\alpha_k(i)rl(r)+\epsilon\theta_k^r(i)/4\rceil} \xi_k^r(i,n) \leq (1+\tilde{\epsilon})l(r)\Big) \\
= \ & \mathbf{P}\Big(\sum_{n=2}^{\lceil\alpha_k(i)rl(r)+\epsilon\theta_k^r(i)/4\rceil} \xi_k(i,n) \leq (1+\tilde{\epsilon})rl(r)\Big) \\
\leq \ & \mathbf{P}\Big(\sum_{n=2}^{\lceil\alpha_k(i)rl(r)+\epsilon\theta_k^r(i)/4\rceil} \xi_k(i,n) - \frac{\lceil\alpha_k(i)rl(r)+\epsilon\theta_k^r(i)/4\rceil - 1}{\alpha_k(i)} \\
& \qquad \leq (1+\tilde{\epsilon})rl(r) - \frac{\lceil\alpha_k(i)rl(r)+\epsilon\theta_k^r(i)/4\rceil - 1}{\alpha_k(i)}\Big) \\
\leq \ & \mathbf{P}\Big(\sum_{n=2}^{\lceil\alpha_k(i)rl(r)+\epsilon\theta_k^r(i)/4\rceil} \xi_k(i,n) - \frac{\lceil\alpha_k(i)rl(r)+\epsilon\theta_k^r(i)/4\rceil - 1}{\alpha_k(i)} \qquad (206) \\
& \qquad \leq (\tilde{\epsilon} - \frac{\epsilon\beta_k(i)}{4\alpha_k(i)})rl(r) + \frac{1}{\alpha_k(i)}\Big). \qquad (207)
\end{aligned}
$$

For any $\epsilon > 0$, choose $\tilde{\epsilon}$ such that $\tilde{\epsilon} \leq \epsilon\underline{\beta}_k/(16\bar{\alpha}_k)$. Hence, $\tilde{\epsilon} - \epsilon\beta_k(i)/(4\alpha_k(i)) < 0$. Since $rl(r) \to \infty$ if $r \to \infty$, we can choose $r_k(\epsilon) > 0$ large enough such that if $r > r_k(\epsilon)$, then

$$
\frac{1}{\alpha_k(i)} \leq \frac{1}{\underline{\alpha}_k} \leq \frac{1}{2}(\frac{\epsilon\beta_k}{4\bar{\alpha}} - \tilde{\epsilon})rl(r) \leq \frac{1}{2}(\frac{\epsilon\beta_k(i)}{4\alpha_k(i)} - \tilde{\epsilon})rl(r).
$$

Hence, for all $r > r_k(\epsilon)$,

$$(\tilde{\epsilon} - \frac{\epsilon\beta_k(i)}{4\alpha_k(i)})rl(r) + \frac{1}{\alpha_k(i)} \leq \frac{1}{2}(\tilde{\epsilon} - \frac{\epsilon\beta_k(i)}{4\alpha_k(i)})rl(r) \leq 0.$$

Therefore,

$$\begin{aligned}
& \mathbf{P}\Big(\sum_{n=2}^{\lceil\alpha_k(i)rl(r)+\epsilon\theta_k^r(i)/4\rceil} \xi_k(i,n) - \frac{\lceil\alpha_k(i)rl(r)+\epsilon\theta_k^r(i)/4\rceil - 1}{\alpha_k(i)} \leq (\tilde{\epsilon} - \frac{\epsilon\beta_k(i)}{4\alpha_k(i)})rl(r) + \frac{1}{\alpha_k(i)}\Big) \\
\leq\ & \mathbf{P}\Big(\Big|\sum_{n=2}^{\lceil\alpha_k(i)rl(r)+\epsilon\theta_k^r(i)/4\rceil} \xi_k(i,n) - \frac{\lceil\alpha_k(i)rl(r)+\epsilon\theta_k^r(i)/4\rceil - 1}{\alpha_k(i)}\Big| \geq \frac{1}{2}(\frac{\epsilon\beta_k(i)}{4\alpha_k(i)} - \tilde{\epsilon})rl(r)\Big) \\
=\ & \mathbf{P}\Big(\Big|\sum_{n=1}^{\lceil\alpha_k(i)rl(r)+\epsilon\theta_k^r(i)/4\rceil-1} \xi_k(i,n) - \frac{\lceil\alpha_k(i)rl(r)+\epsilon\theta_k^r(i)/4\rceil - 1}{\alpha_k(i)}\Big| \geq \frac{1}{2}(\frac{\epsilon\beta_k(i)}{4\alpha_k(i)} - \tilde{\epsilon})rl(r)\Big) \\
\leq\ & \hat{g}_k(\frac{1}{2}(\frac{\epsilon\beta_k(i)}{4\alpha_k(i)} - \tilde{\epsilon})(\alpha_k(i) + \frac{\epsilon\beta_k(i)}{4})^{-1}, (\alpha_k(i) + \frac{\epsilon\beta_k(i)}{4})rl(r)) \\
\leq\ & \hat{g}_k((\frac{\epsilon\underline{\beta}_k}{32\bar{\alpha}_k})(\bar{\alpha}_k + \frac{\epsilon\bar{\beta}_k}{4})^{-1}, \underline{\alpha}_k rl(r)),
\end{aligned} \tag{208}$$

where the second inequality follows from Lemma 30 and the last inequality is from the fact that $\hat{g}_k(x,y)$ is decreasing in x and y.

Second, using the analysis above, we have

$$\begin{aligned}
& \mathbf{P}\Big(\sum_{n=2}^{\lfloor\alpha_k(i)rl(r)-\epsilon\theta_k^r(i)/4\rfloor} \xi_k^r(i,n) \geq (1-2\tilde{\epsilon})l(r)\Big) \\
\leq\ & \mathbf{P}\Big(\sum_{n=2}^{\lfloor\alpha_k(i)rl(r)-\epsilon\theta_k^r(i)/4\rfloor} \xi_k(i,n) - \frac{\lfloor\alpha_k(i)rl(r)-\epsilon\theta_k^r(i)/4\rfloor - 1}{\alpha_k(i)} \geq (\frac{\epsilon\beta_k(i)}{4\alpha_k(i)} - 2\tilde{\epsilon})rl(r) + \frac{2}{\alpha_k(i)}\Big) \\
\leq\ & \mathbf{P}\Big(\sum_{n=2}^{\lfloor\alpha_k(i)rl(r)-\epsilon\theta_k^r(i)/4\rfloor} \xi_k(i,n) - \frac{\lfloor\alpha_k(i)rl(r)-\epsilon\theta_k^r(i)/4\rfloor - 1}{\alpha_k(i)} \geq (\frac{\epsilon\beta_k(i)}{4\alpha_k(i)} - 2\tilde{\epsilon})rl(r)\Big).
\end{aligned} \tag{209}$$

Note that $\tilde{\epsilon} \leq \epsilon\underline{\beta}_k/(16\bar{\alpha}_k)$, hence $((\epsilon\beta_k(i))(4\alpha_k(i))^{-1} - 2\tilde{\epsilon})rl(r) \geq 0$. Therefore,

$$\begin{aligned}
& \mathbf{P}\Big(\sum_{n=2}^{\lfloor\alpha_k(i)rl(r)-\epsilon\theta_k^r(i)/4\rfloor} \xi_k(i,n) - \frac{\lfloor\alpha_k(i)rl(r)-\epsilon\theta_k^r(i)/4\rfloor - 1}{\alpha_k(i)} \geq (\frac{\epsilon\beta_k(i)}{4\alpha_k(i)} - 2\tilde{\epsilon})rl(r)\Big) \\
\leq\ & \mathbf{P}\Big(\Big|\sum_{n=2}^{\lfloor\alpha_k(i)rl(r)-\epsilon\theta_k^r(i)/4\rfloor} \xi_k(i,n) - \frac{\lfloor\alpha_k(i)rl(r)-\epsilon\theta_k^r(i)/4\rfloor - 1}{\alpha_k(i)}\Big| \geq (\frac{\epsilon\beta_k(i)}{4\alpha_k(i)} - 2\tilde{\epsilon})rl(r)\Big) \\
=\ & \mathbf{P}\Big(\Big|\sum_{n=1}^{\lfloor\alpha_k(i)rl(r)-\epsilon\theta_k^r(i)/4\rfloor-1} \xi_k(i,n) - \frac{\lfloor\alpha_k(i)rl(r)-\epsilon\theta_k^r(i)/4\rfloor - 1}{\alpha_k(i)}\Big| \geq (\frac{\epsilon\beta_k(i)}{4\alpha_k(i)} - 2\tilde{\epsilon})rl(r)\Big) \\
\leq\ & \hat{g}_k((\frac{\epsilon\beta_k(i)}{4\alpha_k(i)} - 2\tilde{\epsilon})(\alpha_k(i))^{-1}, \alpha_k(i)rl(r)) \\
\leq\ & \hat{g}_k((\frac{\epsilon\underline{\beta}_k}{32\bar{\alpha}_k})(\bar{\alpha}_k + \frac{\epsilon\bar{\beta}_k}{4})^{-1}, \underline{\alpha}_k rl(r))
\end{aligned} \tag{210}$$

where the second inequality is from Lemma 30, and the last inequality is from the fact that $\hat{g}_k(x,y)$ is decreasing in x and y.

From Lemma 34, we have

$$
\begin{aligned}
& \mathbf{P}(\{|t^r(j+1)-t^r(j)-l(r)|>\tilde{\epsilon}l(r)\}\cap\Gamma^r(i,j)\cap\Upsilon(i,j,\tilde{\epsilon})\,|\mathcal{F}^r_j,X^r_j=i) \\
\leq\;& \sum_{k=1}^{K}\hat{h}_k(\tilde{\epsilon}(4K\bar{\mu}_k)^{-1},\lfloor\underline{\mu}_k rl(r)\rfloor)\chi(q^r(j)\geq\theta^r(i)).
\end{aligned}
$$

Note that $\{\xi^r(i,\kappa^r(i,j)+n),n\geq 2\}$ is independent of $\mathcal{F}^r_j$ and $\chi(q^r(j)\geq\theta^r(X^r_j))$ is measurable with respect to $\mathcal{F}^r_j$. Combining this inequality with (205)-(210) and letting $\tilde{E}^r_k(j)=E^r_k(t^r(j+1))-E^r_k(t^r(j))$, we have

$$
\begin{aligned}
& \mathbf{P}(\{|\tilde{E}^r_k(j)-\alpha_k(i)rl(r)|\geq\frac{\epsilon}{4}\theta^r_k(i)\}\cap\Lambda^r(i,j,\tilde{\epsilon})\cap\Upsilon(i,j,\tilde{\epsilon})\cap\Gamma^r(i,j)|\mathcal{F}^r_j,X^r_j=i) \\
\leq & \Big(2\hat{g}_k((\frac{\epsilon\underline{\beta}_k}{32\bar{\alpha}_k})(\bar{\alpha}_k+\frac{\epsilon\bar{\beta}_k}{4})^{-1},\underline{\alpha}_k rl(r))+\sum_{k=1}^{K}\hat{h}_k(\frac{\tilde{\epsilon}}{4K\bar{\mu}_k},\lfloor\underline{\mu}_k rl(r)\rfloor)\Big)\chi(q^r(j)\geq\theta^r(i)).
\end{aligned}
\tag{211}
$$

Combining (211) with (202) and (203), we have

$$
\begin{aligned}
& \mathbf{P}(\{|\bar{q}^r_k(j+1)-z^r_k(j+1)|\geq\epsilon r^{-1}\theta^r_k(i)\}\cap\Gamma^r(i,j)\cap\Upsilon(i,j,\tilde{\epsilon})\cap\Lambda(i,j,\tilde{\epsilon})|\mathcal{F}^r_j,X^r_j=i) \\
\leq\;& \Big(2\hat{g}_k((\frac{\epsilon\underline{\beta}_k}{32\bar{\alpha}_k})(\bar{\alpha}_k+\frac{\epsilon\bar{\beta}_k}{4})^{-1},\underline{\alpha}_k rl(r))+\sum_{k=1}^{K}\hat{h}_k(\frac{\tilde{\epsilon}}{4K\bar{\mu}_k},\lfloor\underline{\mu}_k rl(r)\rfloor) \\
& +\frac{(4K)^4}{\epsilon^4}\frac{1}{(rl(r)\underline{\beta}_k)^2}\Big)\chi(q^r(j)\geq\theta^r(i)).
\end{aligned}
$$

The above inequality holds for every $i\in I$. Therefore,

$$
\begin{aligned}
& \mathbf{P}(\{|\bar{q}^r_k(j+1)-z^r_k(j+1)|\geq\epsilon r^{-1}\theta^r_k(X^r_j)\}\cap\Gamma^r(j)\cap\Upsilon(X^r_j,j,\tilde{\epsilon})\cap\Lambda(X^r_j,j,\tilde{\epsilon})|\mathcal{F}^r_j) \\
\leq\;& \Big(2\hat{g}_k((\frac{\epsilon\underline{\beta}_k}{32\bar{\alpha}_k})(\bar{\alpha}_k+\frac{\epsilon\bar{\beta}_k}{4})^{-1},\underline{\alpha}_k rl(r))+\sum_{k=1}^{K}\hat{h}_k(\frac{\tilde{\epsilon}}{4K\bar{\mu}_k},\lfloor\underline{\mu}_k rl(r)\rfloor) \\
& +\frac{(4K)^4}{\epsilon^4}\frac{1}{(rl(r)\underline{\beta}_k)^2}\Big)\chi(q^r(j)\geq\theta^r(X^r_j)).
\end{aligned}
$$

Setting $\tilde{\epsilon}=\epsilon\underline{\beta}_k/(16\bar{\alpha}_k)$, we have the desired result. $\square$

Proof of Lemma 47. Note that $z^r(j+1)\geq 0$ and $\bar{q}^r_k(j)=r^{-1}(q^r_k(j)-\theta^r_k(j))$, then

$$
\{q^r(j+1)\not\geq(1-\epsilon)\theta^r(X^r_{j+1})\}\quad\subset\quad\cup_{1\leq k\leq K}\{|\bar{q}^r_k(j+1)-z^r_k(j+1)|\geq\epsilon r^{-1}\theta^r_k(X^r_{j+1})\}.
$$

Therefore,

$$\{q^r(j+1) \not\geq (1-\epsilon)\theta^r(X^r_{j+1})\} \cap \Gamma^r(j) \cap \Upsilon(X^r_j, j, \frac{\epsilon\underline{\beta}}{16\bar{\alpha}}) \cap \Lambda(X^r_j, j, \frac{\epsilon\underline{\beta}}{16\bar{\alpha}})$$
$$\subset \bigcup_{1\le k\le K} \Big(\{|\bar{q}^r_k(j+1) - z^r_k(j+1)| \geq \frac{\epsilon\theta^r_k(X^r_{j+1})}{r}\} \cap \Gamma^r(j) \cap \Upsilon(X^r_j, j, \frac{\epsilon\underline{\beta}_k}{16\bar{\alpha}_k}) \cap \Lambda(X^r_j, j, \frac{\epsilon\underline{\beta}_k}{16\bar{\alpha}_k})\Big).$$

Applying Lemma 45 with slight adaptation, we get the desired result. □

Proof of Lemma 49. We consider a review period during which the target idle policy is implemented, i.e we assume $\tilde{\Gamma}^r(j,\epsilon)$. Without loss of generality, we choose ϵ such that $0 < \epsilon \leq 1$.

Without loss of generality, we assume that $X^r_j = i$. Recall that we have provided an upper bound on the total duration of each target idle review period by (131) and (132) in Section 3.4. That is, if the target idle policy is implemented during the jth review period and $X^r_j = i$, we have

$$\begin{aligned} t^r(j+1) - t^r(j) \quad &\leq \quad \max_{k\in\mathcal{R}(i)} \sum_{n=\kappa^r_k(i,j)+2}^{\kappa^r_k(i,j)+e^r_k(j)} \xi^r_k(i,n) + \tilde{\xi}^r_k(\kappa^r_k(i,j)+1) \\ &\quad + \sum_{k=1}^{K} \sum_{n=\iota^r_k(i,j)+2}^{\iota^r_k(i,j)+\tilde{p}^r_k(j)} \eta^r_k(i,n) + \tilde{\eta}^r_k(i, \iota^r_k(i,j)+1). \end{aligned}$$

Next, we will provide a probabilistic upper bound on $t^r(j+1) - t^r(j)$ based on the inequality above. Recall that we preselected a directed tree $\mathcal{T}(i)$ in order to avoid ambiguity in the description of the policy, and we have also defined $\mathcal{C}_k(i)$ to be the set of child classes of class k at the environment state i. We also define a set of constants for each class k and each environment state i, $1 \leq k \leq K$, $i \in \mathcal{I}$ as follows. Throughout the rest of the proof, we make the convention that the maximum over an empty set is 0. Let

$$\begin{aligned} \tilde{c}_k(i) \quad &= \quad \max\{\frac{1+\tilde{c}_l(i)}{p_{kl}(i)} : \forall l \in \mathcal{C}_k(i)\}, \quad 1 \leq k \leq K, \\ \tilde{C}(i) \quad &= \quad 1 + 2\max_k \beta_k(i)\Big(\max\{\max_{l\in\mathcal{R}(i)}\{\frac{1+2\tilde{c}_l(i)}{\alpha_l(i)}\}, \max_{1\le k\le K}\{\frac{2\tilde{c}_k(i)}{\mu_k(i)}\}\}\Big), \end{aligned}$$

where we let $\tilde{c}_k(i) = 1$ if $\mathcal{C}_k(i) = \emptyset$.

Assuming $\tilde{\Gamma}^r(j,\epsilon)$, $X^r_j = i$, $\Lambda^r(X^r_j, j, \epsilon)$, and $\Upsilon^r(X^r_j, j, \epsilon)$, then we have

$$\{t^r(j+1) - t^r(j) > 2K\tilde{C}(i)\epsilon l(r)\} \cap \Lambda^r(X^r_j, j, \epsilon) \cap \Upsilon^r(X^r_j, j, \epsilon) \cap \tilde{\Gamma}^r(j,\epsilon) \cap \{X^r_j = i\}$$

$$
\begin{aligned}
&\subset \quad (\{\max_{k\in\mathcal{R}(i)}(\sum_{n=\kappa_k^r(i,j)+2}^{\kappa_k^r(i,j)+e_k^r(j)}\xi_k^r(i,n)+\tilde{\xi}_k^r(\kappa_k^r(i,j)+1))>K\tilde{C}(i)\epsilon l(r)\}\cap\Lambda^r(i,j,\epsilon)\cap\tilde{\Gamma}^r(j,\epsilon))\\
&\quad\ \cup(\{\sum_{k=1}^{K}(\sum_{n=\iota_k^r(i,j)+2}^{\iota_k^r(i,j)+\tilde{p}_k^r(j)}\eta_k^r(i,n)+\tilde{\eta}_k^r(i,\iota_k^r(i,j)+1))>K\tilde{C}(i)\epsilon l(r)\}\cap\Upsilon^r(i,j,\epsilon)\cap\tilde{\Gamma}^r(j,\epsilon))\\
&\subset \quad \bigcup_{k\in\mathcal{R}(i)}(\{\sum_{n=\kappa_k^r(i,j)+2}^{\kappa_k^r(i,j)+e_k^r(j)}\xi_k^r(i,n)>(\tilde{C}(i)-(2K)^{-1})\epsilon l(r)\}\cap\tilde{\Gamma}^r(j,\epsilon))\\
&\quad\ \bigcup_{k=1}^{K}(\{\sum_{n=\iota_k^r(i,j)+2}^{\iota_k^r(i,j)+\tilde{p}_k^r(j)}\eta_k^r(i,n)>(\tilde{C}(i)-(2K)^{-1})\epsilon l(r)\}\cap\tilde{\Gamma}^r(j,\epsilon)).
\end{aligned}
\tag{212}
$$

Let $\bar{\theta}^r(i)=\max_{1\le k\le K}\theta_k^r(i)$. From the definition of $e_k^r(j)$ given in Section 3.4 and noting that $X_j^r=i$, we have

$$
e_k^r(j)\le\epsilon\theta_k^r(i)+\tilde{p}_k^r(j)\le\epsilon\bar{\theta}^r(i)+\tilde{p}_k^r(j).
\tag{213}
$$

Since $\xi_k^r(i,n)\ge 0$ for all $n\ge 1$, we have

$$
\begin{aligned}
&\{\sum_{n=\kappa_k^r(i,j)+2}^{\kappa_k^r(i,j)+e_k^r(j)}\xi_k^r(i,n)>(\tilde{C}(i)-\frac{1}{2K})\epsilon l(r)\}\cap\tilde{\Gamma}^r(j,\epsilon)\\
&\subset\{\sum_{n=\kappa_k^r(i,j)+2}^{\kappa_k^r(i,j)+\epsilon\bar{\theta}^r(i)+\tilde{p}_k^r(j)}\xi_k^r(i,n)>(\tilde{C}(i)-\frac{1}{2K})\epsilon l(r)\}\cap\tilde{\Gamma}^r(j,\epsilon)\\
&\subset\{\sum_{n=\kappa_k^r(i,j)+2}^{\kappa_k^r(i,j)+\epsilon\bar{\theta}^r(i)+\tilde{p}_k^r(j)}\xi_k^r(i,n)>(\tilde{C}(i)-\frac{1}{2K})\epsilon l(r),\ \tilde{p}_k^r(j)\le 2\tilde{c}_k(i)\epsilon\bar{\theta}^r(i)\}\\
&\quad\cup(\{\tilde{p}_k^r(j)>2\tilde{c}_k(i)\epsilon\bar{\theta}^r(i)\}\cap\tilde{\Gamma}^r(j,\epsilon))\\
&\subset\{\sum_{n=\kappa_k^r(i,j)+2}^{\kappa_k^r(i,j)+\epsilon\bar{\theta}^r(i)(1+2\tilde{c}_k(i))}\xi_k^r(i,n)>(\tilde{C}(i)-\frac{1}{2K})\epsilon l(r)\}\\
&\quad\cup(\{\tilde{p}_k^r(j)>2\tilde{c}_k(i)\epsilon\bar{\theta}^r(i)\}\cap\tilde{\Gamma}^r(j,\epsilon)).
\end{aligned}
\tag{214}
$$

