DYNAMIC SCHEDULING OF A TWO-CLASS QUEUE WITH SETUPS

MARTIN I. REIMAN

Bell Laboratories, Lucent Technologies, Murray Hill, New Jersey

LAWRENCE M. WEIN

Massachusetts Institute of Technology, Cambridge, Massachusetts

(Received April 1994; revision received June 1995; accepted February 1996)

We analyze two scheduling problems for a queueing system with a single server and two customer classes. Each class has its own renewal arrival process, general service time distribution, and holding cost rate. In the first problem, a setup cost is incurred when the server switches from one class to the other, and the objective is to minimize the long-run expected average cost of holding customers and incurring setups. The setup cost is replaced by a setup time in the second problem, where the objective is to minimize the average holding cost. By assuming that a recently derived heavy traffic principle holds not only for the exhaustive policy but for nonexhaustive policies, we approximate (under standard heavy traffic conditions) the dynamic scheduling problems by diffusion control problems. The diffusion control problem for the setup cost problem is solved exactly, and asymptotics are used to analyze the corresponding setup time problem. Computational results show that the proposed scheduling policies are within several percent of optimal over a broad range of problem parameters.

We consider two dynamic scheduling problems for a single-server queueing system with two classes of customers. In both problems, each class possesses its own renewal arrival process, general service time distribution, and holding cost rate, and the server incurs a setup when switching from one class to the other. In the *setup cost* problem, a setup cost is incurred, and the objective is to minimize the long-run expected average setup and holding cost. In the *setup time* problem, a random setup time is incurred when the server switches class, and the objective is to minimize the long-run expected average holding cost. In both problems, the server has three options at each point in time: (1) serve a customer from the class that is currently set up, (2) switch to the other class (and immediately begin service in the setup cost problem), or (3) sit idle.

These scheduling problems have numerous applications, most notably for manufacturing systems and *polling systems* in computer communication networks. The setup time problem is more realistic than the setup cost problem in most situations, but is also more difficult to analyze. However, the setup cost problem is relevant for some manufacturing systems because, motivated by Just-In-Time (JIT) manufacturing, many facilities have *internalized* their setup times; i.e., they have essentially eliminated their setup times at the expense of incurring significant material, labor, and/or capital costs.

Although many studies have analyzed the performance of polling systems under various scheduling policies (see Takagi 1986, Boxma and Takagi 1992, and references therein), relatively few papers have considered the optimal scheduling of polling systems. The seminal paper in this research area is Hofri and Ross (1987), who analyze a two-class system with setup costs and times. Let c_i and μ_i denote the holding cost rate and service rate, respectively, for class i customers. When $c_1\mu_1 = c_2\mu_2$, they show that a double threshold policy, where the server serves each class until its queue is exhausted and the length of the other queue achieves a certain threshold level, minimizes the cost of setups and holding customers under both the discounted and average cost criteria. When $c_1\mu_1 \neq c_2\mu_2$, several authors have shown that the class with the larger $c\mu$ index should be served to exhaustion, and Koole (1994) has determined the asymptotic behavior of the switching curve for the setup cost problem in the discounted case; his result is described in more detail at the end of Section 1.

Several authors have studied the setup time problem in which more than two classes are present. Structural results for symmetric systems are derived by Liu et al. (1992) and references therein. Browne and Yechiali (1989) derive quasi-dynamic index policies, which allow the server to choose the sequence of classes to visit at the beginning of each cycle, that minimize or maximize the mean cycle length. Boxma et al. (1991) derive an efficient polling table (a predetermined fixed visit sequence) for minimizing the mean waiting cost. Bertsimas and Xu (1993) derive lower bounds and construct static policies that perform close to the bound when all classes have identical $c\mu$ indices. Van Oyen and Duenyas (1996) develop a dynamic scheduling heuristic based on myopic reward rates. Duenyas and Van Oyen (1995) also construct a dynamic policy for the setup cost problem.

Since the two-class asymmetric problem appears to be analytically intractable, heavy traffic approximations are employed in an attempt to make further headway, i.e., we

Subject classifications: Production/scheduling: stochastic scheduling with setups. Queues: diffusion models of control problems.
Area of review: Stochastic Models.

0030-364X/98/4604-0532 $05.00

make the *heavy traffic* assumption that the server must be busy the great majority of the time to satisfy demand. In the setup cost problem, we also need to assume that the setup costs are very large, roughly two orders of magnitude larger than the holding cost rate. Following in the tradition of Foschini (1977) and Harrison (1988), we study the diffusion control problem that arises as a heavy traffic limit of a sequence of queueing scheduling problems. These limiting control problems tend to be more tractable than their queueing counterparts and have led to network scheduling policies (see, for example, Harrison and Wein 1990 and Wein 1990b) that have a surprisingly simple form and appear to perform well.

For two-class queues operating under the exhaustive polling mechanism, Coffman et al. (1995, 1998) derive a heavy traffic averaging principle for systems with and without setup times, respectively. These averaging principles are due to a time scale decomposition that is *inherent in the heavy traffic scaling*, and so there is every reason to believe that these principles hold for more general (i.e., nonexhaustive) classes of policies. In this paper the approximating limiting control problems are obtained by assuming that the averaging principles hold for more general policies. Under this assumption, we show in Section 1 that the setup cost problem simplifies rather dramatically in the limiting heavy traffic regime: the dimension of the state space collapses from three (queue length of each class and the position of the server) to one (total workload). This result also allows our analysis to naturally decompose onto two different time scales. On the very fast time scale over which individual queue lengths change, we myopically optimize a control that specifies the amount of low-priority work to serve as a function of the total workload. This state-dependent control is derived in closed form and offers considerable insight. On the slower time scale over which the total workload varies, a singular control problem is solved that specifies a busy/idle policy. The solution to this control problem leads to a rather complex equation for one variable, which represents a threshold level, that can easily be solved numerically.

The setup time problem is addressed in Section 2, and the limiting control problem again is one-dimensional, although here we obtain an explicit diffusion control problem. The control, which represents the amount of low-priority work to serve as a function of the total workload, appears in the drift term of the diffusion process in a nonlinear fashion; consequently, the optimality equation leads to a nonlinear ordinary differential equation (ODE) that cannot be solved explicitly. However, we use asymptotics to obtain a scheduling policy; the asymptotics also reveal a substantial qualitative difference between the optimal policies in the setup cost and setup time cases.

For both problems, we use the value iteration algorithm to obtain "exact" optimal policies for a variety of test cases, and we show in Section 3 that the proposed policies perform within several percent of optimal over a broad range of problem parameters.

Our presentation of the analysis, and indeed the analysis itself, is rather informal throughout. For example, we do not prove that the limiting control problems are the heavy traffic limit of a sequence of queueing scheduling problems. Also, several of our claims regarding the nature of the limiting control problems and their optimal solutions are not proved. Providing a rigorous presentation of our results would be extremely demanding and would take us far afield from our two main objectives: (1) to obtain fundamental insights into the nature of the optimal policies under heavy traffic conditions, and (2) to develop effective scheduling policies for these systems. Much of our analysis, however, relies upon observations that have been rigorously proven for simpler systems, and we have no doubt that our results are essentially correct; consequently, we view (and refer to) insights gained from our analysis as insights into the nature of the optimal policy in heavy traffic. We hope that this approach increases the accessibility of the paper without sacrificing the persuasiveness of our arguments.

1. THE SETUP COST PROBLEM

1.1. Problem Description

Customers of class $i = 1, 2$ arrive according to independent renewal processes, where λ_i and c_{ai}^2 denote, respectively, the arrival rate and squared coefficient of variation (variance divided by the square of the mean) of the interarrival times. Each class has its own general service time distribution with service rate μ_i and squared coefficient of variation c_{si}^2, and we define the system's traffic intensity by $\rho = \sum_{i=1}^{2} (\lambda_i/\mu_i)$. A cost c_i is incurred per unit time for holding a class i customer in the system. A setup cost $K/2$ is imposed whenever the server switches from one class to the other, so that K is the setup cost per cycle.

The server has three scheduling options at each point in time: (1) serve the class that is currently set up, (2) switch to the other class and initiate service, or (3) sit idle. Since a switchover is instantaneous and costly, the option of switching to the other class and idling need not be considered. We assume that the server works in a preemptive-resume fashion, although the heavy traffic analysis is too crude to capture the effects of the nonpreemptive discipline as an alternative assumption. Let $Q_i(t)$ be the number of class i customers in queue or in service at time t, and let $J(t)$ denote the number of times the server sets up in the time interval $[0, t]$. Then our objective is to find a nonanticipating (with respect to the queue length process) scheduling policy to minimize

$$\limsup_{T\to\infty} \frac{1}{T} E\left[\int_0^T \sum_{i=1}^{2} c_i Q_i(t)\, dt + \frac{K}{2} J(T) \right]. \tag{1.1}$$

1.2. The Heavy Traffic Normalizations

A precise formulation of the approximating diffusion control problem requires much notation that would not be subsequently used. In addition, the limiting control problem will not be explicitly solved; rather, the averaging principle in (1.3) allows us to optimize over a specific form of policy that is introduced in subsection 1.4. Hence, the heavy traffic control problem will not be precisely formulated, and a description of the heavy traffic conditions and normalizations will suffice for our purposes.

The approximating control problem is the limit of a sequence of scheduling problems indexed by the heavy traffic scaling parameter n, where $n \to \infty$. Since a heavy traffic limit theorem will not be proved here, we avoid unnecessary notation by considering a single large integer n satisfying $\sqrt{n}(1 - \rho) = c$, where c is positive and of moderate size (that is, $O(1)$); this standard heavy traffic condition requires the server to be busy the great majority of the time over the long run. As we will see later, the scheduling policy that arises out of our heavy traffic analysis is independent of the system parameter n. Let V_i be the *unfinished workload* process for class i; $V_i(t)$ is the amount of time a continuously busy server requires to clear all of the class i customers who are present in the system at time t. The *normalized*, or scaled, queue length process is defined by $Z_i(t) = Q_i(nt)/\sqrt{n}$; similarly, $W_i(t) = V_i(nt)/\sqrt{n}$ denotes the normalized workload process. We approximate these normalized processes by the appropriate, and yet to be defined, limiting processes. Although $V_i(t)$ is not directly observable by the scheduler at time t, the normalized workload process is more convenient to employ than the normalized queue length process in the approximating heavy traffic control problem. However, we use the linear identity $Z_i = \mu_i W_i$ to translate the solution of the approximating control problem into a scheduling policy that is expressed in terms of the original queue length process (Q_1, Q_2). This linear identity is justified by extant heavy traffic limit theorems for many queueing systems.

