Queueing Systems 48, 263–307, 2004

Dynamic Scheduling of a Multiclass Fluid Model with Transient Overload

JUNXIA CHANG, HAYRIYE AYHAN * and J.G. DAI ** {junxiach;hayhan;dai}@isye.gatech.edu
School of ISyE, Georgia Institute of Technology, Atlanta, GA 30332-0205, USA

CATHY H. XIA cathyx@us.ibm.com
IBM T.J. Watson Research Center, P.O. Box 704, Yorktown Heights, NY 10598, USA

Received 15 May 2003; Revised 4 May 2004

Abstract. We study the optimal dynamic scheduling of different requests of service in a multiclass stochastic fluid model that is motivated by recent and emerging computing paradigms for Internet services and applications. In particular, our focus is on environments with specific performance guarantees for each class under a profit model in which revenues are gained when performance guarantees are satisfied and penalties are incurred otherwise. Within the context of the corresponding fluid model, we investigate the dynamic scheduling of different classes of service under conditions where the workload of certain classes may be overloaded for a transient period of time. Specifically, we consider the case with 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 varies with time and the class 1 fluid can more efficiently reduce the holding cost. Under these assumptions, we characterize the optimal server allocation policy that minimizes the holding cost in the fluid model when the arrival rate function for class 1 is known. Using the insights gained from this deterministic case, we study the stochastic fluid system when the arrival rate function for class 1 is random and develop 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. Numerical examples are also provided to demonstrate further that these policies yield good results in terms of minimizing the expected holding cost.

Keywords: stochastic fluid model, transient overload, e-commerce, quality-of-service, service-level-agreement

AMS subject classification: 60K25, 90B22, 90B36

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

* Research supported in part by National Science Foundation grants DMI-9908161 and DMI-9984352.
** Research supported in part by National Science Foundation grants DMI-9813345 and DMI-0300599, and by TLI-AP, a partnership between National University of Singapore and Georgia Institute of Technology.

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 paper 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 [2]. 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.

Our 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., [9]) and direct modifications to the Internet server scheduling mechanism (e.g., [2,6]). On the contrary, our focus in this paper 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 [12], 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 our research, but it differs from the present study in several important aspects. Our focus in this paper 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 [12].

Our primary concern in this paper is to investigate the preceding optimal server resource control problem as a dynamic scheduling problem. Our motivation behind considering a fixed time horizon is that in reality many Web sites exhibit [13] regular daily access patterns, 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. Our approach is based on formulating the problem as a multiclass stochastic fluid model and employing optimal control theory [14,15] 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 [1,3,7,18], to the best our knowledge, no optimal scheduling policy is known for the general problem considered herein. As mentioned above, we focus 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. Preliminary results on this problem can be found in an earlier work [4]. In this paper, 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 has also been considered in [16,19], 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 the paper is as follows. We define our multiclass fluid model in section 2. Deterministic instance of the model is analyzed in section 3 where we provide the optimal control policy. Sections 4 and 5 consider the stochastic instance of the model. In section 4, 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 5. Our concluding remarks are provided in section 7. Throughout, proofs are relegated to the appendix.

2. The stochastic fluid model

This paper 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 $\mu_i,\ i = 1, 2$.

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, $\lambda_1^{\mathrm{l}} \leqslant \lambda_1^{\mathrm{h}}$. 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 in the paper. 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

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

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

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

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 $T = \{(T_1(t), T_2(t)),\ t \geqslant 0\}$ 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 $\theta_i \geqslant 0$ be constants, $i = 1, 2$. For a real number x, define $x^+ = \max(x, 0)$. Consider the integral

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

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 paper 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, [8,10,17]). Details of such an optimal policy are discussed in appendix A.

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 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 $\rho_1^{\mathrm{h}} = \lambda_1^{\mathrm{h}}/\mu_1$ for the high load period and $\rho_1^{\mathrm{l}} = \lambda_1^{\mathrm{l}}/\mu_1$ 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 $\rho^{\mathrm{h}} = \rho_1^{\mathrm{h}} + \rho_2$ for the high load period and $\rho^{\mathrm{l}} = \rho_1^{\mathrm{l}} + \rho_2$ for the low load period. When $\rho^{\mathrm{h}} > 1$ and $\rho^{\mathrm{l}} < 1$, 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 $\rho^{\mathrm{h}} > 1$ and $\rho^{\mathrm{l}} < 1$ the system experiences an overload period followed by an under-load period, a phenomenon known as transient overload in literature; see, for example, [2]. Although understanding transient overload is the primary motivation of this paper, except explicitly stated otherwise, we do not assume $\rho^{\mathrm{h}} > 1$.

3. 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 appendix B. For convenience, we first define the following policy.

Definition 1. The following policy referred as the *low-period-policy* is implemented in the low period, i.e., when $H < t \leqslant H + L$.

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

The optimal policy depends on the system load. In the next three subsections, we will describe the optimal policy under all load conditions. In the first case, $\rho_1^{\mathrm{h}} > 1$, $\rho^{\mathrm{l}} \leqslant 1$; in the second case, $\rho^{\mathrm{h}} > 1$, $\rho_1^{\mathrm{h}} \leqslant 1$, $\rho^{\mathrm{l}} \leqslant 1$; and in the last case, $\rho^{\mathrm{h}} \leqslant 1,\ \rho^{\mathrm{l}} \leqslant 1$.

3.1. The case $\rho_1^{\mathrm{h}} > 1,\ \rho^{\mathrm{l}} \leqslant 1$

Suppose that $\rho_1^{\mathrm{h}} > 1$ and $\rho^{\mathrm{l}} \leqslant 1$. Then the optimal policy has the following structure:

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

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:

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

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

The quantities $\psi_1, \psi_2, \tilde{\psi}_1$ 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

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

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

where

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

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

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

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 1–3. 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: $Z_1(0) \geqslant \theta_1$. 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

$$a_1 \leqslant H \leqslant B, \qquad L \geqslant \gamma_1(H - a_1), \tag{10}$$

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

$$Z_2(0) + (\lambda_2 - \mu_2)s_1 = \theta_2, \tag{11}$$
$$Z_1(0) + \lambda_1^{\mathrm{h}} s_1 = Z_1(s_1), \tag{12}$$
$$Z_2(s_1) + (\lambda_2 - \mu_2 u_2)(s_2 - s_1) = \theta_2, \tag{13}$$
$$Z_1(s_1) + \big(\lambda_1^{\mathrm{h}} - \mu_1 u_1\big)(s_2 - s_1) = Z_1(s_2), \tag{14}$$
$$u_1 + u_2 = 1, \tag{15}$$
$$Z_1(s_2) + \big(\lambda_1^{\mathrm{h}} - \mu_1\big)(t_1 - s_2) = \theta_1, \tag{16}$$
$$Z_1(t_1) + \big(\lambda_1^{\mathrm{h}} - \mu_1\big)(H - t_1) = Z_1(H), \tag{17}$$
$$Z_1(H) + \big(\lambda_1^{\mathrm{l}} - \mu_1\big)(t_2 - H) = \theta_1, \tag{18}$$
$$\mu_1 h_1(t_2 - t_1) = \mu_2 h_2(t_2 - s_2). \tag{19}$$

Note that equations (11)–(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)–(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)–(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 B that equations (11)–(19) have a unique solution with $0 \leqslant s_1 \leqslant s_2 \leqslant t_1 \leqslant H \leqslant t_2 \leqslant H + L$ and $u_1, u_2 \geqslant 0$.

– Case 2.2: If

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

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

$$\max\big\{B, \tilde{\psi}_1\big\} \leqslant H \leqslant a_2, \qquad L \geqslant \gamma_2\big(H - \tilde{\psi}_1\big),$$

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

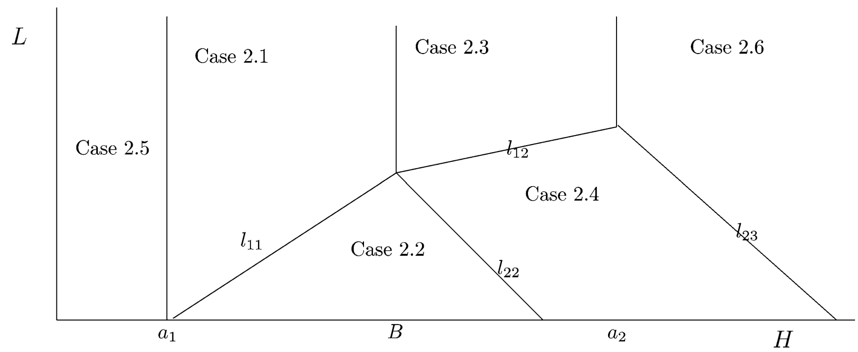

Figure 1. Pictorial representation for the case $Z_1(0) \leqslant \theta_1, Z_2(0) \geqslant \theta_2$ and $\tilde{\psi}_1 \geqslant \tilde{\psi}_2$, where l_{11}: $L = \gamma_1(H - a_1)$, l_{12}: $L = \gamma_2(H - \tilde{\psi}_1)$, l_{22}: $H + L = \tilde{\psi}_1 + ((1 + \eta(\rho_1^{\mathrm{h}} - 1))/((1 - \eta)(\rho_1^{\mathrm{h}} - 1)))(\tilde{\psi}_1 - \tilde{\psi}_2)$, l_{23}: $H + L = \psi_1/(1 - \eta)$.