Similarly, we have

$$
\begin{aligned}
&\{\sum_{n=\iota_k^r(i,j)+2}^{\iota_k^r(i,j)+\tilde{p}_k^r(j)}\eta_k^r(i,n)>(\tilde{C}(i)-\frac{1}{2K})\epsilon l(r)\}\cap\tilde{\Gamma}^r(j,\epsilon)\\
&\subset\{\sum_{n=\iota_k^r(i,j)+2}^{\iota_k^r(i,j)+\tilde{p}_k^r(j)}\eta_k^r(i,n)>(\tilde{C}(i)-\frac{1}{2K})\epsilon l(r),\ \tilde{p}_k^r(j)\le 2\tilde{c}_k(i)\epsilon\bar{\theta}^r(i)\}\\
&\quad\cup(\{\tilde{p}_k^r(j)>2\tilde{c}_k(i)\epsilon\bar{\theta}^r(i)\}\cap\tilde{\Gamma}^r(j,\epsilon))
\end{aligned}
$$

$$\subset \{ \sum_{n=\iota_k^r(i,j)+2}^{\iota_k^r(i,j)+2\tilde{c}_k(i)\epsilon\bar{\theta}^r(i)} \eta_k^r(i,n) > (\tilde{C}(i) - \frac{1}{2K})\epsilon l(r)\}$$
$$\cup(\{\tilde{p}_k^r(j) > 2\tilde{c}_k(i)\epsilon\bar{\theta}^r(i)\} \cap \tilde{\Gamma}^r(j,\epsilon)). \tag{215}$$

Combining the result of (212)-(215), we have

$$\{t^r(j+1) - t^r(j) > 2K\tilde{C}(i)\epsilon l(r)\} \cap \Lambda^r(X_j^r, j, \epsilon) \cap \Upsilon^r(X_j^r, j, \epsilon) \cap \tilde{\Gamma}^r(j,\epsilon) \cap \{X_j^r = i\}$$
$$\subset \Big(\cup_{k\in\mathcal{R}(i)} \{ \sum_{n=\kappa_k^r(i,j)+2}^{\kappa_k^r(i,j)+\epsilon\bar{\theta}^r(i)(1+2\tilde{c}_k(i))} \xi_k^r(i,n) > (\tilde{C}(i) - (2K)^{-1})\epsilon l(r)\}\Big)$$
$$\cup\Big(\cup_{1\le k\le K} \{ \sum_{n=\iota_k^r(i,j)+2}^{\iota_k^r(i,j)+2\tilde{c}_k(i)\epsilon\bar{\theta}^r(i)} \eta_k^r(i,n) > (\tilde{C}(i) - (2K)^{-1})\epsilon l(r)\}\Big)$$
$$\cup\Big(\cup_{1\le k\le K} \{\tilde{p}_k^r(j) > 2\tilde{c}_k(i)\epsilon\bar{\theta}^r(i)\} \cap \tilde{\Gamma}^r(j,\epsilon)\Big). \tag{216}$$

Next, we provide a probabilistic upper bound on the number of jobs processed during the target-idle review period, i.e $\tilde{p}_k^r(j)$ for each k, $1 \le k \le K$. Recall that $\mathcal{C}_k(i)$ denotes the classes that succeed class k immediately, $\mathcal{T}_k(i)$ denotes all the nodes that succeed class k and k itself, i.e the subtree with root node being k, and they are defined in Section 3.4. Note $\tilde{p}_k^r(j) > 2\tilde{c}_k(i)\bar{\theta}^r(i)$ means that we have to see more than $2\tilde{c}_k(i)\bar{\theta}^r(i)$ number of class k completions during the target-idle period in order to fulfill the policy. So if only $2\tilde{c}_k(i)\bar{\theta}^r(i)$ number of class k jobs are processed, there exists a child node of class k, say l, such that it can not reach its safety stock level after processing $\tilde{p}_l^r(j)$ number of jobs or there are less than $\tilde{p}_l^r(j)$ number of jobs for it to process, i.e

$$\{\tilde{p}_k^r(j) > 2\tilde{c}_k(i)\epsilon\bar{\theta}^r(i)\} \cap \tilde{\Gamma}^r(j,\epsilon)$$
$$\subset \Big(\{\exists l \in \mathcal{C}_k(i), \Phi_l^k(\iota_k^r(i,j) + 2\tilde{c}_k(i)\epsilon\bar{\theta}^r(i)) + q_l^r(j) - \tilde{p}_l^r(j) < \theta_l^r(i)\}$$
$$\cup\{\Phi_l^k(\iota_k^r(i,j) + 2\tilde{c}_k(i)\epsilon\bar{\theta}^r(i)) + q_l^r(j) < \tilde{p}_l^r(j)\}\Big) \cap \tilde{\Gamma}^r(j,\epsilon)$$
$$\subset \{\exists l \in \mathcal{C}_k(i), \Phi_l^k(\iota_k^r(i,j) + 2\tilde{c}_k(i)\epsilon\bar{\theta}^r(i)) + q_l^r(j) - \tilde{p}_l^r(j) < \theta_l^r(i)\} \cap \tilde{\Gamma}^r(j,\epsilon)$$
$$\subset (\cup_{l\in\mathcal{C}_k(i)}\{\Phi_l^k(\iota_k^r(i,j) + 2\tilde{c}_k(i)\epsilon\bar{\theta}^r(i)) - \tilde{p}_l^r(j) < \epsilon\bar{\theta}^r(i)\}) \cap \tilde{\Gamma}^r(j,\epsilon)$$

where the last $\subset$ follows from the fact that $q^r(j) \ge (1-\epsilon)\theta^r(i)$ when $\tilde{\Gamma}^r(j,\epsilon)$ holds. We also have

$$\{\Phi_l^k(\iota_k^r(i,j) + 2\tilde{c}_k(i)\epsilon\bar{\theta}^r(i)) - \tilde{p}_l^r(j) < \epsilon\bar{\theta}^r(i)\} \cap \tilde{\Gamma}^r(j,\epsilon)$$

$$\subset \{\Phi_l^k(\iota_k^r(i,j)+2\tilde{c}_k(i)\epsilon\bar{\theta}^r(i))-\tilde{p}_l^r(j)<\epsilon\bar{\theta}^r(i),\, \tilde{p}_l^r(j)\leq 2\tilde{c}_k(i)\epsilon\bar{\theta}^r(i)\}$$
$$\cup(\{\tilde{p}_l^r(j)>2\tilde{c}_k(i)\epsilon\bar{\theta}^r(i)\}\cap\tilde{\Gamma}^r(j,\epsilon))$$
$$\subset \{\Phi_l^k(\iota_k^r(i,j)+2\tilde{c}_k(i)\epsilon\bar{\theta}^r(i))<\epsilon(1+2\tilde{c}_k(i))\bar{\theta}^r(i)\}\cup(\{\tilde{p}_l^r(j)>2\tilde{c}_k(i)\epsilon\bar{\theta}^r(i)\}\cap\tilde{\Gamma}^r(j,\epsilon)).$$

Combining these two results, we have

$$\{\tilde{p}_k^r(j)>2\tilde{c}_k(i)\epsilon\bar{\theta}^r(i)\}\cap\tilde{\Gamma}^r(j,\epsilon)$$
$$\subset \bigcup_{l\in\mathcal{C}_k(i)}\Big(\{\Phi_l^k(\iota_k^r(i,j)+2\tilde{c}_k(i)\epsilon\bar{\theta}^r(i))<\epsilon(1+2\tilde{c}_k(i))\bar{\theta}^r(i)\}\cup(\{\tilde{p}_l^r(j)>2\tilde{c}_k(i)\epsilon\bar{\theta}^r(i)\}\cap\tilde{\Gamma}^r(j,\epsilon))\Big).$$

Going through the same procedure as above for all the nodes that succeed class k, we have

$$\begin{aligned}&\{\tilde{p}_k^r(j)>2\tilde{c}_k(i)\epsilon\bar{\theta}^r(i)\}\cap\tilde{\Gamma}^r(j,\epsilon)\\ \subset\ &\bigcup_{l\in\mathcal{T}_k(i)}\bigcup_{l'\in\mathcal{C}_l(i)}\{\Phi_{l'}^l(\iota_l^r(i,j)+2\tilde{c}_l(i)\epsilon\bar{\theta}^r(i))<\epsilon(1+2\tilde{c}_{l'}(i))\bar{\theta}^r(i)\},\end{aligned} \tag{217}$$

where we let $\cup_{l'\in\mathcal{C}_l(i)}\{\Phi_{l'}^l(\iota_l^r(i,j)+2\tilde{c}_l(i)\epsilon\bar{\theta}^r(i))<\epsilon(1+2\tilde{c}_{l'}(i))\bar{\theta}^r(i)\}=\emptyset$ if $\mathcal{C}_l(i)=\emptyset$, i.e when l is a leaf.

Combining (216) and (217), we have

$$\begin{aligned}&\{t^r(j+1)-t^r(j)>2K\tilde{C}(i)\epsilon l(r)\}\cap\Lambda^r(X_j^r,j,\epsilon)\cap\Upsilon^r(X_j^r,j,\epsilon)\cap\tilde{\Gamma}^r(j,\epsilon)\cap\{X_j^r=i\}\\ \subset\ &\Big(\bigcup_{k\in\mathcal{R}(i)}\{\sum_{n=\kappa_k^r(i,j)+2}^{\kappa_k^r(i,j)+\epsilon\bar{\theta}^r(i)(1+2\tilde{c}_k(i))}\xi_k^r(i,n)>(\tilde{C}(i)-(2K)^{-1})\epsilon l(r)\}\Big)\\ &\cup\Big(\bigcup_{1\leq k\leq K}\{\sum_{n=\iota_k^r(i,j)+2}^{\iota_k^r(i,j)+2\tilde{c}_k(i)\epsilon\bar{\theta}^r(i)}\eta_k^r(i,n)>(\tilde{C}(i)-(2K)^{-1})\epsilon l(r)\}\Big)\\ &\cup\Big(\bigcup_{1\leq k\leq K}\bigcup_{l\in\mathcal{T}_k(i)}\bigcup_{l'\in\mathcal{C}_l(i)}\{\Phi_{l'}^l(\iota_l^r(i,j)+2\tilde{c}_l(i)\epsilon\bar{\theta}^r(i))<\epsilon(1+2\tilde{c}_{l'}(i))\bar{\theta}^r(i)\}\Big).\end{aligned}$$

Therefore,

$$\begin{aligned}&\mathbf{P}\Big(\{t^r(j+1)-t^r(j)>2K\tilde{C}(i)\epsilon l(r)\}\cap\Lambda^r(X_j^r,j,\epsilon)\cap\Upsilon^r(X_j^r,j,\epsilon)\cap\tilde{\Gamma}^r(j,\epsilon)|\mathcal{F}_j^r,X_j^r=i\Big)\\ &\leq\chi(\tilde{\Gamma}^r(j,\epsilon))\Big(\sum_{k\in\mathcal{R}(i)}\mathbf{P}\Big(\sum_{n=\kappa_k^r(i,j)+2}^{\kappa_k^r(i,j)+\epsilon\bar{\theta}^r(i)(1+2\tilde{c}_k(i))}\xi_k^r(i,n)>(\tilde{C}(i)-(2K)^{-1})\epsilon l(r)|\mathcal{F}_j^r,X_j^r=i\Big)\\ &\quad+\sum_{1\leq k\leq K}\mathbf{P}\Big(\sum_{n=\iota_k^r(i,j)+2}^{\iota_k^r(i,j)+2\tilde{c}_k(i)\epsilon\bar{\theta}^r(i)}\eta_k^r(i,n)>(\tilde{C}(i)-(2K)^{-1})\epsilon l(r)|\mathcal{F}_j^r,X_j^r=i\Big)\end{aligned}$$

$$
\begin{aligned}
&+\sum_{1\le k\le K}\sum_{l\in\mathcal{T}_k(i)}\sum_{l'\in\mathcal{C}_l(i)}\mathbf{P}\Big(\Phi^l_{l'}(\iota^r_l(i,j)+2\tilde{c}_l(i)\epsilon\bar{\theta}^r(i))<\epsilon(1+2\tilde{c}_{l'}(i))\bar{\theta}^r(i)|\mathcal{F}^r_j,X^r_j=i\Big)\Big)\\
=\ &\chi(\tilde{\Gamma}^r(j,\epsilon))\Big(\sum_{k\in\mathcal{R}(i)}\mathbf{P}\Big(\sum_{n=1}^{\epsilon\bar{\theta}^r(i)(1+2\tilde{c}_k(i))-1}\xi^r_k(i,n)>(\tilde{C}(i)-(2K)^{-1})\epsilon l(r)\Big)\\
&+\sum_{1\le k\le K}\mathbf{P}\Big(\sum_{n=1}^{2\tilde{c}_k(i)\epsilon\bar{\theta}^r(i)-1}\eta^r_k(i,n)>(\tilde{C}(i)-(2K)^{-1})\epsilon l(r)\Big)\\
&+\sum_{1\le k\le K}\sum_{l\in\mathcal{T}_k(i)}\sum_{l'\in\mathcal{C}_l(i)}\mathbf{P}\Big(\Phi^l_{l'}(2\tilde{c}_l(i)\epsilon\bar{\theta}^r(i))<\epsilon(1+2\tilde{c}_{l'}(i))\bar{\theta}^r(i)\Big)\Big).
\end{aligned}\tag{218}
$$

Note that for all $k\in\mathcal{R}(i)$,

$$
\tilde{C}(i)-\frac{1}{2K}\ge 2\beta_k(i)\frac{1+2\tilde{c}_k(i)}{\alpha_k(i)}
$$

and $\theta^r_k(i)=\beta_k(i)rl(r)$, $\bar{\theta}^r(i)=\max_{1\le k\le K}\beta_k(i)rl(r)$. Thus we have

$$
\begin{aligned}
&\sum_{k\in\mathcal{R}(i)}\mathbf{P}\Big(\sum_{n=1}^{\epsilon\bar{\theta}^r(i)(1+2\tilde{c}_k(i))-1}\xi^r_k(i,n)>(\tilde{C}(i)-(2K)^{-1})\epsilon l(r)\Big)\\
=\ &\sum_{k\in\mathcal{R}(i)}\mathbf{P}\Big(\sum_{n=1}^{\epsilon\bar{\theta}^r(i)(1+2\tilde{c}_k(i))-1}(\xi_k(i,n)-\frac{1}{\alpha_k(i)})>(\tilde{C}(i)-\frac{1}{2K})\epsilon rl(r)-\frac{\epsilon\bar{\theta}^r(i)(1+2\tilde{c}_k(i))-1}{\alpha_k(i)}\Big)\\
\le\ &\sum_{k\in\mathcal{R}(i)}\mathbf{P}\Big(\sum_{n=1}^{\epsilon\bar{\theta}^r(i)(1+2\tilde{c}_k(i))-1}(\xi_k(i,n)-\frac{1}{\alpha_k(i)})>\frac{\epsilon\bar{\theta}^r(i)(1+2\tilde{c}_k(i))}{\alpha_k(i)}\Big)\\
\le\ &\sum_{k\in\mathcal{R}(i)}\hat{g}_k(\frac{\epsilon}{\alpha_k(i)},\bar{\theta}^r(i)(1+2\tilde{c}_k(i)))\\
\le\ &\sum_{k\in\mathcal{R}(i)}\hat{g}_k(\frac{\epsilon}{\bar{\alpha}_k},\underline{\mu}_k rl(r))\\
\le\ &\sum_{1\le k\le K}\hat{g}_k(\frac{\epsilon}{\bar{\alpha}},\underline{\mu}rl(r)),
\end{aligned}\tag{219}
$$

where the last two inequalities are from Lemma 30 and that $\hat{g}_k(x,y)$ is decreasing in x and y.

Similarly, we have

$$
\begin{aligned}
&\sum_{1\le k\le K}\mathbf{P}\Big(\sum_{n=1}^{2\tilde{c}_k(i)\epsilon\bar{\theta}^r(i)-1}\eta^r_k(i,n)>(\tilde{C}(i)-(2K)^{-1})\epsilon l(r)\Big)\\
=\ &\sum_{1\le k\le K}\mathbf{P}\Big(\sum_{n=1}^{2\tilde{c}_k(i)\epsilon\bar{\theta}^r(i)-1}(\eta_k(i,n)-\frac{1}{\mu_k(i)})>(\tilde{C}(i)-\frac{1}{2K})\epsilon rl(r)-\frac{2\tilde{c}_k(i)\epsilon\bar{\theta}^r(i)-1}{\mu_k(i)}\Big)\\
\le\ &\sum_{1\le k\le K}\mathbf{P}\Big(\sum_{n=1}^{2\tilde{c}_k(i)\epsilon\bar{\theta}^r(i)-1}(\eta_k(i,n)-\frac{1}{\mu_k(i)})>\frac{2\tilde{c}_k(i)\epsilon\bar{\theta}^r(i)}{\mu_k(i)}\Big)
\end{aligned}
$$

$$
\begin{aligned}
&\leq \sum_{1\leq k\leq K} \hat{h}_k(\frac{\epsilon}{\mu_k(i)}, 2\tilde{c}_k(i)\beta_k(i)rl(r)) \\
&\leq \sum_{1\leq k\leq K} \hat{h}_k(\frac{\epsilon}{\bar{\mu}}, 2\underline{\mu}rl(r)), \qquad (220)
\end{aligned}
$$

where the last two inequalities are from Lemma 30 and the fact that $\hat{h}_k(x,y)$ is decreasing in x and y.

From Lemma 33 and $2\tilde{c}_l(i)p_{ll'} \geq 2+2\tilde{c}_{l'}(i)$ for all $l' \in \mathcal{C}_l(i)$ and $\bar{\theta}^r(i) = \max_{1\leq k} \beta_k(i)rl(r)$, we have

$$
\begin{aligned}
&\sum_{1\leq k\leq K} \sum_{l\in\mathcal{T}_k(i)} \sum_{l'\in\mathcal{C}_l(i)} \mathbf{P}\Big(\Phi^l_{l'}(2\tilde{c}_l(i)\epsilon\bar{\theta}^r(i)) < \epsilon(1+2\tilde{c}_{l'}(i))\bar{\theta}^r(i)\Big) \\
= &\sum_{1\leq k\leq K} \sum_{l\in\mathcal{T}_k(i)} \sum_{l'\in\mathcal{C}_l(i)} \mathbf{P}\Big(\sum_{n=1}^{2\tilde{c}_l(i)\epsilon\bar{\theta}^r(i)}(\phi^l_{l'}(n) - p_{ll'}(i)) < \epsilon(1+2\tilde{c}_{l'}(i))\bar{\theta}^r(i) - 2\tilde{c}_l(i)\epsilon\bar{\theta}^r(i)p_{ll'}(i)\Big) \\
\leq &\sum_{1\leq k\leq K} \sum_{l\in\mathcal{T}_k(i)} \sum_{l'\in\mathcal{C}_l(i)} \mathbf{P}\Big(\sum_{n=1}^{2\tilde{c}_l(i)\epsilon\bar{\theta}^r(i)}(\phi^l_{l'}(n) - p_{ll'}(i)) < -\epsilon\bar{\theta}^r(i)\Big) \\
\leq &\sum_{1\leq k\leq K} \sum_{l\in\mathcal{T}_k(i)} \sum_{l'\in\mathcal{C}_l(i)} (\frac{\epsilon}{2\tilde{c}_l(i)})^{-4}(\bar{\theta}^r(i))^{-2} \\
\leq &K^3(\frac{\epsilon}{2\tilde{c}_{\max}})^{-4}(\underline{\beta}rl(r))^{-2}, \qquad (221)
\end{aligned}
$$

where

$$
\tilde{c}_{\max} = \max_{i\in\mathcal{I}} \max_{1\leq k\leq K} \tilde{c}_k(i). \qquad (222)
$$

Combining (218)-(221), we have

$$
\begin{aligned}
&\mathbf{P}\Big(\{t^r(j+1) - t^r(j) > 2K\tilde{C}(i)\epsilon l(r)\}\Lambda^r(X^r_j, j, \epsilon)\Upsilon^r(X^r_j, j, \epsilon)\tilde{\Gamma}^r(j,\epsilon)|\mathcal{F}^r_j, X^r_j = i\Big) \\
\leq\ &\chi(\tilde{\Gamma}^r(j,\epsilon))\Big(\sum_{1\leq k\leq K} \hat{g}_k(\frac{\epsilon}{\bar{\alpha}}, \underline{\mu}rl(r)) + \sum_{1\leq k\leq K} \hat{h}_k(\frac{\epsilon}{\bar{\mu}}, 2\underline{\mu}rl(r)) + K^3(\frac{\epsilon}{2\tilde{c}_{\max}})^{-4}(\underline{\beta}rl(r))^{-2}\Big).
\end{aligned}
$$

Since the above bound holds for any i such that $X^r_j = i$, we have

$$
\begin{aligned}
&\mathbf{P}\Big(\{t^r(j+1) - t^r(j) > 2K\tilde{C}(i)\epsilon l(r)\} \cap \Lambda^r(X^r_j, j, \epsilon) \cap \Upsilon^r(X^r_j, j, \epsilon) \cap \tilde{\Gamma}^r(j,\epsilon)|\mathcal{F}^r_j\Big) \\
\leq\ &\chi(\tilde{\Gamma}^r(j,\epsilon))\Big(\sum_{1\leq k\leq K} \hat{g}_k(\frac{\epsilon}{\bar{\alpha}}, \underline{\mu}rl(r)) + \sum_{1\leq k\leq K} \hat{h}_k(\frac{\epsilon}{\bar{\mu}}, 2\underline{\mu}rl(r)) + K^3(\frac{\epsilon}{2\tilde{c}_{\max}})^{-4}(\underline{\beta}rl(r))^{-2}\Big).
\end{aligned}
$$

Letting

$$
\tilde{C}_1 = \max_{i\in\mathcal{I}} \tilde{C}(i), \qquad (223)
$$

and replacing $\tilde{C}(i)$ by $\tilde{C}_1$ in the above inequality, we have the desired result.

