

Control of Polling in Presence of Vacations in Heavy Traffic with Applications to Satellite and Mobile Radio Systems

Eitan Altman
INRIA
2004 Route des Lucioles
B.P. 93, 06902 Sophia-Antipolis Cedex
FRANCE

Harold J. Kushner*
Division of Applied Mathematics
and Lefschetz Center for Dynamical Systems
Brown University
Providence RI 02912
USA

June 30, 1999
Revised June, 2000

Abstract

Consider a queueing system with many queues, each with its own input stream, but with only one server. The server must allocate its time among the queues to minimize or nearly minimize some cost criterion. The allocation of time among the queues is often called polling and is the subject of a large literature. Usually, it is assumed that the queues are always available, and the server can allocate at will. We consider the case where the queues are not always available due to disruption of the connection between them and the server. Such occurrences are common in wireless communications, where any of the mobile sources might become unavailable to the server from time to time due to obstacles, atmospheric or other effects. The possibility of such “vacations” complicates the polling problem enormously. Due to the complexity of the basic problem we analyze it in the heavy traffic regime where the server has little idle time over the average requirements. It is shown that the suitable scaled total workloads converge to a controlled limit diffusion process with jumps.

*The work was supported by contracts DAAD 19-99-1-00223 from the Army Research Office and NSF grant ECS9703895.

The jumps are due to the effects of the vacations. The control enters the dynamics only via its value just before a vacation begins; hence it is only via the the jump value that the control affects the dynamics. This type of model has not received much attention. The individual queued workloads and job numbers can be recovered (asymptotically) from the limit scaled workload. This state space collapse is critical for the effective numerical and analytical work, since the limit process is one dimensional. It is also shown, under appropriate conditions, that the arrival process during a vacation can be approximated by the scaled “fluid” process. With a suitable nonlinear discounted cost rate, it is shown that the optimal costs for the physical problems converge to that for the limit problem as the traffic intensity approaches its heavy traffic limit. Explicit solutions are obtained in some simple but important cases, and it is shown that the $c\mu$ -rule is asymptotically optimal if there are no vacations. The stability of the queues is analyzed via a perturbed Liapunov function method. Finally, we extend the results to unreliable channels where the data might be received with errors and need to be retransmitted.

Key words: heavy traffic analysis, queueing networks, scheduling queues, communication networks, wireless communications, mobile communications, polling, optimal stochastic control

AMS classification numbers: 90B22, 60K25, 60K30, 60F17, 93E20, 93E25

1 Introduction

Consider a queueing system with several queues and a single server. The problem of assigning service among the competing queues in an optimal way has been studied extensively in the last half of the century, starting with [9] (p. 84-85). The assignment is often called “polling.” For linear holding costs, the fixed-priority policy known as the $c\mu$ rule (and other rules closely related to it) has been shown to be an optimal policy under a variety of statistical assumptions and cost structures (see e.g. [1, 3, 4, 9] and references therein). Due to Little’s rule, this policy turns out to minimize also the overall average expected waiting time in the system. The problem can be considered to be one in optimal stochastic control.

We are concerned with this assignment or polling problem when the connections between each queue and the server is broken at random times and for random durations. Such intervals during which a queue is not available to the server (even if it wished to poll it) are often called vacations in the queueing literature. The possibility of such vacations complicates the control problem considerably, since the possibility that any queue might not be available to the server at any future time needs to be accounted for in choosing the current server allocation.

This problem is of considerable importance in contemporary wireless communications, where the queues are in the mobile sources which generate data to be transmitted and the server is the channel or antenna of the base station. At each time, one assigns the channel to one of the sources (i.e., points the antenna

in the direction of that source). In more complex cases with so-called “smart antennas,” the channel can be shared among the sources in a controlled way, but that possibility will not be considered here.

Recently, Tassiulas et al. [28] have considered this problem of how to assign service to competing queues in the presence of random connectivity. The motivation concerned the dynamic assignment of transmission access to a channel between mobile terminals, any of which might be unavailable from time to time due to physical obstacles or to propagation problems (atmospheric attenuation, interference, noise, fading etc). In the context of satellite communications, a survey of such problems can be found in [11].