In addition to speeding up time by a factor of n and reducing the queue lengths by a factor of $\sqrt{n}$, we also need to rescale the cost parameters c_i and K. The crux of problem (1.1) is the tradeoff between setup costs and holding costs, and to obtain a nontrivial solution to the approximating control problem, these two costs need to be of the same order of magnitude. Since only the ratio of these two costs matters, without loss of generality we leave the holding cost rates c_1 and c_2 unscaled at $O(1)$, and we only scale the setup cost K. The appropriate normalization is to divide the setup cost K by the heavy traffic scaling parameter n; readers are referred to subsection 1.2 of an earlier version of this paper (Reiman and Wein 1994) for details. Consequently, let $\kappa = K/n$ denote the normalized setup cost. Thus, heavy traffic conditions for the setup cost problem imply that the traffic intensity should be near one, and the setup cost should be large. A canonical example is to set $n = 100$ and set c, c_1, c_2, and κ all equal to one, so that $\rho = 0.9$ and the setup cost $K = 100$.

1.3. A Preliminary Heavy Traffic Result

The starting point for the setup cost problem is a recent heavy traffic result due to Coffman et al. (1995), which will be referred to hereafter as the CPR result. We present an informal statement of a special case of this heavy traffic limit theorem that will suffice for our purposes. As in problem (1.1), consider a queueing system with a single server and two customer classes. The CPR result is derived under a specific queue discipline: the server serves each class to exhaustion, and then switches class; although the averaging principle has been proven only under the exhaustive policy, we make the crucial but well-founded assumption that it also holds for nonexhaustive policies. The work conserving nature of the exhaustive discipline implies that the total workload process $W = W_1 + W_2$ is identical to the corresponding process under the FCFS policy. It follows from the heavy traffic limit theorem of Iglehart and Whitt (1970) that this process is well approximated under heavy traffic conditions by RBM$(-c, \sigma^2)$, which is a reflected Brownian motion (see Harrison 1985 for a definition) on $[0, \infty)$ with drift $-c$ and variance

$$\sigma^2 = \sum_{i=1}^{2} \frac{\lambda_i}{\mu_i^2} (c_{ai}^2 + c_{si}^2). \tag{1.2}$$

It turns out to be impossible to obtain a limit process for (W_1, W_2) in the usual sense, because in the heavy traffic limit the two-dimensional process moves back and forth along the constant workload line at an infinite rate, the direction being determined by which of the two queues is being served. However, this *time scale decomposition* is used to derive the following *averaging principle*: for any continuous function f and any $T > 0$,

$$\int_0^T f(W_i(t))\, dt \xrightarrow{\text{d}} \int_0^T \left(\int_0^1 f(uW(t))\, du \right) dt \tag{1.3}$$

$$\text{for } i = 1, 2 \quad \text{as } \rho \to 1.$$

Hence, given the normalized total workload W, the two-dimensional workload (W_1, W_2) can be treated as if it is uniformly distributed along the constant workload line from $(0, W)$ to $(W, 0)$. That is, the two-dimensional distribution is $(UW, (1 - U)W)$, where U is a uniform $[0, 1]$ random variable that is independent of W. This averaging principle allows us to *collapse the state space* of the control problem from three dimensions (the number of customers of each class in the system and the location of the server) to one dimension (the total workload).

1.4. The Form of the Optimal Policy

The traditional heavy traffic approach to scheduling problems is to precisely formulate the queueing system scheduling problem, find the limiting control problem that approximates the scheduling problem under heavy traffic conditions, and solve the latter problem. The averaging principle (1.3) allows us to take a slightly different approach: we first argue that the optimal policy should be of a specific form in the heavy traffic limit, and then we optimize the approximating system over this class of policies.

Without loss of generality, we assume that $c_1\mu_1 \geqslant c_2\mu_2$ and sometimes refer to classes 1 and 2 as the high- and low-priority classes, respectively. Existing results (Hofri and Ross for Poisson arrivals and exponential service times, and Duenyas and Van Oyen for Poisson arrivals and general service times) as well as intuition suggest that class 1 should be served to exhaustion. (It is possible to construct examples where this policy is not optimal. Our contention is that it is *asymptotically* optimal in the heavy traffic limit, where fine details, such as the exact nature of the service time distributions and the assumptions regarding preemption, are washed away. A policy is asymptotically optimal if the heavy traffic limit of its associated normalized cost coincides with the limit of optimal normalized costs.) When the server is set up for class 1, the only other decision is to specify whether the server should idle or switch to class 2 when no class 1 customers are present. Since we work with the normalized workload process (W_1, W_2), the only reasonable form of the optimal policy is to switch when $W_2(t) \geqslant w_2$ for some scaled threshold level w_2.

Since switching is instantaneous, $W_1(t) = 0$ and $W_2(t) = x$ at the moment of switching, where x must be greater than or equal to the threshold w_2. Because preemption is allowed, the server should never idle at class 2 when class 2 customers are present. The CPR result implies that the total workload $W = W_1 + W_2$ remains constant in the heavy traffic time scale while the server is serving class 2 customers. Hence, our decision can be expressed as the amount $u(x)$ by which the server depletes class 2's original work. That is, class 2 is served until $W_1(t) = u(x)$ and $W_2(t) = x - u(x)$. The control $u(x)$ must be between zero and x, where $u = x$ is the exhaustive policy. Since a different amount u can be chosen for each value of the total workload x, the control $u(x)$ can generate any possible switching curve in the nonnegative orthant, and so is without loss of generality.

Finally, since the server should never idle at class 2 when $W_2(t) > 0$, if $u(x) < x$, then the server immediately switches back to class 1 when $W_1(t) = u(x)$ and $W_2(t) = x - u(x)$. However, if $u(x) = x$, and hence class 2 is served exhaustively, then the server must decide whether to idle or to switch back to class 1. Once again, the obvious form of the optimal policy in this case is to idle until $W_1(t)$ is greater than or equal to w_1. Notice that if the threshold levels w_1 and w_2 were both zero, then infinite setup costs would be incurred.

In summary, the controls are the function $u(x)$, which specifies the amount of class 2's work to serve, and the threshold levels w_1 and w_2, which dictate the server's busy/idle policy. The form of the optimal policy in heavy traffic is: *serve class 1 until* $W_1(t) = 0$ *and* $W_2(t) \geqslant w_2$; *switch to class 2. If* $W_2(t) = x$ *at the moment of switching, then serve class 2 until* $W_1(t) = u(x)$ *and* $W_2(t) = x - u(x)$. *If* $u(x) < x$, *then switch to class 1; if* $u(x) = x$, *then do not switch until* $W_1(t) \geqslant w_1$.

1.5. An Overview of the Analysis

The analysis hinges on the following crucial observation: *since setups are instantaneous, the total workload process is only affected by the server's busy/idle policy, not by how often the server switches class*. Hence, the control $u(x)$ only influences the total workload indirectly via the idling. However, $u(x)$ does affect the rate at which holding costs and setup costs are incurred when the total workload is x. Therefore, a two-step procedure is employed to find the optimal policy $(u(x), w_1, w_2)$ within the specified form. In the first step, the control $u(x)$ is chosen to minimize the cost rate for each state x; this minimization is performed independently for each state x. In the second step, we attempt to find the optimal threshold levels w_1 and w_2, and hence the optimal total workload process. Our heavy traffic analysis will show that the optimal total workload process is a RBM$(-c, \sigma^2)$ on $[w, \infty)$, where w is a parameter that is chosen to minimize the total expected cost. Hence, the Brownian model is too crude to distinguish between the two thresholds w_1 and w_2, and so we set both w_1 and w_2 equal to the derived value of w.

1.6. The Optimal *u*(*x*)

The control $u(x)$ is chosen to minimize the cost rate that is incurred when the normalized total workload process is x. Under the policy characterized by $u(x)$, class 2's work is depleted by the amount $u(x)$ if the total workload when the server arrives to class 2 is x. The CPR result implies that, for our purposes, it is as if W_1 is uniformly distributed between 0 and $u(x)$, and W_2 is uniformly distributed between $x - u(x)$ and x. Since $Z_i = \mu_i W_i$, the holding cost rate when in state x can be expressed as

$$\sum_{i=1}^{2} c_i \mu_i E[W_i] = c_2 \mu_2 x + \frac{\Delta u(x)}{2}, \tag{1.4}$$

where $\Delta = c_1\mu_1 - c_2\mu_2$.

To find the setup cost rate when in state x, we need to find the cycle length. For a fixed total unfinished workload x, the two-dimensional workload process (W_1, W_2) moves back and forth at an asymptotically infinite rate along the line segment from $(0, x)$ to $(u(x), x - u(x))$. We determine the cycle length, and hence the setup cost rate, as a function of the normalized workload by slowing down time so that the two-dimensional workload moves at a finite and positive rate, the total workload stays fixed, and the movement of the two-dimensional workload is deterministic. If the server finds x units of work in class 2 upon arrival, then this work will be depleted at rate $1 - \rho_2$. The server works until $W_1(t) = u(x)$ and $W_2(t) = x - u(x)$, which occurs after $u(x)/(1 - \rho_2)$ time units. As we will see later, the normalized total workload process W never spends any time below $\max(w_1, w_2)$, and so we need not include any unnecessary inserted idle time into the cycle length calculation. Therefore, it takes $u(x)/(1 - \rho_1)$ time units to deplete class 1 and complete the cycle, resulting in a cycle of length $u(x)/(1 - \rho_2) + u(x)/(1 - \rho_1)$. Since the holding costs are estimated using a heavy traffic approximation, and the scheduling problem essentially trades off the setup and holding costs, a more accurate analysis results if we assume that $\rho = 1$ in our cycle length expression, which simplifies the cycle length to $u(x)/\rho_1\rho_2$. Because two setups are incurred in each cycle, the setup cost rate when in state x is $\rho_1\rho_2\kappa/u(x)$.

Now we find the optimal $u(x)$ by solving:

$$\min_{u(x)\in[0,x]} c_2\mu_2 x + \frac{\Delta u(x)}{2} + \frac{\rho_1\rho_2\kappa}{u(x)}. \tag{1.5}$$

If we define

$$\hat{w} = \sqrt{\frac{2\rho_1\rho_2\kappa}{\Delta}}, \tag{1.6}$$

then straightforward calculus leads to

$$u^*(x) = \min(x, \hat{w}). \tag{1.7}$$

Hence, $\hat{w}$ is the largest value of the total workload for which class 2 is served exhaustively. Notice that $\hat{w} = \infty$ when $\Delta = 0$, and so the optimal control in the balanced case is $u^*(x) = x$ for all x, which corresponds to exhaustive service for class 2.