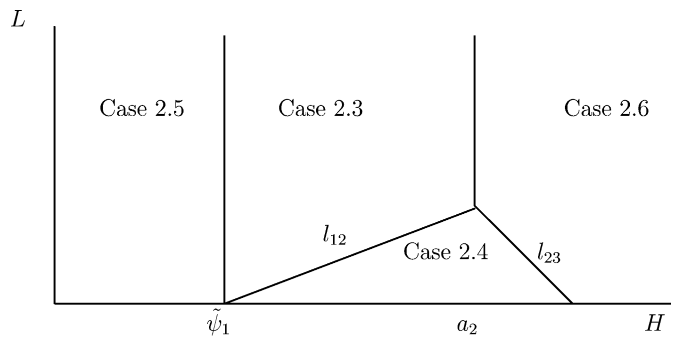

Figure 2. Pictorial representation for the case $Z_1(0) \leqslant \theta_1$, $Z_2(0) \geqslant \theta_2$ and $\tilde{\psi}_1 \leqslant \tilde{\psi}_2$, where l_{12}: $L = \gamma_2(H - \tilde{\psi}_1)$, l_{23}: $H + L = \psi_1/(1-\eta)$.

- Case 2.4: If

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

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 $H \leqslant \max\{a_1, \tilde{\psi}_1\}$, 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 $H \geqslant a_2$ and $H + L \geqslant (1-\eta)^{-1}\psi_1$, then the optimal policy is given by (OPT) with $s_1 = s_2 = 0$.

• Case 3: $Z_1(0) < \theta_1, Z_2(0) \leqslant \theta_2, \psi_1 \leqslant \psi_2$. In this case, the optimal policy is given by (OPT) with $s_1 = s_2 = 0$.
• Case 4: $Z_1(0) < \theta_1, Z_2(0) \leqslant \theta_2, \psi_1 \geqslant \psi_2$. 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 $a_1 \leqslant H \leqslant a_2, L \geqslant \gamma_1(H - a_1)$, 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

$$H \geqslant a_1, \qquad H + L \leqslant \psi_1 + \frac{\eta}{1-\eta}(\psi_1 - \psi_2), \qquad L \leqslant \gamma_1(H - a_1),$$

then we set $t_2 = H + L$, and solve equations (13)–(17) and (19) with $s_1 = 0$ to compute s_2, u_1, u_2 and t_1.
- Case 4.3: If $H \leqslant a_1$, 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 $[(\rho_2 - d_2(\mu_2 H)^{-1})^+, d_1(\mu_1 H)^{-1} - (\rho_1^{\mathrm{h}} - 1)]$ and setting $u_1 = 1 - u_2$.

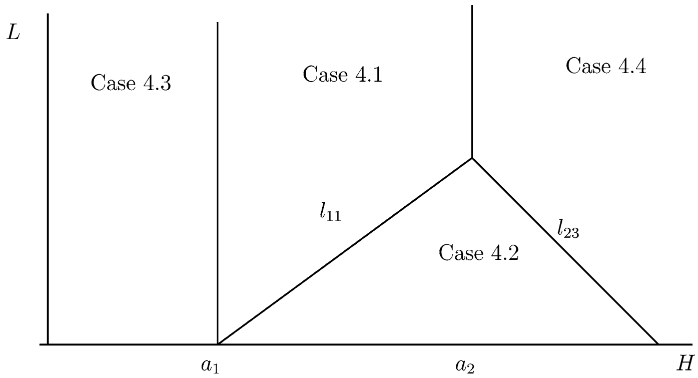

Figure 3. Pictorial representation for the case $Z_1(0) \leqslant \theta_1, Z_2(0) \leqslant \theta_2$ and $\psi_2 \leqslant \psi_1$, where l_{11}: $L = \gamma_1(H - a_1), l_{23}$: $H + L = \psi_1 + (\eta/(1-\eta))(\psi_1 - \psi_2)$.

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

As mentioned above, we prove the optimality of this policy in appendix B. 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 $Z_1(0) < \theta_1$, $Z_2(0) \leqslant \theta_2$, $\psi_1 \leqslant \psi_2$, then the optimal policy is given by (OPT) with $s_1 = s_2 = 0$. In order to see this, first consider the case $H \geqslant \psi_1$. 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 $t \geqslant \psi_1$, we have

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

since $\psi_2 \geqslant \psi_1 \geqslant 0$ 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) $Z_1(0) \geqslant \theta_1$ or,

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

then the policy with

$$\begin{array}{lll} \forall t \in (0, H) & \dot{T}_1(t) = 1, & \dot{T}_2(t) = 0; \\ \forall t \in (H, H+L) & \text{low-period-policy} & \end{array}$$

is optimal for all $H \geqslant 0$ and $L \geqslant 0$.

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.

3.2. *The case* $\rho^{\mathrm{h}} > 1$, $\rho_1^{\mathrm{h}} \leqslant 1$, $\rho^{\mathrm{l}} \leqslant 1$

In this case the optimal policy has the following structure:

$$
\begin{array}{lll}
\forall t \in (0, s_1): & \dot{T}_2(t) = 1, & \dot{T}_1(t) = 0; \\
\forall t \in (s_1, s_2): & \dot{T}_2(t) = \rho_2 - \dfrac{(\theta_2 - Z_2(s_1))/\mu_2}{a_1(s_1)}, & \dot{T}_1(t) = 1 - \dot{T}_2(t); \\
\forall t \in (s_2, s_3): & \dot{T}_2(t) = 0, & \dot{T}_1(t) = 1; \\
\forall t \in (s_3, H): & \dot{T}_2(t) = 1 - \rho_1^{\mathrm{h}}, & \dot{T}_1(t) = \rho_1^{\mathrm{h}}; \\
\forall t \in (H, H+L): & \text{low-period-policy;} &
\end{array}
$$

where

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

and s_1, s_2, s_3 are given as

$$
\begin{aligned}
s_1 &= \max\bigl\{t\colon\ 0 \leqslant t \leqslant H,\ Z_2(t) \geqslant \theta_2,\ Z_1(t) \leqslant \theta_1\bigr\}, \\
s_2 &= \max\bigl\{t\colon\ s_1 \leqslant t \leqslant H,\ Z_1(t) \leqslant \theta_1\bigr\}, \\
s_3 &= \max\bigl\{t\colon\ s_2 \leqslant t \leqslant H,\ Z_1(t) \geqslant \theta_1\bigr\}.
\end{aligned}
$$

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

3.3. *The case* $\rho^{\mathrm{h}} \leqslant 1$, $\rho^{\mathrm{l}} \leqslant 1$

In this case the optimal policy has the following structure:

$$
\begin{array}{ll}
\forall t \in (0, H) & \text{low-period-policy except replace } \rho_1^{\mathrm{l}} \text{ by } \rho_1^{\mathrm{h}}; \\
\forall t \in (H, H+L) & \text{low-period-policy.}
\end{array}
$$

Remark 3. The policies described in sections 3.2 and 3.3 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.

4. Discrete review policies in the stochastic case

Throughout the rest of this paper, we shall consider the stochastic instance of the fluid model described in section 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 $\rho^{\mathrm{h}} > 1$, $\rho_1^{\mathrm{h}} \leqslant 1$, $\rho^{\mathrm{l}} \leqslant 1$ is given in section 3.2 and the optimal policy when $\rho^{\mathrm{h}} \leqslant 1$, $\rho^{\mathrm{l}} \leqslant 1$ is given in section 3.3 (see remark 3). We therefore focus only on the case when

$$\rho_1^{\mathrm{h}} > 1, \qquad \rho^{\mathrm{l}} \leqslant 1.$$

To specify the control policy in this case, we shall always consider the following four subcases which were first introduced in section 3 and are summarized below:

$$\text{case 1:} \quad Z_1(0) \geqslant \theta_1, \tag{20}$$

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

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

$$\text{case 4:} \quad Z_1(0) < \theta_1, \qquad Z_2(0) \leqslant \theta_2, \qquad \psi_1 \geqslant \psi_2. \tag{23}$$

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

$$M = \min\{n \in \mathbb{N}\colon n\tau \geqslant H_0\}.$$

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 $\widetilde{H}(t)$ and the low period $\widetilde{L}(t)$ using one of the methods described below. If $t < M\tau$, we implement the policy described in section 3 from t to $t + \tau$ using $\widetilde{H}(t)$ as the length of the high period, $\widetilde{L}(t)$ 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

$$\widetilde{H}(t) = \mathbb{E}[H \mid H > t] - t, \tag{24}$$

or

$$\widetilde{H}(t) = \min\bigl\{x \geqslant 0\colon \mathbb{P}(H > x + t \mid H > t) = p\bigr\}, \tag{25}$$

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 6, 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, $\widetilde{L}(t) = \mathbb{E}[L]$.

We now show that our discrete review policy is asymptotically optimal as the expected length of the high period tends to infinity. 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 3. The closed form expression for $c(H_0, L_0)$ is given in [5, appendix C] (which is a longer version of the current paper). 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 $\beta_1 \geqslant 0$ (which depend on the arrival rates, service rates, initial fluid levels, threshold values and holding costs per unit time) such that if

$$\widetilde{H}(0) \geqslant D,$$

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

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

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

Proof. We provide the proof for the discrete review policy where $\widetilde{H}(t)$ 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

$$D = \max\left\{a_2(0), \psi_1(0) + \frac{\eta}{1-\eta}\big(\psi_1(0) - \psi_2(0)\big)\right\}. \tag{27}$$

We first show by induction that for all $0 \leqslant n \leqslant M - 1$, 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 cases 1 and 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 $\psi_2 = \psi_2(0) \leqslant 0$, then $D \geqslant a_2$ and $D \geqslant \psi_1 + \eta(1-\eta)^{-1}\psi_1 = (1-\eta)^{-1}\psi_1$. Hence, $\widetilde{H}(t) \geqslant D$ (where D is given in (27)) which implies that the condition of case 2.6 in section 3.1 is satisfied, where the discrete review policy gives fixed priority to class 1 in the interval $[0, \tau)$.