The classical $c\mu$ rule turns to be not only far from optimal for this system, but in fact the system may be unstable when any fixed priority policy is used. Tassiulas et al. [28, 29] considered the problem of obtaining a dynamic assignment policy that maximizes the throughput. The solution methodology is based on stability analysis using Liapunov functions; first a necessary condition for stability is identified, which holds under any policy. Then a particular policy is identified for which a sufficient stability condition coincides with the above necessary stability condition. It then turns out that this policy stabilizes the system under the largest range of input rates and is thus shown to maximize the throughput that the system can handle. Such a policy has a very simple form [28]: assign a transmission opportunity to the longest connected queue.

It turns out that there is a very large class of policies other than the above one which also achieve that maximum stability region and maximum throughput. (For example, if we first multiply the length of each queue by some (queue dependent) positive constant and then assign transmission to the queue with longest *weighted* length, still an optimal throughput is achieved; this is suggested by our stability analysis in Section 6.) The aim of this paper is thus to consider control and optimal control under more sensitive cost criteria, which can not only maximize the throughput but can also minimize some expected holding costs (or in particular, discounted mean values of a broad class of functions of the queue lengths or expected workload in the system).

Due to the complexity of the system and the generality of the statistical assumptions, we consider the optimal control problem only in an asymptotic sense: the one obtained by an appropriate scaling, corresponding to the heavy traffic regime, where the server has little spare capacity over the mean requirements. As usual with heavy traffic analysis, the limit system is substantially simpler than the original physical system. The aim is to use the limit model to get nearly optimal controls and approximations to the optimal value functions for the actual physical system under heavy traffic conditions. For some large class of policies, we establish the convergence of the total workload processes to a one dimensional diffusion process with jumps, where the jumps are due to the possibility that the server will have no work to do during part of a vacation of some source. The individual workloads and queue lengths can be approximated in terms of this limit process; hence there is a substantial reduction in dimension; from that of the original problem to unity. This limit result is then used to obtain a closed form solution for the asymptotic problem for several types

of cost functions. In particular, a closed form solution is obtained for the case when the cost corresponds to the total workload in the system.

A problem related to the one we solve here has been treated in [8] and references therein. There too, optimal scheduling of service opportunities is considered. However, the problem of random (unpredictable) disconnectivity is not considered there; instead, there are predictable instances in which service opportunities appear. These correspond to transmission opportunities between adjacent satellites which use inter-satellite links within a satellite constellation. As in [28], the criterion is to maximize the throughput. Several policies are proposed there and their performance is compared.

The structure of the paper is as follows. Section 2 describes the problem and lists the assumptions which are needed to get the weak convergence results of Section 3. Jobs (batches of data) arrive at the queues of the individual sources at random times, and in random amounts. It is assumed that the vacations are relatively “rare,” and that the ratio of the vacation intervals to the intervacation intervals is small. Nevertheless they have a very important effect on the performance. For notational simplicity, only two sources are considered in the details. Yet all of the results hold irrespective of the number of sources, and the extensions are commented on. The analysis contains results which are of broader interest. Examples are the proof that the arrival (and, indeed, the (suitably time scaled) workload and queue processes) process *during a vacation* can be well approximated by a “fluid” process under heavy traffic, and that the individual queue sizes can be approximated by linear functions of the individual queued workloads. The policy affects the limit process only via the magnitude of the jumps. In particular, the jumps depend only on the “control” values just before a vacation begins. If we restrict the policy to being a member of a large class of piecewise continuous feedback policies, then Theorem 3.2 shows that the individual workloads can be well approximated by linear functions of the total workload, under heavy traffic.

The discounted cost function is introduced in Section 4. The limit control problem appears to be nonstandard, owing to the special way that the control appears in the dynamics, even though it appears in the cost function in a standard way. To get weak convergence of the cost or optimal cost values, one needs a uniform integrability condition as well as weak convergence, and this is dealt with as well.