1.7. The Optimal Threshold Level

In this subsection we analyze the normalized total workload process under the form of the proposed policy, using the control $u^*(x)$ in (1.7). This analysis shows that the total workload process W is a RBM$(-c, \sigma^2)$ on $[w, \infty)$, where w is a parameter that will be optimized over.

In the balanced case, the control $u^*(x)$ implies that the form of the optimal policy is to switch from class 1 to class 2 when $W_1(t) = 0$ and $W_2(t) \geqslant w_2$, and switch from class 2 to class 1 when $W_2(t) = 0$ and $W_1(t) \geqslant w_1$. Let us begin by assuming that $w_1 < w_2$. When the two-dimensional workload process hits the point $(x, 0)$, where $x \in [w_1, w_2)$, then the server will switch to class 1 and the process instantaneously moves to the point $(0, x)$. Since $x < w_2$, the server will not immediately switch back to class 2. Rather, the server serves newly arriving class 1 customers or sits idle until class 2's workload reaches w_2. In the heavy traffic limit, time is sped up by a factor of n, and the two-dimensional workload process instantaneously moves from the point $(0, x)$ to the point $(0, w_2)$; consequently, the total workload process never spends any time below the value of w_2. A similar argument when $w_1 > w_2$ implies that the total workload process is a RBM$(-c, \sigma^2)$ on $[\max(w_1, w_2), \infty)$. Thus, the heavy traffic analysis is too crude to distinguish between the thresholds w_1 and w_2, and we follow the convention of setting them both equal to w. Later in this subsection the cost minimizing value of w will be derived. Hence, the setup cost problem *decomposes* in the balanced case, and we can optimize over a single threshold parameter w independently of $u^*(x)$.

For the imbalanced case, the total workload process needs to be investigated under four different cases, depending upon the relative values of the normalized threshold levels w_1, w_2 and $\hat{w}$.

Case 1. $0 \leqslant w_1, w_2 \leqslant \hat{w}$. The curves for switching from class 2 to class 1 for all four cases are shown in Figure 1, where the vertical portion of the switching curve follows from (1.7). The argument put forth in the balanced case implies that the total workload process in this case is an RBM$(-c, \sigma^2)$ on $[\max(w_1, w_2), \infty)$. We again set w_1 and w_2 equal to the parameter w, and we model the optimal total workload process as an RBM$(-c, \sigma^2)$ on $[w, \infty)$. In this case, the parameter w is optimized over the region $0 \leqslant w \leqslant \hat{w}$.

Case 2. $\hat{w} \leqslant w_1, w_2$. The state $(w_1, 0)$ is never reached, and hence the parameter w_1 does not play a role here. By a similar argument as above, W is an RBM$(-c, \sigma^2)$ on $[w_2, \infty)$. Thus, once again, we set w_1 and w_2 equal to a parameter w, let W be an RBM$(-c, \sigma^2)$ on $[w, \infty)$, and optimize w over the region $w \geqslant \hat{w}$.

Case 3. $0 \leqslant w_1 \leqslant \hat{w} \leqslant w_2$. The total workload W is an RBM$(-c, \sigma^2)$ on $[w_2, \infty)$, and so we set w_1 and w_2 equal to w and optimize over $w \geqslant \hat{w}$. Thus, Case 3 reduces to Case 2.

Case 4. $0 \leqslant w_2 \leqslant \hat{w} \leqslant w_1$. The parameter w_1 is not a factor, and W is an RBM$(-c, \sigma^2)$ on $[w_2, \infty)$. Hence, Case 4 reduces to Case 1.

In summary, it suffices to restrict our attention to Cases 1 and 2. Thus, as in the balanced case, the single threshold parameter $w \geqslant 0$ can be optimized independently of $u^*(x)$.

We now derive the optimal value of the parameter w. Substituting the optimal control $u^*(x)$ from (1.7) into the cost rate function in (1.5) yields the optimal cost rate when the normalized workload is x, which is

$$c_2\mu_2 x + \frac{\Delta x}{2} + \frac{\rho_1\rho_2\kappa}{x} \quad \text{when } x \leqslant \hat{w}, \text{ and} \tag{1.8}$$

$$c_2\mu_2 x + \sqrt{2\rho_1\rho_2\Delta\kappa} \quad \text{when } x \geqslant \hat{w}. \tag{1.9}$$

To find the total expected average cost, the optimal cost rate is integrated over the steady-state distribution of the total workload process. The normalized workload process is approximated by an RBM$(-c, \sigma^2)$ on $[w, \infty)$, which has stationary density function $\alpha e^{-\alpha(x-w)}$ for $x \geqslant w$, where $\alpha = 2c/\sigma^2$.

If $w \geqslant \hat{w}$, then the total expected cost is $C(w) = c_2\mu_2(w + \alpha^{-1}) + \sqrt{2\rho_1\rho_2\Delta\kappa}$, which is increasing in w. Therefore, the optimal value of w is less than or equal to $\hat{w}$, and Case 1 of the previous subsection holds. Define the aggregate cost parameter $\bar{C} = (c_1\mu_1 + c_2\mu_2)/2$. Then the total expected cost equals

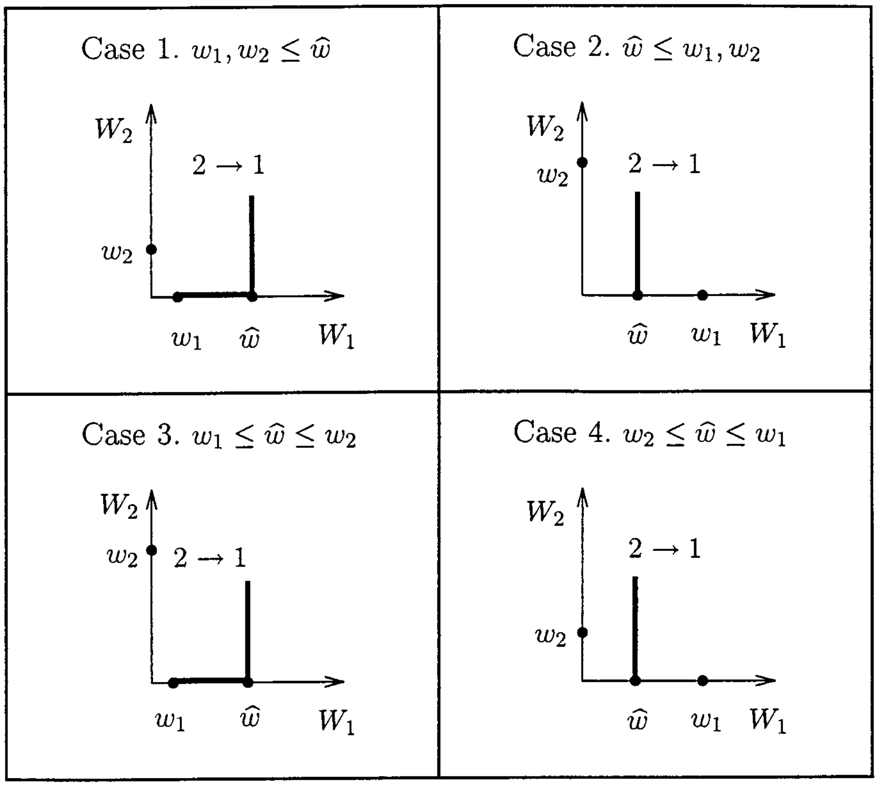

Figure 1. The total workload process for various values of w_1, w_2, and $\hat{w}$.

$$C(w) = \alpha e^{\alpha w}\bigg(\bar{C}\int_w^{\hat{w}} x e^{-\alpha x}\,dx + \rho_1\rho_2\kappa\int_w^{\hat{w}} \frac{e^{-\alpha x}}{x}\,dx + c_2\mu_2\int_{\hat{w}}^{\infty} x e^{-\alpha x}\,dx + \sqrt{2\rho_1\rho_2\Delta\kappa}\int_{\hat{w}}^{\infty} e^{-\alpha x}\,dx\bigg). \qquad (1.10)$$

Setting the derivative of the total expected cost with respect to w equal to zero yields

$$-\bar{C} = e^{\alpha(w-\hat{w})}\bigg(\alpha\sqrt{2\rho_1\rho_2\Delta\kappa} - \frac{\Delta(\alpha\hat{w}+1)}{2}\bigg) + \alpha^2\rho_1\rho_2\kappa\bigg(e^{\alpha w}(E_1(\alpha w) - E_1(\alpha\hat{w})) - \frac{1}{\alpha w}\bigg), \qquad (1.11)$$

where $E_1(x) = \int_x^\infty t^{-1}e^{-t}\,dt$, $x > 0$ is the exponential integral. It turns out that $C(w)$ is not convex; however, the solution to (1.11) is well-behaved numerically, and yields the global minimum of $C(w)$ for the cases we consider. We denote this solution by w^* and refer to it as the optimal threshold level. Since $C'(w) = \alpha C(w) - \alpha(\bar{C}w + \rho_1\rho_2\kappa w^{-1})$, it follows that the optimal total expected cost is $C(w^*) = \bar{C}w^* + \rho_1\rho_2\kappa w^{*-1}$.

In the balanced case where $\Delta = 0$, the first order condition (1.11) reduces to

$$\frac{1}{\alpha w} - e^{\alpha w}E_1(\alpha w) = \frac{c_1\mu_1}{\alpha^2\rho_1\rho_2\kappa}. \qquad (1.12)$$

Moreover, $C''(w) = \alpha^3\rho_1\rho_2\kappa(e^{\alpha w}E_1(\alpha w) + (\alpha w)^{-2} - (\alpha w)^{-1})$, and the convexity of $C(w)$ follows from the bound $e^x E_1(x) > 1/(x+1)$.