For case 4, $\widetilde{H}(t) \geqslant D$ (where D is given in (27)) which implies that the condition of case 4.4 in section (3.1) 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 $1 \leqslant n \leqslant M - 1$. Then the fluid levels of the two classes at time $t = n\tau$ are $Z_1(n\tau) = Z_1(0) + n\tau(\lambda_1^{\mathrm{h}} - \mu_1)$, and $Z_2(n\tau) = Z_2(0) + n\tau\lambda_2$, respectively. It is easily checked from (6), (7) and (8) that

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

To specify the discrete review policy at time $t = n\tau$, we again consider cases 1–4 given in (20)–(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 cases 1 and 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 $\widetilde{H}(0) = \mathbb{E}[H]$ and

$$\begin{aligned} \widetilde{H}(n\tau) &\geqslant \mathbb{E}[H] - n\tau \geqslant D - n\tau \\ &= \max\left\{a_2(n\tau),\, \psi_1(n\tau) + \frac{\eta}{1-\eta}\big(\psi_1(n\tau) - \psi_2(n\tau)\big)\right\}, \end{aligned} \tag{28}$$

it follows from case 2.6 in section 3.1 that the discrete review policy gives fixed priority to class 1 in the interval $[n\tau, (n+1)\tau)$.

Similarly, for case 4, (28) implies that conditions of case 4.4 in section 3.1 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 $0 \leqslant t \leqslant H_0$. The result in (26) then follows from proposition 6 in section 5. ☐

Remark 5. The proof for other methods are the same except $\mathbb{E}[H]$ is replaced by $\widetilde{H}(0)$ in (28).

5. Other policies that are asymptotically optimal

Throughout this section, we assume that $\rho_1^{\mathrm{h}} > 1$ and $\rho^{\mathrm{l}} \leqslant 1$. 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 $1 - \rho_1^{\mathrm{l}}$ 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 $c^{\mathrm{FP1}}(H_0, L_0)$ 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 3) 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 paper) are given in [5, appendix C]. We have the following proposition.

Proposition 6. There exists $\beta_2 \geqslant 0$, which does not depend on the duration of the high period and low period, such that

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

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

Proof. We need to consider the holding costs under cases 1–4 separately. Note that for cases 1–3, 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 3) 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 (30) below. The upper bounds on the cost differences for other values of the high and low periods are given in [5].

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 $Z_1(t) = Z_1(0) + (\lambda_1^{\mathrm{h}} - \mu_1)t$, and $Z_1(\psi_1) = \theta_1$ if $H_0 \geqslant \psi_1$. We consider the sample paths such that conditions of case 4.1 and $H_0 \geqslant \psi_1$ are both satisfied. Thus, we have

$$\psi_1 \leqslant H_0 \leqslant a_2, \qquad L_0 \geqslant \gamma_1(H_0 - a_1). \tag{29}$$

For (H_0, L_0) such that (29) 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 $Z_1(\psi_1) + (\lambda_1^{\mathrm{h}} - \mu_1)(H_0 - \psi_1) + (\lambda_1^{\mathrm{l}} - \mu_1)(t_2' - H_0) = \theta_1$. Since $Z_1(\psi_1) = \theta_1$, we obtain

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

Note that the conditions of case 4.1 imply that $t_2' \leqslant H_0 + L_0$. 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 (29) imply that $H_0 \geqslant \psi_2$. 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

$\mu_2(1-\rho_1^{\mathrm{l}})-\lambda_2$. Let $\tilde{t}_2'$ denote the time that class 2 decreases to its threshold value in the low period. Then

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

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

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

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

$$\psi_1\leqslant H_0\leqslant a_2,\qquad L_0\geqslant\gamma_4(H_0-a_1)\tag{30}$$

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

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

and

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

So, we can calculate the holding cost under the FP1 policy for (H_0,L_0) satisfying (30) as

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

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

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

Since $H_0 \leqslant a_2, c^{\mathrm{FP1}}(H_0, L_0) - c(H_0, L_0)$ is bounded when (H_0, L_0) satisfies (30). 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 (30). Expressions in [5] 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 (30)).

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

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

Definition 7. Assume conditions of case 4. We define the π^{a_1} policy as follows:

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

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 10 that when $\rho_2 \to 1 - \rho_1^1$ and $\mathbb{E}[L] \to \infty$, π^{a_1} is strongly asymptotically optimal under the assumptions of case 4. We use a notion of strongly asymptotically optimal (as introduced in [16]) in the following sense:

Definition 8. 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

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

We will also use the following notation.

Definition 9. For $f : \mathbb{R} \to \mathbb{R}$, we write

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

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 [5, appendix C].

Proposition 10. Assume conditions of case 4. Suppose H and L are random variables with $\mathbb{E}[H^2] < \infty$. If $\mathbb{E}[L] \to \infty$ and $\rho_2 \to (1 - \rho_1^{\mathrm{l}})$ (where ρ_1^{l} is fixed), then

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

and π^{a_1} is strongly asymptotically optimal.

Proof. Similar to the proof of proposition 6, 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 (35) below. The upper bounds on the cost differences for other values of the high and low periods are given in [5].

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 3.1). 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 $H \geqslant a_2$ and $\psi_1 \geqslant \psi_2$ imply that $H \geqslant \psi_1$. 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

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

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

$$t_2 = \big(\rho_1^{\mathrm{h}} - 1\big)\big(1 - \rho_1^{\mathrm{l}}\big)^{-1}(H - \psi_1) + H. \tag{31}$$

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

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

which is equivalent to

$$H \geqslant a_2, \qquad L \geqslant \gamma_3(H - \psi_1). \tag{32}$$

If H and L satisfy (32), the evolution of class 1 fluid under the optimal policy is as follows:

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

where t_2 is given in (31). The holding cost incurred by class 1 is given as

$$\int_0^{H+L} h_1\big(Z_1(t) - \theta_1\big)^+ \mathrm{d}t = \frac{1}{2} h_1 \mu_1 \frac{(\rho_1^{\mathrm{h}} - 1)(\rho_1^{\mathrm{h}} - \rho_1^{\mathrm{l}})}{(1 - \rho_1^{\mathrm{l}})} (H - \psi_1)^2. \tag{33}$$

Next we compute the holding cost incurred by class 2 under the optimal policy when H and L satisfy (32). 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 (32) 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 $\mu_2(1 - \rho_1^{\mathrm{l}}) - \lambda_2$ under the low-period-policy. Let $\tilde{t}_2$ denote the time that class 2 decreases to its threshold in the low period. Then

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

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

$$\tilde{t}_2 = t_2 + \frac{\rho_2}{1 - \rho_2 - \rho_1^{\mathrm{l}}}(t_2 - H), \tag{34}$$

where t_2 is given in (31). Note that $\tilde{t}_2 \leqslant H + L$ implies that $L \geqslant \gamma_4(H - a_1)$.

Thus, we consider the sample paths such that H and L satisfy both (32) and $L \geqslant \gamma_4(H - a_1)$, which is equivalent to

$$H \geqslant a_2, \qquad L \geqslant \gamma_4(H - a_1). \tag{35}$$

Now we can specify the evolution of class 2 fluid under the optimal policy if (H, L) satisfies (35). That is

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

where t_2 and $\tilde{t}_2$ are given in (31) and (34), respectively. The holding cost incurred by class 2 under the optimal policy if (H, L) satisfies (35) is given as

$$\int_0^{H+L} h_2\big(Z_2(t) - \theta_2\big)^+ \mathrm{d}t = \frac{1}{2} h_2 \mu_2 \big(\rho_2 (t_2 - \psi_2)^2 + \big(1 - \rho_2 - \rho_1^{\mathrm{l}}\big)\big(\tilde{t}_2 - t_2\big)^2\big). \quad (36)$$

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 (33) and (36) is equal to

$$\begin{aligned} c(H, L) = \frac{1}{2} h_2 \mu_2 \Bigg\{ & \frac{(\rho_1^{\mathrm{h}} - 1)(\rho_1^{\mathrm{h}} - \rho_1^{\mathrm{l}})}{\eta(1 - \rho_1^{\mathrm{l}})} (H - \psi_1)^2 \\ & + \frac{(1 - \rho_2 - \rho_1^{\mathrm{l}})^2}{\rho_2} \left[\frac{\rho_1^{\mathrm{h}} + \rho_2 - 1}{1 - \rho_1^{\mathrm{l}} - \rho_2} (H - a_1) - \frac{\rho_1^{\mathrm{h}} - 1}{1 - \rho_1^{\mathrm{l}}} (H - \psi_1) \right]^2 \\ & + \big(1 - \rho_1^{\mathrm{l}} - \rho_2\big) \left[\frac{\rho_1^{\mathrm{h}} + \rho_2 - 1}{1 - \rho_1^{\mathrm{l}} - \rho_2} (H - a_1) - \frac{\rho_1^{\mathrm{h}} - 1}{1 - \rho_1^{\mathrm{l}}} (H - \psi_1) \right]^2 \Bigg\}. \end{aligned} \quad (37)$$

The expression in (37) yields the lower bound for the holding cost if (H, L) satisfies (35).