$\square$

Proof of Lemma 51. Suppose that the jth review period is a target-idle review period (i.e during this review period, the target-idle policy is implemented). Without loss of generality, we assume that $X_j^r = i$ and $q^r(j) \ngeq \theta^r(i)$. Recall that $\tilde{p}_k^r(j)$ denotes the number of class k service completions during the jth review period. Throughout the rest of the proof, we let $\tilde{E}_k^r(j) = E_k^r(t^r(j+1)) - E_k^r(t^r(j))$, i.e the exogenous arrivals to class k during the jth review period. The queue length of class k at the end of the jth review period is

$$q_k^r(j+1) = q_k^r(j) + \tilde{E}_k^r(j) - \tilde{p}_k^r(j) + \sum_{1\le l\le K} (\Phi(\iota_l^r(i,j) + \tilde{p}_l^r(j)) - \Phi(\iota_l^r(i,j))).$$

Note that $\Phi(\iota_l^r(i,j) + \tilde{p}_l^r(j)) - \Phi(\iota_l^r(i,j)) \le \tilde{p}_l^r(j)$ for all $l = 1, \dots, K$, thus

$$|q_k^r(j+1) - q_k^r(j)| \quad \le \quad E_k^r(t^r(j+1)) - E_k^r(t^r(j)) + \tilde{p}_k^r(j) + \sum_{l=1}^{K} \tilde{p}_l^r(j). \tag{224}$$

For k such that $\alpha_k(i) > 0$, we have

$$\begin{aligned}
&\mathbf{P}(\{\tilde{E}_k^r(j) > 4K\tilde{C}_1\alpha_k(i)\epsilon r l(r)\} \cap \Lambda^r(X_j^r, j, \epsilon)\Upsilon^r(X_j^r, j, \epsilon)\tilde{\Gamma}^r(j,\epsilon) | \mathcal{F}_j^r, X_j^r = i) \\
\le\ & \mathbf{P}(\{\sum_{n=\kappa^r(i,j)+2}^{\kappa^r(i,j)+\lceil 4K\tilde{C}_1\alpha_k(i)\epsilon r l(r)\rceil} \xi_k^r(i,n) + \tilde{\xi}_k^r(i,\kappa^r(i,j)+1) \le 2K\tilde{C}_1\epsilon l(r)\} | \mathcal{F}_j^r, X_j^r = i) \\
&\quad + \mathbf{P}(\{t^r(j+1) - t^r(j) > 2K\tilde{C}_1\epsilon l(r)\} \cap \Lambda^r(X_j^r, j, \epsilon) \cap \Upsilon^r(X_j^r, j, \epsilon) \cap \tilde{\Gamma}^r(j,\epsilon) | \mathcal{F}_j^r, X_j^r = i) \\
\le\ & \mathbf{P}(\{\sum_{n=1}^{\lceil 4K\tilde{C}_1\alpha_k(i)\epsilon r l(r)\rceil - 1} \xi_k^r(i,n) \le 2K\tilde{C}_1\epsilon l(r)\}) + \hat{f}_5(\epsilon, r)
\end{aligned} \tag{225}$$

where the last inequality is from Lemma 49, the fact that $\tilde{\xi}^r(i,n) \ge 0$ for any $n \ge 1$ and any $i \in \mathcal{I}$, and the fact that $\xi^r(\kappa^r(i,j)+n)$ is independent from $\mathcal{F}_j^r$ for $n \ge 2$.

Similar to the proof of (211), choose $r(\epsilon) > 0$ large enough, such that if $r > r(\epsilon)$, then $-2K\tilde{C}_1\epsilon r l(r) + \underline{\alpha}^{-1} \le -K\tilde{C}_1\epsilon r l(r)$. Thus, we have

$$\begin{aligned}
&\mathbf{P}(\sum_{n=1}^{\lceil 4K\tilde{C}_1\alpha_k(i)\epsilon r l(r)\rceil - 1} \xi_k^r(i,n) \le 2K\tilde{C}_1\epsilon l(r)) \\
=\ & \mathbf{P}(\sum_{n=1}^{\lceil 4K\tilde{C}_1\alpha_k(i)\epsilon r l(r)\rceil - 1} (\xi_k(i,n) - \frac{1}{\alpha_k(i)}) \le 2K\tilde{C}_1\epsilon r l(r) - \frac{\lceil 4K\tilde{C}_1\alpha_k(i)\epsilon r l(r)\rceil - 1}{\alpha_k(i)})
\end{aligned}$$

$$
\begin{aligned}
&\leq \mathbf{P}\Big(\sum_{n=1}^{\lceil 4K\tilde{C}_1\alpha_k(i)\epsilon rl(r)\rceil-1}(\xi_k(i,n)-\frac{1}{\alpha_k(i)})\leq -2K\tilde{C}_1\epsilon rl(r)+\frac{1}{\alpha_k(i)}\Big)\\
&\leq \mathbf{P}\Big(\sum_{n=1}^{\lceil 4K\tilde{C}_1\alpha_k(i)\epsilon rl(r)\rceil-1}(\xi_k(i,n)-\frac{1}{\alpha_k(i)})\leq -K\tilde{C}_1\epsilon rl(r)\Big)\\
&\leq \hat{g}_k(\frac{1}{4\bar{\alpha}},4K\tilde{C}_1\bar{\alpha}\epsilon rl(r)),
\end{aligned}
\tag{226}
$$

where the last inequality is from Lemma 30.

Combining (225) and (226) and noting that $\bar{\alpha}\geq\alpha_k(i)$, we have

$$
\begin{aligned}
&\mathbf{P}(\{\tilde{E}_k^r(j)>4K\tilde{C}_1\bar{\alpha}\epsilon rl(r)\}\cap\Lambda^r(X_j^r,j,\epsilon)\Upsilon^r(X_j^r,j,\epsilon)\tilde{\Gamma}^r(j,\epsilon)|\mathcal{F}_j^r)\\
\leq\ &\hat{g}_k(\frac{1}{4\bar{\alpha}},4K\tilde{C}_1\bar{\alpha}rl(r))+\hat{f}_5(\epsilon,r).
\end{aligned}
\tag{227}
$$

From (217) and (221), we have

$$
\mathbf{P}(\{\tilde{p}_k^r(j)>2\tilde{c}_k(i)\bar{\theta}^r(i)\}\cap\tilde{\Gamma}^r(j,1)|\mathcal{F}_j^r,X_j^r=i)\leq K^2(\frac{\epsilon}{2\tilde{c}_{\max}})^{-4}(\underline{\beta}\epsilon rl(r))^{-2}.
$$

Since $\tilde{c}_{\max}\geq\tilde{c}_k(i)$, we can replace $\tilde{c}_k(i)$ by $\tilde{c}_{\max}$ in the above inequality and have

$$
\mathbf{P}(\{\tilde{p}_k^r(j)>2\tilde{c}_{\max}\bar{\beta}\epsilon rl(r)\}\cap\tilde{\Gamma}^r(j,1)|\mathcal{F}_j^r)\leq K^2(\frac{\epsilon}{2\tilde{c}_{\max}})^{-4}(\underline{\beta}\epsilon rl(r))^{-2}. \tag{228}
$$

Let $\tilde{C}_2=4K\tilde{C}_1\bar{\alpha}+2(K+1)\tilde{c}_{\max}\bar{\beta}$. Combining the results of (224), (227) and (228), we have

$$
\begin{aligned}
&\mathbf{P}(\{|q_k^r(j+1)-q_k^r(j)|>\tilde{C}_2\epsilon rl(r)\}\cap\Lambda^r(X_j^r,j,\epsilon)\Upsilon^r(X_j^r,j,\epsilon)\tilde{\Gamma}^r(j,\epsilon)|\mathcal{F}_j^r)\\
\leq\ &\hat{g}_k(\frac{1}{4\bar{\alpha}},4K\tilde{C}_1\bar{\alpha}\epsilon rl(r))+K^2(\frac{\epsilon}{2\tilde{c}_{\max}})^{-4}(\underline{\beta}\epsilon rl(r))^{-2}+\hat{f}_5(\epsilon,r).
\end{aligned}
$$

Since the above result holds for all $k\in\{1,\dots,K\}$, the proof is completed.

$\square$

Proof of Lemma 53. Suppose that the jth review period is a fluid period, i.e during this period the fluid policy is implemented. Thus, it satisfies

$$
q^r(j)\geq\theta^r(X_j^r).
$$

Without loss of generality, we assume that $X_j^r=i$. If the jth review period is not interrupted by an environment transition, then the number of jobs of each class processed during this period is the same as defined by (126). If the jth review period is interrupted by an environment transition, then the actual number of jobs of each class processed during this review period may be less than the value designated by (126). Let $\hat{p}^r_k(j)$ denote the actual number of class k jobs processed during the jth review period, then

$$\hat{p}^r_k(j) \le p^r_k(j) \le r\mu_k(i)l(r). \tag{229}$$

For the rest of the proof, we let $\tilde{E}^r_k(j) = E^r_k(t^r(j+1)) - E^r_k(t^r(j))$. Then the queue length of class k at the end of the jth review period is $q^r_k(j+1)$ and it satisfies

$$q^r_k(j+1) = q^r_k(j) + \tilde{E}^r_k(j) - \hat{p}^r_k(j) + \sum_{l=1}^{K}(\Phi^l_k(i, \iota^r_l(i,j) + \hat{p}^r_l(j)) - \Phi^l_k(i, \iota^r_l(i,j))),$$

where $\iota^r_l(i,j)$ is the number of jobs that has been completed at the environment state i before $t^r(j)$. Note that $0 \le \Phi^l_k(i, \iota^r_l(i,j) + \hat{p}^r_l(j)) - \Phi^l_k(i, \iota^r_l(i,j)) \le \hat{p}^r_l(j)$, therefore,

$$|q^r_k(j+1) - q^r_k(j)| \le \tilde{E}^r_k(j) + \hat{p}^r_k(j) + \sum_{l=1}^{K}\hat{p}^r_l(j) \le \tilde{E}^r_k(j) + (K+1)\bar{\mu}rl(r), \tag{230}$$

where the second inequality is from (229). Following the procedure used in the proof of (225), (226), and (227), except we use Lemma 36 instead of Lemma 49, we know that there exists $r(\epsilon) > 0$ such that if $r > r(\epsilon)$, then

$$\begin{aligned}
&\mathbf{P}(\{\tilde{E}^r_k(j) > 2\bar{\alpha}(1+\epsilon)rl(r)\} \cap \Upsilon^r(X^r_j, j, \epsilon)\chi(q^r(j) \ge \theta^r(X^r_j))|\mathcal{F}^r_j) \\
\le\ & \hat{g}_k(\frac{1}{4\bar{\alpha}}, 2\bar{\alpha}(1+\epsilon)rl(r)) + \sum_{k=1}^{K}\hat{h}_k(\epsilon(4K\bar{\mu}_k)^{-1}, \lfloor \mu_k rl(r) \rfloor).
\end{aligned} \tag{231}$$

Combining the result of (230) and (231), and extending it to the vector form, we have the desired result. □

Proof of Lemma 55. From the definition of $t^r(j)$, $t^r(j+1) - t^r(j)$ is the duration of the jth review period. The duration of a review period depends on the policy implemented during this review period. If the queue length at the beginning of a review period is above its threshold level (i.e $q^r(j) \ge \theta^r(X^r_j)$), then the fluid policy is implemented during this review period; otherwise, the target-idle policy is implemented during this review period. We will consider both of these two cases respectively. Without loss of generality, we assume that the network is staying at the state i of the environment at the beginning of the jth review period, i.e $X_j^r = i$.

First, we consider the case that the fluid policy is implemented during the jth review period, which means that the queue length at the beginning of this review period is above its threshold level, i.e $q^r(j) \geq \theta^r(X_j^r)$. Then from (130), we have

$$(t^r(j+1) - t^r(j))\chi(q^r(j) \geq \theta^r(X_j^r)) \leq (\max_{1\leq s\leq S} e_s^{r,F}(j)),$$

where $e_s^{r,F}(j)$ is defined in (129). From (129), we also see that

$$e_s^{r,F}(j) \leq \max\{b_s^r(j), l(r)\} \leq b_s^r(j) + l(r),$$

where $b_s^r(j)$ is defined in (128). Combining these two inequalities, we have

$$(t^r(j+1) - t^r(j))\chi(q^r(j) \geq \theta^r(X_j^r)) \leq \max_{1\leq s\leq S}\{b_s^r(j) + l(r)\} \leq \sum_{1\leq s\leq S} b_s^r(j) + l(r). \quad (232)$$

From (126) and (128), we have

$$b_s^r(j) \quad \leq \quad \sum_{k\in C_s} \sum_{n=\iota_k^r(i,j)+2}^{\iota_k^r(i,j)+\lfloor r\mu_k(i)l(r)\rfloor} \eta_k^r(i,n) + \tilde{\eta}_k^r(i, \iota_k^r(i,j)+1).$$

Therefore,

$$\begin{aligned}
& \mathbf{E}[\mathbf{E}[b_s^r(j)|\mathcal{F}_j^r, X_j^r = i]] \\
\leq \quad & \sum_{k\in C_s} \mathbf{E}[\mathbf{E}[\sum_{n=\iota_k^r(i,j)+2}^{\iota_k^r(i,j)+\lfloor r\mu_k(i)l(r)\rfloor} \eta_k^r(i,n) + \tilde{\eta}_k^r(i, \iota_k^r(i,j)+1)|\mathcal{F}_j^r, X_j^r = i]] \\
= \quad & \sum_{k\in C_s} \mathbf{E}[\sum_{n=2}^{\lfloor r\mu_k(i)l(r)\rfloor} \eta_k^r(i,n)] + \mathbf{E}[\mathbf{E}[\tilde{\eta}_k^r(i, \iota_k^r(i,j)+1)|\mathcal{F}_j^r, X_j^r = i]] \\
\leq \quad & \sum_{k\in C_s} r\mu_k(i)l(r)(r\mu_k(i))^{-1} + r^{-1}\hat{c}_2(rt) \\
= \quad & \sum_{k\in C_s} l(r) + r^{-1}\hat{c}_2(rt),
\end{aligned}$$

where the second equality is from the fact that $\eta_k^r(\iota_k^r(i,j)+n)$ is independent from $\mathcal{F}_j^r$ for $n \geq 2$, and the second inequality is from Remark 26 and the fact that $t^r(j) \leq t$.

The above inequality holds for any $i \in \mathcal{I}$, hence

$$\mathbf{E}[b_s^r(j)] = \mathbf{E}[\mathbf{E}[b_s^r(j)|\mathcal{F}_j^r]] \quad \leq \quad \sum_{k\in C_s} l(r) + r^{-1}\hat{c}_2(rt). \quad (233)$$

Combining (232) and (233) and noting that $\sum_{1\le s\le S}\sum_{k\in C_s} 1 = K$, we have

$$\mathbf{E}[(t^r(j+1)-t^r(j))\chi(q^r(j)\ge\theta^r(X_j^r))] \quad\le\quad (K+1)l(r)+r^{-1}\hat{c}_2(rt). \tag{234}$$

Second, we consider the case that the jth review period is a target-idle review period, i.e the target-idle policy is implemented during this review period. The jth review period is a target-idle period if and only if the queueing length at the beginning of this review period is not above the safety stock level, i.e $q^r(j)\not\ge\theta^r(X_j^r)$. Then from (132),

$$(t^r(j+1)-t^r(j))\chi(q^r(j)\not\ge\theta^r(X_j^r)) \le \max_{1\le s\le S} e_s^{r,I}(j) \le \sum_{1\le s\le S} e_s^{r,I}(j).$$

From (131), we have

$$\begin{aligned}
&\mathbf{E}[(t^r(j+1)-t^r(j))\chi(q^r(j)\not\ge\theta^r(X_j^r))]\\
\le\;& \sum_{k=1}^{K}\mathbf{E}[\mathbf{E}[\sum_{n=\kappa_k^r(X_j^r,j)+2}^{\kappa_k^r(X_j^r,j)+e_k^r(j)}\xi_k^r(X_j^r,n)+\tilde{\xi}_k^r(\kappa_k^r(X_j^r,j)+1)|\mathcal{F}_j^r]]\\
&+\sum_{k=1}^{K}\mathbf{E}[\mathbf{E}[\sum_{n=\iota_k^r(X_j^r,j)+2}^{\iota_k^r(X_j^r,j)+\tilde{p}_k^r(j)}\eta_k^r(X_j^r,n)+\tilde{\eta}_k^r(X_j^r,\iota_k^r(X_j^r,j)+1)|\mathcal{F}_j^r]]\\
\le\;& \sum_{k=1}^{K}\mathbf{E}[\mathbf{E}[\sum_{n=\kappa_k^r(X_j^r,j)+2}^{\kappa_k^r(X_j^r,j)+e_k^r(j)}\xi_k^r(X_j^r,n)|\mathcal{F}_j^r]]+r^{-1}\hat{c}_1(rt)\\
&+\sum_{k=1}^{K}\mathbf{E}[\mathbf{E}[\sum_{n=\iota_k^r(X_j^r,j)+2}^{\iota_k^r(X_j^r,j)+\tilde{p}_k^r(j)}\eta_k^r(X_j^r,n)|\mathcal{F}_j^r]]+r^{-1}\hat{c}_2(rt),
\end{aligned}\tag{235}$$

where the second inequality is from Remark 26 and the fact that $t^r(j)\le t$.

From the target-idle policy, we know that the number of jobs for each class to process or wait for only depends on the queue length at the beginning of the review period and the routing process. Therefore, $\{e_k^r(j), k\in\mathcal{R}(i)\}$ and $\{\tilde{p}_k^r(j), 1\le k\le K\}$ are independent of $\{\xi_k^r(\kappa_k^r(i,j)+n), \eta_k^r(i,\iota_k^r(i,j)+n),\ n\ge 2\}$. Also note that $\{\eta_k^r(\iota_k^r(i,j)+n), \xi_k^r(\iota_k^r(i,j)+n),\ n\ge 2\}$ is independent of $\mathcal{F}_j^r$. From Wald's equality, we have

$$\begin{aligned}
\mathbf{E}[\mathbf{E}[\sum_{n=\kappa_k^r(i,j)+2}^{\kappa_k^r(i,j)+e_k^r(j)}\xi_k^r(i,n)|\mathcal{F}_j^r,X_j^r=i]] &= \mathbf{E}[\xi_k^r(i,n)]\mathbf{E}[\mathbf{E}[e_k^r(j)|\mathcal{F}_j^r,X_j^r=i]]\\
&= r^{-1}\alpha_k(i)\mathbf{E}[\mathbf{E}[e_k^r(j)|\mathcal{F}_j^r,X_j^r=i]]
\end{aligned}$$

$$
\begin{aligned}
&\leq r^{-1}\bar{\alpha}_k \mathbf{E}[\mathbf{E}[e_k^r(j)|\mathcal{F}_j^r, X_j^r = i]], \\
\mathbf{E}[\mathbf{E}[\sum_{n=\iota_k^r(i,j)+2}^{\iota_k^r(i,j)+\tilde{p}_k^r(j)} \eta_k^r(i,n)|\mathcal{F}_j^r, X_j^r = i]] &= \mathbf{E}[\eta_k^r(i,1)]\mathbf{E}[\mathbf{E}[\tilde{p}_k^r(j)|\mathcal{F}_j^r, X_j^r = i]] \\
&= r^{-1}\mu_k(i)\mathbf{E}[\mathbf{E}[\tilde{p}_k^r(j)|\mathcal{F}_j^r, X_j^r = i]] \\
&\leq r^{-1}\bar{\mu}_k \mathbf{E}[\mathbf{E}[\tilde{p}_k^r(j)|\mathcal{F}_j^r, X_j^r = i]].
\end{aligned}
$$

Therefore,

$$
\mathbf{E}[\mathbf{E}[\sum_{n=\kappa_k^r(X_j^r,j)+2}^{\kappa_k^r(X_j^r,j)+e_k^r(j)} \xi_k^r(X_j^r,n)|\mathcal{F}_j^r]] \leq r^{-1}\bar{\alpha}_k \mathbf{E}[\mathbf{E}[e_k^r(j)|\mathcal{F}_j^r]], \tag{236}
$$

$$
\mathbf{E}[\mathbf{E}[\sum_{n=\iota_k^r(X_j^r,j)+2}^{\iota_k^r(X_j^r,j)+\tilde{p}_k^r(j)} \eta_k^r(X_j^r,n)|\mathcal{F}_j^r]] \leq r^{-1}\bar{\mu}_k \mathbf{E}[\mathbf{E}[\tilde{p}_k^r(j)|\mathcal{F}_j^r]]. \tag{237}
$$

Note that $\tilde{\Gamma}^r(j,1) = \{q^r(j) \ngeq \theta^r(X_j^r)\}$, then going through the same procedure used in the proof of (217), we have

$$
\{\tilde{p}_k^r(j) > 2\tilde{c}_k(i)x\} \cap \tilde{\Gamma}^r(j,1) \subset \bigcup_{l\in\mathcal{T}_k(i)} \bigcup_{l'\in\mathcal{C}_l(i)} \{\Phi_{l'}^l(\iota_l^r(i,j) + 2\tilde{c}_l(i)x) < 2\tilde{c}_{l'}(i)x + \bar{\theta}^r(i)\}. \tag{238}
$$

Note that $\tilde{c}_l(i)p_{ll'}(i) > \tilde{c}_{l'}(i)$ for all l and $l' \in \mathcal{C}_l(i)$, $1 \leq l \leq K$. Hence, there exists $c_3 > 0$ and $0 < c_4 < 1$ such that if $x \geq c_3\bar{\theta}^r(i)$, then $2c_4\tilde{c}_l(i)p_{ll'}(i)x \geq 2\tilde{c}_{l'}(i)x + \bar{\theta}^r(i)$ for all $i \in \mathcal{I}$, $1 \leq l \leq K$ and all $l' \in \mathcal{C}_l(i)$. Hence, for all $x \geq c_3\bar{\theta}^r(i)$,

$$
\begin{aligned}
&\mathbf{P}(\Phi_{l'}^l(\iota_l^r(i,j) + 2\tilde{c}_l(i)x) < 2\tilde{c}_{l'}(i)x + \bar{\theta}^r(i)|\mathcal{F}_j^r, X_j^r = i) \\
= \; &\mathbf{P}(\Phi_{l'}^l(2\tilde{c}_l(i)x) < 2\tilde{c}_{l'}(i)x + \bar{\theta}^r(i)) \\
\leq \; &\mathbf{P}(\Phi_{l'}^l(2\tilde{c}_l(i)x) < 2c_4\tilde{c}_l(i)p_{ll'}(i)x) \\
\leq \; &\mathbf{P}(|\Phi_{l'}^l(2\tilde{c}_l(i)x) - 2\tilde{c}_l(i)p_{ll'}(i)x| > 2(1-c_4)\tilde{c}_l(i)p_{ll'}(i)x) \\
\leq \; &\frac{1}{((1-c_4)p_{ll'}(i))^4(2\tilde{c}_l(i)x)^2},
\end{aligned} \tag{239}
$$

where the last inequality is from Lemma 33.