The optimal control is computed under some assumptions on the costs in Section 5. In Section 6 we establish the stability conditions for a large class of policies. We extend our model and results in Section 7, where we consider the case of unreliable channels which may require retransmissions of erroneous information.

2 The Problem Formulation

There are two sources with inputs which generate data in some random way. Any number of sources can be used, but we stick to two for notational simplicity. The

results for the general case will be apparent from the results for the two source case. The data (or jobs) created by each source are queued there, and the server alternates service (i.e., polling) between them in some way to be determined later. There is no switchover time in going from one source to another. Each of the sources will become unavailable to the server from time to time. The periods of unavailability are called vacations, in accordance with current terminology. Strictly speaking, it is the connection from the source to the server which is “on vacation,” but we simply say that the source itself is on vacation. A source that is on vacation cannot be polled, but the data or job inputs which it creates still arrive to its queue. In that case, the content of the unavailable queue grows, but the server can only work on the available queue. Service is nonpreemptive, and FCFS: I.e., a job once started is completed, assuming that there is no intervening vacation. Also, the system is assumed to be “work conserving” in that the server will not idle if there is work to do on an available queue. Suppose that a vacation starts in the middle of a job. We suppose *either* (a) that the job is allowed to be completed, or (b) that it is stopped, but when that source is next served the job needs only its residual time, or (c) that the entire job needs to be redone. Because of the “rarity” of vacations under the assumptions (A2.3) and (A2.4) to be introduced below, the results will be the same for all cases. For specificity, and without loss of generality in the results, we suppose that both sources are available at time 0.

Our approach is that of heavy traffic analysis, where the “spare capacity” of the system is small. As usual in heavy traffic analysis, the problem is embedded in a sequence of problems, indexed by n . As $n \rightarrow \infty$, the spare capacity of the server goes to zero, and this is quantified in (A2.2). Let $\{\Delta_{i,l}^{a,n}, l < \infty\}$ denote the interarrival times for jobs at source $i = 1, 2$, and $\{\Delta_{i,l}^{d,n}, l < \infty\}$ the corresponding work (real time) requirements.

Comment on weak convergence. Let \mathbb{R}^r denote r -dimensional Euclidean space. The path space for all of the processes will be $D(\mathbb{R}^r; 0, \infty)$, the space of \mathbb{R}^r -valued functions which are right continuous and have left hand limits, for the appropriate values of r . If $r = 1$, we write simply $D(\mathbb{R}; 0, \infty)$. The Skorohod topology will be used on this space. All of the concepts concerning weak convergence which will be used can be found in the standard references [5, 12]. Summaries with applications to stochastic systems can be found in [17, 19]. The following is a convenient criterion for tightness in $D(\mathbb{R}; 0, \infty)$. It will be used implicitly, without specific mention. Let $Y^n(\cdot)$ be a sequence of processes with paths in $D(\mathbb{R}; 0, \infty)$. Let $T^n(t)$ denote the stopping times with respect to the filtration engendered by $Y^n(\cdot)$ and which are no larger than t . If, for each t ,

$$\limsup_{\delta \rightarrow 0} \sup_n \sup_{\tau \in T^n(t)} E(1 \wedge |Y^n(\tau + \delta) - Y^n(\tau)|) = 0, \quad (2.1a)$$

and

$$\{Y^n(s) : n < \infty, s \leq t\} \text{ is tight in } \mathbb{R}, \quad (2.1b)$$

then $\{Y^n(\cdot)\}$ is tight [12].

Notation and assumptions. For some centering constants $\bar{\Delta}_i^{\alpha,n}, \alpha = a, d$, whose properties will be specified below, define the processes

$$w_i^{\alpha,n}(t) = \frac{1}{\sqrt{n}} \sum_{l=1}^{nt} \left[1 - \frac{\Delta_{i,l}^{\alpha,n}}{\bar{\Delta}_i^{\alpha,n}} \right], \quad t \geq 0, \quad \alpha = a, d, \quad i = 1, 2. \quad (2.2)$$