Next we calculate the holding cost under the π^{a_1} policy when (H, L) belongs to the region in (35). Note that $H \geqslant a_2$ and $\psi_1 \geqslant \psi_2 \geqslant 0$ imply that $H \geqslant a_1$. According to the π^{a_1} policy, we know that both classes share the service capacity until a_1 as specified in definition 7. Since $a_1 \leqslant H$, 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 $\rho_1^{\mathrm{h}} > 1$, 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

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

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

$$t_2' = \frac{\rho_1^{\mathrm{h}} - 1}{1 - \rho_1^{\mathrm{l}}}(H - a_1) + H. \quad (38)$$

Since $t_2' \leqslant H + L$, $L \geqslant (\rho_1^{\mathrm{h}} - 1)(1 - \rho_1^{\mathrm{l}})^{-1}(H - a_1)$. Note that for every (H, L) that satisfies (35), this condition is satisfied. That is if (H, L) belongs to the region in (35), 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

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

where t_2' is given in (38). The holding cost incurred by class 1 under the π^{a_1} policy is equal to

$$\int_0^{H+L} h_1\big(Z_1(t) - \theta_1\big)^+ \mathrm{d}t = \frac{1}{2} h_1 \mu_1 \frac{(\rho_1^{\mathrm{h}} - 1)(\rho_1^{\mathrm{h}} - \rho_1^{\mathrm{l}})}{1 - \rho_1^{\mathrm{l}}} (H - a_1)^2. \tag{39}$$

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 $\mu_2(1 - \rho_1^{\mathrm{l}})$ under the low-period-policy. Since $\mu_2(1 - \rho_1^{\mathrm{l}}) > \lambda_2$, 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

$$Z_2(a_1) + \lambda_2\big(t_2' - a_1\big) + \big(\lambda_2 - \mu_2\big(1 - \rho_1^{\mathrm{l}}\big)\big)\big(\tilde{t}_2' - t_2'\big) = \theta_2,$$

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

$$\tilde{t}_2' = t_2' + \frac{\rho_2}{1 - \rho_2 - \rho_1^{\mathrm{l}}}\big(t_2' - a_1\big). \tag{40}$$

Since $\tilde{t}_2' \leqslant H + L$, $L \geqslant \gamma_4(H - a_1)$. For each sample path such that (H, L) satisfies (35), 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

$$\begin{array}{ll} \text{if } t \in [0, a_1], & Z_2(t) = Z_2(0) + (\lambda_2 - \mu_2 a_1)t \leqslant \theta_2, \\ \text{if } t \in \big(a_1, t_2'\big), & Z_2(t) = \theta_2 + \lambda_2(t - a_1) > \theta_2, \\ \text{if } t \in \big(t_2', \tilde{t}_2'\big), & Z_2(t) = Z_2\big(t_2'\big) + \big(\lambda_2 - \mu_2\big(1 - \rho_1^{\mathrm{l}}\big)\big)\big(t - t_2'\big) > \theta_2, \\ \text{if } t \in \big[\tilde{t}_2', H + L\big], & Z_2(t) \leqslant \theta_2, \end{array}$$

where t_2' and $\tilde{t}_2'$ are given in (38) and (40), respectively. The holding cost incurred by class 2 under the π^{a_1} policy when (H, L) satisfies (35) is equal to

$$\begin{aligned} \int_0^{H+L} h_2\big(Z_2(t) - \theta_2\big)^+ \mathrm{d}t &= \frac{1}{2} h_2 \mu_2 \big\{\rho_2\big(t_2' - a_1\big)^2 + \big(1 - \rho_2 - \rho_1^{\mathrm{l}}\big)\big(\tilde{t}_2' - t_2'\big)^2\big\} \\ &= \frac{1}{2} h_2 \mu_2 \frac{\rho_2(\rho_1^{\mathrm{h}} - \rho_1^{\mathrm{l}})^2}{(1 - \rho_1^{\mathrm{l}})(1 - \rho_1^{\mathrm{l}} - \rho_2)} (H - a_1)^2. \end{aligned} \tag{41}$$

Summing (39) and (41), we obtain the total holding cost under π^{a_1} when (H, L) satisfies (35) as

$$c^{a_1}(H,L) = \frac{1}{2}h_2\mu_2\left(\frac{(\rho_1^{\mathrm{h}}-1)(\rho_1^{\mathrm{h}}-\rho_1^{\mathrm{l}})}{\eta(1-\rho_1^{\mathrm{l}})} + \frac{\rho_2(\rho_1^{\mathrm{h}}-\rho_1^{\mathrm{l}})^2}{(1-\rho_1^{\mathrm{l}})(1-\rho_1^{\mathrm{l}}-\rho_2)}\right)(H-a_1)^2. \quad (42)$$

We can now compute the difference between the holding costs of the optimal policy and the π^{a_1} policy. Subtracting (37) from (42), we have

$$\begin{aligned} &c^{a_1}(H,L) - c(H,L) \\ &= \frac{1}{2}\mu_2 h_2\left\{\frac{(\rho_1^{\mathrm{h}}-1)(\rho_1^{\mathrm{h}}-\rho_1^{\mathrm{l}})}{\eta(1-\rho_1^{\mathrm{l}})}\big[(H-a_1)^2-(H-\psi_1)^2\big]\right. \\ &\quad + \frac{\rho_2(\rho_1^{\mathrm{h}}-\rho_1^{\mathrm{l}})^2}{(1-\rho_1^{\mathrm{l}})(1-\rho_1^{\mathrm{l}}-\rho_2)}(H-a_1)^2 - \frac{\rho_2(\rho_1^{\mathrm{h}}-1)^2}{(1-\rho_1^{\mathrm{l}})(1-\rho_1^{\mathrm{l}}-\rho_2)}(H-\psi_1)^2 \\ &\quad \left. - \frac{2\rho_2(\rho_1^{\mathrm{h}}-1)}{1-\rho_2-\rho_1^{\mathrm{l}}}(H-\psi_2)(H-\psi_1) - \frac{\rho_2(1-\rho_1^{\mathrm{l}})}{1-\rho_2-\rho_1^{\mathrm{l}}}(H-\psi_2)^2\right\}. \end{aligned} \quad (43)$$

First consider the last three terms in (43). Factoring out $\rho_2[(1-\rho_1^{\mathrm{l}})(1-\rho_1^{\mathrm{l}}-\rho_2)]^{-1}$, we can combine them into

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

Adding this value to the second term in (43) and taking the common factor $\rho_2[(1-\rho_1^{\mathrm{l}})(1-\rho_1^{\mathrm{l}}-\rho_2)]^{-1}$ out, we can combine all the terms with $(1-\rho_2-\rho_1^{\mathrm{l}})$ in the denominator into

$$\begin{aligned} &\frac{\rho_2}{(1-\rho_1^{\mathrm{l}})(1-\rho_1^{\mathrm{l}}-\rho_2)} \\ &\quad \times\big\{\big[(\rho_1^{\mathrm{h}}-\rho_1^{\mathrm{l}})(H-a_1)+(\rho_1^{\mathrm{h}}-1)(H-\psi_1)+(1-\rho_1^{\mathrm{l}})(H-\psi_2)\big] \\ &\quad \times\big[(\rho_1^{\mathrm{h}}-\rho_1^{\mathrm{l}})(H-a_1)-(\rho_1^{\mathrm{h}}-1)(H-\psi_1)-(1-\rho_1^{\mathrm{l}})(H-\psi_2)\big]\big\}. \end{aligned} \quad (44)$$

From the definitions of a_1, ψ_1, and ψ_2, we know that $a_1 = ((\rho_1^{\mathrm{h}}-1)\psi_1+\rho_2\psi_2)(\rho_1^{\mathrm{h}}+\rho_2-1)^{-1}$. Plugging in this expression of a_1, we can further simplify the expression in the second line of (44) as

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

Thus, we have

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

$$
- \frac{\rho_2(\rho_1^{\mathrm{h}} - \rho_1^{\mathrm{l}})(\rho_1^{\mathrm{h}} - 1)}{(1 - \rho_1^{\mathrm{l}})(\rho_1^{\mathrm{h}} + \rho_2 - 1)}\left(2H - a_1 - \frac{(\rho_1^{\mathrm{h}} - 1)\psi_1 + (1 - \rho_1^{\mathrm{l}})\psi_2}{\rho_1^{\mathrm{h}} - \rho_1^{\mathrm{l}}}\right)(\psi_1 - \psi_2)\Bigg\}
$$

$$
\leqslant \frac{1}{2}h_2\mu_2\left\{\frac{(\rho_1^{\mathrm{h}} - 1)(\rho_1^{\mathrm{h}} - \rho_1^{\mathrm{l}})}{\eta(1 - \rho_1^{\mathrm{l}})}(2H - a_1 - \psi_1)(\psi_1 - a_1)\right\},
$$

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

$$
\begin{aligned}
c^{a_1}(H, L) - c(H, L) &\leqslant \frac{1}{2}h_2\mu_2\left(\frac{(\rho_1^{\mathrm{h}} - 1)(\rho_1^{\mathrm{h}} - \rho_1^{\mathrm{l}})}{\eta(1 - \rho_1^{\mathrm{l}})}\right)2H^2 \\
&= \frac{1}{2}h_1\mu_1\left(\frac{(\rho_1^{\mathrm{h}} - 1)(\rho_1^{\mathrm{h}} - \rho_1^{\mathrm{l}})}{1 - \rho_1^{\mathrm{l}}}\right)2H^2, \qquad (45)
\end{aligned}
$$

where the equality follows from the definition of η. In [5], upper bounds similar to the one in (45) are also be obtained for other values of the high and the low periods (i.e. when H and L do not satisfy (35)). Since $E[H^2] \leqslant \infty$, we have the desired result. □

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

Definition 11. Assume conditions of case 2. We define the FP2-FP1 policy as follows:

$\forall t \in (0, H)$,	if $Z_2(t) > \theta_2$, $Z_1(t) < \theta_1$	then $\dot{T}_2(t) = 1$,	$\dot{T}_1(t) = 0$;
$\forall t \in (0, H)$,	if $Z_2(t) = \theta_2$, $Z_1(t) < \theta_1$	then $\dot{T}_2(t) = \rho_2$,	$\dot{T}_1(t) = 1 - \rho_2$;
$\forall t \in (0, H)$,	if $Z_1(t) \geqslant \theta_1$	then $\dot{T}_2(t) = 0$,	$\dot{T}_1(t) = 1$;
$\forall t \in (H, H + L)$		low-period-policy.	

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 $c^{\mathrm{FP2\text{-}FP1}}(H, L)$ 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 $c^{\mathrm{FP2\text{-}FP1}}(H, L)$ is given in [5, appendix C].

Proposition 12. Assume conditions of case 2. Suppose H and L are random variables with $\mathbb{E}[H^2] < \infty$. If $\mathbb{E}[L] \to \infty$ and $\rho_2 \to 1 - \rho_1^{\mathrm{l}}$ (where ρ_1^{l} is fixed), then

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

and FP2-FP1 policy is strongly asymptotically optimal.