Let $\beta_{\max} = \max\{\beta_k(i) : 1 \leq k \leq K, i \in \mathcal{I}\}$, $\beta_{\min} = \min\{\beta_k(i) : 1 \leq k \leq K, i \in \mathcal{I}\}$, and $p_{\min} = \min\{p_{kl}(i) : p_{kl}(i) > 0, 1 \leq k, l \leq K, i \in \mathcal{I}\}$. From (238), (239) and $\bar{\theta}^r(i) = \max_{1\leq k\leq K} \beta_k(i) r l(r)$, we have

$$
\mathbf{E}[\mathbf{E}[\tilde{p}_k^r(j)|\mathcal{F}_j^r, X_j^r = i]]
$$

$$
\begin{aligned}
&\leq \mathbf{E}[1+\int_0^\infty \mathbf{P}(\tilde{p}_k^r(j) > y|\mathcal{F}_j^r, X_j^r = i)\, dy] \\
&= 1+\mathbf{E}[\int_0^\infty 2\tilde{c}_k(i)\mathbf{P}(\tilde{p}_k^r(j) > 2\tilde{c}_k(i)x|\mathcal{F}_j^r, X_j^r = i)\, dx] \\
&\leq 1+2\tilde{c}_{\max}\mathbf{E}[\sum_{l\in\mathcal{T}_k(i)}\sum_{l'\in\mathcal{C}_l(i)}\int_0^\infty \mathbf{P}(\Phi_{l'}^l(\iota_l^r(i,j)+2\tilde{c}_l(i)x) < 2\tilde{c}_{l'}(i)x+\bar{\theta}^r(i)|\mathcal{F}_j^r, X_j^r = i)\, dx] \\
&\leq 1+2\tilde{c}_{\max}\mathbf{E}[\sum_{l\in\mathcal{T}_k(i)}\sum_{l'\in\mathcal{C}_l(i)}\Big(c_3\bar{\theta}^r(i)+\int_{c_3\bar{\theta}^r(i)}^\infty \frac{1}{((1-c_4)p_{ll'}(i))^4(2\tilde{c}_l(i)x)^2}\, dx\Big)] \\
&\leq 1+2\tilde{c}_{\max}\mathbf{E}[\sum_{l\in\mathcal{T}_k(i)}\sum_{l'\in\mathcal{C}_l(i)}\Big(c_3\bar{\theta}^r(i)+\frac{1}{((1-c_4)p_{ll'}(i))^4(2\tilde{c}_l(i))^2c_3\bar{\theta}^r(i)}\Big)] \\
&\leq 1+2\tilde{c}_{\max}K^2\Big(c_3\beta_{\max}rl(r)+\frac{1}{4((1-c_4)p_{\min})^4c_3\beta_{\min}rl(r)}\Big),
\end{aligned}
$$

where the last inequality is from the fact that $\tilde{c}_k(i) \geq 1$ for all $k = 1, \ldots K$. Note that $rl(r) \to \infty$ as $r \to \infty$. We choose r_1 large enough so that if $r > r_1$, then $(4((1-c_4)p_{\min})^4c_3\beta_{\min}rl(r))^{-1} \leq 1$. Let

$$
f_1(r) = 2+2\tilde{c}_{\max}K^2c_3\beta_{\max}rl(r),
$$

then the above inequality implies that for any $i \in \mathcal{I}$ and $r > r_1$,

$$
\mathbf{E}[\mathbf{E}[\tilde{p}_k^r(j)|\mathcal{F}_j^r, X_j^r = i]] \leq f_1(r).
$$

Therefore,

$$
\mathbf{E}[\mathbf{E}[\tilde{p}_k^r(j)|\mathcal{F}_j^r]] \leq f_1(r).
$$

Note that we are considering the case that the event $\tilde{\Gamma}(1,j)$ happens. Taking $\epsilon = 1$, then from (213) we have $e_k^r(j) \leq \bar{\theta}^r(X_j^r) + \tilde{p}_k^r(j)$ and

$$
\mathbf{E}[\mathbf{E}[e_k^r(j)|\mathcal{F}_j^r]] \leq \beta_{\max}rl(r) + f_1(r). \tag{240}
$$

Combining the results of (235), (236), (237), (240), and (240), we have

$$
\begin{aligned}
&\mathbf{E}[(t^r(j+1)-t^r(j))\chi(q^r(j) \ngeq \theta^r(X_j^r))] \\
\leq\ &\sum_{k=1}^K r^{-1}\bar{\alpha}_k(\beta_{\max}rl(r)+f_1(r)) + r^{-1}\hat{c}_1(rt) + \sum_{k=1}^K r^{-1}\bar{\mu}_k f_1(r)) + r^{-1}\hat{c}_2(rt).
\end{aligned}
$$

Denote the right hand side of the above inequality as $f_2(r)$, and combining this result with (234), we have

$$
\mathbf{E}(t^r(j+1)-t^r(j)) \leq \max\{(K+1)l(r) + r^{-1}\hat{c}_2(rt), f_2(r)\}
$$

for any $0 \le j \le j^r(t) - 1$. Let $f(r)$ denote the right hand side of the above inequality and simplify its expression, we have the conclusion of the lemma.

□

Proof of Lemma 57. From the definition of $j^r(t)$, we know that

$$\sum_{j=1}^{j^r(t)} (t^r(j) - t^r(j-1)) \le t + (t^r(j^r(t)) - t^r(j^r(t) - 1)). \tag{241}$$

Let $X_0^r = 0$ and $X_n^r = \sum_{j=1}^{n} ((t^r(j) - t^r(j-1)) - \mathbf{E}(t^r(j) - t^r(j-1)|\mathcal{F}_{j-1}^r))$, for $n > 0$. From Lemma 55, $\mathbf{E}|X_n^r| \le 2nf(r) < \infty$. Note that $X_n^r \in \mathcal{F}_n^r$ and $\mathbf{E}[X_{n+1}|\mathcal{F}_n] = X_n$. Therefore, $\{X_n^r, \mathcal{F}_n^r,\ n \ge 1\}$ is a martingale. Note that $j^r(t)$ is optional relative to the filtration $\{\mathcal{F}_n^r, n \ge 1\}$, and $\mathbf{E}[X_n^r] = 0$. By the optional sampling theorem, we know that $\mathbf{E}[X_{j^r(t)\wedge n}^r] = 0$, i.e

$$\mathbf{E}\Big[\sum_{j=1}^{j^r(t)\wedge n} ((t^r(j) - t^r(j-1)) - \mathbf{E}(t^r(j) - t^r(j-1)|\mathcal{F}_{j-1}^r))\Big] = 0. \tag{242}$$

Recall that $\Gamma^r(j) = \{X_{j+1}^r = X_j^r,\ q_j^r \ge \theta^r(X_j^r)\}$ and note that $\chi(q^r(j-1) \ge \theta^r(X_{j-1}^r))$ is measurable with respect to $\mathcal{F}_{j-1}^r$, then

$$\begin{aligned}
&\mathbf{E}[\sum_{j=1}^{j^r(t)\wedge n} (t^r(j) - t^r(j-1))] = \mathbf{E}[\sum_{j=1}^{j^r(t)\wedge n} \mathbf{E}[t^r(j) - t^r(j-1)|\mathcal{F}_{j-1}^r]] \\
\ge\ & (1-\epsilon)l(r)\mathbf{E}[\sum_{j=1}^{j^r(t)\wedge n} \mathbf{E}[\chi(t^r(j) - t^r(j-1) > (1-\epsilon)l(r))\chi(\Gamma^r(j-1))|\mathcal{F}_{j-1}^r]] \\
=\ & (1-\epsilon)l(r)\mathbf{E}\Big[\sum_{j=1}^{j^r(t)\wedge n} \chi(q^r(j-1) \ge \theta^r(X_{j-1}^r))\Big(1 - \mathbf{E}[\chi(X_j^r \ne X_{j-1}^r)|\mathcal{F}_{j-1}^r] \\
& - \mathbf{E}[\chi(t^r(j) - t^r(j-1) < (1-\epsilon)l(r))\chi(X_j^r = X_{j-1}^r)|\mathcal{F}_{j-1}^r]\Big)\Big].
\end{aligned} \tag{243}$$

Applying the same technique as we use to prove (242), we have

$$\begin{aligned}
&\mathbf{E}[\sum_{j=1}^{j^r(t)\wedge n} \chi(q^r(j-1) \ge \theta^r(X_{j-1}^r))\mathbf{E}[\chi(X_j^r \ne X_{j-1}^r)|\mathcal{F}_{j-1}^r]] \\
\le\ & \mathbf{E}[\sum_{j=1}^{j^r(t)\wedge n} \mathbf{E}[\chi(X_j^r \ne X_{j-1}^r)|\mathcal{F}_{j-1}^r]] \\
=\ & \mathbf{E}[\sum_{j=1}^{j^r(t)\wedge n} \chi(X_j^r \ne X_{j-1}^r)] \le \mathbf{E}[N^r(t)].
\end{aligned} \tag{244}$$

Let $\tilde{t}^r(j-1) = t^r(j) - t^r(j-1)$ for the rest of the proof. Since $\chi(q^r(j-1) \geq \theta^r(X^r_{j-1}))$ is measurable with respect to $\mathcal{F}^r_{j-1}$, we have

$$\begin{aligned}
& \mathbf{E}[\sum_{j=1}^{j^r(t)\wedge n} \chi(q^r(j-1) \geq \theta^r(X^r_{j-1}))\mathbf{E}[\chi(\tilde{t}^r(j-1) < (1-\epsilon)l(r))\chi(X^r_j = X^r_{j-1})|\mathcal{F}^r_{j-1}]] \\
= \;& \mathbf{E}[\sum_{j=1}^{j^r(t)\wedge n} \mathbf{E}[\chi(\tilde{t}^r(j-1) < (1-\epsilon)l(r))\chi(q^r(j-1) \geq \theta^r(X^r_{j-1}))\chi(X^r_j = X^r_{j-1})|\mathcal{F}^r_{j-1}]] \\
\leq \;& \mathbf{E}[\sum_{j=1}^{j^r(t)\wedge n} \sum_{k=1}^{K} \hat{h}_k(\epsilon(2K\bar{\mu}_k)^{-1}, \lfloor \underline{\mu}_k rl(r)\rfloor)\chi(q^r(j-1) \geq \theta^r(X^r_{j-1}))] \\
\leq \;& \sum_{k=1}^{K} \hat{h}_k(\epsilon(2K\bar{\mu}_k)^{-1}, \lfloor \underline{\mu}_k rl(r)\rfloor)\mathbf{E}[\sum_{j=1}^{j^r(t)\wedge n} \chi(q^r(j-1) \geq \theta^r(X^r_{j-1}))],
\end{aligned}$$

where the first inequality is from Lemma 35 and it holds if $> r_1(\epsilon)$ for some $r_1(\epsilon) > 0$. Note that $\hat{h}_k(x, y) \to 0$ if $y \to \infty$. Since $rl(r) \to \infty$ if $r \to \infty$, then there exists $r_2(\epsilon) > r_1(\epsilon)$ such that if $r > r_2(\epsilon)$, then $\sum_{k=1}^{K} \hat{h}_k(\epsilon(2K\bar{\mu}_k)^{-1}, \lfloor \underline{\mu}_k rl(r)\rfloor) \leq \epsilon$. Therefore, if $r > r_2(\epsilon)$, then the above inequality implies

$$\begin{aligned}
& \mathbf{E}[\sum_{j=1}^{j^r(t)\wedge n} \chi(q^r(j-1) \geq \theta^r(X^r_{j-1}))\mathbf{E}[\chi(\tilde{t}^r(j-1) < (1-\epsilon)l(r))\chi(X^r_j = X^r_{j-1})|\mathcal{F}^r_{j-1}]] \\
\leq \;& \epsilon\mathbf{E}[\sum_{j=1}^{j^r(t)\wedge n} \chi(q^r(j-1) \geq \theta^r(X^r_{j-1}))]. \qquad (245)
\end{aligned}$$

Combining (243), (244), and (245), for $r > r_2(\epsilon)$, we have

$$\mathbf{E}[\sum_{j=1}^{j^r(t)\wedge n} \tilde{t}^r(j-1)] \geq (1-\epsilon)l(r)\Big(\mathbf{E}[\sum_{j=1}^{j^r(t)\wedge n} \chi(q^r(j-1) \geq \theta^r(X^r_{j-1}))(1-\epsilon) - \mathbf{E}[N^r(t)]\Big).$$

Note that the target idle policy ensures that the queue length of each class is above the safety stock level at the end of this target idle period. Therefore, there will be at least half of the review periods (except those initiated by an environment transition) such that the queue length at the beginning of this review period is above the chosen safety stock level. That is,

$$\sum_{j=1}^{j^r(t)\wedge n} \chi(q^r(j-1) \geq \theta^r(X^r_{j-1})) \geq \frac{1}{2}(j^r(t) \wedge n - N^r(t)).$$

Combining this result with the last inequality and noting that $\tilde{t}^r(j-1) = t^r(j) - t^r(j-1)$, we have

$$\mathbf{E}[\sum_{j=1}^{j^r(t)\wedge n}(t^r(j) - t^r(j-1))] \geq (1-\epsilon)l(r)\Big(\frac{(1-\epsilon)\mathbf{E}[j^r(t)\wedge n]}{2} - \frac{3\mathbf{E}[N^r(t)]}{2}\Big).$$

Let n go to infinity at both sides of the above inequality. Then from the monotone convergence theorem, we have

$$\mathbf{E}[\sum_{j=1}^{j^r(t)}(t^r(j) - t^r(j-1))] \geq (1-\epsilon)l(r)\Big(\frac{(1-\epsilon)\mathbf{E}[j^r(t)]}{2} - \frac{3\mathbf{E}[N^r(t)]}{2}\Big).$$

From (241) and Lemma 55, we have

$$\mathbf{E}[\sum_{j=1}^{j^r(t)}(t^r(j) - t^r(j-1))] \leq t + f(r).$$

Therefore,

$$(1-\epsilon)^2 l(r)\frac{E[j^r(t)]}{2} - \frac{3}{2}(1-\epsilon)l(r)E[N^r(t)] \leq t + f(r),$$

which implies the conclusion of the lemma. □

Proof of Lemma 59. Let

$$B^r(t,\epsilon) \quad = \quad \cup_{1\leq j\leq j^r(t)}\{X_j^r = X_{j-1}^r,\, q^r(j) \ngeq (1-\epsilon)\theta^r(X_j^r)\}.$$

Then

$$B^r(t,\epsilon) \subset \Big(B^r(t,\epsilon)\cap\Lambda^r(t,\frac{\epsilon\underline{\beta}}{32K\bar{\alpha}})\cap\Upsilon^r(t,\frac{\epsilon\underline{\beta}}{32K\bar{\alpha}})\Big)\cup\Big((\Lambda^r(t,\frac{\epsilon\underline{\beta}}{32K\bar{\alpha}}))^c\cup(\Upsilon^r(t,\frac{\epsilon\underline{\beta}}{32K\bar{\alpha}}))^c\Big).$$

Let

$$\tilde{B}^r(j,t,\epsilon) = \{X_j^r = X_{j-1}^r,\, q^r(j) \ngeq (1-\epsilon)\theta^r(X_j^r)\}\Lambda(X_{j-1}^r, j-1, \frac{\epsilon\underline{\beta}}{16\bar{\alpha}})\Upsilon(X_{j-1}^r, j-1, \frac{\epsilon\underline{\beta}}{16\bar{\alpha}}),$$

then from (153), we have

$$B^r(t,\epsilon) \quad \subset \quad \Big(\cup_{1\leq j\leq j^r(t)}\tilde{B}^r(j,t,\epsilon)\Big)\cup\Big((\Lambda^r(t,\frac{\epsilon\underline{\beta}}{32K\bar{\alpha}}))^c\cup(\Upsilon^r(t,\frac{\epsilon\underline{\beta}}{32K\bar{\alpha}}))^c\Big).$$

Therefore, the indicator function satisfies

$$\chi(B^r(t,\epsilon)) \quad \leq \quad \sum_{j=1}^{j^r(t)}\chi(\tilde{B}^r(j,t,\epsilon)) + \chi((\Lambda^r(t,\frac{\epsilon\underline{\beta}}{32K\bar{\alpha}}))^c) + \chi(\Upsilon^r(t,\frac{\epsilon\underline{\beta}}{32K\bar{\alpha}}))^c)$$

hence,

$$\mathbf{P}(B^r(t,\epsilon)) \leq \mathbf{E}[\sum_{j=1}^{j^r(t)} \chi(\tilde{B}^r(j,t,\epsilon))] + \mathbf{P}((\Lambda^r(t, \frac{\epsilon\underline{\beta}}{32K\bar{\alpha}}))^c) + \mathbf{P}((\Upsilon^r(t, \frac{\epsilon\underline{\beta}}{32K\bar{\alpha}}))^c). \quad (246)$$

From Lemma 41 and Lemma 43, we have

$$\mathbf{P}((\Lambda^r(t, \frac{\epsilon\underline{\beta}}{32K\bar{\alpha}}))^c) \leq \hat{f}_1(\frac{\epsilon\underline{\beta}}{32K\bar{\alpha}}, t, r), \qquad \mathbf{P}((\Upsilon^r(t, \frac{\epsilon\underline{\beta}}{32K\bar{\alpha}}))^c) \leq \hat{f}_2(\frac{\epsilon\underline{\beta}}{32K\bar{\alpha}}, t, r). \quad (247)$$

To estimate the first term of the right hand side of the inequality (246), we construct a martingale. Let

$$Y_n^r \;=\; \sum_{j=1}^{n} \Big(\chi(\tilde{B}^r(j,t,\epsilon)) - \mathbf{E}[\chi(\tilde{B}^r(j,t,\epsilon))|\mathcal{F}_{j-1}^r]\Big),$$

then Y_n^r is measurable with respect to $\mathcal{F}_n^r$, $\mathbf{E}|Y_n^r| \leq 2n$ and $\mathbf{E}[Y_{n+1}^r|\mathcal{F}_n^r] = Y_n^r$. Therefore, $\{Y_n^r, n \geq 1\}$ is a martingale with respect to the filtration $\{\mathcal{F}_n^r, n \geq 1\}$. Note that $j^r(t)$ is a stopping time with respect to the filtration $\{\mathcal{F}_n^r, n \geq 1\}$. From the optional stopping theorem and from the fact that $\mathbf{E}[Y_n^r] = 0$, we have $\mathbf{E}[Y_{j^r(t)\wedge n}^r] = 0$. Therefore,

$$\begin{aligned}
& \mathbf{E}[\sum_{j=1}^{j^r(t)\wedge n} \chi(\tilde{B}^r(j,t,\epsilon))] \\
= \;& \mathbf{E}[\sum_{j=1}^{j^r(t)\wedge n} \mathbf{E}[\chi(\tilde{B}^r(j,t,\epsilon))|\mathcal{F}_{j-1}^r]] \\
= \;& \mathbf{E}[\sum_{j=1}^{j^r(t)\wedge n} \mathbf{E}[\chi(\tilde{B}^r(j,t,\epsilon))\chi(q^r(j-1) \geq \theta^r(X_{j-1}^r))|\mathcal{F}_{j-1}^r]] \qquad (248) \\
\leq \;& \mathbf{E}[\sum_{j=1}^{j^r(t)\wedge n} \hat{f}_4(\epsilon, r)]. \qquad (249)
\end{aligned}$$

Note that if $q^r(j-1) \ngeq \theta^r(X_{j-1}^r)$, then we implement the target-idle policy such that $q^r(j) \geq \theta^r(X_j^r)$ if the $(j-1)$th review period is not interrupted by an environment transition, i.e $X_j^r = X_{j-1}^r$. Therefore (248) holds. The inequality (249) is from Lemma 47. Let $n \to \infty$. From the monotone convergence theorem, we have

$$\mathbf{E}[\sum_{j=1}^{j^r(t)} \chi(\tilde{B}^r(j,t,\epsilon))] \;\leq\; \mathbf{E}[\sum_{j=1}^{j^r(t)} \hat{f}_4(\epsilon, r)] \;=\; \hat{f}_4(\epsilon, r)\mathbf{E}[j^r(t)]. \quad (250)$$

Combining the result of (246), (247) and (250), we have

$$\mathbf{P}(B^r(t,\epsilon)) \;\leq\; \hat{f}_1(\frac{\epsilon\underline{\beta}}{32K\bar{\alpha}}, t, r) + \hat{f}_2(\frac{\epsilon\underline{\beta}}{32K\bar{\alpha}}, t, r) + \hat{f}_4(\epsilon, r)\mathbf{E}[j^r(t)].$$

From Remark 58, for any fixed $t > 0$ and $\epsilon > 0$, we know that $\mathbf{E}[j^r(t)] \leq \tilde{f}_7(\epsilon, t) < \infty$. From Remark 42 and Remark 44, for any fixed $\epsilon > 0$ (w.l.o.g, $\epsilon < 1$) and $t > 0$, we have $\hat{f}_n(\epsilon, t, r) = \mathcal{O}(r^{-(1+\gamma')})$ for $n = 1, 2$. From Remark 48, we know that $\hat{f}_4(\epsilon, r) = o(r^{-(1+\gamma/9)})$. Therefore, for any fixed $t > 0$ and $\epsilon > 0$, $\mathbf{P}(B^r(t,\epsilon)) = o(r^{-(1+\gamma/9)})$. Therefore, for any sequence $\{r_n, n \geq 1\}$ such that $r_n \to \infty$ if $n \to \infty$, we have

$$\sum_{n=1}^{\infty} \mathbf{P}(B^{r_n}(t,\epsilon)) < \infty.$$

Applying the Broel-Cantelli Lemma, we know that

$$\mathbf{P}(\cap_{m\geq 1} \cup_{n \geq m} B^{r_n}(t,\epsilon)) = 0,$$

which implies the conclusion of Lemma 59 since it holds for any sequence $\{r_n, n \geq 1\}$. $\square$

Proof of Lemma 60. The proof is similar to that of Lemma 59. We let

$$B^r(t,\epsilon) = \bigcup_{1 \leq j \leq j^r(t)} (\{X_j^r = X_{j-1}^r,\, q^r(j-1) \geq \theta^r(X_{j-1}^r)\} \cap \{|(t^r(j) - t^r(j-1)) - l(r)| \geq \epsilon l(r)\}).$$

Then

$$B^r(t,\epsilon) \quad \subset \quad (B^r(t,\epsilon) \cap \Upsilon^r(t,\epsilon)) \cup (\Upsilon^r(t,\epsilon))^c.$$

Recall that $\Gamma^r(i,j) = \{X_j^r = X_{j+1}^r,\, q^r(j) \geq \theta^r(i)\}$. Let $\tilde{t}^r(j-1) = t^r(j) - t^r(j-1)$, then

$$\begin{aligned} B^r(t,\epsilon) \quad &\subset \Big(\bigcup_{1 \leq j \leq j^r(t)} (\{|\tilde{t}^r(j-1) - l(r)| \geq \epsilon l(r)\} \cap \Gamma^r(j-1) \cap \Upsilon^r(t,\epsilon)) \Big) \cup (\Upsilon^r(t,\epsilon))^c \\ &\subset \Big(\bigcup_{1 \leq j \leq j^r(t)} (\{|\tilde{t}^r(j-1) - l(r)| \geq \epsilon l(r)\} \cap \Gamma^r(j-1) \cap \Upsilon^r(X_j^r, j, \epsilon)) \Big) \cup (\Upsilon^r(t,\epsilon))^c. \end{aligned}$$

The rest of the proof is similar to that of Lemma 59, except that we apply the result of Lemma 34 instead of Lemma 47 in a similar inequality to (249). This concludes the proof of Lemma 60. $\square$

Proof of Lemma 61. The proof is similar to that of Lemma 60 except that we apply the result of Lemma 36 instead of Lemma 34. $\square$

Proof of Lemma 62. The proof is similar to that of Lemma 59 except that we apply the result of Lemma 49 instead of Lemma 47. $\square$

Proof of Lemma 63. Taking $\epsilon = 1$ in Lemmas 61 and 62, and noting that $K\tilde{C}_1 \geq 1$, we have the result of Lemma 63.