When an index of summation is nt we mean the *integer* part $[nt]$. Let $x_i^n(t)$ denote $1/\sqrt{n}$ times the number of jobs in queue i at *real time* nt , including the one in service, if any. Let $WL_i^n(t)$ (called the workload at queue i) denote $1/\sqrt{n}$ times the real time that the server must work to complete all of the jobs which are in queue i at *real time* nt . Thus, time is scaled by $1/n$ and the state by $1/\sqrt{n}$. Define the *total workload* $WL^n(t) = \sum_i WL_i^n(t)$. By *scaled work* we mean $1/\sqrt{n}$ times the actual physical work in question. The expression *scaled time* of some event always refers to the real time of that event divided by n .

Define the index of a job at queue i as one plus the number of jobs that arrived or were there before it, starting with the ordered $\sqrt{n}x_i^n(0)$ jobs which are in that queue at time zero. Let $L_i^n(t) \geq 0$ denote the index of the last customer to enter service in queue i at or before real time nt . For future use, note that $x_i^n(\cdot)$ and $WL_i^n(\cdot)$ are related by:

$$WL_i^n(t) \in \left[\frac{1}{\sqrt{n}} \sum_{l=L_i^n(t)}^{L_i^n(t)+\sqrt{n}x_i^n(t)-1} \Delta_{i,l}^{d,n}, \frac{1}{\sqrt{n}} \sum_{l=L_i^n(t)+1}^{L_i^n(t)+\sqrt{n}x_i^n(t)-1} \Delta_{i,l}^{d,n} \right]. \quad (2.3)$$

We will use the following conditions. By “driving random variables,” we mean the set of initial conditions, the arrival times and service requirements, and the starting and stopping times of the vacations.

A2.0. For each n , $x_i^n(0), WL_i^n(0), i = 1, 2$, are independent of all of the “future” driving random variables. None of the sources is on vacation at $t = 0$. [The last sentence is used only to simplify the notation.]

A2.1. For $\alpha = a, d; i = 1, 2$, there are constants $\bar{\Delta}_i^\alpha$ such that

$$\bar{\Delta}_i^{\alpha,n} \rightarrow \bar{\Delta}_i^\alpha.$$

As $n \rightarrow \infty$, the sequences $w_i^{\alpha,n}(\cdot), n < \infty, \alpha = a, d; i = 1, 2$, converge weakly to mutually independent Wiener processes $w_i^\alpha(\cdot), \alpha = a, d, i = 1, 2$, with variance parameters $\sigma_{\alpha,i}^2$, resp.

Define $\bar{\lambda}_i^{\alpha,n} = 1/\bar{\Delta}_i^{\alpha,n}$, which will be used interchangeably, and similarly for $\bar{\lambda}_i^\alpha = 1/\bar{\Delta}_i^\alpha$. Define the traffic intensities

$$\rho_i^n = \bar{\Delta}_i^{d,n} / \bar{\Delta}_i^{a,n} = \bar{\Delta}_i^{d,n} \bar{\lambda}_i^{a,n}, \quad \rho_i = \bar{\Delta}_i^d / \bar{\Delta}_i^a = \bar{\Delta}_i^d \bar{\lambda}_i^a.$$

A2.2. There is a real number b such that

$$\lim_n \sqrt{n} \left[\sum_i \rho_i^n - 1 \right] = b.$$

Note that (A2.2) implies that $\sum_i \rho_i = 1$.

A2.3. For each n, i , the intervals between the end of the l -th vacation and the start of the next one for source i are denoted by $n\tau_{i,l}^{s,n}$, $l = 1, \dots$. They are mutually independent, exponentially distributed, independent of all the other “driving” random variables and have rate $\bar{\lambda}_i^{s,n}/n$, where $\bar{\lambda}_i^{s,n}$ converges to $\bar{\lambda}_i^s > 0$ as $n \rightarrow \infty$. The intervals for the different sources are mutually independent. By convention, we define $\tau_{i,0}^{s,n} = 0$.