Proof. 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 10, 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 (49) below and $\tilde{\psi}_1 \leqslant \tilde{\psi}_2$. The upper bounds on the cost differences for other values of the high and low periods are given in [5].

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 3.1). Note that in this case $Z_2(0) \geqslant \theta_2$ and $Z_1(0) \leqslant \theta_1$, hence $\psi_2 \leqslant 0 \leqslant \psi_1$. 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 $Z_1(t_1) = Z_1(0) + (\lambda_1^{\mathrm{h}} - \mu_1)t_1 = \theta_1$. Hence, $t_1 = \psi_1$, and the condition of case 2.6, in particular $H \geqslant a_2$, guarantees $\psi_1 \leqslant H$. Similar to the analysis in the proof of proposition 10, 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 $Z_1(H) \geqslant \theta_1$. 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

$$Z_1(t_1) + \left(\lambda_1^{\mathrm{h}} - \mu_1\right)(H - t_1) + \left(\lambda_1^{\mathrm{l}} - \mu_1\right)(t_2 - H) = \theta_1.$$

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

$$t_2 = H + \frac{\rho_1^{\mathrm{h}} - 1}{1 - \rho_1^{\mathrm{l}}}(H - \psi_1) = H + \gamma_3(H - \psi_1). \tag{46}$$

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

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

which is equivalent to

$$H \geqslant a_2, \qquad L \geqslant \gamma_3(H - \psi_1). \tag{47}$$

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

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

where t_2 is given in (46). We can calculate the holding cost incurred by class 1 under the optimal policy for each sample path such that (H, L) satisfies (47) and it is, in fact, the same as the one given in (39).

Now we analyze the evolution of class 2 fluid under the optimal policy when (H, L) satisfies (47). 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 $\mu_2(1-\rho_1^{\mathrm{l}})$ 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

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

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

$$\tilde{t}_2=t_2+\rho_2\big(1-\rho_2-\rho_1^{\mathrm{l}}\big)^{-1}(t_2-\psi_2), \tag{48}$$

where t_2 is given in (46). Since $\tilde{t}_2\leqslant H+L$, $L\geqslant\gamma_4(H-a_1)$.

Now we consider sample paths such that (H,L) satisfies both (47) and $L\geqslant\gamma_4(H-a_1)$. Thus, (H,L) satisfies

$$H\geqslant a_2,\qquad L\geqslant\gamma_4(H-a_1). \tag{49}$$

For each sample path such that (H,L) satisfies (49), the evolution of class 2 fluid can be specified according to the optimal policy as follows

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

where t_2 and $\tilde{t}_2$ are given in (46) and (48), respectively. Then the holding cost incurred by class 2 under the optimal policy for each sample path with (H,L) satisfying (49) is equal to

$$\int_0^{H+L}h_2\big(Z_2(t)-\theta_2\big)^+\mathrm{d}t=\frac{1}{2}h_2\mu_2\big\{\rho_2(t_2-\psi_2)^2-\rho_2\psi_2^2+\big(1-\rho_2-\rho_1^{\mathrm{l}}\big)(\tilde{t}_2-t_2)^2\big\}. \tag{50}$$

Summing (39) and (50) and plugging in the expressions of t_2 and $\tilde{t}_2$, we have

$$\begin{aligned}
c(H,L)=\frac{1}{2}h_2\mu_2\Bigg\{&\frac{(\rho_1^{\mathrm{h}}-1)(\rho_1^{\mathrm{h}}-\rho_1^{\mathrm{l}})}{\eta(1-\rho_1^{\mathrm{l}})}(H-\psi_1)^2-\rho_2(\psi_2^-)^2\\
&+\frac{(1-\rho_2-\rho_1^{\mathrm{l}})^2}{\rho_2}\left[\frac{\rho_1^{\mathrm{h}}+\rho_2-1}{1-\rho_1^{\mathrm{l}}-\rho_2}(H-a_1)-\frac{\rho_1^{\mathrm{h}}-1}{1-\rho_1^{\mathrm{l}}}(H-\psi_1)\right]^2\\
&+\big(1-\rho_1^{\mathrm{l}}-\rho_2\big)\left[\frac{\rho_1^{\mathrm{h}}+\rho_2-1}{1-\rho_1^{\mathrm{l}}-\rho_2}(H-a_1)-\frac{\rho_1^{\mathrm{h}}-1}{1-\rho_1^{\mathrm{l}}}(H-\psi_1)\right]^2\Bigg\}.
\end{aligned} \tag{51}$$

Now we analyze the fluid level evolution under FP2-FP1 policy when (H,L) satisfies (49). 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

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

and we have $t_1' = \tilde{\psi}_1$ and $\tilde{t}_1' = \tilde{\psi}_2$. We first consider the case that $\tilde{\psi}_1 \leqslant \tilde{\psi}_2$, 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 $H \geqslant t_1'$, i.e. $H \geqslant \tilde{\psi}_1$. One can verify that if (H, L) satisfies (49), then $H \geqslant \tilde{\psi}_1$. Therefore, for any (H, L) that satisfies (49) and $\tilde{\psi}_1 \leqslant \tilde{\psi}_2$, 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

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

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

$$t_2' = H + (\rho_1^{\mathrm{h}} - 1)(1 - \rho_1^{\mathrm{l}})^{-1}(H - \tilde{\psi}_1) = H + \gamma_3(H - \tilde{\psi}_1). \tag{52}$$

If the fluid level of class 1 decreases to its threshold before the low period is over, $t_2' \leqslant H + L$. Thus, we need $L \geqslant \gamma_3(H - \tilde{\psi}_1)$. But for any sample path with (H, L) satisfying (49) and $\tilde{\psi}_1 \leqslant \tilde{\psi}_2$, $L \geqslant \gamma_3(H - \tilde{\psi}_1)$ holds.

If (H, L) satisfies (49) and $\tilde{\psi}_1 \leqslant \tilde{\psi}_2$, the evolution of class 1 fluid is given as

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

where t_2' is in by (52). The holding cost incurred by class 1 is equal to

$$\int_0^{H+L} h_1(Z_1(t) - \theta_1)^+ \mathrm{d}t = \frac{(\rho_1^{\mathrm{h}} - 1)(\rho_1^{\mathrm{h}} - \rho_1^{\mathrm{l}})}{1 - \rho_1^{\mathrm{l}}}(H - \tilde{\psi}_1)^2. \tag{53}$$

We now consider the evolution of class 2 fluid when (H, L) satisfies (49) and $\tilde{\psi}_1 \leqslant \tilde{\psi}_2$. 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 $\mu_2(1 - \rho_1^{\mathrm{l}})$. 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

$$Z_2(0) + (\lambda_2 - \mu_2)\tilde{\psi}_1 + \lambda_2(t_2' - \tilde{\psi}_1) + (\lambda_2 - \mu_2(1 - \rho_1^{\mathrm{l}}))(\tilde{t}_2' - t_2') = \theta_2.$$

From this equation, we can get

$$\tilde{t}_2' = t_2' + \big((1-\rho_2)\big(\tilde{\psi}_2 - \tilde{\psi}_1\big) + \rho_2\big(t_2' - \tilde{\psi}_1\big)\big)\big(1-\rho_2-\rho_1^{\mathrm{l}}\big)^{-1}, \tag{54}$$

where t_2' is given in (52). For class 2 fluid to decrease to its threshold level in the low period, we need to have $\tilde{t}_2' \leqslant H + L$, which requires that $L \geqslant \gamma_4(H - a_1)$.

Then the evolution of class 2 fluid under FP2-FP1 policy with (H, L) satisfying (49) and $\tilde{\psi}_1 \leqslant \tilde{\psi}_2$ is given as

$$\begin{aligned}
&\text{if } t \in \big[0, \tilde{\psi}_1\big], && Z_2(t) = Z_2(0) + (\lambda_2 - \mu_2)t \geqslant \theta_2,\\
&\text{if } t \in \big(\tilde{\psi}_1, t_2'\big), && Z_2(t) = Z_2\big(\tilde{\psi}_1\big) + \lambda_2\big(t - \tilde{\psi}_1\big) > \theta_2,\\
&\text{if } t \in \big(t_2', \tilde{t}_2'\big), && Z_2(t) = Z_2\big(t_2'\big) + \big(\lambda_2 - \mu_2\big(1-\rho_1^{\mathrm{l}}\big)\big)\big(t - t_2'\big) > \theta_2,\\
&\text{if } t \in \big[\tilde{t}_2', H+L\big], && Z_2(t) \leqslant \theta_2,
\end{aligned}$$

where t_2' and $\tilde{t}_2'$ are given in (52) and (54), respectively. The holding cost incurred by class 2 under the FP2-FP1 policy can be computed as

$$\begin{aligned}
&\int_0^{H+L} h_2\big(Z_2(t) - \theta_2\big)^+ \mathrm{d}t\\
&\quad = \frac{1}{2}h_2\mu_2\Bigg\{2\frac{(1-\rho_2)(\rho_1^{\mathrm{h}} - \rho_1^{\mathrm{l}})}{1-\rho_1^{\mathrm{l}}}\big(\tilde{\psi}_2 - \tilde{\psi}_1\big)\big(H - \tilde{\psi}_1\big) + \rho_2\bigg[\frac{\rho_1^{\mathrm{h}} - \rho_1^{\mathrm{l}}}{1-\rho_1^{\mathrm{l}}}\big(H - \tilde{\psi}_1\big)\bigg]^2\\
&\qquad + \big(1-\rho_2-\rho_1^{\mathrm{l}}\big)\bigg[\frac{\rho_1^{\mathrm{h}} + \rho_2 - 1}{1-\rho_2-\rho_1^{\mathrm{l}}}(H - a_1) - \frac{\rho_1^{\mathrm{h}} - 1}{1-\rho_1^{\mathrm{l}}}\big(H - \tilde{\psi}_1\big)\bigg]^2\\
&\qquad + (1-\rho_2)\big(2\tilde{\psi}_2 - \tilde{\psi}_1\big)\tilde{\psi}_1\Bigg\}.
\end{aligned} \tag{55}$$