$\square$

Proof of Lemma 64. The proof is similar to that of Lemma 59 except that we apply the result of Lemma 45 instead of Lemma 47 to prove an inequality similar to (249). $\square$

Proof of Lemma 65. The proof is similar to that of Lemma 59 except that we apply the result of Lemma 51 instead of Lemma 47 to prove an inequality similar to (249). $\square$

Proof of Lemma 67. The rest of the proof is similar to that of Lemma 59, except that we apply the result of Lemma 53 instead of Lemma 47 to obtain an inequality similar to (249). $\square$

Proof of Lemma 68. Take $\epsilon = 1$, then the result of Lemma 68 follows immediately from that of Lemma 65 and Lemma 67.

$\square$

Proof of Lemma 70. For any $s \geq 0$, recall that $j^r(s)$th is the index of the first review period after s, hence

$$t^r(j^r(s) - 1) \leq s \leq t^r(j^r(s)).$$

From Lemma 63, we know that for any $0 \leq s \leq t$,

$$|t^r(j^r(s)) - s| \leq |t^r(j^r(s)) - t^r(j^r(s) - 1)| \leq 2K\tilde{C}_1 l(r).$$

Note that $l(r) \to 0$ as $r \to \infty$. Thus, we have

$$\lim_{r\to\infty} \sup_{0\leq s\leq t} |t^r(j^r(s)) - s| = 0$$

which concludes the proof of the lemma. $\square$

Proof of Lemma 71. We consider $t > 0$ and $s \in [0, t]$, then from the definition of $n^r(\cdot, \cdot)$,

$$n^r(0, s)l(r) - t^r(j^r(s))$$

$$
\begin{aligned}
&= \sum_{j=0}^{j^r(s)-1} \Big(\chi(q^r(j) \geq \theta^r(X_j^r),\, X_{j+1}^r = X_j^r) l(r) - (t^r(j+1) - t^r(j))\Big) \\
&= \sum_{j=0}^{j^r(s)-1} \chi(q^r(j) \geq \theta^r(X_j^r),\, X_{j+1}^r = X_j^r) \Big(l(r) - (t^r(j+1) - t^r(j))\Big) \\
&\quad - \sum_{j=0}^{j^r(s)-1} \chi(q^r(j) \not\geq \theta^r(X_j^r),\, X_{j+1}^r = X_j^r)(t^r(j+1) - t^r(j)) \\
&\quad - \sum_{j=0}^{j^r(s)-1} \chi(X_{j+1}^r \neq X_j^r)(t^r(j+1) - t^r(j)).
\end{aligned}
$$

Therefore,

$$
\begin{aligned}
&|n^r(0,s)l(r) - t^r(j^r(s))| \\
\leq & \sum_{j=0}^{j^r(s)-1} \chi(q^r(j) \geq \theta^r(X_j^r),\, X_{j+1}^r = X_j^r)|t^r(j+1) - t^r(j) - l(r)| \\
&+ \sum_{j=0}^{j^r(s)-1} \chi(q^r(j) \not\geq \theta^r(X_j^r),\, X_{j+1}^r = X_j^r)(t^r(j+1) - t^r(j)) \qquad (251) \\
&+ \sum_{j=0}^{j^r(s)-1} \chi(X_{j+1}^r \neq X_j^r)(t^r(j+1) - t^r(j)).
\end{aligned}
$$

From Lemma 60, for any $\epsilon > 0$ and $t > 0$ and almost any sample path ω, there exists $r(\omega, t, \epsilon) > 0$ such that for all $0 \leq j \leq j^r(t)$, if $r > r(\omega, t, \epsilon)$,

$$
\begin{aligned}
&\chi(q^r(j) \geq \theta^r(X_j^r),\, X_{j+1}^r = X_j^r)|t^r(j+1) - t^r(j) - l(r)| \\
\leq\ & \chi(q^r(j) \geq \theta^r(X_j^r),\, X_{j+1}^r = X_j^r)\epsilon l(r). \qquad (252)
\end{aligned}
$$

For the second term of the right hand side of the above inequality, we have

$$
\begin{aligned}
&\sum_{j=0}^{j^r(s)-1} \chi(q^r(j) \not\geq \theta^r(X_j^r),\, X_{j+1}^r = X_j^r)(t^r(j+1) - t^r(j)) \\
\leq & \sum_{j=1}^{j^r(s)-1} \chi(q^r(j) \not\geq \theta^r(X_j^r),\, X_{j+1}^r = X_j^r, X_j^r = X_{j-1}^r)(t^r(j+1) - t^r(j)) \qquad (253) \\
&+ \sum_{j=1}^{j^r(s)-1} \chi(X_j^r \neq X_{j-1}^r)(t^r(j+1) - t^r(j)) + (t^r(1) + t^r(0)).
\end{aligned}
$$

From Lemma 59, for any $\epsilon > 0$, there exists $r(\omega, t, \epsilon) > 0$ such that if $r > r(\omega, t, \epsilon)$, then for all $1 \leq j \leq j^r(t)$, we have $\{X_j^r = X_{j-1}^r\} = \{X_j^r = X_{j-1}^r, q^r(j) \geq (1-\epsilon)\theta^r(X_j^r)\}$. Without loss of generality, we assume $0 < \epsilon < 1$ and we assume $r > r(\omega, t, \epsilon)$ for the rest of the proof. Then,

$$\begin{aligned}
& \sum_{j=1}^{j^r(s)-1} \chi(q^r(j) \ngeq \theta^r(X_j^r),\, X_{j+1}^r = X_j^r, X_j^r = X_{j-1}^r)(t^r(j+1) - t^r(j)) \\
= & \sum_{j=1}^{j^r(s)-1} \chi(q^r(j) \ngeq \theta^r(X_j^r), X_{j+1}^r = X_j^r = X_{j-1}^r, q^r(j) \geq (1-\epsilon)\theta^r(X_j^r))(t^r(j+1) - t^r(j)) \\
\leq & \sum_{j=1}^{j^r(s)-1} \chi(q^r(j) \ngeq \theta^r(X_j^r), q^r(j) \geq (1-\epsilon)\theta^r(X_j^r),\, X_{j+1}^r = X_j^r, X_j^r = X_{j-1}^r) 2K\tilde{C}_1 \epsilon l(r) \\
\leq & \sum_{j=1}^{j^r(s)-1} \chi(q^r(j-1) \geq \theta^r(X_{j-1}^r),\, X_{j-1}^r = X_j^r) 2K\tilde{C}_1 \epsilon l(r), \qquad (254)
\end{aligned}$$

where the first inequality is from Lemma 62. Note that from the designated policy, the queue length at the end of an uninterrupted target-idle period will be above the safety stock level, i.e $\{q^r(j) \ngeq \theta^r(X_j^r),\, X_{j-1}^r = X_j^r\} \subset \{q^r(j-1) \geq \theta^r(X_{j-1}^r)\, X_{j-1}^r = X_j^r\}$. This implies the second inequality above.

Combining (251)-(254), we have

$$\begin{aligned}
|n^r(0,s)l(r) - t^r(j^r(s))| \leq & \sum_{j=0}^{j^r(s)-1} \chi(q^r(j) \geq \theta^r(X_j^r),\, X_{j+1}^r = X_j^r)(1 + 2K\tilde{C}_1)\epsilon l(r) \quad (255) \\
& +2 \sum_{j=0}^{j^r(s)-1} \chi(X_{j+1}^r \neq X_j^r)(t^r(j+1) - t^r(j)) + (t^r(1) - t^r(0)).
\end{aligned}$$

From Lemma 60, we have

$$\chi(q^r(j) \geq \theta^r(X_j^r),\, X_{j+1}^r = X_j^r)l(r) \leq \chi(q^r(j) \geq \theta^r(X_j^r),\, X_{j+1}^r = X_j^r)\frac{(t^r(j+1) - t^r(j))}{1-\epsilon}$$

Therefore,

$$\begin{aligned}
& \sum_{j=0}^{j^r(s)-1} \chi(q^r(j) \geq \theta^r(X_j^r),\, X_{j+1}^r = X_j^r)(1 + 2K\tilde{C}_1)\epsilon l(r) \\
\leq & \sum_{j=0}^{j^r(s)-1} \chi(q^r(j) \geq \theta^r(X_j^r),\, X_{j+1}^r = X_j^r)(1 + 2K\tilde{C}_1)\epsilon \frac{(t^r(j+1) - t^r(j))}{1-\epsilon} \\
\leq & \frac{(1 + 2K\tilde{C}_1)\epsilon}{1-\epsilon}(s + t^r(j^r(s)) - t^r(j^r(s)-1)) \\
\leq & \frac{(1 + 2K\tilde{C}_1)\epsilon}{1-\epsilon}(s + 2K\tilde{C}_1 l(r)), \qquad (256)
\end{aligned}$$

where the last inequality is from Lemma 63. And from Lemma 63, we also have

$$\begin{aligned} & 2\sum_{j=0}^{j^r(s)-1}\chi(X^r_{j+1}\neq X^r_j)(t^r(j+1)-t^r(j))+(t^r(1)-t^r(0)) \\ \leq\ & 2\sum_{j=0}^{j^r(s)-1}\chi(X^r_{j+1}\neq X^r_j)2K\tilde{C}_1 l(r) \\ =\ & 4K\tilde{C}_1(N^r(s)+1)l(r), \end{aligned} \tag{257}$$

where $N^r(s)$ denotes the number of environment transitions until time s. Without loss of generality, we choose $0<\epsilon<2^{-1}$, then combining the results of (255)-(257), we have

$$\begin{aligned} |n^r(0,s)l(r)-t^r(j^r(s))| &\leq 2(1+2K\tilde{C}_1)\epsilon s+2K\tilde{C}_1(2N^r(s)+2K\tilde{C}_1+3)l(r) \\ &\leq 2(1+2K\tilde{C}_1)\epsilon t+2K\tilde{C}_1(2N^r(t)+2K\tilde{C}_1+3)l(r). \end{aligned}$$

From Lemma 15, we know that there exists $r_2(\omega,t,\epsilon)>0$, such that $N^r(t)=N(t)$ if $r>r_2(\omega,t,\epsilon)$. Since $l(r)\to 0$ when $r\to\infty$, we can choose $r(\omega,t,\epsilon)>\max\{r_1(\omega,t,\epsilon),r_2(\omega,t,\epsilon)\}$ such that if $r>r(\omega,t,\epsilon)$, then $2K\tilde{C}_1(2N^r(t)+2K\tilde{C}_1+3)l(r)l(r)<\epsilon$ and thus

$$|n^r(0,s)l(r)-t^r(j^r(s))| \quad\leq\quad (2(1+2K\tilde{C}_1)t+1)\epsilon$$

for all $s\in[0,t]$. Combining this result and Lemma 70, we have the conclusion of Lemma 71.

□

Proof of Lemma 73. First, since $X(\cdot)$ satisfies the regularity condition, we know that for fixed $t\geq 0$, there exists a finite $m\geq 0$ and $\epsilon_0>0$, such that $\tau_m\leq t<t+\epsilon_0<\tau_{m+1}$. From Lemma 15, there exists $r_1>0$ such that if $r>r_1$, $X^r(\tau^r_n)=X(\tau_n)$ for all $n=1,\ldots,m$, and $\tau^r_{m+1}>t+\epsilon_0$. Let $i_n=X(\tau_n)$, $n=0,\ldots,m$, if $r>r_1$, then for all $s\in[0,t+\epsilon_0]$,

$$X^r(s)\in\{i_0,\ldots,i_m\}. \tag{258}$$

Throughout the rest of the proof, we consider only those r such that $r>r_1$. Hence, up to time t there are only finite environment transitions for all networks. From Lemma 70, we know that

$$t^r(j^r(\cdot))\to\mathbf{1}(\cdot)\quad\text{u.o.c}\quad\text{as } r\to\infty.$$

From Corollary 22,

$$\lim_{r\to\infty} \sup_{0\le s\le t} |r^{-1}Z^r(s) - Z^r(t^r(j^r(s)))/r| = 0.$$

Note that $Z^r(t^r(j^r(t)))$ is also denoted by $q^r(j^r(t))$, we see that

$$\lim_{r\to\infty} \sup_{0\le s\le t} |r^{-1}Z^r(s) - q^r(j^r(s))/r| = 0. \tag{259}$$

From the definition of $\bar{q}^r(j)$, for any $0 \le s \le t$,

$$|\bar{q}^r(j^r(s)) - q^r(j^r(s))/r| \le \theta^r(X^r_{j^r(s)})/r.$$

From Lemma 70, there exists $r_2 > 0$ such that if $r > r_2$, for all $0 \le s \le t$,

$$t^r(j^r(s)) \le t + \epsilon_0.$$

Note $X^r_{j^r(s)} = X^r(t^r(j^r(s)))$, from (258), and recall the definition of $\theta^r(i)$, $i \in \mathcal{I}$, we know that for all $0 \le s \le t$, if $r > \max\{r_1, r_2\}$,

$$|\bar{q}^r(j^r(s)) - q^r(j^r(s))/r| \le \max\{\theta^r(i_n)/r : 0 \le n \le m\} = \max\{\beta(i_n)l(r) : 0 \le n \le m\}.$$

Since $l(r) \to 0$ as $r \to \infty$, we have

$$\lim_{r\to\infty} \sup_{0\le s\le t} |\bar{q}^r(j^r(s)) - q^r(j^r(s))/r| = 0. \tag{260}$$

From (259) and (260), we see that the conclusion of Lemma 73 holds. □

3.8 Summary

In this chapter, we presented a study of the dynamic scheduling of computer communication networks with time varying characteristics. In particular, we model such networks as multiclass queueing networks in a slowly changing environment and we provided a hierarchy decision frame work for such networks.

We consider that the network is operating in slowly changing environment. The changing environment is modelled as a general stochastic process which takes only discrete values, where each value represent an environment state or a network operating state. The arrival processes, service processes and routing matrices are marked renewal processes for each environment state. Our focus in this chapter is to establish a frame work to facilitate the searching for a nearly optimal scheduling policy for such networks.

We first show that a general mutliclass open queueing network in a slowly changing environment can be approximated by a stochastic fluid model when the dynamics of the network tend to change much more frequently than the environment changes states. Next we provide a general method to derive a scheduling policy from any given solution of the stochastic fluid model. We also show that if implementing the derived policy, the dynamics of the network captures the fluid level evolution of the given stochastic fluid model solution. This result holds under very general conditions.

Through this study, we have established a general approach to searching for an nearly optimal scheduling policy for multiclass queueing networks in a slowly changing environment. This is a three step approach. The first step is to approximate the queueing network by a stochastic fluid model; the second step is to solve the stochastic fluid model which is much more tractable than the original queueing network; the third step is to derive a scheduling policy for the original network through the method we provided in this study. If the solution of the stochastic fluid model is optimal, then the derived scheduling policy is asymptotically optimal in the fluid scale.

CHAPTER 4

SUMMARY AND CONCLUSIONS

In this study, we investigate the dynamic scheduling of computer communication networks that can be periodically overloaded. We model such networks as mutliclass queueing networks in a slowly changing environment. We establish a hierarchical framework to search for a suitable scheduling policy for such networks through its connection with stochastic fluid models. We first study the dynamic scheduling of a multiclass stochastic fluid model where the server is under the quality of service contract. Then, we unveil a relationship between the scheduling of stochastic fluid models and that of the queueing networks in a changing environment.

In the multiclass stochastic fluid model that we study, we focus on a system with two fluid classes and a single server whose capacity can be shared arbitrarily among these two classes. The server is under a quality of service contract which is indicated by a threshold value of each class. Whenever the fluid level of a certain class is above the designated threshold value, penalty cost is incurred to the server. We also allow that the server may be overloaded transiently. We specify the optimal and asymptotically optimal resource allocation policies for such a stochastic fluid model.

Afterwards, we relate the problem of optimizing the queueing network to that of optimizing the stochastic fluid model. We connect them by providing a general and successful interpretation of the fluid model solution in order to construct a scheduling policy of the queueing network. The connection we establish facilitates the process of searching for a nearly optimal scheduling policy for the queueing network.

To establish the connection between the queueing networks in a changing environment and the stochastic fluid models, we take a two step approach. The first step is to approximate such networks by their corresponding stochastic fluid models with a proper scaling method. The second step is to provide a general scheme to interpret the stochastic fluid model solution and construct a suitable policy for the queueing network.

In the first step, we scale the space and the rate of arrival and service processes in a similar fashion as the scaling method of the law of large numbers. With this scaling method and assuming that the changing environment can be captured by a limiting stochastic process as the scaler increases, we prove that all the limiting points of queueing processes satisfy a stochastic fluid model. The stochastic fluid model captures the stochastic pattern of the operating environment of the network, but replaces the discrete events that changes highly frequently by their average values. In other words, in a stochastic fluid model, the operating environment of network still randomly transits from one state to another state. However, at each state of the environment, the discrete customers of the network are replaced by fluid units; and the dynamics of the network at each particular environment state is deterministic, which is determined by the associated arrival rates, service rates, and routing proportions matrices.

The stochastic fluid model has a much simpler structure than the queueing network model and is easier to study than the original queueing network model. Assuming that the fluid trajectory, i.e the evolution of the fluid levels, of the stochastic fluid model is given, we provide a method to construct a scheduling policy for the original queueing network. With the derived scheduling policy, the dynamics of the queueing network tracks the fluid level evolution of the given stochastic fluid solution almost surely. Therefore, if the optimal fluid trajectory of the stochastic fluid model is given, then the original network is controlled in a nearly optimal way.

This two step approach provides us with a general hierarchical scheme to search for an asymptotically optimal scheduling policy for the queueing networks in a slowly changing environment. It is important to note that although our research is motivated by the computer computing paradigms for Internet services, it also applies to other type of networks with similar characteristics.

APPENDIX A

HOLDING COST EXPRESSIONS

In this section, we provide expressions for the holding cost under various policies when the length of the high period is H and the length of the low period is L. These expressions are used extensively in Section 2.6. We only consider the cases given in Section 2.4, i.e., we assume $\rho_1^h > 1$ and $\rho_1^l + \rho_2 < 1$.

We know from Corollary 2 that for Case 1 and Case 3, specified in (20) and (22) respectively, FP1 policy is optimal. Hence, we focus only on Cases 2 and 4 given in (21) and (23) respectively. In order to see the performance of the policies considered in Section 2.6, we first provide the holding cost expressions under the optimal policy. These expressions serve as the lower bound for all the other policies. Then we also provide the holding cost expressions under FP1 policy for Cases 2 and 4. In addition, we also provide the holding cost expressions under FP2-FP1 policy for Case 2, and the holding cost expressions under π^{a_1} policy for Case 4. These expressions help evaluate the performance of these two policies when $\rho_2 + \rho_1^l \to 1$.

When H and L are known, the optimal policy is given in Section 3.1. In order to compute the holding cost expression under a given policy, we observe the evolution of the fluid levels of both classes under this policy. Given the fluid levels, holding cost incurred by class 1 and class 2 can be computed easily. For example, for Case 2, when H and L satisfy the conditions of Case 2.6 (in Section 2.4.1), i.e., $H > a_2$, $H + L > \psi_1(1-\eta)^{-1}$, the optimal policy is to set $s_1 = s_2 = 0$ which is equivalent to the FP1 policy. We know that fluid levels of both classes will increase, and at $t_1 = \psi_1$, class 1 fluid reaches its threshold from below and starts to incur cost. Fluid levels of both classes continue to increase linearly until the beginning of the low period. In the low period, fluid level of class 1 begins to decrease and class 2 fluid continues to increase until class 1 fluid decreases to its threshold, which happens at t_2. After t_2, class 1 fluid is kept at its threshold and class 2 fluid begins to decrease and reaches its threshold at $\tilde{t}_2$. We know that after $\tilde{t}_2$, both classes will be kept below their thresholds. Note that when $H > a_2$, under the optimal policy, $L \geq \tilde{t}_2 - H$ is equivalent to $L \geq \gamma_4(H - a_1)$ (which is OPT:1 below) and $H \geq a_2$, $L \geq \gamma_4(H - a_1)$ imply that the conditions of Case 2.6, i.e., $H \geq a_2$, $H + L > \psi_1(1-\eta)^{-1}$ are satisfied. So, we can compute the holding cost when $H \geq a_2$ and $L \geq \gamma_4(H - a_1)$. We obtain the holding cost expressions for the other cases in a similar way.

Next, we provide the lower bound of the holding cost, i.e th holding cost under the optimal policy for each sample path (H, L) in Section A.1. The cost expressions under the FP1 policy is provided in Section A.2, that of π^{a_1} policy is provided in Section A.3, and that of FP2-FP1 policy is provided in Section A.4.

A.1 Cost under the optimal policy

While computing the holding cost under the optimal policy, we combine Cases 2 and 4 whenever $\psi_2^- = 0$ (in Case 4), where $a^- = \max\{-a, 0\}$. However, we have to divide each case into several subcases in order to obtain closed form expressions for the holding cost. As a result, we have 17 subcases labeled (OPT:1) to (OPT:17). Recall that t_1 is the time that class 1 increases to its threshold from below in the high period, and t_2 is the time that class 1 decreases to its threshold from above in the low period if the low period is long enough, and $\tilde{t}_2$ is the time that class 2 decreases to its threshold from above if the low period is long enough. Also, recall that $\tilde{\psi}_1 \geq \tilde{\psi}_2$ is equivalent to $B \geq a_1 \geq \tilde{\psi}_1 \geq \tilde{\psi}_2$.