A2.4. For each n, i , there are mutually independent and identically distributed random variables $\tau_{i,l}^{v,n}$, $l = 1, \dots$, such that the duration of the l -th vacation interval for source i is $\sqrt{n}\tau_{i,l}^{v,n}$. Also, $\tau_{i,l}^{v,n}$ converges weakly to a random variable τ_i^v as $n \rightarrow \infty$. For each i , the $\tau_{i,l}^{v,n}$, $l = 1, \dots$, are independent of all other “driving” random variables. The intervals for the different sources are mutually independent. $\sup_{i,n} E\tau_{i,l}^{v,n} < \infty$. Define $\tau_{i,0}^{v,n} = 0$.

The convergence to the Wiener process in (A2.1) is a convenient way of covering many common models, the simplest being the iid cases. The extension of (A2.4) which covers simultaneous or overlapping vacation intervals is discussed at the end of Section 3. The added difficulties are only algebraic. We can work with different information structures, in that the server controller can know either the numbers queued or the work queued. We will confine ourselves to the first case, but it will be seen that the results are asymptotically (as $n \rightarrow \infty$) equivalent in that the minimum costs are the same and a good policy for one is equivalent to a good policy for the other. This holds since (Theorem 3.1) the scaled number queued and scaled work queued are (asymptotically) linearly related except on an arbitrarily small (scaled) time interval. Thus, unless mentioned otherwise, the server does not know the work in the queues, only the number of jobs in each queue. It also knows the entire past history, which is the set of past polling decisions, the work done for each job already served and the timing, as well as the starting and ending times of the vacations to date, for each source. The *admissible control* (or polling) policy is defined in the following way.

A2.5. The server can select the queue served in any nonanticipative way at all, provided that it does not switch while a job is being processed. By nonanticipative we mean the following. Suppose that a job has been completed at real time nt and both sources are available. Then, the next source to be polled is determined by the value (0 or 1) of a measurable function of the initial queue sizes, all arrival and service data and the record of vacation starts and completions up to real time nt for each queue.

Later we will also deal with the set of policies which satisfy either of the following special but important conditions.

A2.6a. [In terms of queued numbers.] *There is a real-valued function $\phi(\cdot)$, which is continuous and nondecreasing, such that between vacations the server polls source 1 at real time nt if $x_2^n(t) < \phi(x_1^n(t))$, and polls source 2 otherwise. The server does not switch while a job is being processed.*

A2.6b. [In terms of queued workload, a less restrictive function.] *There is a real-valued function $\theta(\cdot)$ which is continuous at all but a finite number of values such that between vacations the server polls source 1 if $WL_1^n(t) \geq \theta(WL^n(t))$, and polls source 2 otherwise. The server does not switch while a job is being processed.*

3 Weak Convergence of the Workload and Content Processes: Arbitrary Controls

For $l > 0$, define

$$\nu_{i,l}^n = \sum_{k=1}^l \left[\tau_{i,k}^{s,n} + \tau_{i,k-1}^{v,n} / \sqrt{n} \right],$$

which is $1/n$ times the real time of the start of the l -th vacation at source i . I.e., it is the *scaled* time of the start of the l -th vacation at source i . The (scaled) l -th vacation interval for source i is the half open (scaled) interval $[\nu_{i,l}^n, \nu_{i,l}^n + \tau_{i,l}^{v,n} / \sqrt{n})$.

Discussion of the control problem. The weak convergence result will be in terms of the total workload (rather than in terms of the workload in each queue), which will be seen to be enough to get the desired results. In fact, except under special policies such as (A2.6), in general there is no weak convergence result for the $(WL_i^n(\cdot), i = 1, 2)$ or $(x_i^n(\cdot), i = 1, 2)$. Also, the use of $WL(\cdot)$ yields a one dimensional limit control problem, a considerable advantage. Next, we introduce some needed notation, and set the problem up and discuss some of its features in a way that facilitates the proof of Theorem 3.1.