Summing (53) and (55), we get the total holding cost under FP2-FP1 policy for each (H, L) that satisfies (49) and $\tilde{\psi}_1 \leqslant \tilde{\psi}_2$ as

$$\begin{aligned}
c^{\text{FP2-FP1}}(H, L) = \frac{1}{2}h_2\mu_2\Bigg\{&\frac{(\rho_1^{\mathrm{h}} - 1)(\rho_1^{\mathrm{h}} - \rho_1^{\mathrm{l}})}{\eta(1-\rho_1^{\mathrm{l}})}\big(H - \tilde{\psi}_1\big)^2 + (1-\rho_2)\big(2\tilde{\psi}_2 - \tilde{\psi}_1\big)\tilde{\psi}_1\\
&+ 2\frac{(1-\rho_2)(\rho_1^{\mathrm{h}} - \rho_1^{\mathrm{l}})}{1-\rho_1^{\mathrm{l}}}\big(\tilde{\psi}_2 - \tilde{\psi}_1\big)\big(H - \tilde{\psi}_1\big)\\
&+ \rho_2\bigg[\frac{\rho_1^{\mathrm{h}} - \rho_1^{\mathrm{l}}}{1-\rho_1^{\mathrm{l}}}\big(H - \tilde{\psi}_1\big)\bigg]^2 + \big(1-\rho_2-\rho_1^{\mathrm{l}}\big)\bigg[\frac{\rho_1^{\mathrm{h}} + \rho_2 - 1}{1-\rho_2-\rho_1^{\mathrm{l}}}(H - a_1)\\
&- \frac{\rho_1^{\mathrm{h}} - 1}{1-\rho_1^{\mathrm{l}}}\big(H - \tilde{\psi}_1\big)\bigg]^2\Bigg\}.
\end{aligned} \tag{56}$$

Subtracting (51) from (56), with some algebra we have

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

$$
\begin{aligned}
&- \rho_2(H-\psi_2)^2 - \frac{2\rho_2(\rho_1^{\mathrm{h}}-1)}{1-\rho_1^{\mathrm{l}}}(H-\psi_2)\big(H-\tilde{\psi}_1\big) \\
&- \frac{2(\rho_1^{\mathrm{h}}-1)^2}{1-\rho_1^{\mathrm{l}}}(H-\psi_1)\big(\psi_1-\tilde{\psi}_1\big) + (1-\rho_2)\big(2\tilde{\psi}_2-\tilde{\psi}_1\big)\tilde{\psi}_1 \\
&+ \frac{2(1-\rho_2)(\rho_1^{\mathrm{h}}-\rho_1^{\mathrm{l}})}{1-\rho_1^{\mathrm{l}}}\big(\tilde{\psi}_2-\tilde{\psi}_1\big)\big(H-\tilde{\psi}_1\big) + \rho_2\psi_2^2\bigg\} \\
\leqslant{}& \frac{1}{2}h_2\mu_2\bigg\{\bigg(\frac{(\rho_1^{\mathrm{h}}-1)(\rho_1^{\mathrm{h}}-\rho_1^{\mathrm{l}})}{\eta(1-\rho_1^{\mathrm{l}})} + \frac{(\rho_1^{\mathrm{h}}-1)^2}{1-\rho_1^{\mathrm{l}}}\bigg)\big(2H-\psi_1-\tilde{\psi}_1\big)\big(\psi_1-\tilde{\psi}_1\big) \\
&- \rho_2(H-\psi_2)^2 - \frac{2\rho_2(\rho_1^{\mathrm{h}}-1)}{1-\rho_1^{\mathrm{l}}}(H-\psi_2)\big(H-\tilde{\psi}_1\big) \\
&- \frac{2(\rho_1^{\mathrm{h}}-1)^2}{1-\rho_1^{\mathrm{l}}}(H-\psi_1)\big(\psi_1-\tilde{\psi}_1\big) + 2(1-\rho_2)\tilde{\psi}_2\tilde{\psi}_1 \\
&+ \frac{2(1-\rho_2)(\rho_1^{\mathrm{h}}-\rho_1^{\mathrm{l}})}{1-\rho_1^{\mathrm{l}}}\tilde{\psi}_2\big(H-\tilde{\psi}_1\big) + \rho_2\psi_2^2\bigg\}.
\end{aligned}
$$

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

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

Thus,

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

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

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

Then we have

$$\begin{aligned}
&c^{\mathrm{FP2\text{-}FP1}}(H,L)-c(H,L)\\
&\quad\leqslant \frac{1}{2}h_2\mu_2\left\{\left(\frac{(\rho_1^{\mathrm{h}}-1)(\rho_1^{\mathrm{h}}-\rho_1^{\mathrm{l}})}{\eta(1-\rho_1^{\mathrm{l}})}+\frac{(\rho_1^{\mathrm{h}}-1)^2}{1-\rho_1^{\mathrm{l}}}\right)\left(2H-\psi_1-\tilde{\psi}_1\right)\left(\psi_1-\tilde{\psi}_1\right)\right\}\\
&\quad\leqslant \frac{1}{2}h_1\mu_1\left\{\left(\frac{(\rho_1^{\mathrm{h}}-1)(\rho_1^{\mathrm{h}}-\rho_1^{\mathrm{l}})}{1-\rho_1^{\mathrm{l}}}+\frac{(\rho_1^{\mathrm{h}}-1)^2}{1-\rho_1^{\mathrm{l}}}\right)\left(2H-\psi_1-\tilde{\psi}_1\right)\left(\psi_1-\tilde{\psi}_1\right)\right\}, \quad (57)
\end{aligned}$$

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. In [5], we also obtain upper bounds similar to the one in (57) for other values of the high and the low periods (i.e. when H and L do not satisfy (49)). Since $E[H^2] \leqslant \infty$, we have the desired result. □

6. Numerical results

In this section, we provide numerical examples to demonstrate the performance of the discrete review policy described in section 4 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 3. 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 $\mathbb{E}[c(H, L)]$ 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 DSview1, and the discrete review policies implemented with the method given in (25) with $p = 0.25, 0.5$ and 0.75 will be called DSview2, DSview3, and DSview4, 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 as system I and system II, respectively. In system I, parameters are set as follow: $\theta_1 = 50$, $\theta_2 = 100$, $h_1 = 2$, $h_2 = 1$, $Z_1(0) = 0$, $Z_2(0) = 90$, $\rho_1^{\mathrm{h}} = 2$, $\rho_1^{\mathrm{l}} = 0.1$ and $\rho_2 = 0.4$. In system II, $\rho_1^{\mathrm{l}} = 0$ and $\rho_2 = 0.95$

and the remaining parameters remain the same. We consider four different distributions (referred 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 cases A and B are $1/2$ and 1, respectively. In our experiments, $\mathbb{E}[H]$ attains the values: 5, 12.5, 25, 37.5 and 50 and $\mathbb{E}[L]$ attains the values: 12.5, 25, 50 and 1000.

Under a specified distribution with fixed values of $\mathbb{E}[H]$ and $\mathbb{E}[L]$, we generate 500,000 sets of H and L values. For each set of H and L values, we compute $\mathrm{c}(H, L)$ (lower bound), $c^{\mathrm{FP1}}(H, L)$, $c^{a_1}(H, L)$ 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 1–4 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}, DSview1, DSview2, DSview3 and DSview4 policies.

As tables 1–4 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 $\mathbb{E}[H]$ is small to moderate. The same result also holds for case C when $\rho_1^1 = 0$ 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 $\rho_1^1 = 0.1$ and $\rho_2 = 0.4$, the discrete review policies outperform the π^{a_1} policy. For systems with $\rho_1^1 = 0$ and $\rho_2 = 0.95$, the same observation holds for the DSview1, DSview3 and DSview4 policies. If $\rho_1^1 = 0$ and $\rho_2 = 0.95$, DSview2 has higher holding cost than π^{a_1} policy in cases A and B when $\mathbb{E}[H]$ is small and $\mathbb{E}[L]$ is not large or when $\mathbb{E}[L]$ is large.

In systems with $\rho_1^1 = 0.1$ 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 $\mathbb{E}[H]$ in case A. On the other hand, DSview2 significantly outperforms DSview1 and DSview3 policies in cases A and B and in case C when $\mathbb{E}[H]$ is not large. In case C, as $\mathbb{E}[H]$ increases, DSview1 policy starts dominating the other discrete review policies. On the other hand, in case D, DSviewl policy always outperforms the other discrete review policies in systems with $\rho_1^1 = 0.1$ and $\rho_2 = 0.4$. The same assertion holds for systems with $\rho_1^1 = 0$ and $\rho_2 = 0.95$ except when $E[L]$ and $E[H]$ are both large (see table 4).

Table 1
Average holding costs when $\mathbb{E}[L] = 12.5$. * indicates the actual value of the average holding cost for the FP1 policy.

System	Case	$\mathbb{E}[H]$	LB	Percentage differences off the lower bound					
				FP1	π^{a_1}	DSview1	DSview2	DSview3	DSview4
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

Table 2
Average holding costs when $\mathbb{E}[L] = 25$. * indicates the actual value of the average holding cost for the FP1 policy.

System	Case	$\mathbb{E}[H]$	LB	Percentage differences off the lower bound					
				FP1	π^{a_1}	DSview1	DSview2	DSview3	DSview4
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

Table 3
Average holding costs when $\mathbb{E}[L] = 50$. * indicates the actual value of the average holding cost for the FP1 policy.

System	Case	$\mathbb{E}[H]$	LB	Percentage differences off the lower bound					
				FP1	π^{a_1}	DSview1	DSview2	DSview3	DSview4
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

Table 4
Average holding costs when $\mathbb{E}[L] = 1000$. * indicates the actual value of the average holding cost for the FP1 policy.