1. Assume that the conditions of Case 2.6 (or Case 4.4) are satisfied and $L \geq \tilde{t}_2$. In Case 2.6 (and Case 4.4), the optimal policy sets $s_1 = s_2 = 0$, i.e. implements the FP1 policy. If $L \geq \tilde{t}_2$ is also satisfied, then the low period is long enough so that the fluid levels of both classes reach their thresholds. This is equivalent to

$$\text{(OPT:1)} \qquad H \geq a_2,\ L \geq \gamma_4(H - a_1),$$

where γ_4 is given in the proof of Proposition 5 in Section 2.6, and the holding cost is

$$c(H, L) \;=\; \frac{1}{2} h_2 \mu_2 \left\{ \frac{(\rho_1^h - 1)(\rho_1^h - \rho_1^l)}{\eta(1 - \rho_1^l)} (H - \psi_1)^2 - \frac{(\rho_1^h - 1)}{\eta} (\psi_1^-)^2 - \rho_2 (\psi_2^-)^2 \right.$$

$$+\frac{(1-\rho_2-\rho_1^l)^2}{\rho_2}\Big[\frac{\rho_1^h+\rho_2-1}{1-\rho_1^l-\rho_2}(H-a_1)-\frac{(\rho_1^h-1)}{(1-\rho_1^l)}(H-\psi_1)\Big]^2$$
$$+(1-\rho_1^l-\rho_2)\Big[\frac{\rho_1^h+\rho_2-1}{1-\rho_1^l-\rho_2}(H-a_1)-\frac{(\rho_1^h-1)}{(1-\rho_1^l)}(H-\psi_1)\Big]^2\Big\}.$$

2. Assume that the conditions of Case 2.6 (or Case 4.4) are satisfied and $t_2 \le L \le \tilde{t}_2$. As mentioned above, in Case 2.6 (and Case 4.4), the optimal policy implements the FP1 policy. If $t_2 \le L \le \tilde{t}_2$, then the low period is long enough such that class 1 fluid level reaches its threshold, but class 2 fluid is still above its threshold when the low period is over. This is equivalent to

$$\text{(OPT:2)} \qquad H \ge a_2,\ \gamma_3(H-\psi_1) \le L \le \gamma_4(H-a_1),$$

and the holding cost is

$$\begin{aligned}c(H,L) \;=\; & \frac{1}{2}h_2\mu_2\Big\{\frac{(\rho_1^h-1)(\rho_1^h-\rho_1^l)}{\eta(1-\rho_1^l)}(H-\psi_1)^2-\frac{(\rho_1^h-1)}{\eta}(\psi_1^-)^2-\rho_2(\psi_2^-)^2\\ & +\frac{(1-\rho_2-\rho_1^l)^2}{\rho_2}\Big[\frac{\rho_1^h+\rho_2-1}{1-\rho_1^l-\rho_2}(H-a_1)-\frac{(\rho_1^h-1)}{(1-\rho_1^l)}(H-\psi_1)\Big]^2\\ & +(1-\rho_1^l-\rho_2)\Big[\frac{\rho_1^h+\rho_2-1}{1-\rho_1^l-\rho_2}(H-a_1)-\frac{(\rho_1^h-1)}{(1-\rho_1^l)}(H-\psi_1)\Big]^2\Big\}\\ & -\frac{1}{2}h_2\mu_2(1-\rho_2-\rho_1^l)\Big[\frac{(\rho_1^h+\rho_2-1)}{(1-\rho_2-\rho_1^l)}(H-a_1)-L\Big]^2.\end{aligned}$$

3. Assume that the conditions of Case 2.6 (or Case 4.4) are satisfied and $L \le t_2$. The optimal policy sets $s_1 = s_2 = 0$, i.e. implements the FP1 policy. Since $L \le t_2$, at the end of the low period, both classes will be above their thresholds. This is equivalent to

$$\text{(OPT:3)} \qquad H \ge a_2,\ L \le \gamma_3(H-\psi_1),\ H+L \ge \psi_1^+ + \frac{\eta}{1-\eta}(\psi_1^+-\psi_2^+),$$

and the optimal cost is

$$\begin{aligned}c(H,L) \;=\; & \frac{1}{2}h_2\mu_2\Big\{\frac{(\rho_1^h-1)(\rho_1^h-\rho_1^l)}{\eta(1-\rho_1^l)}(H-\psi_1)^2+\rho_2(H+L-\psi_2)^2\\ & -\frac{1-\rho_1^l}{\eta}\Big[\frac{\rho_1^h-1}{1-\rho_1^l}(H-\psi_1)-L\Big]^2-\frac{(\rho_1^h-1)}{\eta}(\psi_1^-)^2-\rho_2(\psi_2^-)^2\Big\}.\end{aligned}$$

4. Assume that the conditions of Case 2.3 are satisfied and $L \ge \tilde{t}_2 - H$. In Case 2.3 optimal policy sets $s_1 = s_2$. Let $s_1 = s_2 = s$. Note that $L \ge \tilde{t}_2 - H$ is equivalent to

$L > \gamma_4(H - a_1)$, which means that the low period is long enough so that fluid levels of both classes reach their thresholds. Thus, if

$$\text{(OPT:4)} \qquad \max(\tilde{\psi}_1, B) \leq H \leq a_2,\ L \geq \gamma_4(H - a_1),$$

then the holding cost is

$$\begin{aligned} c(H, L) &= \frac{1}{2} h_2 \mu_2 \Bigg\{ \frac{(1-\rho_1^l)(\rho_1^h - \rho_1^l)}{\eta(\rho_1^h - 1)} (t_2 - H)^2 + \rho_2 (t_2 - s)^2 \\ &\quad + (1-\rho_2)(2\tilde{\psi}_2 - s)s + 2(1-\rho_2)(\tilde{\psi}_2 - s)(t_2 - s) \\ &\quad + (1 - \rho_2 - \rho_1^l) \Big[\frac{\rho_1^h + \rho_2 - 1}{1 - \rho_2 - \rho_1^l} (H - a_1) - (t_2 - H) \Big]^2 \Bigg\}, \end{aligned}$$

where

$$\begin{aligned} s &= \frac{d_1/\mu_1 - (\rho_1^h - 1)(1-\eta) t_2}{1 + \eta(\rho_1^h - 1)}, \\ t_1 &= \frac{(1-\eta)t_2 + \eta d_1/\mu_1}{1 + \eta(\rho_1^h - 1)}, \\ t_2 &= \frac{(\rho_1^h - \rho_1^l) H - \eta(\rho_1^h - 1) d_1 (\mu_1 (1 + \eta(\rho_1^h - 1)))^{-1}}{(1 - \rho_1^l) + (1-\eta)(\rho_1^h - 1)(1 + \eta(\rho_1^h - 1))^{-1}}, \end{aligned}$$

5. Assume that the conditions of Case 2.3 are satisfied and $L \leq \tilde{t}_2 - H$. When H and L satisfy the conditions of Case 2.3, it implies that $L \geq t_2$, i.e., the low period is long enough so that class 1 reaches its threshold. However, since $L \leq \tilde{t}_2 - H$, the low period is not long enough for class 2 to reach its threshold. At the end of the low period, class 2 fluid is still above its threshold, but class 1 is below its threshold. Hence, if

$$\text{(OPT:5)} \qquad \max(\tilde{\psi}_1, B) \leq H \leq a_2,\ \gamma_2(H - \tilde{\psi}_1) \leq L \leq \gamma_4(H - a_1)$$

then the holding cost is

$$\begin{aligned} c(H, L) &= \frac{1}{2} h_2 \mu_2 \Bigg\{ \frac{(1-\rho_1^l)(\rho_1^h - \rho_1^l)}{\eta(\rho_1^h - 1)} (t_2 - H)^2 + \rho_2 (t_2 - s)^2 \\ &\quad + (1-\rho_2)(2\tilde{\psi}_2 - s)s + 2(1-\rho_2)(\tilde{\psi}_2 - s)(t_2 - s) \\ &\quad + (1 - \rho_2 - \rho_1^l) \Big[\frac{\rho_1^h + \rho_2 - 1}{1 - \rho_2 - \rho_1^l} (H - a_1) - (t_2 - H) \Big]^2 \Bigg\} \\ &\quad - \frac{1}{2} h_2 \mu_2 (1 - \rho_2 - \rho_1^l) \Big[\frac{(\rho_1^h + \rho_2 - 1)}{(1 - \rho_2 - \rho_1^l)} (H - a_1) - L \Big]^2, \end{aligned}$$

where t_2, t_1, s are the same as in the previous case.

6. Assume that the conditions of Case 2.4 are satisfied. In Case 2.4, the optimal policy sets $s_1 = s_2 = s$ and $t_2 = H + L$. In this case, at the end of the low period the fluid levels of both classes are above their thresholds. Thus, if

$$\text{(OPT:6)} \qquad L \le \gamma_2(H - \tilde{\psi}_1),$$
$$\max\{\tilde{\psi}_1,\, \tilde{\psi}_1 + \frac{1 + \eta(\rho_1^h - 1)}{(1-\eta)(\rho_1^h - 1)}(\tilde{\psi}_1 - \tilde{\psi}_2)\} \le H + L \le \frac{\psi_1}{1-\eta},$$

then the holding cost is

$$\begin{aligned} c(H, L) &= \frac{1}{2} h_2 \mu_2 \Bigg\{ \frac{(\rho_1^h - 1)(\rho_1^h - \rho_1^l)}{\eta(1 - \rho_1^l)}(H - t_1)^2 + (1 - \rho_2)(2\tilde{\psi}_2 - s)s \\ &\quad + 2(1 - \rho_2)(\tilde{\psi}_2 - s)(H + L - s) + \rho_2(H + L - s)^2 \\ &\quad - \frac{(1 - \rho_1^l)}{\eta}\Big[\frac{(\rho_1^h - 1)}{(1 - \rho_1^l)}(H - t_1) - L\Big]^2 \Bigg\}, \end{aligned}$$

where

$$\begin{aligned} s &= \frac{d_1/\mu_1 - (1-\eta)(\rho_1^h - 1)(H + L)}{1 + \eta(\rho_1^h - 1)}, \\ t_1 &= \frac{\eta d_1/\mu_1 + (1 - \eta)(H + L)}{1 + \eta(\rho_1^h - 1)}. \end{aligned}$$

7. Assume that the conditions of Case 2.5 are satisfied, $\tilde{\psi}_1 \le \tilde{\psi}_2$, class 2 decreases to its threshold before class 1 increases to its threshold, and the low period is long enough to decrease class 2 fluid to its threshold. When conditions of Case 2.5 are satisfied, and $\tilde{\psi}_1 \le \tilde{\psi}_2$, we have $H \le \tilde{\psi}_1 \le \tilde{\psi}_2$ and the optimal policy sets $s_1 = s_2 = H$. In the high period, class 2 has higher priority and in the low period, class 2 has higher priority until class 1 fluid increases to its threshold or class 2 fluid decreases to its threshold. Under this policy, let t_1' be the time that class 1 fluid increases to its threshold in the low period if the low period is long enough. Then, $t_1' = H + \rho_1^h/\rho_1^l(\tilde{\psi}_1 - H)$. If $\tilde{\psi}_2 \le t_1'$, then class 2 fluid decreases to its threshold in the low period, at $\tilde{\psi}_2$, before class 1 fluid increases to its threshold. In this case, after $\tilde{\psi}_2$, no class will incur cost under the Low-period-policy. Hence, if

$$\text{(OPT:7)} \qquad H \le \tilde{\psi}_1 \le \tilde{\psi}_2 \le \frac{\rho_1^h}{\rho_1^l}(\tilde{\psi}_1 - H) + H, \quad H + L > \tilde{\psi}_2,$$

then the holding cost is

$$c(H, L) = \frac{1}{2} h_2 \mu_2 (1 - \rho_2) \tilde{\psi}_2^2.$$

8. Assume that the conditions of Case 2.5 are satisfied, $\tilde{\psi}_1 \le \tilde{\psi}_2$, class 2 fluid decreases to its threshold before class 1 fluid increases to its threshold, but the low period is not long enough for class 2 fluid to reach its threshold. Hence, $H + L \le \tilde{\psi}_2$. If

$$\text{(OPT:8)} \qquad H \le \tilde{\psi}_1 \le \tilde{\psi}_2 \le \frac{\rho_1^h}{\rho_1^l}(\tilde{\psi}_1 - H) + H, \quad H + L \le > \tilde{\psi}_2,$$

then the holding cost is

$$c(H, L) = \frac{1}{2} h_2 \mu_2 (1 - \rho_2)(2\tilde{\psi}_2 - H - L)(H + L).$$

9. Assume that the conditions of Case 2.5 are satisfied, $\tilde{\psi}_1 \le \tilde{\psi}_2$, class 1 fluid level increases to its threshold before class 2 fluid level decreases to its threshold, and the low period is long enough for class 2 fluid to reach its threshold. Since the conditions of Case 2.5 are satisfied and $\tilde{\psi}_1 \le \tilde{\psi}_2$, we have $H \le \tilde{\psi}_1 \le \tilde{\psi}_2$. Following the optimal policy, we set $s_1 = s_2 = H$. Class 2 has higher priority in the high period and also in the low period before class 1 fluid reaches its threshold at t_1'. So, if $\tilde{\psi}_2 \ge t_1'$, it means that class 2 is still above its threshold when class 1 increases to its threshold in the low period. Based on the Low-period-policy after t_1', server will spend just enough effort ($u_1 = \rho_1^l$) to keep class 1 at its threshold, and use the remaining effort ($u_2 = 1 - \rho_1^l > \rho_2$) to serve class 2. Let $\tilde{t}_2$ be the time that class 2 fluid level decreases to its threshold, then $L + H \ge \tilde{t}_2$, which is equivalent to $L \ge \gamma_4(H - a_1)$. So, if

$$\text{(OPT:9)} \qquad H \le \tilde{\psi}_1 \le \frac{\rho_1^h}{\rho_1^l}(\tilde{\psi}_1 - H) + H \le \tilde{\psi}_2, \quad L \ge \gamma_4(H - a_1),$$

then the holding cost is

$$\begin{aligned} c(H, L) = {} & \frac{1}{2} h_2 \mu_2 \Big\{ (1 - \rho_2) \Big[2\tilde{\psi}_2 - \frac{\rho_1^h}{\rho_1^l}(\tilde{\psi}_1 - H) - H \Big] \Big[\frac{\rho_1^h}{\rho_1^l}(\tilde{\psi}_1 - H) + H \Big] \\ & + (1 - \rho_2 - \rho_1^l) \Big[\frac{\rho_1^h + \rho_2 - 1}{1 - \rho_2 - \rho_1^l}(H - a_1) - \frac{\rho_1^h}{\rho_1^l}(\tilde{\psi}_1 - H) \Big]^2 \Big\}. \end{aligned}$$

10. Assume that the conditions of Case 2.5 are satisfied, $\tilde{\psi}_1 \leq \tilde{\psi}_2$, class 1 fluid level increases to its threshold before class 2 fluid level decreases to its threshold, and the low period is not long enough for class 2 fluid to reach its threshold, but long enough for class 1 fluid to reach its threshold. Hence, $H + L \geq t_1'$, which is equivalent to $L \geq \rho_1^h/\rho_1^l(\tilde{\psi}_1 - H)$. If

$$\text{(OPT:10)} \qquad H \leq \tilde{\psi}_1 \leq \frac{\rho_1^h}{\rho_1^l}(\tilde{\psi}_1 - H) + H \leq \tilde{\psi}_2, \quad \frac{\rho_1^h}{\rho_1^l}(\tilde{\psi}_1 - H) \leq L \leq \gamma_4(H - a_1),$$

then the holding cost is

$$\begin{aligned} c(H,L) &= \frac{1}{2}h_2\mu_2\Big\{(1-\rho_2)\Big[2\tilde{\psi}_2 - \frac{\rho_1^h}{\rho_1^l}(\tilde{\psi}_1 - H) - H\Big]\Big[\frac{\rho_1^h}{\rho_1^l}(\tilde{\psi}_1 - H) + H\Big] \\ &\quad +(1-\rho_2-\rho_1^l)\Big[\frac{\rho_1^h+\rho_2-1}{1-\rho_2-\rho_1^l}(H-a_1) - \frac{\rho_1^h}{\rho_1^l}(\tilde{\psi}_1 - H)\Big]^2\Big\} \\ &\quad -\frac{1}{2}h_2\mu_2(1-\rho_2-\rho_1^l)\Big[\frac{\rho_1^h+\rho_2-1}{1-\rho_2-\rho_1^l}(H-a_1) - L\Big]^2. \end{aligned}$$

11. Assume that the conditions of Case 2.5 are satisfied, $\tilde{\psi}_1 \leq \tilde{\psi}_2$, the low period is neither long enough for class 1 fluid to increase to its threshold nor long enough for class 2 fluid to decrease to its threshold. However, if the low period were long enough class 1 fluid would increase to its threshold before class 2 would decrease to its threshold. Hence, $H + L \leq t_1'$. If

$$\text{(OPT:11)} \qquad H \leq \tilde{\psi}_1 \leq \frac{\rho_1^h}{\rho_1^l}(\tilde{\psi}_1 - H) + H \leq \tilde{\psi}_2, \quad L \leq \frac{\rho_1^h}{\rho_1^l}(\tilde{\psi}_1 - H),$$

then the holding cost is

$$c(H,L) = \frac{1}{2}h_2\mu_2(1-\rho_2)(2\tilde{\psi}_2 - H - L)(H+L).$$

12. Assume that the conditions of Case 2.1 (or Case 4.1) are satisfied and the low period is long enough for class 2 fluid to decrease to its threshold. Recall that when conditions of Case 2.1 (or Case 4.1) are satisfied, we denote the time that class 2 decreases to its threshold as $\tilde{\psi}_2$. Moreover, $H + L \geq \tilde{\psi}_2$ is equivalent to $L > \gamma_4(H - a_1)$ which implies that $L \geq \gamma_1(H - a_1)$. In this case, at the end of the low period, fluid levels of both classes will be below their thresholds. Notice that $a_1 \leq B$ implies that $\tilde{\psi}_2 < \tilde{\psi}_1$.

Hence, if

$$\text{(OPT:12)} \qquad a_1 \le H \le B, \quad L \ge \gamma_4(H - a_1), \text{ for Case 2}$$

$$a_1 \le H \le a_2, \quad L \ge \gamma_4(H - a_1), \text{ for Case 4}$$

then the holding cost is

$$\begin{aligned} c(H, L) &= \frac{1}{2}h_2\mu_2 \left\{ \frac{(1-\rho_1^l)(\rho_1^h - \rho_1^l)}{\eta(\rho_1^h - 1)}(t_2 - H)^2 + (1-\rho_2)s_1^2 + \rho_2(t_2 - s_2)^2 \right. \\ & \left. \quad + (1 - \rho_2 - \rho_1^l)\left[\frac{\rho_1^h + \rho_2 - 1}{1 - \rho_2 - \rho_1^l}(H - a_1) - (t_2 - H)\right]^2 \right\}, \end{aligned}$$

where

$$\begin{aligned} s_1 &= \tilde{\psi}_2^+, \\ s_2 &= \frac{(d_1/\mu_1 + d_2/\mu_2) - (\rho_1^h - 1)(1-\eta)t_2}{\rho_2 + \eta(\rho_1^h - 1)}, \\ t_1 &= \frac{\eta(d_1/\mu_1 + d_2/\mu_2) + \rho_2(1-\eta)t_2}{\rho_2 + \eta(\rho_1^h - 1)}, \\ t_2 &= \frac{(\rho_1^h - \rho_1^l)H - (\rho_1^h - 1)\eta(d_1/\mu_1 + d_2/\mu_2)(\rho_2 + \eta(\rho_1^h - 1))^{-1}}{(1 - \rho_1^l) + (\rho_1^h - 1)(1-\eta)\rho_2(\rho_2 + \eta(\rho_1^h - 1))^{-1}}. \end{aligned}$$

13. Assume that the conditions of Case 2.1 (or Case 4.1) are satisfied but the low period is not long enough for class 2 fluid to decrease to its threshold. Hence, $L \le \gamma_4(H - a_1)$. At the end of the low period, class 2 is still above its threshold but class 1 is below its threshold. Hence, if

$$\text{(OPT:13)} \qquad a_1 \le H \le B, \quad \gamma_1(H - a_1) \le L \le \gamma_4(H - a_1), \text{ for Case 2}$$

$$a_1 \le H \le a_2, \quad \gamma_1(H - a_1) \le L \le \gamma_4(H - a_1), \text{ for Case 4}$$

then

$$\begin{aligned} c(H, L) &= \frac{1}{2}h_2\mu_2 \left\{ \frac{(1-\rho_1^l)(\rho_1^h - \rho_1^l)}{\eta(\rho_1^h - 1)}(t_2 - H)^2 + (1-\rho_2)s_1^2 + \rho_2(t_2 - s_2)^2 \right. \\ & \left. \quad + (1 - \rho_2 - \rho_1^l)\left[\frac{\rho_1^h + \rho_2 - 1}{1 - \rho_2 - \rho_1^l}(H - a_1) - (t_2 - H)\right]^2 \right\} \\ &- \frac{1}{2}h_2\mu_2(1 - \rho_2 - \rho_1^l)\left[\frac{\rho_1^h + \rho_2 - 1}{1 - \rho_2 - \rho_1^l}(H - a_1) - L\right]^2 \end{aligned}$$

where s_1, s_2, t_1, t_2 are the same as given in Case 2.1 (Case 4.1).