1.8. The Proposed Scheduling Policy

The heavy traffic solution is given by the control $u^*(x)$ defined in (1.7), which specifies a switching curve, and the threshold level w^* satisfying (1.11) or (1.12). We use this solution to propose a scheduling policy in terms of the three-dimensional state of the original problem, which is the two-dimensional queue length process (Q_1, Q_2), and the server location. Since both $u^*(x)$ and w^* are expressed in terms of the normalized workload W, several steps are required to translate this heavy traffic solution into a proposed policy. First, we reverse the heavy traffic scaling to express the quantities $u^*(x)$ and w^* in terms of the unscaled workload V. Since $W(t) = V(nt)/\sqrt{n}$, when the normalized workload W equals x, then the original workload V equals y, where $y = \sqrt{n}x$. The control $u^*(x)$ requires the server to serve class 2 until $W_1 = u^*(x)$, or equivalently, until $V_1/\sqrt{n} = u^*(y/\sqrt{n})$. If we substitute K/n for the normalized setup cost κ in (1.6), then when the total workload V equals y, class 2 is served until

$$V_1 = \min\bigg(y, \sqrt{\frac{2\rho_1\rho_2 K}{\Delta}}\bigg). \qquad (1.13)$$

By (1.13), class 2 is served exhaustively as long as the total workload V is less than or equal to

$$\hat{v} = \sqrt{\frac{2\rho_1\rho_2 K}{\Delta}}, \qquad (1.14)$$

which, not surprisingly, equals $\sqrt{n}\hat{w}$. Similarly, if we define the parameter $\theta = 2(1-\rho)/\sigma^2$ and the unscaled threshold

$v^* = \sqrt{n}w^*$, then substitution of $v/\sqrt{n}$ for w, K/n for κ, and $\sqrt{n}\theta$ for α in (1.11) and (1.12) yields, respectively,

$$-\bar{C} = e^{\theta(v-\hat{v})}\left(\theta\sqrt{2\rho_1\rho_2\Delta K} - \frac{\Delta(\theta\hat{v}+1)}{2}\right) + \theta^2\rho_1\rho_2 K\left(e^{\theta v}(E_1(\theta v) - E_1(\theta\hat{v})) - \frac{1}{\theta v}\right), \tag{1.15}$$

and

$$\frac{1}{\theta v} - e^{\theta v}E_1(\theta v) = \frac{c_1\mu_1}{\theta^2\rho_1\rho_2 K}. \tag{1.16}$$

Finally, the predicted optimal average cost for the original scheduling problem is:

$$\sqrt{n}C(w^*) = \bar{C}v^* + \frac{\rho_1\rho_2 K}{v^*}. \tag{1.17}$$

Notice that the quantities in (1.13)–(1.17) *are independent of the heavy traffic scaling parameter n*, and are expressed solely in terms of the primitive problem parameters.

Now that the optimal control has been translated into unscaled workloads, we use the simple heavy traffic relationship $\mu_i W_i = Z_i$ between workloads and queue lengths to express the switching curve and threshold level in terms of queue lengths. The only remaining hurdle is that the resulting quantities are continuous, whereas the two-dimensional queue length process resides on a lattice. We naively ignore this difference between our continuous solution and the discrete state space, which essentially amounts to rounding the threshold level up to the next highest integer and rounding the switching curve out to the next largest lattice points. In addition to being the most natural translation of the continuous solution, it also prevents us from rounding a threshold level down to zero, where infinite setup costs would be incurred.

In the balanced case, the critical value $\hat{v}$ in (1.14) equals infinity, which corresponds to exhaustive service. The proposed policy is: *when $Q_1(t) = 0$ and $Q_2(t) \geqslant \mu_2 v^*$, then switch from class 1 to class 2; when $Q_2(t) = 0$ and $Q_1(t) \geqslant \mu_1 v^*$, then switch from class 2 to class 1.* The parameter v^* is the solution to (1.16). This policy is a special case of the double threshold policy introduced by Hofri and Ross, who prove that the optimal policy is of this form in the balanced case when arrivals are Poisson.

By (1.13), the proposed policy for the imbalanced case has a particularly simple form, and is pictured in Figure 2: *when $Q_1(t) = 0$ and $Q_2(t) \geqslant \mu_2 v^*$, then switch from class 1 to class 2. When $Q_1(t) \geqslant \mu_1\hat{v}$ or ($Q_2(t) = 0$ and $Q_1(t) \geqslant \mu_1 v^*$), then switch from class 2 to class 1.* The parameters $\hat{v}$ and v^* are defined in (1.14) and (1.15), respectively. Hence, the server switches to the high-priority class as soon as the queue length of that class grows to the level $\mu_1\hat{v}$. By (1.14), this critical level increases with the setup cost K and decreases as Δ, the $c\mu$ differential between the two classes, gets larger. It is worth noting that the heuristic policy of Duenyas and Van Oyen has the same general form as our proposed policy.

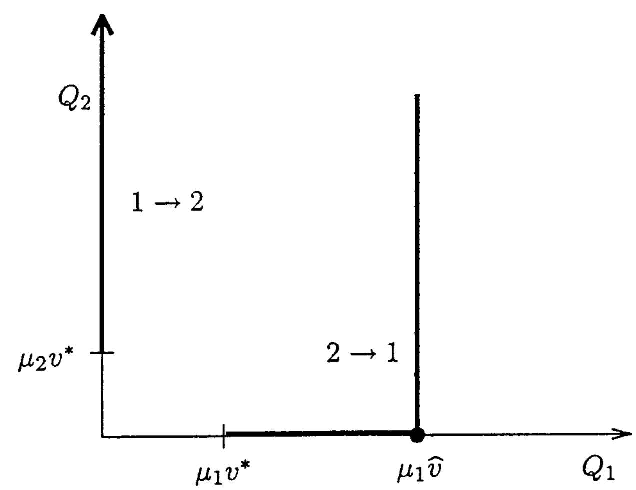

Figure 2. The proposed scheduling policy when $c_1\mu_1 > c_2\mu_2$.

The most intriguing aspect of Figure 2 is the switching curve from class 2 to class 1. In the Markovian discounted case, Koole derives the interesting asymptotic result that the upper part of this switching curve is, in fact, vertical; that is, the curve satisfies Q_1 equals a constant for all values of Q_2 larger than some finite quantity. Moreover, the numerically computed optimal switching curves in our computational study (see also Table 1 of Koole and the end of Section 3 of Duenyas and Van Oyen) behave as in Figure 2, *except* that the vertical part of the curve veers slightly to the right during its approach to the horizontal axis. The point of departure from a straight vertical line occurs relatively close to the horizontal axis (e.g., at $Q_2 = 3$ in Duenyas and Van Oyen's example, at $Q_2 = 2$ or 3 in most of our examples, and $Q_2 = 6$ in Koole's discounted example). It is important to note that existing numerical results are consistent with the conjecture that our proposed policy is asymptotically optimal (of course, not the unique optimum) in heavy traffic: the discrepancy between the form of the optimal switching curve and the form pictured in Figure 2, which typically affects only several states in the positive quadrant, vanishes in the heavy traffic limit because the point of departure from the vertical line (which is quantified in the parenthetic remark in the last sentence) becomes zero under the heavy traffic normalization. Hence, as in previous heavy traffic scheduling work, the heavy traffic analysis leads to a simple policy that *captures the key features* of the optimal policy and *ignores fine details* (such as the departure from the vertical switching curve near the horizontal axis) that do not have a significant effect on system performance when the system is congested.

2. THE SETUP TIME PROBLEM

2.1. Problem Description

The only difference between the setup time problem considered in this section and the setup cost problem is that a random setup time, rather than a setup cost, is incurred when the server switches from one class to the other. All relevant notation from the setup cost problem will be retained. By Coffman et al. (1998), the performance of this system in heavy traffic depends upon the setup time distributions only through the mean setup time per cycle, which we denote by s. The server has three scheduling options at each point in time: (1) serve a customer from the class that is currently set up, (2) initiate a setup, or (3) sit idle. The objective is to find a preemptive-resume, nonanticipating scheduling policy to minimize

$$\limsup_{T\to\infty} \frac{1}{T} E\left[\int_0^T \sum_{i=1}^2 c_i Q_i(t)\, dt \right]. \tag{2.1}$$

2.2. The Approximating Diffusion Control Problem

The setup times are not rescaled as the heavy traffic limit is approached; that is, we assume that the setup times are $O(1)$. The lack of setup costs has eliminated the incentive to insert unnecessary idleness in heavy traffic. Inserted idleness increases the workload, which in turn increases the holding costs. Hence, the proposed form of the optimal policy is simpler than in the setup cost problem: *serve class 1 to exhaustion and then set up for class 2. If class 2's normalized unfinished workload $W_2(t) = x$ at the setup completion epoch, then serve class 2 until $W_1(t) = u(x)$ and $W_2(t) = x - u(x)$, and immediately switch back to class 1.* As in the setup cost problem, the control $\{u(x), x \geq 0\}$ can generate any arbitrary switching curve in the nonnegative orthant.

Since the setup times are $O(1)$, switchovers occur instantaneously in the heavy traffic limit. Hence, the relationship between the two-dimensional normalized unfinished workload process (W_1, W_2) and the total normalized workload W still satisfies (1.3), just as in the setup cost problem.

Under the exhaustive service policy, Coffman et al. (1998) show that as $\rho \to 1$ the normalized total unfinished workload process W converges weakly to a diffusion process with drift $\rho_1\rho_2 s x^{-1} - c$ and variance σ^2 given by (1.2). They calculate the state-dependent drift by taking the limit as $n \to \infty$ of $\sqrt{n}(\rho - f(x))$, where $f(x)$ is the fraction of time that the server spends doing useful work when the normalized unfinished workload W equals x. Mimicking their calculations, we find that the corresponding drift for the controlled policy $u(x)$ is

$$\mu(x) = \frac{\rho_1 \rho_2 s}{u(x)} - c, \tag{2.2}$$

which agrees with their drift in the special case of exhaustive service (that is, $u(x) = x$ for all x). In summary, we approximate the normalized total unfinished workload process W by a $(\mu(x), \sigma^2)$ diffusion.

As we mentioned earlier, given $W(t) = x$, the two-dimensional process (W_1, W_2) behaves the same with or without setup times; hence, the holding cost rate when in state x is given by (1.4). Therefore, the approximating diffusion control problem is to choose $\{u(x), x \geq 0\}$ to minimize

$$\limsup_{T\to\infty} \frac{1}{T} E\left[\int_0^T \left(c_2 \mu_2 X(t) + \frac{\Delta u(X(t))}{2} \right) dt \right], \tag{2.3}$$

where X is a $(\mu(x), \sigma^2)$ diffusion process and $u(x) \in [0, x]$ for all $x \geq 0$.

The previous literature on heavy traffic approximations of queueing scheduling problems assumes zero setup times, and the time scale decomposition described in Section 1 leads to a deterministic pathwise optimization for the optimal queue length process, and a singular control problem for the optimal cumulative idleness process. The presence of setup times destroys this simplifying structure, and (2.3) provides the first example of a scheduling problem for a queueing system that is approximated in heavy traffic by a drift control problem.