System	Case	$\mathbb{E}[H]$	LB	Percentage differences off the lower bound					
				FP1	π^{a_1}	DSview1	DSview2	DSview3	DSview4
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

In systems with $\rho_1^1 = 0$ and $\rho_2 = 0.95$, the performances of DSview2 and DSview4 policies depend on the expected length of the low period. Even though the DSview4 policy shows poor performance (compared to the other discrete review policies) when $\mathbb{E}[L]$ is small, its performance improves (in particular, in cases A and B) as $\mathbb{E}[L]$ gets large. On the other hand, even though DSview2 policy has one of the best performances among the discrete review policies when $\mathbb{E}[L]$ is small, its performance deteriorates in cases A and B as $\mathbb{E}[L]$ gets large. However, in cases C and D, DSview1 and DSview2 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 DSview2 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, DSview1 policy seems to outperform the other discrete review policies. On the other hand, if class 2 is heavily loaded, DSview1 policy has a good overall policy.

7. 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 paper 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 paper. 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.

Appendix A. An optimal policy if $h_1\mu_1 \leqslant h_2\mu_2$

Under the assumption that the class 2 has constant arrival rate λ_2 and $\rho_2 < 1$, if $h_1\mu_1 \leqslant h_2\mu_2$, 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 appendix B as is done when the assumption in (5) holds.

- 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 $Z_1(t) \geqslant \theta_1$, 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 $Z_2(t) \leqslant \theta_2$ and $Z_1(t) \leqslant \theta_1$, 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 $\dot{T}_1(t) \geqslant \rho_1^1$, $\dot{T}_2(t) \geqslant \rho_2$ such that $\dot{T}_1(t) + \dot{T}_2(t) \leqslant 1$.

Appendix B. Proofs for the optimality of the policies in section 3

Before we prove the optimality of the policies given in section 3, we provide a lemma related to the Pontryagin maximum principle. Originally, this lemma was given in [15] 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:

$$\max \int_{B_0}^{B_1} f^0\big(x(t), u(t), t\big)\,\mathrm{d}t \tag{B.1}$$

such that

$$\dot{x}(t) = f\big(x(t), u(t), t\big), \tag{B.2}$$

$$x(B_0) = x_0, \tag{B.3}$$

$$x(B_1) \geqslant x_1, \tag{B.4}$$

$$u(t) \in \mathsf{U}, \quad \text{where } \mathsf{U} \subset \mathbb{R}^r \text{ and } \big(x(t), u(t)\big) \in \mathbb{R}^n \times \mathbb{R}^r, \tag{B.5}$$

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, $u(t)$ is piecewise continuous, and they satisfy (B.2)–(B.5). We want to find an optimal admissible pair $(x(t), u(t))$ that maximizes integral in (B.1). In the following lemma, for vectors a and b, $a \cdot b$ denotes the usual inner product of a and b.

Lemma 13. Let $(\bar{x}(t), \bar{u}(t))$ be an admissible pair for the problem given in (B.1)–(B.5). Suppose there exists a continuous function $p(t) = (p_1(t), p_2(t), \dots, p_n(t))$ 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

$$p_i(B_1) \geqslant 0 \quad \text{and} \quad p_i(B_1) = 0 \quad \text{if } \bar{x}_i(B_1) > x_1^i,\ \forall i = 1, \dots, n. \tag{B.6}$$

In addition, the Hamiltonian function

$$H\bigl(x(t), u(t), p(t), t\bigr) = f^0\bigl(x(t), u(t), t\bigr) + p(t) \cdot f\bigl(x(t), u(t), t\bigr) \tag{B.7}$$

satisfies the following

$$H\bigl(\bar{x}(t), \bar{u}(t), p(t), t\bigr) - H\bigl(x(t), u(t), p(x), t\bigr) \geqslant \dot{p}(t) \cdot \bigl(x(t) - \bar{x}(t)\bigr) \tag{B.8}$$

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 (B.1)–(B.5).

Proof. We use Δ to denote the following

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

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

According to (B.7) we have

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

It then follows from (B.2) and (B.8) that

$$\Delta \geqslant \int_{B_0}^{B_1} \dot{p}(t)\bigl[x(t) - \bar{x}(t)\bigr]\,\mathrm{d}t + \int_{B_0}^{B_1} p(t)\bigl[\dot{x}(t) - \dot{\bar{x}}(t)\bigr]\,\mathrm{d}t.$$

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

$$\begin{aligned}&\sum_{i=0}^{k} \left\{ \int_{\xi_i}^{\xi_{i+1}} \dot{p}(t)\bigl[x(t) - \bar{x}(t)\bigr]\,\mathrm{d}t + \int_{\xi_i}^{\xi_{i+1}} p(t)\bigl[\dot{x}(t) - \dot{\bar{x}}(t)\,\mathrm{d}t\bigr] \right\} \\ &\quad = \sum_{i=0}^{k} \int_{\xi_i}^{\xi_{i+1}} \frac{\mathrm{d}}{\mathrm{d}t}\bigl[p(t)\bigl(x(t) - \bar{x}(t)\bigr)\bigr]\end{aligned}$$

$$\begin{aligned}
&= \sum_{i=0}^{k} \left[p(\xi_{i+1})\big(x(\xi_{i+1}) - \bar{x}(\xi_{i+1})\big) - p(\xi_i)\big(x(\xi_i) - \bar{x}(\xi_i)\big)\right] \\
&= p(B_1)\big(x(B_1) - \bar{x}(B_1)\big) \\
&\geqslant 0,
\end{aligned}$$

where the last equality is due to the continuity of $p(t)$, $x(t)$, $\bar{x}(t)$ and (B.3), and the last inequality is based on (B.4) and (B.6). Hence, $\Delta \geqslant 0$, and the optimality of $(\bar{x}(t), \bar{u}(t))$ is proven. □

We next prove that the policy specified in section 3 is optimal for our original problem described in section 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

$$\max \int_0^{H+L} \sum_{i=1}^{2} -h_i\big(Z_i(t) - \theta_i\big)^+ \,\mathrm{d}t. \tag{B.9}$$

such that

$$\begin{aligned}
&\dot{Z}_i(t) = \lambda_i(t) - \mu_i u_i(t), && i = 1, 2, && \text{(B.10)}\\
&Z_i(t) \geqslant 0 && \forall t \in [0, H + L],\ i = 1, 2, && \text{(B.11)}\\
&u_i(t) \geqslant 0 && \forall t \in [0, H + L],\ i = 1, 2, && \text{(B.12)}\\
&u_1(t) + u_2(t) \leqslant 1 && \forall t \in [0, H + L], && \text{(B.13)}
\end{aligned}$$

where $\lambda_1(t) = \lambda_1^{\mathrm{h}}$, $\forall t \in (0, H)$, and $\lambda_1(t) = \lambda_1^{\mathrm{l}}$, $\forall t \in (H, H+L)$, 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 13, 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 (B.6) and (B.8). 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.

B.1. Proof for the optimality of the policy in section 3.1

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 $Z_1^*(t) \leqslant \theta_1$. Corresponding to this policy, we define two critical time instances for class 1 as follow

$$t_1 = \max\big\{t\colon\ s_2 \leqslant t \leqslant H,\ Z_1^*(t) \leqslant \theta_1\big\}, \tag{B.14}$$

$$t_2 = \max\big\{t\colon\ H \leqslant t \leqslant H + L,\ Z_1^*(t) \geqslant \theta_1\big\}, \tag{B.15}$$

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

$$\tilde{s}_2 = \max\{t\colon\ s_2 \leqslant t \leqslant t_2,\, Z_2^*(t) \leqslant \theta_2\}, \tag{B.16}$$

$$\tilde{t}_2 = \max\{t\colon\ t_2 \leqslant t \leqslant H + L,\, Z_2^*(t) \geqslant \theta_2\}, \tag{B.17}$$

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 3) and the definition of $t_1, t_2, \tilde{s}_2, \tilde{t}_2$, we claim the following holds:

Claim 1.

$$s_1 \leqslant s_2 \leqslant t_1 \leqslant H \leqslant t_2 \leqslant H + L,$$
$$s_1 \leqslant s_2 \leqslant \tilde{s}_2 \leqslant t_2 \leqslant \tilde{t}_2 \leqslant H + L.$$

Claim 2.

$$\begin{array}{lll}
\forall t \in (0, s_1) & Z_1^*(t) < \theta_1, & Z_2^*(t) > \theta_2,\\
\forall t \in (s_1, s_2) & Z_1^*(t) < \theta_1, & Z_2^*(t) \leqslant \theta_2,\\
\forall t \in (s_1, t_1) & Z_1^*(t) < \theta_1, & \\
\forall t \in (t_1, t_2) & Z_1^*(t) > \theta_1, & \\
\forall t \in (t_2, H+L) & Z_1^*(t) \leqslant \theta_1, & \\
\forall t \in (s_2, \tilde{s}_2) & Z_2^*(t) < \theta_2, & \\
\forall t \in (\tilde{s}_2, \tilde{t}_2) & Z_2^*(t) > \theta_2, & \\
\forall t \in (\tilde{t}_2, H+L) & Z_2^*(t) \leqslant \theta_2. &
\end{array}$$

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 $\dot{p}_i(t) = (\partial/\partial Z_i)H(Z(t), p(t), t)$ at the differentiable points, where the Hamiltonian function is given by

$$H\big(Z(t), u(t), p(t), t\big) = \sum_{i=1}^{2} \big(-h_i\big(Z_i(t) - \theta_i\big)^+ + p_i(t)\big(\lambda_i(t) - \mu_i u_i(t)\big)\big). \tag{B.18}$$