14. Assume that conditions of Case 2.2 (or Case 4.2) are satisfied. Note that

$$\tilde{\psi}_1 + (1 + \eta(\rho_1^h - 1))((1 - \eta)(\tilde{\psi}_1 - \tilde{\psi}_2^+)(\rho_1^h - 1))^{-1} \geq a_1$$

implies that $\tilde{\psi}_1 \geq \tilde{\psi}_2$. Since the low period is not long enough for class 1 fluid to decrease to its threshold, $t_2 = H + L$. With some algebra we have

$$\text{(OPT:14)} \qquad \begin{aligned} & a_1 \leq H, \quad L \leq \gamma_1(H - a_1), \\ & H + L \leq \tilde{\psi}_1 + \frac{1 + \eta(\rho_1^h - 1)}{(1 - \eta)(\rho_1^h - 1)}(\tilde{\psi}_1 - \tilde{\psi}_2^+) - \frac{\eta}{1 - \eta}\psi^+, \end{aligned}$$

and the holding cost is

$$\begin{aligned} c(H, L) \quad = \quad & \left\{ \frac{(\rho_1^h - 1)(\rho_1^h - \rho_1^l)}{\eta(1 - \rho_1^l)}(H - t_1)^2 + (1 - \rho_2)s_1^2 \right. \\ & \left. + \rho_2(H + L - s_2)^2 - \frac{(1 - \rho_1^l)}{\eta}\left[\frac{\rho_1^h - 1}{1 - \rho_1^l}(H - t_1) - L\right]^2 \right\} \end{aligned}$$

where

$$\begin{aligned} s_1 &= \tilde{\psi}_2^+, \\ s_2 &= \frac{(d_1/\mu_1 + d_2/\mu_2) - (\rho_1^h - 1)(1 - \eta)t_2}{\rho_2 + \eta(\rho_1^h - 1)}, \\ t_1 &= \frac{\eta(d_1/\mu_1 + d_2/\mu_2) + \rho_2(1 - \eta)t_2}{\rho_2 + \eta(\rho_1^h - 1)}, \\ t_2 &= H + L. \end{aligned}$$

15. Assume that conditions of Case 2.5 are satisfied, $\tilde{\psi}_2 \leq \tilde{\psi}_1$ and class 2 reaches its threshold from below in the high period or conditions of Case 4.3 are satisfied. Recall that $\tilde{\psi}_2 \leq \tilde{\psi}_1$ implies that $\tilde{\psi}_2 \leq \tilde{\psi}_1 \leq a_1$. Since conditions of Case 2.5 and $\tilde{\psi}_2 \leq \tilde{\psi}_1$ are satisfied, we have $H \leq a_1$. According to the optimal policy, class 2 has higher priority in the high period as long as its fluid level is above its threshold, and class 2 fluid reaches its threshold at $\tilde{\psi}_2$. After $\tilde{\psi}_2$, server will allocate enough capacity to keep class 2 fluid level below its threshold and the remaining capacity will be allocated to class 1. In this case, class 1 fluid will never reach its threshold in the high period.

Similarly, for Case 4.3, under the optimal policy class 1 and class 2 fluids will stay below their thresholds in the high period. Hence, if

$$\text{(OPT:15)} \qquad \tilde{\psi}_2 \leq \tilde{\psi}_1, \quad \tilde{\psi}_2 \leq H \leq a_1,$$

then the holding cost is

$$c(H, L) \;=\; \frac{1}{2} h_2 \mu_2 (1 - \rho_2) s_1^2.$$

16. Assume that conditions of Case 2.5 are satisfied, $\tilde{\psi}_2 \leq \tilde{\psi}_1$, and class 2 fluid does not reach its threshold in the high period, but it reaches its threshold in the low period. If

$$\text{(OPT:16)} \qquad \tilde{\psi}_2 \leq \tilde{\psi}_1, \quad H \leq \tilde{\psi}_2 \leq H + L,$$

then the holding cost is

$$c(H, L) \;=\; \frac{1}{2} h_2 \mu_2 (1 - \rho_2) \tilde{\psi}_2^2.$$

17. Assume that conditions of Case 2.5 are satisfied $\tilde{\psi}_2 \leq \tilde{\psi}_1$, and class 2 fluid does not reach its threshold in the low period. Hence, if

$$\text{(OPT:17)} \qquad \tilde{\psi}_2 \leq \tilde{\psi}_1, \quad H + L \leq \tilde{\psi}_2, \tag{261}$$

then the holding cost is

$$c(H, L) \;=\; \frac{1}{2} h_2 \mu_2 (2\tilde{\psi}_2 - H - L)(H + L).$$

A.2 Cost under the FP1 policy for Case 2 and Case 4

Note that when $\psi_2 < \psi_1$, under the FP1 policy, class 2 fluid increases to its threshold before class 1 fluid increases to its threshold in the high period if the high period is long enough, i.e. if $H \geq \psi_2$. In order to compute the holding cost under the FP1 policy, we consider 9 different cases labeled (FP1:1) to (FP1:9).

1. Assume that class 1 fluid increases to its threshold in the high period and decreases to its threshold in the low period and class 2 fluid also decreases to its threshold in the low period. Hence, if

$$\text{(FP1:1)} \qquad H \geq \psi_1, \quad L \geq \gamma_4(H - a_1)$$

then the holding cost under FP1 policy is

$$\begin{aligned} c^{FP1}(H,L) &= \frac{1}{2}h_2\mu_2 \Big\{ \frac{(\rho_1^h - 1)(\rho_1^h - \rho_1^l)}{\eta(1-\rho_1^l)}(H-\psi_1)^2 - \frac{(\rho_1^h - 1)}{\eta}(\psi_1^-)^2 - \rho_2(\psi_2^-)^2 \\ &+ \frac{(1-\rho_2-\rho_1^l)^2}{\rho_2}\Big[\frac{\rho_1^h+\rho_2-1}{1-\rho_1^l-\rho_2}(H-a_1) - \frac{(\rho_1^h-1)}{(1-\rho_1^l)}(H-\psi_1)\Big]^2 \\ &+ (1-\rho_1^l-\rho_2)\Big[\frac{\rho_1^h+\rho_2-1}{1-\rho_1^l-\rho_2}(H-a_1) - \frac{(\rho_1^h-1)}{(1-\rho_1^l)}(H-\psi_1)\Big]^2\Big\}. \end{aligned}$$

2. Assume that class 1 fluid increases to its threshold in the high period and decreases to its threshold in the low period but class 2 fluid does not decrease to its threshold at the end of the low period. Hence, if

$$\text{(FP1:2)} \qquad H \geq \psi_1, \qquad \gamma_3(H-\psi_1) \leq L \leq \gamma_4(H-a_1)$$

the holding cost under FP1 policy is

$$\begin{aligned} c^{FP1}(H,L) &= \frac{1}{2}h_2\mu_2 \Big\{ \frac{(\rho_1^h - 1)(\rho_1^h - \rho_1^l)}{\eta(1-\rho_1^l)}(H-\psi_1)^2 - \frac{(\rho_1^h - 1)}{\eta}(\psi_1^-)^2 - \rho_2(\psi_2^-)^2 \\ &+ \frac{(1-\rho_2-\rho_1^l)^2}{\rho_2}\Big[\frac{\rho_1^h+\rho_2-1}{1-\rho_1^l-\rho_2}(H-a_1) - \frac{(\rho_1^h-1)}{(1-\rho_1^l)}(H-\psi_1)\Big]^2 \\ &+ (1-\rho_1^l-\rho_2)\Big[\frac{\rho_1^h+\rho_2-1}{1-\rho_1^l-\rho_2}(H-a_1) - \frac{(\rho_1^h-1)}{(1-\rho_1^l)}(H-\psi_1)\Big]^2\Big\} \\ &- \frac{1}{2}h_2\mu_2(1-\rho_2-\rho_1^l)\Big[\frac{(\rho_1^h+\rho_2-1)}{(1-\rho_2-\rho_1^l)}(H-a_1) - L\Big]^2. \end{aligned}$$

3. Assume that class 1 fluid increases to its threshold in the high period but does not decrease to its threshold in the low period. Hence, if

$$\text{(FP1:3)} \qquad H \geq \psi_1, \qquad L \leq \gamma_3(H-\psi_1), \tag{262}$$

then the holding cost under FP1 policy is

$$\begin{aligned} c^{FP1}(H,L) &= \frac{1}{2}h_2\mu_2 \Big\{ \frac{(\rho_1^h - 1)(\rho_1^h - \rho_1^l)}{\eta(1-\rho_1^l)}(H-\psi_1)^2 + \rho_2(H+L-\psi_2)^2 \\ &- \frac{(\rho_1^h - 1)}{\eta}(\psi_1^-)^2 - \rho_2(\psi_2^-)^2 - \frac{1-\rho_1^l}{\eta}\Big[\frac{\rho_1^h-1}{1-\rho_1^l}(H-\psi_1) - L\Big]^2\Big\}. \end{aligned}$$

4. Assume that class 1 fluid does not increase to its threshold in the high period, but class 2 fluid increases to its threshold in the high period. In the low period, before class 2 fluid decreases to its threshold, class 1 fluid increases to its threshold, and the server allocates enough capacity to maintain class 1 fluid at its threshold until class 2 fluid decreases to its threshold. The low period is long enough for class 2 fluid to decrease to its threshold. At the end of the low period, fluid levels of both classes are below their thresholds. Hence, if

$$\text{(FP1:4)} \qquad \hat{\psi} \leq H \leq \psi_1, \qquad L \leq \gamma_4(H - a_1)$$

where

$$\hat{\psi} = \frac{(\rho_1^h - 1)(1-\rho_2)}{\rho_2\rho_1^l + (\rho_1^h - 1)(1-\rho_2)}\psi_1 + \frac{\rho_2\rho_1^l}{\rho_2\rho_1^l + (\rho_1^h - 1)(1-\rho_2)}\psi_2,$$

then the holding cost under FP1 policy is

$$\begin{aligned} c^{FP1}(H,L) &= \frac{1}{2}h_2\mu_2\left\{-\rho_2(\psi_2^-)^2 + \frac{\rho_2(1-\rho_1^l)}{1-\rho_1^l-\rho_2}(H-\psi_2)^2\right. \\ &\quad + \frac{2\rho_2(\rho_1^h-1)}{1-\rho_2-\rho_1^l}(H-\psi_1)(H-\psi_2) \\ &\quad \left. + \frac{(1-\rho_2)(\rho_1^h-1)^2}{\rho_1^l(1-\rho_2-\rho_1^l)}(H-\psi_1)^2\right\}. \end{aligned}$$

5. Assume that class 1 fluid does not increase to its threshold in the high period, but class 2 fluid increases to its threshold in the high period. In the low period, before class 2 fluid decreases to its threshold, class 1 fluid increases to its threshold, and the server allocates just enough capacity to maintain class 1 fluid at its threshold until class 2 fluid decreases to its threshold. The low period is not long enough for class 2 fluid to decrease to its threshold. At the end of the low period, class 1 fluid is below its threshold and class 2 fluid is still above its threshold. Hence, if

$$\text{(FP1:5)} \qquad \hat{\psi} \leq H \leq \psi_1, \qquad -\gamma_5(\psi_1 - H) \leq L \leq \gamma_4(H - a_1)$$

where $\gamma_5 = -(\rho_1^h - 1)(\rho_1^l)^{-1}$ then the holding cost under FP1 policy is

$$c^{FP1}(H,L) = \frac{1}{2}h_2\mu_2\left\{-\rho_2(\psi_2^-)^2 + \frac{\rho_2(1-\rho_1^l)}{1-\rho_1^l-\rho_2}(H-\psi_2)^2\right.$$

$$
\begin{aligned}
&+\frac{2\rho_2(\rho_1^h-1)}{1-\rho_2-\rho_1^l}(H-\psi_1)(H-\psi_2)\\
&+\frac{(1-\rho_2)(\rho_1^h-1)^2}{\rho_1^l(1-\rho_2-\rho_1^l)}(H-\psi_1)^2\Big\}\\
&-\frac{1}{2}h_2\mu_2(1-\rho_2-\rho_1^l)\Big[\frac{\rho_1^h+\rho_2-1}{1-\rho_2-\rho_1^l}(H-a_1)-L\Big]^2.
\end{aligned}
$$

6. Assume that class 1 does not increase to its threshold in the high period, but class 2 increases to its threshold in the high period. If the low period were long enough, class 1 fluid would increase to its threshold before class 2 fluid decreases to its threshold. However, the low period is not long enough and class 1 fluid is still below its threshold and class 2 is above its threshold at the end of the low period. Hence, if

$$
\text{(FP1:6)} \qquad \hat{\psi} \le H \le \psi_1, \qquad L \le -\gamma_5(\psi_1 - H)
$$

then the holding cost under FP1 policy is

$$
\begin{aligned}
c^{FP1}(H,L) &= \frac{1}{2}h_2\mu_2\Big\{\frac{\rho_2}{1-\rho_2}(H-\psi_2)^2-\rho_2(\psi_2^-)^2\Big\}\\
&\quad-\frac{1}{2}h_2\mu_2(1-\rho_2)\Big[\frac{\rho_2}{1-\rho_2}(H-\psi_2)-L\Big]^2.
\end{aligned}
$$

7. Assume that class 1 fluid does not increase to its threshold in the high period, but class 2 fluid increases to its threshold in the high period. In the low period, class 2 fluid decreases to its threshold before class 1 fluid increases to its threshold. The low period is long enough such that at the end of the low period, both class 1 and class 2 fluids are below their thresholds. Hence, if

$$
\text{(FP1:7)} \qquad \psi_2 \le H \le \hat{\psi}, \qquad L \ge \gamma_6(H - \psi_2),
$$

where $\gamma_6 = \rho_2(1-\rho_2)^{-1}$ then the holding cost under FP1 policy is

$$
c^{FP1}(H,L) = \frac{1}{2}h_2\mu_2\Big\{\frac{\rho_2}{1-\rho_2}(H-\psi_2)^2-\rho_2(\psi_2^-)^2\Big\}.
$$

8. Assume that class 1 fluid does not increase to its threshold in the high period, but class 2 fluid increases to its threshold in the high period. If the low period were long enough, class 2 fluid would decrease to its threshold before class 1 fluid increases to

its threshold. However, the low period is not long enough. So, at the end of the low period, class 2 fluid is still above its threshold and class 1 fluid is below its threshold. Hence, if

$$\text{(FP1:8)} \qquad \psi_2 \leq H \leq \hat{\psi}, \qquad L \leq \gamma_6(H - \psi_2),$$

then the holding cost under FP1 policy is

$$\begin{aligned} c^{FP1}(H,L) &= \frac{1}{2}h_2\mu_2\left\{\frac{\rho_2}{1-\rho_2}(H-\psi_2)^2 - \rho_2(\psi_2^-)^2\right\} \\ &\quad -\frac{1}{2}h_2\mu_2(1-\rho_2)\Big[\frac{\rho_2}{1-\rho_2}(H-\psi_2) - L\Big]^2. \end{aligned}$$

9. Assume that neither class 1 fluid nor class 2 fluid reaches its threshold in the high period. Hence, if

$$\text{(FP1:9)} \qquad H \leq \psi_2,$$

then the cost under FP1 policy is

$$c^{FP1}(H,L) \;=\; 0.$$

A.3 Cost under the pi-a1 policy for Case 4

Note that in this case if the high period is long enough (i.e. if $H \geq a_1$), under the π^{a_1} policy, class 1 and class 2 fluids reach their thresholds at the same time, namely, at a_1. In order to compute the holding cost under the π^{a_1} policy, we consider 4 different cases labeled (a_1:1) to (a_1:4).

1. Assume that fluid levels of both classes increase to their thresholds at the same time in the high period, and the low period is long enough to decrease fluid levels of both classes below their thresholds. Thus, at the end of the low period, both class 1 and class 2 fluids are below their thresholds. Hence, if

$$(a_1{:}1) \qquad H \geq a_1, \qquad L \geq \gamma_4(H - a_1),$$

then the holding cost under π^{a_1} policy is

$$c^{a_1}(H,L) \;=\; \frac{1}{2}h_2\mu_2\left(\frac{(\rho_1^h-1)(\rho_1^h-\rho_1^l)}{\eta(1-\rho_1^l)} + \frac{\rho_2(\rho_1^h-\rho_1^l)^2}{(1-\rho_1^l)(1-\rho_1^l-\rho_2)}\right)(H-a_1)^2.$$

2. Assume that both classes increase to their thresholds at the same time in the high period but the low period is not long enough for class 2 fluid to decrease to its threshold. At the end of the low period, class 1 fluid is below its threshold but class 2 fluid is still above its threshold. Hence, if

$$(a_1{:}2) \qquad H \geq a_1, \qquad \gamma_3(H-a_1) \leq L \leq \gamma_4(H-a_1),$$

then the holding cost under π^{a_1} policy is

$$\begin{aligned} c^{a_1}(H,L) &= \frac{1}{2}h_2\mu_2\left(\frac{(\rho_1^h-1)(\rho_1^h-\rho_1^l)}{\eta(1-\rho_1^l)} + \frac{\rho_2(\rho_1^h-\rho_1^l)^2}{(1-\rho_1^l)(1-\rho_1^l-\rho_2)}\right)(H-a_1)^2 \\ &\quad -\frac{1}{2}h_2\mu_2(1-\rho_2-\rho_1^l)\left[\frac{\rho_1^h+\rho_2-1}{1-\rho_2-\rho_1^l}(H-a_1)-L\right]^2. \end{aligned}$$

3. Assume that fluid levels of both classes increase to their thresholds at the same time in the high period but the low period is not long enough for either class 1 or class 2 fluid to decrease to its threshold. At the end of the low period, both class 1 and class 2 fluids are above their thresholds. Hence, if

$$(a_1{:}3) \qquad H \geq a_1, \qquad L \leq \gamma_3(H-a_1),$$

then the cost under π^{a_1} policy is

$$\begin{aligned} c^{a_1}(H,L) &= \frac{1}{2}h_2\mu_2\left\{\frac{(\rho_1^h-1)(\rho_1^h-\rho_1^l)}{\eta(1-\rho_1^l)}(H-a_1)^2 + \rho_2(H+L-a_1)^2\right. \\ &\quad \left. -\frac{(1-\rho_1^l)}{\eta}\left[\frac{\rho_1^h-1}{1-\rho_1^l}(H-a_1)-L\right]^2\right\}. \end{aligned}$$

4. Assume that fluid levels of both classes are still below their thresholds at the end of the high period. Hence, if

$$(a_1{:}4) \qquad H \leq a_1,$$

then the holding cost under π^{a_1} policy is

$$c^{a_1}(H,L) = 0.$$

A.4 Cost under the FP2-FP1 policy for Case 2

Case 2 has two subcases: $\tilde{\psi}_1 \leq \tilde{\psi}_2$ and $\tilde{\psi}_1 \geq \tilde{\psi}_2$. Recall that $\tilde{\psi}_1$ ($\tilde{\psi}_2$) is the time that class 1 fluid increases (class 2 fluid decreases) to its threshold from below (from above) in the high period if class 2 has higher priority and if the high period is long enough. So, if $\tilde{\psi}_1 \leq \tilde{\psi}_2 \leq H$, class 1 fluid increases to its threshold before class 2 fluid decreases to its threshold. However, if $\tilde{\psi}_2 \leq \tilde{\psi}_1 \leq H$, then class 1 fluid is still below its threshold when class 2 fluid reaches its threshold in the high period.

1. Assume that in the high period class 1 fluid increases to its threshold (at $\tilde{\psi}_1$) before class 2 fluid decreases to its threshold, after $\tilde{\psi}_1$, class 1 has higher priority in the high period. Suppose that the low period is long enough to reduce fluid levels of both classes below their thresholds. Thus, at the end of the low period, fluid levels of both classes are below their thresholds. Hence, if

$$\text{(FP2-FP1:1)} \qquad \tilde{\psi}_1 \leq \tilde{\psi}_2, \quad H \geq \tilde{\psi}_1, \quad L \geq \gamma_4(H - a_1),$$

then the holding cost under FP2-FP1 policy is

$$\begin{aligned} c^{\mathrm{FP2-FP1}}(H,L) \;=\;& \frac{1}{2}h_2\mu_2\Bigg\{\frac{(\rho_1^h - 1)(\rho_1^h - \rho_1^l)}{\eta(1-\rho_1^l)}(H-\tilde{\psi}_1)^2 + (1-\rho_2)(2\tilde{\psi}_2 - \tilde{\psi}_1)\tilde{\psi}_1 \\ & +2\frac{(1-\rho_2)(\rho_1^h - \rho_1^l)}{1-\rho_1^l}(\tilde{\psi}_2 - \tilde{\psi}_1)(H-\tilde{\psi}_1) + \rho_2\Big[\frac{\rho_1^h - \rho_1^l}{1-\rho_1^l}(H-\tilde{\psi}_1)\Big]^2 \\ & +(1-\rho_2-\rho_1^l)\Big[\frac{\rho_1^h + \rho_2 - 1}{1-\rho_2-\rho_1^l}(H-a_1) - \frac{\rho_1^h - 1}{1-\rho_1^l}(H-\tilde{\psi}_1)\Big]^2\Bigg\}. \end{aligned}$$

2. Assume that in the high period class 1 fluid increases to its threshold (at $\tilde{\psi}_1$) before class 2 fluid decreases to its threshold, after $\tilde{\psi}_1$, class 1 has higher priority in the high period. Suppose that the low period is long enough for class 1 fluid to decrease below its threshold, but not long enough for class 2 fluid to decrease to its threshold. Thus, at the end of the low period, class 1 fluid level is at its threshold but class 2 fluid is still above its threshold. Hence, if

$$\text{(FP2-FP1:2)} \qquad \tilde{\psi}_1 \leq \tilde{\psi}_2, \quad H \geq \tilde{\psi}_1,\ \gamma_3(H-\tilde{\psi}_1) \leq L \leq \gamma_4(H-a_1),$$

then the holding cost under FP2-FP1 policy is

$$
\begin{aligned}
c^{\mathrm{FP2-FP1}}(H,L) &= \frac{1}{2}h_2\mu_2\Bigg\{\frac{(\rho_1^h-1)(\rho_1^h-\rho_1^l)}{\eta(1-\rho_1^l)}(H-\tilde{\psi}_1)^2+(1-\rho_2)(2\tilde{\psi}_2-\tilde{\psi}_1)\tilde{\psi}_1 \\
&\quad +2\frac{(1-\rho_2)(\rho_1^h-\rho_1^l)}{1-\rho_1^l}(\tilde{\psi}_2-\tilde{\psi}_1)(H-\tilde{\psi}_1)+\rho_2\Big[\frac{\rho_1^h-\rho_1^l}{1-\rho_1^l}(H-\tilde{\psi}_1)\Big]^2 \\
&\quad +(1-\rho_2-\rho_1^l)\Big[\frac{\rho_1^h+\rho_2-1}{1-\rho_2-\rho_1^l}(H-a_1)-\frac{\rho_1^h-1}{1-\rho_1^l}(H-\tilde{\psi}_1)\Big]^2\Bigg\} \\
&\quad -\frac{1}{2}h_2\mu_2(1-\rho_2-\rho_1^l)\Big[\frac{\rho_1^h+\rho_2-1}{1-\rho_2-\rho_1^l}(H-a_1)-L\Big]^2.
\end{aligned}
$$