2.3. Analysis of the Diffusion Control Problem

We begin by briefly discussing the balanced case where each class has the same $c\mu$ index. (Details of the analysis can be found in subsection 2.3 of Reiman and Wein.) Setting Δ equal to zero in (2.3) shows that the problem reduces to choosing $u(x)$ to minimize the mean of the stationary distribution of the diffusion process X. This goal is achieved by minimizing the drift $\mu(x)$ in (2.2), and hence the optimal control is $u(x) = x$ for all x; therefore, *the proposed scheduling policy for the balanced case is to serve each class to exhaustion, and immediately switch class*. The resulting diffusion process is a Bessel process with an additive drift, and the expected average cost is $\bar{C}(2\rho_1\rho_2 s + \sigma^2)/(2(1 - \rho))$, where the cost parameter $\bar{C}$ was defined earlier as $(c_1\mu_1 + c_2\mu_2)/2$.

For the balanced case, we can also introduce setup costs into the setup time problem without sacrificing tractability. For the balanced case with setup times and setup costs, the proposed scheduling policy is: *when $Q_1(t) = 0$ and $Q_2(t) \geq \mu_2 v^*$, switch from class 1 to class 2; when $Q_2(t) = 0$ and $Q_1(t) \geq \mu_1 v^*$, then switch from class 2 to class 1.* The threshold v^* is found by solving

$$\frac{(\theta v)^{\beta} e^{-\theta v}}{\Gamma(\beta + 1, \theta v)} = \left(1 - \frac{\beta}{\theta v}\right) + \frac{\beta \bar{C}}{\theta(\rho_1 \rho_2 \theta K - \beta \bar{C} v)}, \tag{2.4}$$

where $\beta = 2\rho_1\rho_2 s/\sigma^2$. Although we have not been able to prove the existence of a unique positive root v^* to (2.4), the numerical solution to this equation was well behaved for our test examples. The averages suboptimality of this policy on 45 test cases was 3.4%. (See subsection 3.2 of Reiman and Wein.)

For the remainder of this paper, we assume $c_1\mu_1 > c_2\mu_2$ in the setup time problem. Notice that (2.3) is nonstandard, in the sense that the drift is unbounded at zero and will be unbounded whenever the control $u(x) = 0$. Nonetheless, we proceed as if standard arguments apply (see, for example, Mandl 1968), and write the Hamilton-Jacobi-Bellman optimality equation for problem (2.3) as

$$\min_{u(x)\in[0,x]}\left\{c_2\mu_2 x+\frac{\Delta u(x)}{2}-g\right.$$
$$\left.+\left(\frac{\rho_1\rho_2 s}{u(x)}-c\right)V'(x)+\frac{\sigma^2}{2}V''(x)=0.\right. \tag{2.5}$$

Hence, if we can find a constant g, which is referred to as the *gain*, and a *potential* (relative value) function $V(x)$ that solves (2.5), then the control $u^*(x)$ that minimizes the expression in brackets in (2.5) is optimal, and g is the minimal average cost per unit time (independent of initial state). The resulting potential function $V(x)$ represents the cost incurred under the optimal policy when the initial state is x minus the cost incurred under the optimal policy when the initial state is zero. We assume that $V\in\mathbf{C}^2$ and, to avoid notational confusion between the potential function and the unscaled workload process, we employ the first derivative of the potential function, which is denoted by $p(x)=V'(x)$.

Rewriting (2.5) as

$$\min_{u(x)\in[0,x]}\left\{\frac{\Delta u(x)}{2}+\frac{\rho_1\rho_2 sp(x)}{u(x)}\right\}$$
$$+c_2\mu_2 x-g-cp(x)+\frac{\sigma^2}{2}p'(x)=0, \tag{2.6}$$

we obtain $u(x)=\sqrt{2\rho_1\rho_2 sp(x)/\Delta}$ from the first order optimality condition. Since greater initial workload implies greater cost, we have $p(x)>0$ and the function in brackets in (2.6) is convex with respect to $u(x)$. Hence, the optimal control is given by

$$u^*(x)=\min\left\{x,\ \sqrt{\frac{2\rho_1\rho_2 sp(x)}{\Delta}}\right\}. \tag{2.7}$$

It is interesting to compare (2.7) with the corresponding solution (1.6)–(1.7) in the setup cost problem. The solutions are identical, except that the normalized setup cost per cycle κ in (1.6) is replaced by the expected setup time per cycle s multiplied by $p(x)$. Hence, the two optimal controls will be qualitatively similar if the potential function $V(x)$ is linear, which will turn out not to be the case. Thus, solutions for the two problems lead to fundamentally different qualitative behavior.

We assume that $2\rho_1\rho_2 sp(x)/\Delta$ is monotone enough (e.g., p is nondecreasing) and is greater than x^2 as $x\to 0$, so that

$$u^*(x)=\begin{cases} x & \text{if } x\leq\hat{w},\\ \sqrt{\dfrac{2\rho_1\rho_2 sp(x)}{\Delta}} & \text{if } x\geq\hat{w}, \end{cases} \tag{2.8}$$

where the normalized threshold level $\hat{w}$ is unknown at this point and satisfies the fixed point equation

$$\hat{w}=\sqrt{\frac{2\rho_1\rho_2 sp(\hat{w})}{\Delta}}. \tag{2.9}$$

If we substitute (2.8) into (2.5), then the optimality equation reduces to two ordinary differential equations (ODE's) for $p(x)$

$$\frac{\sigma^2}{2}p'(x)+\left(\frac{\rho_1\rho_2 s}{x}-c\right)p(x)=g-\bar{C}x\quad\text{for } x\in[0,\hat{w}] \tag{2.10}$$

and

$$\frac{\sigma^2}{2}p'(x)-cp(x)+\sqrt{2\rho_1\rho_2\Delta sp(x)}$$
$$=g-c_2\mu_2 x\quad\text{for } x\geq\hat{w}. \tag{2.11}$$

The ODE in (2.10) is linear and possesses an explicit solution (that satisfies the properties assumed above). Unfortunately, the ODE in (2.11) is nonlinear and does not appear to admit an analytical solution. At this point, we can resort to approximate analytical methods or numerical methods. In the remainder of this section we derive an approximate analytical solution. However, we also employed the *Markov chain approximation* technique (see Kushner and Dupuis 1992) to numerically compute an optimal solution to the diffusion control problem. The resulting scheduling policy performed well in a computational study (it averaged 1.9% suboptimality over the 27 test cases in Table VI). However, the numerical procedure has three shortcomings when compared to the approximate analytical approach: (1) it did not perform quite as well in our computational study, (2) it is more difficult to implement, and (3) it is not independent of the heavy traffic parameter n. Consequently, for brevity's sake, we omit the numerical approach and refer readers to Section 2.6 of Reiman and Wein.

We conclude this subsection with an asymptotic result that forms the basis of our approximate analytical solution. Although (2.11) cannot be solved analytically, first hitting time arguments can be employed to obtain the asymptotic value of $p(x)$ as $x\to\infty$. A derivation in the appendix shows that the derivative of the potential function satisfies

$$p(x)=\frac{c_2\mu_2 x}{c}+o(x)\quad\text{as } x\to\infty. \tag{2.12}$$

This asymptotic result allows us to see how the control $u^*(x)$ behaves as $x\to\infty$. More specifically, (2.7) and (2.12) imply that

$$\frac{u^*(x)}{\sqrt{x}}\to\sqrt{\frac{2c_2\mu_2\rho_1\rho_2 s}{c\Delta}}\quad\text{as } x\to\infty. \tag{2.13}$$

This result is in direct contrast to the solution (1.6)–(1.7) of the setup cost problem, which implies that

$$u^*(x)\to\sqrt{\frac{2\rho_1\rho_2\kappa}{\Delta}}\quad\text{as } x\to\infty. \tag{2.14}$$

Equations (2.13)–(2.14) summarize the contrasting qualitative behavior between the solutions to the two problems: $u^*(x)$ grows as $\sqrt{x}$ in the setup time problem and is a constant for large x in the setup cost problem.

2.4. An Approximate Analytical Solution

One of our goals is to find a scheduling policy that performs well and is relatively easy to derive. Consequently, we investigate a simple class of policies, which we refer to as *asymptotic policies*; these policies can be constructed by patching together the asymptotic result (2.12) with the first part of solution (2.8). In particular, we assume that $u(x) = x$ for x less than or equal to some unknown threshold $\hat{w}$, and $u(x)/\sqrt{x}$ equals a constant thereafter; hence, we are assuming that the asymptotic result holds not only for very large x, but for all $x \geq \hat{w}$. Continuity at $\hat{w}$ gives

$$u(x) = \begin{cases} x & \text{if } x \leq \hat{w}, \\ \sqrt{\hat{w}x} & \text{if } x \geq \hat{w}. \end{cases} \quad (2.15)$$

This control, and hence the resulting scheduling policy, is characterized by a single parameter, the threshold level $\hat{w}$. Although the stability of the actual queue length process under the asymptotic policy may be difficult to ascertain, by (2.2) and (2.15) the corresponding diffusion process is positive recurrent for all $\hat{w} > 0$, because the drift is negative and monotone decreasing for all sufficiently large values of the normalized total workload.

We estimate $\hat{w}$ by assuming that this parameter satisfies the fixed point equation (2.9), and by approximating the unknown function $p(x)$ in this equation. In an attempt to refine (2.12), we *assume* that $p(x) = ax + b\sqrt{x} + o(\sqrt{x})$ as $x \to \infty$ for unknown constants a and b. Substituting this expression into the nonlinear ODE (2.11) and ignoring all $o(\sqrt{x})$ terms leads to

$$p(x) = \frac{c_2\mu_2 x}{c} + \sqrt{\frac{2c_2\mu_2\rho_1\rho_2\Delta s x}{c^3}} + o(\sqrt{x}) \quad \text{as } x \to \infty. \quad (2.16)$$

Substituting this expression into the fixed point Equation (2.9) yields

$$\frac{2c_2\mu_2\rho_1\rho_2 s}{\Delta c x} + \sqrt{\frac{8c_2\mu_2\rho_1^3\rho_2^3 s^3}{\Delta c^3 x^3}} = 1. \quad (2.17)$$

If we set $z = \sqrt{cx}$, then (2.17) becomes the cubic equation

$$\Delta z^3 - 2c_2\mu_2\rho_1\rho_2 s z - (2\rho_1\rho_2 s)^{3/2}\sqrt{c_2\mu_2\Delta} = 0. \quad (2.18)$$

Since $\sqrt{c\hat{w}} = \sqrt{(1-\rho)\hat{v}}$, it follows that the optimal unscaled threshold level $\hat{v}$ is $z^2/(1-\rho)$, where z solves (2.18).