Let $S_i^{a,n}(t)$ (resp., $S_i^{d,n}(t)$) denote $1/n$ times the number of jobs that arrived to (resp., were completely served at) queue i by real time nt . Let $Z^n(t)$ denote $1/\sqrt{n}$ times the total real time that both queues are empty *and* neither source is on vacation up to real time nt . Let $T^{v,n}(t)$ denote $1/\sqrt{n}$ times the total time up to real time nt that the server could not work due to a vacation (i.e., where the contents of the available queue, if any, is zero, or where there are no available queues). Then we can write (remaining work = arrived work - work done)

$$\begin{aligned} WL^n(t) = & WL^n(0) + \frac{1}{\sqrt{n}} \sum_i \sum_{l=1}^{nS_i^{a,n}(t)} \Delta_{i,l}^{d,n} \\ & - \frac{1}{\sqrt{n}} [\text{real time of all service by real time } nt], \end{aligned} \tag{3.1}$$

where the last term on the right is

$$- [\sqrt{n}t - Z^n(t) - T^{v,n}(t)]. \quad (3.2)$$

The effect of the vacations: A heuristic discussion. We need to examine $T^{v,n}(t)$ more carefully, since (as will be seen) it is through this term that the control affects the paths of the process in the limit. By (A2.4), in scaled time the vacations last $\tau_{i,l}^{v,n}/\sqrt{n}$ units of time, an amount which vanishes as $n \rightarrow \infty$. By (A2.3), the intervacation times are $\tau_{i,l}^{v,n}$ in scaled time. Owing to the mutual independence of the intervacation times for the same queue and for the different queues, for any $T > 0$ the probability that vacations will overlap at some point on the scaled time interval $[0, T]$ is of the order of $1/\sqrt{n}$. Thus, in the following discussion, which evaluates the effects of the vacations on the paths for arbitrarily large n , we will suppose (without loss of generality) that only one source can be on vacation at a time. More particularly, if one or more vacations overlap, ignore all but the first. Since this modification alters the paths on each interval $[0, T]$ with a probability of the order $O(1/\sqrt{n})$, it does not affect the distribution of the limit quantities. While the possibility of overlapping vacations is not important for the purely weak convergence aspects in this section, it will have to be taken into account when dealing with the convergence of the costs in Section 4, owing to the uniform integrability issues which arise. Define $u^n(t) = WL_1^n(t)$. Hence, $WL_2^n(t) = WL^n(t) - u^n(t)$. It will turn out that $u^n(\cdot)$ has the effect of a control. Its value will be seen to be the mechanism for controlling the values of the jumps.

Consider the l -th vacation of source i . It starts at (scaled) time $\nu_{i,l}^n$, and the total workload just before that is $WL^n(\nu_{i,l}^n -)$. Define $\bar{A}_{i,l}^{j,n}$ to be $1/\sqrt{n}$ times the work arriving at queue i during the l -th vacation of source j . Define

$$\bar{A}_{i,l}^{j,n}(t) = \frac{1}{\sqrt{n}} \sum_{l=nS_i^{a,n}(\nu_{j,l}^n -) + 1}^{nS_i^{a,n}((\nu_{j,l}^n + (\tau_{j,l}^{v,n} \wedge t)/\sqrt{n}) -)} \Delta_{i,l}^{d,n}. \quad (3.3)$$

The minus's ($-$) in the lower and upper limits of summation in (3.3) are due to the fact that (recall that the interval is half open) we count $\nu_{i,l}^n$ as part of the (scaled) vacation period, but not $\nu_{1,l}^n + \tau_{1,l}^{v,n}/\sqrt{n}$, for specificity, although the exact accounting procedure used at the end points is asymptotically (as $n \rightarrow \infty$) irrelevant.

The time scale used in (3.3) will be called the *local fluid scale*. In this scale, t denotes an interval of real time of length $\sqrt{n}t$, or, equivalently, an interval in scaled time of length t/\sqrt{n} .

Until further notice, for notational simplicity in the motivational discussion, let us fix attention on the l -th vacation of source 1. Thus, $\bar{A}_{i,l}^{1,n} = \bar{A}_{i,l}^{1,n}(\infty)$, the scaled arriving work at queue i during this vacation. Also, $\bar{A}_{i,l}^{1,n}(t)$ is $1/\sqrt{n}$ times the work arriving at queue i in the *real* time interval $[n\nu_{1,l}^n, n\nu_{1,l}^n + \sqrt{n}t)$ for $t < \tau_{1,l}^{v,n}$.