3. Assume that in the high period class 1 fluid increases to its threshold (at $\tilde{\psi}_1$) before class 2 fluid decreases to its threshold, after $\tilde{\psi}_1$, class 1 has higher priority in the high period. Suppose that low period is not long enough for class 1 or class 2 fluid to reach its threshold. Thus, at the end of the low period, both class 1 and class 2 fluid levels are above their thresholds. Hence, if

$$
\text{(FP2-FP1:3)} \qquad \tilde{\psi}_1 \le \tilde{\psi}_2, \quad H \ge \tilde{\psi}_1, \quad L \le \gamma_3(H-\tilde{\psi}_1),
$$

then the holding cost under FP2-FP1 policy is

$$
\begin{aligned}
c^{\mathrm{FP2-FP1}}(H,L) &= \frac{1}{2}h_2\mu_2\Bigg\{\frac{(\rho_1^h-1)(\rho_1^h-\rho_1^l)}{\eta(1-\rho_1^l)}(H-\tilde{\psi}_1)^2+(1-\rho_2)(2\tilde{\psi}_2-\tilde{\psi}_1)\tilde{\psi}_1 \\
&\quad +\rho_2(H+L-\tilde{\psi}_1)^2+2(1-\rho_2)(\tilde{\psi}_2-\tilde{\psi}_1)(H+L-\tilde{\psi}_1) \\
&\quad -\frac{1-\rho_1^l}{\eta}\Big[\frac{\rho_1^h-1}{1-\rho_1^l}(H-\tilde{\psi}_1)-L\Big]^2\Bigg\}.
\end{aligned}
$$

4. Assume that if the high period were long enough, class 1 fluid would increase to its threshold (at $\tilde{\psi}_1$) before class 2 fluid decreases to its threshold. However, the length of the high period is shorter than $\tilde{\psi}_1$. Thus, at the end of the high period, class 2 fluid is above its threshold but class 1 fluid is still below its threshold. In the low period, according to FP2-FP1 policy, class 1 fluid increases to its threshold at $\rho_1^h(\tilde{\psi}_1-H)(\rho_1^l)^{-1}+H$. Suppose that this happens before class 2 fluid decreases to its threshold. Then the server allocates enough capacity to class 1 to maintain class 1 fluid it at its threshold level and the remaining capacity is allocated to serving class 2. Moreover, assume that $L \ge \gamma_4(H-a_1)$, i.e. the low period is long enough for class

2 fluid to reach its threshold. Hence, if

$$\text{(FP2-FP1:4)} \qquad H \le \tilde{\psi}_1 \le \frac{\rho_1^h}{\rho_1^l}(\tilde{\psi}_1 - H) + H \le \tilde{\psi}_2, \quad L \ge \gamma_4(H - a_1),$$

then the holding cost under FP2-FP1 policy is

$$\begin{aligned} c^{\mathrm{FP2-FP1}}(H,L) &= \frac{1}{2}h_2\mu_2\Big\{(1-\rho_2)\Big[2\tilde{\psi}_2 - \frac{\rho_1^h}{\rho_1^l}(\tilde{\psi}_1 - H) - H\Big]\Big[\frac{\rho_1^h}{\rho_1^l}(\tilde{\psi}_1 - H) + H\Big] \\ &\quad + (1-\rho_2-\rho_1^l)\Big[\frac{\rho_1^h + \rho_2 - 1}{1-\rho_2-\rho_1^l}(H - a_1) - \frac{\rho_1^h}{\rho_1^l}(\tilde{\psi}_1 - H)\Big]^2\Big\}. \end{aligned}$$

5. Assume that all the assumptions of (FP2-FP1:4) hold except $L \le \gamma_4(H - a_1)$, i.e. the low period is not long enough for class 2 fluid to reach its threshold. Thus, at the end of the low period, class 1 fluid is at its threshold and class 2 fluid is still above its threshold. Hence, if

$$\text{(FP2-FP1:5)} \qquad H \le \tilde{\psi}_1 \le \frac{\rho_1^h}{\rho_1^l}(\tilde{\psi}_1 - H) + H \le \tilde{\psi}_2, \quad \frac{\rho_1^h}{\rho_1^l}(\tilde{\psi}_1 - H) \le L \le \gamma_4(H - a_1),$$

then the holding cost under FP2-FP1 policy is

$$\begin{aligned} c^{\mathrm{FP2-FP1}}(H,L) &= \frac{1}{2}h_2\mu_2\Big\{(1-\rho_2)\Big[2\tilde{\psi}_2 - \frac{\rho_1^h}{\rho_1^l}(\tilde{\psi}_1 - H) - H\Big]\Big[\frac{\rho_1^h}{\rho_1^l}(\tilde{\psi}_1 - H) + H\Big] \\ &\quad + (1-\rho_2-\rho_1^l)\Big[\frac{\rho_1^h + \rho_2 - 1}{1-\rho_2-\rho_1^l}(H - a_1) - \frac{\rho_1^h}{\rho_1^l}(\tilde{\psi}_1 - H)\Big]^2\Big\} \\ &\quad - \frac{1}{2}h_2\mu_2(1-\rho_2-\rho_1^l)\Big[\frac{\rho_1^h + \rho_2 - 1}{1-\rho_2-\rho_1^l}(H - a_1) - L\Big]^2. \end{aligned}$$

6. Assume that the assumptions of (FP2-FP1:5) hold except the low period is not long enough for either class 1 fluid or class 2 fluid to reach its threshold. Thus, at the end of the low period, class 1 fluid is below its threshold and class 2 fluid is above its threshold. Hence, if

$$\text{(FP2-FP1:6)} \qquad H \le \tilde{\psi}_1 \le \frac{\rho_1^h}{\rho_1^l}(\tilde{\psi}_1 - H) + H \le \tilde{\psi}_2, \quad L \le \frac{\rho_1^h}{\rho_1^l}(\tilde{\psi}_1 - H),$$

then the cost under FP2-FP1 policy is

$$c^{\mathrm{FP2-FP1}}(H,L) = \frac{1}{2}h_2\mu_2(1-\rho_2)(2\tilde{\psi}_2 - H - L)(H + L).$$

7. Assume that if the high period were long enough, class 1 fluid would increase to its threshold before class 2 fluid decreases to its threshold. However, the high period is not long enough for class 1 fluid to increase to its threshold. Suppose that at the end of the high period, class 1 fluid is still below its threshold, and class 2 fluid is still above its threshold. Moreover, assume that the low period is long enough for class 2 fluid to decrease to its threshold and in the low period, class 2 fluid decreases to its threshold earlier than class 1 fluid increases to its threshold. Hence, if

$$\text{(FP2-FP1:7)} \qquad H \leq \tilde{\psi}_1 \leq \tilde{\psi}_2 \leq \frac{\rho_1^h}{\rho_1^l}(\tilde{\psi}_1 - H) + H, \quad L + H \geq \tilde{\psi}_2,$$

then the holding cost under FP2-FP1 policy is

$$c^{\mathrm{FP2-FP1}}(H, L) = \frac{1}{2} h_2 \mu_2 (1 - \rho_2) \tilde{\psi}_2^2.$$

8. Assume that all the assumptions of (FP2-FP1:7) hold except that the low period is not long enough for class 2 fluid to decrease to its threshold. Hence, if

$$\text{(FP2-FP1:8)} \qquad H \leq \tilde{\psi}_1 \leq \tilde{\psi}_2 \leq \frac{\rho_1^h}{\rho_1^l}(\tilde{\psi}_1 - H) + H, \quad L + H \leq \tilde{\psi}_2$$

then the holding cost under FP2-FP1 policy is

$$c^{\mathrm{FP2-FP1}}(H, L) = \frac{1}{2} h_2 \mu_2 (1 - \rho_2)(2\tilde{\psi}_2 - H - L)(H + L).$$

9. Assume that class 2 fluid decreases to its threshold before class 1 fluid increases to its threshold. Hence, $\tilde{\psi}_1 \geq \tilde{\psi}_2$. After $\tilde{\psi}_2$, the server allocates just enough capacity to keep class 2 fluid at its threshold, i.e. $u_2 = \rho_2$, and the remaining capacity is allocated to class 1, i.e $u_1 = 1 - \rho_2$. If $H \geq a_1$, then class 1 fluid reaches its threshold at a_1 and after a_1, class 1 has higher priority until class 1 fluid decreases to its threshold again in the low period. After class 1 fluid decreases to its threshold in the low period, we have $u_1 = \rho_1^l$ and $u_2 = 1 - \rho_1^l$. Moreover, assume that $L \geq \gamma_4(H - a_1)$. Thus, at the end of the low period, fluid levels of both classes are below their thresholds. Hence, if

$$\text{(FP2-FP1:9)} \qquad \tilde{\psi}_1 \geq \tilde{\psi}_2, \quad H \geq a_1, \quad L \geq \gamma_4(H - a_1),$$

then the holding cost under FP2-FP1 policy is

$$\begin{aligned}
& c^{\mathrm{FP2-FP1}}(H,L) \\
= \; & \frac{1}{2}h_2\mu_2\left\{\frac{(\rho_1^h-1)(\rho_1^h-\rho_1^l)}{\eta(1-\rho_1^l)}(H-a_1)^2+\rho_2\Big[\frac{\rho_1^h-\rho_1^l}{1-\rho_1^l}(H-a_1)\Big]^2\right. \\
& \left.+(1-\rho_2)\tilde{\psi}_2^2+(1-\rho_2-\rho_1^l)\Big[\frac{\rho_1^h+\rho_2-1}{1-\rho_2-\rho_1^l}(H-a_1)-\frac{\rho_1^h-1}{1-\rho_1^l}(H-a_1)\Big]^2\right\}.
\end{aligned}$$

10. Assume that all the assumptions of (FP2-FP1:9) hold except $L \leq \gamma_4(H-a_1)$. Thus, at the end of the low period, class 1 fluid is below its threshold, but class 2 fluid is above its threshold. Hence, if

$$\text{(FP2-FP1:10)} \qquad \tilde{\psi}_1 \geq \tilde{\psi}_2, \quad H \geq a_1, \quad \gamma_3(H-a_1) \leq L \leq \gamma_4(H-a_1),$$

then the holding cost under FP2-FP1 policy is

$$\begin{aligned}
& c^{\mathrm{FP2-FP1}}(H,L) \\
= \; & \frac{1}{2}h_2\mu_2\left\{\frac{(\rho_1^h-1)(\rho_1^h-\rho_1^l)}{\eta(1-\rho_1^l)}(H-a_1)^2+\rho_2\Big[\frac{\rho_1^h-\rho_1^l}{1-\rho_1^l}(H-a_1)\Big]^2\right. \\
& \left.+(1-\rho_2)\tilde{\psi}_2^2+(1-\rho_2-\rho_1^l)\Big[\frac{\rho_1^h+\rho_2-1}{1-\rho_2-\rho_1^l}(H-a_1)-\frac{\rho_1^h-1}{1-\rho_1^l}(H-a_1)\Big]^2\right\} \\
& -\frac{1}{2}h_2\mu_2(1-\rho_2-\rho_1^l)\Big[\frac{\rho_1^h+\rho_2-1}{1-\rho_2-\rho_1^l}(H-a_1)-L\Big]^2.
\end{aligned}$$

11. Assume that all the assumptions of (FP2-FP1:10) hold except that the low period is not long enough for class 1 fluid to decrease to its threshold, i.e. $L \leq \gamma_3(H-a_1)$. Thus, at the end of the low period, class 1 and class 2 fluids are above their thresholds. Hence, if

$$\text{(FP2-FP1:11)} \qquad \tilde{\psi}_1 \geq \tilde{\psi}_2, \quad H \geq a_1, \quad L \leq \gamma_3(H-a_1),$$

then the holding cost under FP2-FP1 policy is

$$\begin{aligned}
c^{\mathrm{FP2-FP1}}(H,L) \;=\; & \frac{1}{2}h_2\mu_2\left\{\frac{(\rho_1^h-1)(\rho_1^h-\rho_1^l)}{\eta(1-\rho_1^l)}(H-a_1)^2+\rho_2(H+L-a_1)^2\right. \\
& \left.+(1-\rho_2)\tilde{\psi}_2^2-\frac{1-\rho_1^l}{\eta}\Big[\frac{\rho_1^h-1}{1-\rho_1^l}(H-a_1)-L\Big]^2\right\}.
\end{aligned}$$

12. Assume that all the assumptions of (FP2-FP1:9) hold except $H \leq a_1$, i.e. high period is not long enough for class 1 fluid to increase to its threshold. Thus, at the end of the high period, class 2 and class 1 fluids are below their thresholds. Hence, if

$$\text{(FP2-FP1:12)} \qquad \tilde{\psi}_1 \geq \tilde{\psi}_2, \quad \tilde{\psi}_2 \leq H \leq a_1,$$

then the holding cost under FP2-FP1 policy is

$$c^{\mathrm{FP2-FP1}}(H, L) \;=\; \frac{1}{2} h_2 \mu_2 (1 - \rho_2) \tilde{\psi}_2^2.$$

13. Assume that class 2 fluid decreases to its threshold before class 1 fluid increases to its threshold but $H \leq \tilde{\psi}_2$. Then, at the end of the high period, class 2 fluid is above its threshold and class 1 fluid is below its threshold. Suppose that in the low period class 2 has higher priority and class 2 fluid decreases to its threshold at $\tilde{\psi}_2$ and class 1 fluid remains below its threshold. Thus, at the end of the low period both classes are below their thresholds. Hence, if

$$\text{(FP2-FP1:13)} \qquad \tilde{\psi}_1 \geq \tilde{\psi}_2, \quad H \leq \tilde{\psi}_2, \quad H + L \geq \tilde{\psi}_2,$$

then the holding cost under FP2-FP1 policy is

$$c^{\mathrm{FP2-FP1}}(H, L) \;=\; \frac{1}{2} h_2 \mu_2 (1 - \rho_2) \tilde{\psi}_2^2.$$

14. Assume that all assumptions of (FP2-FP1:13) hold except that the low period is not long enough for class 2 fluid to decrease to its threshold. Hence, if

$$\text{(FP2-FP1:14)} \qquad \tilde{\psi}_1 \geq \tilde{\psi}_2, \quad H \leq \tilde{\psi}_2, \quad H + L \leq \tilde{\psi}_2,$$

then the holding cost under the FP2-FP1 policy is

$$c^{\mathrm{FP2-FP1}}(H, L) \;=\; \frac{1}{2} h_2 \mu_2 (1 - \rho_2)(2\tilde{\psi}_2 - H - L)(H + L).$$

REFERENCES

[1] Arlitt, M. and Jin, T., "Workload characterization study of the 1998 world cup web site.," Tech. Rep. HPL-99-35RPL, Hewlett-Packard Labs, HP Laboratories, Palo Alto, Sept 1999.

[2] Arlitt, M. and Williamson, C., "Internet web servers: workload characterization and performance implications," *IEEE/ACM Transactions on Network*, vol. 5, pp. 631–645, Oct 1997.

[3] Ata, B. and Kumar, S., "Heavy traffic analysis of open processing networks with complete resource pooling: asymptotic optimality of discrete review policies," *Annals of Applied Probability*, Preprint, 2004.

[4] Avram, F., Bertsimas, D., and Ricard, M., "Fluid models of sequencing problems in open queueing networks: an optimal control approach," in *Stochastic Networks* (Kelly, F. and Williams, R., eds.), pp. 199–234, 1995.

[5] Baeuerle, N., "Asymptotic optimality of tracking policies in stochastic networks," *Annals of Applied Probability*, vol. 10, no. 4, pp. 1065–1083, 2000.

[6] Baeuerle, N. and Rieder, U., "Optimal control of single-server fluid networks," *Queueing Systems: Theory and Applications*, vol. 35, pp. 185–200, 2000.

[7] Bansal, N. and Harchol-Balter, M., "Scheduling solutions for coping with transient overload," Tech. Rep. CMU-CS-01-134, Department of Computer Science, Carnegie Mellon University, 2001.

[8] Billingsley, P., *Convergence of Probability Measures*. New York: Wiley & Sons, Inc, 1968.

[9] Bramson, M., "State space collapse with application to heavy traffic limits for multiclass queueing networks," *Queueing Systems: Theory and Applications*, vol. 30, pp. 89–148, 1998.

[10] Chen, H. and Mohapatra, P., "Session-based overload control in qos-aware web servers," in *INFOCOM*, 2002.

[11] Chen, H. and Yao, D., "Daymanic scheduling of a multiclass fluid network," *Operations Research*, vol. 41, pp. 1104–1115, 1993.

[12] Chen, R. and Meyn, S., "Value iteration and optimization of multiclass queueing networks," *Queueing Systems: Theory and Applications*, vol. 32, pp. 65–97, 1999.

[13] Cheng, H. and Mandelbaum, A., "Hierarchical modeling of stochastic networks part I: fluid models," in *Stochastic Modeling and Analysis of Manufacturing Systmes* (Yao, D. D., ed.), (New York), pp. 199–234, Springer, 1994.

[14] Choudhury, G. L., Mandelbaum, A., Reiman, M. I., and Whitt, W., "Fluid and diffusion limits for queues in slowly changing environments," *Stochastic Models*, vol. 13, pp. 121–146, 1997.

[15] Dai, J., "On positive harris recurrence of multiclass queueing networks: A unified approach via fluid limt models," *The Annals of Applied Probability*, vol. 5, no. 1, pp. 49–77, 1995.

[16] Ethier, S. and Kurtz, T., *Markov Processes: Characterization and Covergence.* 1986.

[17] Gajrat, A. and Hordijk, A., "Fluid approximation of a controlled multiclass tandem network," *Queueing Systems: Theory and Applications*, vol. 35, pp. 349–380, 2000.

[18] Green, T. and Stidham, S., "Sample-path conservation laws, with applications to scheduling queues and fluid systems," *Queueing Systems: Theory and Applications*, vol. 36, pp. 175–199, 2000.

[19] Harrison, J. M. and Zeevi, A., "A method for staffing large call centers based on stochastic fluid models," Preprint, 2004.

[20] Harrison, J., "The big step approach to flow management in stochastic processing networks," in *Stochastic Networks: Theory and Applications* (Kelly, F., Zachary, S., and Ziedins, I., eds.), pp. 57–90, Oxford University Press, 1996.

[21] Iyengar, A., MacNair, E., and Nguyen, T., "An analysis of web server performance," *Global Telecommunications Conference, GLOBECOM*, vol. 3, pp. 1943–1947, Nov 1997.

[22] Iyengar, A., Squillante, M., and Zhang, L., "Analysis and characterization of large-scale web server access patterns and performance," *World Wide Web*, vol. 2, pp. 85–100, June 1999.

[23] Klimov, G., "Time sharing service systems I," *Theory of Probability and Its Applications*, vol. 19, no. 3, pp. 532–551, 1974.

[24] Liu, Z., Squillante, M., Xia, C., Yu, S.-Z., and Zhang, L., "Profile-based traffic characterization of commercial web sites," in *Proceedings of the 18th International Teletraffic Congress (ITC18)*, (Berlin, Germany), pp. 231–240, 2003.

[25] Liu, Z., Squillante, M., and Wolf, J., "Optimal control of resource allocation in e-business environments with strict quality-of-service performance guarantees," *Proceedings of the IEEE Conference on Decision and Control*, 2002.

[26] Maglaras, C., "Dynamic scheduling in multiclass queuing networks: stability under discrete-review policies," *Queueing Systems: Theory and Applications*, vol. 31, pp. 171–206, 1999.

[27] Maglaras, C., "Dicrete review policies for scheduling stochastic networks: trajectory tracking and fluid scale asymptotic optimality," *The Annals of Applied Probability*, vol. 10, no. 3, pp. 897–929, 2000.

[28] Massey, W., "The analyis of queues with time-varying rates for teleccommunication models," *Telecommunication Systems*, vol. 21, no. 2-4, pp. 173–204, 2002.

[29] Meyn, S. P., "The policy iteration algorithm for average reward markov decision processes with general state space," *IEEE Transactions on Autmatic Control*, vol. 42, pp. 1663–1680, 1997.

[30] Meyn, S. P., "Stability and optimization of queueing networks and their fluid models," in *Mathematics of Stochastic Manufacturing Systems. Lectures in Applied Math* (Yin, G. and Zhang, Q., eds.), (Providence, RI), pp. 175–200, Amer. Math. Soc., 1997.

[31] Pontryagin, L., Boltyanskii, V., Gamkrelidze, R., and Mishchenko, E., *The Mathematical Theory of Optimal Processes*. New York: Interscience Publishers, 1962.

[32] Seierstad, A. and Sydsater, K., "Optimal control theory with economic applications," 1987.

[33] Sethi, S. and Zhang, Q., *Hierarchical Decision Making in Stochastic Manufacturing Systems*. Boston: Birkhauser, 1994.

[34] Shaked, M. and Shanthikuumar, J. G., *Stochastic Orders and Their Applications*. San Diego, CA: Academic Press, 1994.

[35] Shanthikumar, J. G. and Xu, S. H., "Strongly asymptotically optimal design and control of production and service systems," *IIE Transactions*, vol. 32, no. 9, pp. 881–890, 2000.

[36] Smith, W. E., "Various optimizers for single-stage production," *Naval Res. Logist. Quart.*, vol. 3, pp. 59–66, 1956.

[37] Weiss, G., "Scheduling and control of manufacturing systems – a fluid approach," in *Proceedings of the 37 Allerton Conference*, pp. 577–586, 1999.

[38] Xia, C. and Shanthikumar, J., "Asymptotic optimal control of multi-class $g/g/1$ queues with feedback," in *Stochastic Modeling and Optimization of Manufacturing Systems and Supply Chains* (Shanthikumar, J., Yao, D., and Zijm, W., eds.), p. To appear, 2003.

[39] Yeh, D., Dai, J., and Zhou, C., "The qnet method for re-entrent queueing networks with priority disciplines," *Operations Research*, vol. 45, no. 4, pp. 610–623, 1997.

VITA

Junxia Chang was born in a small village of ShanDong province, China. She received her B.Sc degree in Probability and Statistics from Nanjing University in 1996, Nanjing, China. Then she went to pursue the graduate study at the Institute of Applied Mathematics, Chinese Academy of Sciences, Beijing, China. There she received the M.Sc degree in Probability and Statistics in 1999. She attended the doctoral program in the School of Industrial and Systems Engineering, Georgia Institute of Technology since the fall of 1999.