As in Section 1, when the unscaled total workload V equals y, the control $u^*(x)$ requires the server to serve class 2 until the unscaled class 1 workload V_1 equals $\sqrt{n}u^*(y/\sqrt{n})$. Substituting $\hat{v}/\sqrt{n}$ for $\hat{w}$ in (2.15) gives

$$\sqrt{n}u^*\left(\frac{y}{\sqrt{n}}\right) = \begin{cases} y & \text{if } y \leq \hat{v}, \\ \sqrt{\hat{v}y} & \text{if } y \geq \hat{v}. \end{cases} \quad (2.19)$$

Translating workloads into queue lengths gives our proposed scheduling policy: *serve class 1 to exhaustion and then switch to class 2; serve class 2 until*

$$\mu_1^{-1}Q_1(t) \geq \begin{cases} \mu_1^{-1}Q_1(t) + \mu_2^{-1}Q_2(t) \\ \quad \textit{if } \mu_1^{-1}Q_1(t) + \mu_2^{-1}Q_2(t) \leq \hat{v}, \\ \sqrt{\hat{v}(\mu_1^{-1}Q_1(t) + \mu_2^{-1}Q_2(t))} \\ \quad \textit{if } \mu_1^{-1}Q_1(t) + \mu_2^{-1}Q_2(t) \geq \hat{v}, \end{cases} \quad (2.20)$$

and then switch back to class 1. This policy implies that class 2 is served to exhaustion as long as $\mu_1^{-1}Q_1(t) + \mu_2^{-1}Q_2(t) \leq \hat{v}$.

Table I
The 36 Test Cases for the Setup Cost Problem

	Holding Cost c_1	Setup Cost K	Traffic Intensity ρ
Balanced	1.0	—	—
Low	1.5	2	0.5
Medium	5.0	20	0.7
High	10.0	200	0.9

3. COMPUTATIONAL STUDY

A numerical experiment is undertaken in this section to investigate the effectiveness of our proposed policies. For both the setup cost and setup time problems, we compare the performance of the optimal policy, the proposed policy, and a straw policy. The straw policy for the setup cost problem is the *patient exhaustive* policy: *switch out of a class whenever it is exhausted and at least one customer of the other class is present*. The straw policy for the setup time problem is the *exhaustive* policy: *serve each class to exhaustion and then switch class*. These straw policies are studied because they are simple to implement in practice and are commonly found in the literature. The value iteration algorithm is used to derive optimal policies and to evaluate the cost of the proposed and straw policies. To simplify the computational effort, we assume that all interarrival times, service times, and setup times are exponential. The state space was truncated in the value iteration algorithm, and larger and larger state spaces were tested until the results were insensitive to increasing the state space. State spaces up to 90 by 90 and up to 4,000 value iterations were required to achieve three-digit accuracy. We report the *suboptimality* of the proposed and straw policies, where a

$$\text{policy's suboptimality} = \frac{\text{policy's cost} - \text{optimal cost}}{\text{optimal cost}} \times 100\%. \quad (3.1)$$

3.1. The Setup Cost Results

Forty test cases for the setup cost problem are considered, and we begin with the 36 test cases that are generated by all combinations of the parameter values in Table I; results for 12 additional test cases with $K = 10$ can be found in Reiman and Wein. For each test case in Table I, we set the service rates $\mu_1 = \mu_2 = 1$ and the arrival rates $\lambda_1 = \lambda_2 = \rho/2$, and let the holding cost $c_2 = 1$. Hence, each test case is characterized by the holding cost c_1 of the high-priority class, the setup cost per cycle K, and the traffic intensity ρ. This experimental design allows us to isolate the impact of three key parameters: (1) the difference in $c\mu$ values between classes, (2) the setup cost, and (3) the traffic intensity. Notice that nine cases in Table I are *balanced*, that is, $c_1\mu_1 = c_2\mu_2$, and 27 cases are imbalanced. Although our

Table II
Results for the Setup Cost Problem

Holding Cost c_1	Setup Cost K	Traffic Intensity ρ	Cost of Optimal Policy	Suboptimality of Proposed Policy	Suboptimality of Straw Policy
1.0	2	0.5	1.171	0.0%	0.0%
1.0	2	0.7	2.510	0.0%	0.0%
1.0	2	0.9	9.073	0.4%	0.4%
1.0	20	0.5	2.558	12.4%	5.9%
1.0	20	0.7	4.040	0.0%	1.6%
1.0	20	0.9	10.088	0.3%	0.3%
1.0	200	0.5	7.929	5.0%	128.2%
1.0	200	0.7	10.524	1.1%	90.3%
1.0	200	0.9	16.614	0.0%	21.9%
1.5	2	0.5	1.382	0.6%	2.8%
1.5	2	0.7	2.880	0.2%	7.4%
1.5	2	0.9	9.691	0.7%	17.2%
1.5	20	0.5	2.886	3.1%	3.4%
1.5	20	0.7	4.611	0.6%	1.6%
1.5	20	0.9	11.319	0.4%	9.3%
1.5	200	0.5	8.719	13.1%	110.4%
1.5	200	0.7	11.752	3.0%	75.4%
1.5	200	0.9	19.042	0.7%	18.2%
5.0	2	0.5	2.557	0.0%	24.0%
5.0	2	0.7	4.771	0.0%	50.4%
5.0	2	0.9	12.557	0.3%	115.4%
5.0	20	0.5	4.368	4.6%	7.8%
5.0	20	0.7	7.200	0.4%	21.8%
5.0	20	0.9	15.408	1.0%	82.1%
5.0	200	0.5	11.789	27.8%	70.4%
5.0	200	0.7	17.331	14.3%	42.5%
5.0	200	0.9	28.555	1.5%	34.0%
10.0	2	0.5	4.224	0.0%	34.2%
10.0	2	0.7	7.463	0.0%	74.3%
10.0	2	0.9	16.647	0.2%	197.1%
10.0	20	0.5	6.144	1.5%	17.3%
10.0	20	0.7	10.008	0.1%	45.9%
10.0	20	0.9	19.548	0.2%	158.2%
10.0	200	0.5	14.945	30.7%	51.2%
10.0	200	0.7	22.815	11.9%	33.8%
10.0	200	0.9	37.144	1.6%	63.6%

proposed policy, which is described in subsection 1.8, was derived under heavy traffic conditions, the policy is tested with traffic intensities as low as 0.5, and with setup costs as small as one-tenth of the holding cost c_1. Table II provides the long-run average cost under the optimal policy, and the suboptimalities of the proposed policy and the patient exhaustive policy for the 36 test cases. These results are summarized in Tables III and IV to isolate the effects of the three key parameters. Each entry in Tables III and IV represents the average suboptimality of the 12 test cases (nine cases for the holding cost) that have a particular parameter equal to a particular value.

The proposed policy performs remarkably well: its overall average suboptimality is 3.2%, and its suboptimality is less than 1% for 23 of the 36 test cases, and is less than 3.1% for 28 of the 36 test cases. Under these 28 cases, comparison of the optimal switching curves (not displayed here) with the proposed switching curves shows that the two curves differ on at most several states in the state space. In particular, the vertical boundary in Figure 2 is very close to optimal in the imbalanced case. For the test cases in Table I, the quantity $\lceil v^* \rceil$ in (1.15)–(1.16) ranges from one to three, and $\lceil \hat{v} \rceil$ in (1.14) has a mean of 3 and varies from 1 to 13.

Table III
Average Suboptimality of the Proposed Policy: Setup Cost Problem

	Holding Cost c_1	Setup Cost K	Traffic Intensity ρ
Balanced	2.1%	—	—
Low	2.5%	0.2%	8.2%
Medium	5.5%	2.0%	2.6%
High	5.1%	9.2%	0.6%
Overall Average Suboptimality = 3.8%			

Recall that many of the 36 test cases grossly violate the heavy traffic conditions stated in subsection 1.2, which requires heavy loading and much larger setup costs than holding costs. Perhaps the case that comes closest to satisfying these conditions is $c_1 = 1$, $K = 200$ and $\rho = 0.9$, where the proposed policy is optimal. As in previous heavy traffic work (see, for example, Chevalier and Wein 1993), the performance of the proposed policy is relatively insensitive to the heavy traffic assumptions underlying the analysis: the average suboptimality of the proposed policy is 2.3% over the 16 cases in Table I that have the setup cost and the traffic intensity set at their medium or high levels. However, the suboptimality deteriorates to as high as 30% when the setup cost is large, the traffic intensity is low, and the holding cost is high. In fact, most of the suboptimality in the 48 cases occurs when $K = 200$: the average suboptimality for the 24 cases in which $K < 200$ is 1.1%. In summary, the proposed policy performs very well over a broad range of parameter values, and then deteriorates outside of this range.

The patient exhaustive policy, with an average suboptimality of 45.0%, is clearly outperformed by the proposed policy. Not surprisingly, its performance degrades significantly as the holding cost c_1 and the traffic intensity ρ increase. Its suboptimality appears to be convex in the setup cost K. As K initially increases, holding costs play less of a role, and its suboptimality decreases; however, for very large K, the optimal policy idles much more than the

Table IV
Average Suboptimality of the Straw Policy: Setup Cost Problem

	Holding Cost c_1	Setup Cost K	Traffic Intensity ρ
Balanced	27.6%	—	—
Low	27.3%	43.6%	38.0%
Medium	49.8%	29.6%	37.1%
High	75.1%	61.7%	59.8%
Overall Average Suboptimality = 45.0%			

Table V
Results for the Setup Cost Problem: Miscellaneous Cases

Holding Cost c_2	Arrival Rate λ_1	Arrival Rate λ_2	Cost of Optimal Policy	Suboptimality of Proposed Policy	Suboptimality of Straw Policy
5	4.5	0.45	31.410	9.2%	49.7%
5	9/11	9/11	44.986	4.4%	37.2%
10	4.5	0.45	51.020	0.9%	18.9%
10	9/11	9/11	82.590	0.9%	8.6%

patient exhaustive policy, particularly when the traffic intensity is low.

Finally, to assess the proposed policy's performance on severely imbalanced problems, we consider four test cases that have $\mu_1 = 10$, $\mu_2 = 1$, $c_1 = 1$, $\rho = 0.9$, and $K = 200$. The results for these four test cases are displayed in Table V. In all four cases, the proposed policy performs much better than the straw policy.

3.2. The Setup Time Results

Table VI enumerates the 27 test cases for the imbalanced (that is, $c_1\mu_1 > c_2\mu_2$) problem with setup times; recall that Reiman and Wein contains results for 45 test cases for the balanced problem with setup costs and setup times. The results for the proposed policy and the exhaustive policy, which is the straw policy for these test cases, are given in Table VII and summarized in Tables VIII and IX.