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

$$\begin{array}{lll} & p_i(H+L) = 0; & i = 1, 2, \\ \forall t \in (\tilde{t}_2, H+L): & \dot{p}_1(t) = 0, & \dot{p}_2(t) = 0, \\ \forall t \in (t_2, \tilde{t}_2): & \dot{p}_1(t) = \dfrac{\mu_2 h_2}{\mu_1}, & \dot{p}(t) = h_2, \\ \forall t \in (t_1, t_2): & \dot{p}_1(t) = h_1, & \\ \forall t \in (s_2, t_1): & \dot{p}_1(t) = 0, & \\ \forall t \in (\tilde{s}_2, t_2): & \dot{p}_2(t) = h_2, & \\ \forall t \in (s_2, \tilde{s}_2): & \dot{p}_2(t) = 0, & \\ \forall t \in (s_1, s_2): & \dot{p}_1(t) = 0; & \dot{p}_2(t) = 0, \\ \forall t \in (0, s_1): & \dot{p}_1(t) = 0; & \dot{p}_2(t) = h_2. \end{array}$$

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

Claim 3.

$$\begin{array}{ll} \forall t \in (t_2, H+L): & \mu_1 p_1(t) = \mu_2 p_2(t) \leqslant 0; \\ \forall t \in (s_2, t_2): & \mu_1 p_1(t) < \mu_2 p_2(t) \leqslant 0; \\ \forall t \in (s_1, s_2): & \mu_1 p_1(t) = \mu_2 p_2(t) \leqslant 0; \\ \forall t \in (0, s_1): & 0 \geqslant \mu_1 p_1(t) > \mu_2 p_2(t). \end{array}$$

Based on lemma 13, the optimality follows once we show that $(Z^*(t), u^*(t), p(t))$ satisfies (B.6) and (B.8). From the construction of $p_i(t)$, (B.6) holds immediately. It remains to show that (B.8) holds in each time interval throughout $(0, H+L)$ under all four cases given in (20)–(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 (B.18), we have the left-hand side of (B.8) equal to

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

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

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

Note that $u_1^*(t) + u_2^*(t) = 1$, and the admissible $u_i(t), i = 1, 2$, satisfies $u_1(t) + u_2(t) \leqslant 1$, so the right-hand side of the above inequality is non-negative. It follows immediately that (B.8) holds for all time t in the interval $(0, s_1)$.

Repeating this procedure for the remaining intervals, we can similarly prove that (B.8) 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)–(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 $\lambda_1^{\mathrm{h}}, \lambda_1^{\mathrm{l}}$ 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 $t_2 \leqslant H + L$ is equivalent to $L \geqslant \gamma_1(H - a_1)$. In addition, $t_1 \leqslant H \leqslant t_2$ is equivalent to $a_1 \leqslant H$, and $s_1 \leqslant s_2$ is equivalent to $H \leqslant B$. So, in case 2.1 of section 3.1, condition (10) guarantees that we have $0 \leqslant s_1 \leqslant s_2 \leqslant t_1 \leqslant H \leqslant t_2 \leqslant H + L$ and $u_1 > 0, u_2 > 0$.

Under the proposed policy, we know that $\lambda_1^{\mathrm{h}} > \mu_1$. 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 (B.14) and (B.15). 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,

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

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 $u_1^*(t) = \rho_1^{\mathrm{l}}$, and class 2 fluid is going to decrease by holding service capacity at $u_2^*(t) = 1 - \rho_1^{\mathrm{l}} > \rho_2$ until class 2 fluid reaches its threshold from above at $\tilde{t}_2$ (see the definition of $\tilde{t}_2$ given in (B.17)). After $\tilde{t}_2, u_1^*(t) > \rho_1^{\mathrm{l}}$ and

$u_2^*(t) > \rho_2$. So, fluid levels of both classes are going to decrease and are maintained below their thresholds. Hence,

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

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 $\dot{p}_i(t) \geqslant 0$ at all differentiable points, we know $p_i(t) \leqslant 0,\ i = 1, 2$, 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

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

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

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

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

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

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 $\mu_1 p_1(t_1) \leqslant \mu_2 p_2(t_1)$, we have

$$\forall t \in (s_2, t_1), \quad \mu_1 p_1(t) \leqslant \mu_2 p_2(t).$$

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

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

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

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

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

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

This completes the proof of claim 3. □

B.2. Proof for the optimality of the policy in section 3.2

We will only construct the auxiliary function $p_i(t)$, $i = 1, 2$. To complete the proof of (B.8), one only needs to go through the routine procedure as described in appendix B.1. We define t_2, $\tilde{s}_2$ and $\tilde{t}_2$ in the same way as in (B.15), (B.16) and (B.17) 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

$$\begin{array}{lllll}
\forall t \in (0, s_1): & Z_1^*(t) < \theta_1, & Z_2^*(t) > \theta_2, & \dot{p}_1(t) = 0, & \dot{p}_2(t) = h_2, \\
\forall t \in (s_1, s_2): & Z_1^*(t) < \theta_1, & Z_2^*(t) \leqslant \theta_2, & \dot{p}_1(t) = 0, & \dot{p}_2(t) = 0, \\
\forall t \in (s_2, \tilde{s}_2): & Z_2^*(t) < \theta_2, & \dot{p}_2(t) = 0, & & \\
\forall t \in (\tilde{s}_2, \tilde{t}_2), & Z_2^*(t) > \theta_2, & \dot{p}_2(t) = h_2, & & \\
\forall t \in (s_2, s_3): & Z_1^*(t) > \theta_1, & \dot{p}_1(t) = h_1, & & \\
\forall t \in (s_3, H): & Z_1^*(t) = \theta_1, & \dot{p}_1(t) = \dfrac{\mu_2 \dot{p}_2(t)}{\mu_1}, & & \\
\forall t \in (H, t_2): & Z_1^*(t) > \theta_1, & \dot{p}_1(t) = h_1, & & \\
\forall t \in (t_2, \tilde{t}_2): & Z_1^*(t) = \theta_1, & \dot{p}_1(t) = \dfrac{\mu_2 \dot{p}_2(t)}{\mu_1}, & & \\
\forall t \in (\tilde{t}_2, H + L): & Z_1^*(t) \leqslant \theta_1, & Z_2^*(t) \leqslant \theta_2, & \dot{p}_i(t) = 0, & i = 1, 2,
\end{array}$$

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 (B.6). □

B.3. Proof for the optimality of the policy in section 3.3

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.

References

[1] F. Avram, D. Bertsimas and M. Ricard, Fluid models of sequencing problems in open queueing networks: An optimal control approach, in: *Stochastic Networks*, eds. F.P. Kelly and R.J. Williams (Springer, New York, 1995) pp. 199–234.

[2] N. Bansal and M. Harchol-Balter, Scheduling solutions for coping with transient overload, Technical Report, CMU-CS-01-134, Department of Computer Science, Carnegie Mellon University (2001).

[3] N. Bauerle and U. Rieder, Optimal control of single-server fluid networks, Queueing Systems 35 (2000) 185–200.

[4] J.X. Chang, H. Ayhan, J. Dai, Z. Liu, M.S. Squillante and C.H. Xia, Optimal dynamic scheduling in a multiclass fluid model of Internet servers with transient overload, in: *Proc. of the 42nd IEEE Conf. on Decision and Control* (2003) pp. 721–727.

[5] J.X. Chang, H. Ayhan, J. Dai and C.H. Xia, Dynamic scheduling of a multiclass fluid model with transient overload, Technical Report, School of ISyE, Georgia Institute of Technology, `www.isye.gatech.edu/people/faculty/Hayriye_Ayhan` (2003).

[6] H. Chen and P. Mohapatra, Session-based overload control in QoS-aware Web servers, in: *Proc. of INFOCOM*, Vol. 2 (2002) pp. 516–524.

[7] A. Gajrat and A. Hordijk, Fluid approximation of a controlled multiclass tandem network, Queueing Systems 35 (2000) 349–380.

[8] T.C. Green and S. Stidham, Sample-path conservation laws, with applications to scheduling queues and fluid systems, Queueing Systems 36 (2000) 175–199.

[9] A. Iyengar, E. MacNair and T. Nguyen, An analysis of Web server performance, in: *Proc. of IEEE Global Telecommunications Conf.*, Vol. 3 (1997) pp. 1943–1947.

[10] G.P. Klimov, Time sharing service systems I, Theory Probab. Appl. 19(3) (1974) 532–551.

[11] Z. Liu, N. Niclausse and C. Jalpa-Villanueva, Traffic model and performance evaluation of Web servers, Performance Evaluation 46 (2001) 77–100.

[12] Z. Liu, M.S. Squillante and J.L. Wolf, Optimal control of resource allocation in e-business environments with strict quality-of-service performance guarantees, in: *Proc. of the 41st IEEE Conf. on Decision and Control*, Vol. 4 (2002) pp. 4431–4439.

[13] Z. Liu, M.S. Squillante, C.H. Xia, S.-Z. Yu and L. Zhang, Profile-based traffic characterization of commercial Web sites, in: *Proc. of the 18th Internat. Teletraffic Congress (ITC18)*, Berlin, Germany (2003) pp. 231–240.

[14] L.S. Pontryagin, V.G. Boltyanskii, R.V. Gamkrelidze and E.F. Mishchenko, *The Mathematical Theory of Optimal Processes* (Interscience, New York, 1962).

[15] A. Seierstad and K. Sydsater, Sufficient conditions in optimal control theory, *Internat. Economic Rev.* 18(2) (1977) 367–391.

[16] J.G. Shanthikumar and S.H. Xu, Strongly asymptotically optimal design and control of production and service systems, IIE Trans. 32(9) (2000) 881–890.

[17] W.E. Smith, Various optimizers for single-stage production, Naval Res. Logist. Quart. 3 (1956) 59–66.

[18] G. Weiss, Scheduling and control of manufacturing systems – a fluid approach, in: *Proc. of the 37th Allerton Conf.* (1999) pp. 577–586.

[19] C.H. Xia and J.G. Shanthikumar, Asymptotic optimal control of multi-class $G/G/1$ queues with feedback, in: *Stochastic Modeling and Optimization of Manufacturing Systems and Supply Chains*, eds. J.G. Shanthikumar, D.D. Yao and W.H.M. Zijm (2003) pp. 127–142.