Table VI
The 27 Test Cases for the Setup Times Problem

	Holding Cost c_1	Setup Time s	Traffic Intensity ρ
Low	1.5	2	0.5
Medium	5	10	0.7
High	10	20	0.9

The proposed policy performs very impressively on these test cases. The suboptimality is never above 5% and the average suboptimality over the 27 test cases is 1.5%. In contrast, the suboptimality for the exhaustive policy averages 8.7%. Not surprisingly, the policy's performance degrades when the holding cost c_1 is large and the setup times are small.

Table X gives results for four test cases that are identical to those in Table V, except that setup times (with $s = 10$) are incurred instead of setup costs. In contrast to Table V, where the optimal policy requires significant idling, both the proposed and straw policies are very close to optimal in all four cases.

Table VII
Results for the Setup Time Problem

Holding Cost c_1	Setup Time s	Traffic Intensity ρ	Cost of Optimal Policy	Suboptimality of Proposed Policy	Suboptimality of Straw Policy
1.5	2	0.5	1.845	1.9%	0.1%
1.5	2	0.7	4.348	0.3%	0.2%
1.5	2	0.9	16.814	0.0%	0.3%
1.5	10	0.5	4.336	2.1%	0.9%
1.5	10	0.7	10.195	0.2%	0.1%
1.5	10	0.9	39.342	0.0%	0.1%
1.5	20	0.5	7.470	1.2%	0.4%
1.5	20	0.7	17.494	0.1%	0.0%
1.5	20	0.9	67.062	0.0%	0.6%
5.0	2	0.5	3.918	1.4%	8.9%
5.0	2	0.7	9.047	0.2%	14.4%
5.0	2	0.9	33.873	1.6%	19.4%
5.0	10	0.5	9.946	4.9%	5.6%
5.0	10	0.7	23.671	0.5%	3.5%
5.0	10	0.9	90.943	0.2%	3.9%
5.0	20	0.5	17.292	5.0%	4.1%
5.0	20	0.7	41.195	0.5%	2.0%
5.0	20	0.9	154.934	0.1%	4.4%
10.0	2	0.5	6.422	1.4%	21.5%
10.0	2	0.7	14.532	1.8%	29.9%
10.0	2	0.9	52.846	4.5%	40.3%
10.0	10	0.5	17.144	4.7%	12.3%
10.0	10	0.7	40.895	0.8%	9.8%
10.0	10	0.9	154.388	0.7%	12.2%
10.0	20	0.5	30.095	4.8%	9.7%
10.0	20	0.7	72.431	0.4%	6.3%
10.0	20	0.9	240.781	0.6%	22.9%

Table VIII
Average Suboptimality of the Proposed Policy: Setup Time Problem

	Holding Cost c_1	Setup Time s	Traffic Intensity ρ
Low	0.6%	1.5%	3.1%
Medium	1.6%	1.6%	0.5%
High	2.2%	1.4%	0.9%
Overall Average Suboptimality = 1.5%			

Table IX
Average Suboptimality of the Straw Policy: Setup Time Problem

	Holding Cost c_1	Setup Time s	Traffic Intensity ρ
Low	0.3%	15.0%	7.0%
Medium	7.4%	5.4%	7.4%
High	18.3%	5.6%	11.6%
Overall Average Suboptimality = 8.7%			

4. CONCLUDING REMARKS

Using heavy traffic approximations, we analyze a dynamic scheduling problem for a two-class queue with either setup costs or setup times. As in previous heavy traffic scheduling studies, these approximations yield control problems that are more amenable to analysis than the original queueing control problems. Our analysis yields a simple two-parameter policy for the setup cost problem, where one parameter is found in closed form and the other is a solution to a specified equation; we conjecture that this policy is asymptotically optimal in heavy traffic. Although the diffusion control problem that approximates the setup time problem in heavy traffic is not explicitly solvable, a scheduling policy is constructed from an asymptotic result. Computational results indicate that our proposed policies are close to optimal over a broad range of parameter values, including some cases where the heavy traffic conditions are severely violated. An interesting implication of our analysis is that setup cost and setup time problems lead to fundamentally different qualitative solutions; see (2.13)–(2.14). Setup times eat into capacity in a nonlinear fashion, and hence setup costs cannot be used as a surrogate for setup times, as is sometimes done in deterministic scheduling problems with setups. (See, for example, the survey paper by Elmaghraby 1978.)

Markowitz (1996) has generalized our heavy traffic analysis from the two-class case to the multiclass setting. For both the setup cost and setup time problems, he restricts himself to cyclic policies (i.e., the server can idle, continue serving class k customers, or switch to class $k+1$ customers) and shows that *at most one class* (*the smallest* $c\mu$ *class*) *is served nonexhaustively in heavy traffic*. Notice that in the two-class setting, the focus of the setup cost problem is on determining when to serve nonexhaustively and when to idle, and the focus of the setup time problem is on nonexhaustion. Markowitz's results suggest that for setup time problems with many classes and in sufficiently heavy traffic, good noncyclic exhaustive policies, such as the polling tables developed by Boxma et al., should perform reasonably close to optimal.

APPENDIX

The goal in this appendix is to show that:

$$\lim_{x\to\infty} x^{-1} p(x) = \frac{c_2\mu_2}{c},$$

which is equivalent to (2.12). Since

$$p(x) = \lim_{\delta\to 0} \delta^{-1}[V(x+\delta) - V(x)], \tag{A.1}$$

we want to show that:

$$\lim_{x\to\infty}\lim_{\delta\to 0} \frac{V(x+\delta)-V(x)}{x\delta} = \frac{c_2\mu_2}{c}. \tag{A.2}$$

Thus we consider the quantity $V(x+\delta) - V(x)$. We can write

$$V(x+\delta) - V(x) = E_{x+\delta}\left[\int_0^{T_x}\left(c_2\mu_2 X(t) + \frac{\Delta u^*(X(t))}{2} - g\right) dt\right], \tag{A.3}$$

where T_x is the first hitting time of x for the $(\rho_1\rho_2 s/u^*(x) - c, \sigma^2)$ diffusion process X, and the expectation is with respect to the initial state $x+\delta$. Combining (A.1) and (A.3) yields

$$p(x) = \lim_{\delta\to 0}\frac{1}{\delta} E_{x+\delta}\left[\int_0^{T_x}\left(c_2\mu_2 X(t) + \frac{\Delta u^*(X(t))}{2} - g\right) dt\right]. \tag{A.4}$$

Table X
Results for the Setup Time Problem: Miscellaneous Cases

Holding Cost c_2	Arrival Rate λ_1	Arrival Rate λ_2	Cost of Optimal Policy	Suboptimality of Proposed Policy	Suboptimality of Straw Policy
5	4.5	0.45	191.924	0.1%	0.1%
5	9/11	9/11	118.720	0.6%	0.6%
10	4.5	0.45	256.383	0.3%	0.3%
10	9/11	9/11	170.706	0.7%	0.7%

To obtain the desired result, we need to first show that

$$\frac{u^*(x)}{x} \to 0 \quad \text{as } x \to \infty, \tag{A.5}$$

and

$$u^*(x) \to \infty \quad \text{as } x \to \infty. \tag{A.6}$$

These two asymptotic results will be derived in turn. Throughout this appendix we make the intuitively reasonable assumption that $u(x)$ is nondecreasing in x.

We prove (A.5) by contradiction, and hence initially assume that $\overline{\lim}_{x\to\infty} x^{-1}u^*(x) > 0$. Since $u^*(x) \in [0, x]$ for all $x \geq 0$, it follows that

$$p(x) \leq \overline{\lim_{\delta\to 0}}\left\{ \frac{\left(c_2\mu_2 + \frac{\Delta}{2}\right)}{\delta} E_{x+\delta}\left[\int_0^{T_x} X(t)\,dt\right] - \frac{g}{\delta} E_{x+\delta}[T_x] \right\}. \tag{A.7}$$

The assumed monotonicity of $u^*(x)$ yields $u^* \to \infty$, so that the drift of $X(t)$ satisfies:

$$\mu(x) = \frac{\rho_1\rho_2 s}{u^*(x)} - c \to -c \quad \text{as } x \to \infty. \tag{A.8}$$

Take x_0 large enough so that $\mu(x_0) \leq -c/2$. Note that $\mu(x) \leq -c/2$ for $x \geq x_0$. Let $\tilde{X}$ denote a $(-c/2, \sigma^2)$ Brownian motion, and $\tilde{T}_x$ its first passage time. For $x \geq x_0$, it follows that the integral in (A.7) has the bound

$$E_{x+\delta}\left[\int_0^{T_x} X(t)\,dt\right] \leq xE_{x+\delta}[\tilde{T}_x] + E_\delta\left[\int_0^{\tilde{T}_0} \tilde{X}(t)\,dt\right], \tag{A.9}$$

where $\tilde{T}_0$ is the first passage time to zero for a $(-c/2, \sigma^2)$ Brownian motion.

To evaluate the last term in (A.9), let

$$h(\delta) = E_\delta\left[\int_0^{\tilde{T}_{0\wedge b}} \tilde{X}(t)\,dt\right], \tag{A.10}$$

where $\tilde{T}_{0\wedge b}$ denotes the first hitting time for $\tilde{X}$ to either 0 or b. This function satisfies the ordinary differential equation (c.f. Karlin and Taylor)

$$-\frac{c}{2}h'(\delta) + \frac{\sigma^2}{2}h''(\delta) = -\delta, \tag{A.11}$$

subject to the boundary conditions $h(0) = h(b) = 0$, which yields

$$h(\delta) = \frac{2\sigma^2\delta}{c^2} + \frac{\delta^2}{c} + \frac{2(\sigma^2 b + b^2 c)(1 - e^{c\delta/\sigma^2})}{c^2(e^{cb/\sigma^2} - 1)}. \tag{A.12}$$

Therefore,

$$E_\delta\left[\int_0^{\tilde{T}_0} \tilde{X}(t)\,dt\right] = \lim_{b\to\infty} h(\delta) = \frac{2\sigma^2\delta}{c^2} + \frac{\delta^2}{c}. \tag{A.13}$$

Since

$$E_{x+\delta}[\tilde{T}_x] = \frac{2\delta}{c}, \tag{A.14}$$

it follows from (A.7), (A.9), and (A.13) that as $x \to \infty$,

$$p(x) \leq \lim_{\delta\to 0}\left(\frac{c_2\mu_2 + \frac{\Delta}{2}}{\delta}\right)\left(\frac{2x\delta}{c} + \frac{2\sigma^2\delta}{c^2} + \frac{\delta^2}{c}\right) = \left(c_2\mu_2 + \frac{\Delta}{2}\right)\left(\frac{2x}{c} + \frac{2\sigma^2}{c^2}\right). \tag{A.15}$$

Since $p(x)/x^2 \to 0$ as $x \to \infty$, by (2.7), we have $u^*(x)/x \to 0$ as $x \to \infty$, which is a contradiction; hence, (A.5) has been shown. An immediate consequence of (A.5) is

$$\frac{c_2\mu_2 x + \frac{\Delta u^*(x)}{2}}{x} \to c_2\mu_2 \quad \text{as } x \to \infty. \tag{A.16}$$

We next show (A.6), again by contradiction. Since we have assumed $u^*(x)$ nondecreasing, assuming that (A.6) does not hold is equivalent to assuming that $u^*(x)$ approaches some finite constant as $x \to \infty$, which we denote by $u^*(\infty)$. For large x, $X(t)$ behaves as a (μ, σ^2) Brownian motion, where $\mu = \rho_1\rho_2 s/u^*(\infty) - c$ could be of either sign. From (A.4) and the fact that $\rho_1\rho_2 s/u^*(\infty) \leq \rho_1\rho_2 s/u^*(x)$ we obtain

$$p(x) \geq \lim_{\delta\to 0}\frac{c_2\mu_2}{\delta} E_{x+\delta}\left[\int_0^{T_x} X(t)\,dt\right] - \lim_{\delta\to 0}\frac{g}{\delta}E_{x+\delta}[T_x] \tag{A.17}$$

$$\geq \lim_{\delta\to 0}\frac{c_2\mu_2 x - g}{\delta} E_\delta[T_0], \tag{A.18}$$

where T_0 is the first hitting time for a Brownian motion with drift μ and variance σ^2. If $\mu \geq 0$, then $E_\delta[T_0] = \infty$, and if $\mu < 0$, then $E_\delta[T_0] = -\delta/\mu$. Hence,

$$\lim_{x\to\infty} p(x) \geq \lim_{x\to\infty}\frac{c_2\mu_2 x - g}{|\mu|} = \infty. \tag{A.19}$$

Equations (A.19) and (2.7) imply that $u^*(x) \to \infty$, which yields the desired contradiction.

Armed with (A.5) and (A.6), we can now show (A.2). Equation (A.3) can be rewritten as

$$V(x+\delta) - V(x) = E_{x+\delta}\left[\int_0^{T_x} c_2\mu_2 X(t)\,dt\right] + E_{x+\delta}\left[\int_0^{T_x} \frac{\Delta u^*(X(t))}{2}\,dt\right] - gE_{x+\delta}[T_x]. \tag{A.20}$$

Since (A.6) implies (A.8), equations (A.9) and (A.14) implies that for $x \geq x_0$,

$$\frac{gE_{x+\delta}[T_x]}{\delta x} \leq \frac{2g}{cx},$$

which converges to zero as $x \to \infty$. Let

$$\epsilon_x = \sup_{z\geq x}\frac{u^*(z)}{z}.$$

By (A.5), $\epsilon_x \to 0$ as $x \to \infty$. For $x \geq x_0$ we can write

$$E_{x+\delta}\left[\int_0^{T_x} u^*(X(t))\,dt\right] \leq \epsilon_x E_{x+\delta}\left[\int_0^{T_x} X(t)\,dt\right]$$

$$\leq \epsilon_x\left[\frac{2\delta x}{c}+\frac{2\sigma^2\delta}{c^2}+\frac{\delta^2}{c}\right], \tag{A.21}$$

where the last inequality follows from (A.9), (A.13), and (A.14). Since $\epsilon_x \to 0$ as $x \to \infty$, it is clear that

$$\lim_{x\to\infty}\lim_{\delta\to 0}\frac{1}{x\delta}E_{x+\delta}\left[\int_0^{T_x}\frac{\Delta u^*(X(t))}{2}\,dt\right]=0.$$

We are, finally, faced with the first term on the right-hand side of (A.20), which is the only one that does not vanish. Fix x and let $\tilde{X}^{(x)}(t)$ denote a Brownian motion with (constant) drift $\mu(x) = \rho_1\rho_2 s/u^*(x) - c$, and (constant) variance σ^2. Let $\tilde{T}^{(x)}$ denote the first passage times for this process. As in (A.9), the monotonicity of $u^*(x)$ implies that

$$xE_{x+\delta}[\tilde{T}_x^{(\infty)}] + E_\delta\left[\int_0^{\tilde{T}_0^{(\infty)}}\tilde{X}^{(\infty)}(t)\,dt\right]$$

$$\leq E_{x+\delta}\left[\int_0^{T_x}X(t)\,dt\right]$$

$$\leq xE_{x+\delta}[\tilde{T}_x^{(x)}] + E_\delta\left[\int_0^{\tilde{T}_0^{(x)}}\tilde{X}^{(x)}(t)\,dt\right],$$

where $\tilde{X}^{(\infty)}$ is a Brownian motion with drift $-c$ and variance σ^2. Following the analysis that led to (A.13) and (A.14), we obtain

$$\frac{x\delta}{c}+\frac{\sigma^2\delta}{2c^2}+\frac{\delta^2}{2c}$$

$$\leq E_{x+\delta}\left[\int_0^{T_x}X(t)\,dt\right]$$

$$\leq \frac{x\delta}{\mu(x)}+\frac{\sigma^2\delta}{2\mu^2(x)}+\frac{\delta^2}{2\mu(x)}. \tag{A.22}$$

Since $\mu(x) \to -c$ as $x \to \infty$ by (A.6), we have

$$\lim_{x\to\infty}\lim_{\delta\to 0}\frac{E_{x+\delta}\left[\int_0^{T_x}X(t)\,dt\right]}{x\delta}=\frac{1}{c},$$

which yields

$$\lim_{x\to\infty}\lim_{\delta\to 0}\frac{V(x+\delta)-V(x)}{x\delta}=\frac{c_2\mu_2}{c},$$

by (A.20). This is what we set out to show.

ACKNOWLEDGMENT

The second author was supported by a grant from the Leaders for Manufacturing Program at MIT and NSF grant DDM-9057297. We are grateful to David Markowitz for performing the numerical computations reported in Section 3. We would like to thank Mike Harrison for suggesting the diffusion control approach to the setup time problem, and Larry Shepp for helpful discussions related to the diffusion control problem.

REFERENCES

BERTSIMAS, D. AND H. XU. 1993. Optimization of Polling Systems and Dynamic Vehicle Routing Problems on Networks. Sloan School of Management, MIT, Cambridge, MA.

BOXMA, O. J., H. LEVY, AND J. A. WESTRATE. 1991. Efficient Visit Frequencies for Polling Tables: Minimization of Waiting Cost. *Queueing Systems,* **9,** 133–162.

BOXMA, O. J. AND H. TAKAGI. 1992. Editors, Special Issue on Polling Systems. *Queueing Systems,* **11,** 1 and 2.

BROWNE, S. AND U. YECHIALI. 1989. Dynamic Priority Rules for Cyclic-Type Queues. *Adv. in Appl. Probab.* **21,** 432–450.

CHEVALIER, P. B. AND L. M. WEIN. 1993. Scheduling Networks of Queues: Heavy Traffic Analysis of a Multistation Closed Network. *Opns. Res*. **41,** 743–758.

COFFMAN, E. G., JR., A. A. PUHALSKII, AND M. I. REIMAN. 1995. Polling Systems with Zero Switchover Times: A Heavy-Traffic Averaging Principle. *Ann. Appl. Probab*. **5,** 681–719.

COFFMAN, E. G., JR., A. A. PUHALSKII, AND M. I. REIMAN. 1998. Polling Systems in Heavy-Traffic: A Bessel Process Limit. To appear in *Math. Opns. Res.*

DUENYAS, I. AND M. P. VAN OYEN. 1995. Stochastic Scheduling of Parallel Queues with Setup Costs. *Queueing Systems*, **19,** 421–444.

ELMAGHRABY, S. E. 1978. The Economic Lot Scheduling Problem (ELSP): Review and Extensions. *Mgmt. Sci.* **24,** 587–598.

FOSCHINI, G. J. 1977. On Heavy Traffic Diffusion Analysis and Dynamic Routing in Packet Switched Networks. In *Computer Performance*. K. M. Chandy and M. Reiser (eds.). North-Holland, Amsterdam.

HARRISON, J. M. 1985. *Brownian Motion and Stochastic Flow Systems*. John Wiley, New York.

HARRISON, J. M. 1988. Brownian Models of Queueing Networks With Heterogeneous Customer Populations. In *Stochastic Differential Systems, Stochastic Control Theory and Applications*, IMA Volume 10. W. Fleming and P. L. Lions (eds.), Springer-Verlag, New York, 147–186.

HARRISON, J. M. AND L. M. WEIN. 1990. Scheduling Networks of Queues: Heavy Traffic Analysis of a Two-Station Closed Network. *Opns. Res.* **38,** 1052–1064.

HOFRI, M. AND K. W. ROSS. 1987. On the Optimal Control of Two Queues with Server Set-up Times and its Analysis. *SIAM J. Comput.* **16,** 399–420.

IGLEHART, D. L. AND W. WHITT. 1970. Multiple Channel Queues in Heavy Traffic, I and II. *Adv. in Appl. Probab.* **2,** 150–177 and 355–364.

KOOLE, G. 1994. Assigning a Single Server to Inhomogeneous Queues with Switching Costs. To appear in *Theoretical Computer Sci.*

KUSHNER, H. J. AND P. G. DUPUIS. 1992. *Numerical Methods for Stochastic Control Problems in Continuous Time*. Springer-Verlag, New York.

Liu, Z., P. Nain, and D. Towsley. 1992. On Optimal Polling Policies. *Queueing Systems,* **11,** 59–83.

Mandl, P. 1968. *Analytical Treatment of One-Dimensional Markov Processes*. Springer-Verlag, New York.

Markowitz, D. 1996. A Unified Approach to Single Machine Scheduling: Heavy Traffic Analysis of Dynamic Cyclic Policies. Ph.D. thesis, Sloan School of Management, MIT, Cambridge, MA. In preparation.

Reiman, M. I. and L. M. Wein. 1994. Dynamic Scheduling of a Two-Class Queue with Setups. Technical Report, Sloan School of Management, MIT, Cambridge, MA.

Takagi, H. 1986. *Analysis of Polling Systems*. The MIT Press, Cambridge, MA.

Van Oyen, M. P. and I. Duenyas. 1996. Heuristic Scheduling of Parallel Heterogeneous Queues with Set-ups. *Mgmt. Sci.* **42,** 817–829.

Wein, L. M. 1990a. Optimal Control of a Two-Station Brownian Network. *Math. O. R.* **15,** 215–242.

Wein, L. M. 1990b. Scheduling Networks of Queues: Heavy Traffic Analysis of a Two-Station Network With Controllable Inputs. *Opns. Res.* **38,** 1065–1078.