Let $\xi_{i,l}^{v,n}$ denote the change in the *total workload* during the l -th vacation of source i . If

$$\tau_{1,l}^{v,n} < WL_2^n(\nu_{1,l}^n-) + \bar{A}_{2,l}^{1,n} = WL^n(\nu_{1,l}^n-) - u^n(\nu_{1,l}^n-) + \bar{A}_{2,l}^{1,n}, \quad (3.4a)$$

then the vacation ends before queue 2 is emptied, and the vacation would not seem to have an immediate effect on the total idle time and workload and (asymptotically, as $n \rightarrow \infty$) $\xi_{1,l}^{v,n} = 0$. This is not quite obvious, and the point is both important and subtle, since it is possible that the scaled work that arrives during a vacation all arrives close the end in which case there might be idle time. However as seen from Theorem 3.1, it turns out that (asymptotically, as $n \rightarrow \infty$ and in the local fluid time scale) the scaled work can be supposed to arrive “continuously” and at the mean rate during the vacation, analogously to a “fluid.” This implies that, asymptotically as $n \rightarrow \infty$ and under (3.4a), the vacation has no effect on the total workload.

On the other hand, if

$$\tau_{1,l}^{v,n} > WL^n(\nu_{1,l}^n-) - u^n(\nu_{1,l}^n-) + \bar{A}_{2,l}^{1,n} \equiv \hat{\tau}_{1,l}^{v,n}, \quad (3.4b)$$

then queue 2 is emptied before the vacation ends. The proof of Theorem 3.1 allows us to suppose that (asymptotically, as $n \rightarrow \infty$, and in the local fluid time scale) the scaled work arrives continuously, as a fluid, at the mean rate, as noted above. But the heavy traffic condition (A2.2) implies that the service rate is so much faster than the arrival rate of work at queue 2, that (asymptotically, as $n \rightarrow \infty$) the workload in queue 2 is zero for a nonvanishing fraction of the vacation time. Thus, there is (asymptotically) forced idle time and an increase in the total workload due to the vacation. This increase will depend on the value of the workload at queue 2 at the time that the vacation at queue 1 starts. In turn, that value depends on the control policy. This is the *only* way that the control policy affects the workload: via the sizes of the jumps due to the vacations, which (in turn) is determined by the distribution of the total workload just before the vacation starts. It will be seen that the difference between the scaled work that arrives at queue j during this l -th vacation of source 1 and $\bar{\lambda}_j^a \bar{\Delta}_j^d \tau_{1,l}^{v,n} = \rho_j \tau_{1,l}^{v,n}$ converges weakly to zero. Obviously, one can reverse sources 1 and 2 in the above discussion. From the above discussion, we see that the control can be viewed as the division of the total workload among the two queues.

Suppose, formally, that n indexes a weakly convergent subsequence of

$$\left(\tau_{1,l}^{v,n}, u^n(\nu_{1,l}^n-), WL^n(\nu_{1,l}^n-) \right),$$

use the same notation (dropping the n -superscript) for the weak sense limits, and formally use the asymptotic fluid approximation to (3.3). This fluid approximation simply replaces the terms in (3.3) by their asymptotic mean value and is $\rho_i[\tau_{1,l}^v \wedge t]$. We see, formally, that (in the limit) the increase in $T^{v,n}(\cdot)$ (equivalently, in the total workload) during the l -th vacation of source 1 can be written as

$$\xi_{1,l}^v = [(1 - \rho_2) \tau_{1,l}^v - [WL(\nu_{1,l}-) - u(\nu_{1,l}-)]]^+$$

which equals

$$[\rho_1 \tau_{1,l}^v - [\text{WL}(\nu_{1,l}-) - u(\nu_{1,l}-)]]^+. \quad (3.5a)$$

The analog for the l -th vacation of source 2 is

$$\xi_{2,l}^v = [\rho_2 \tau_{2,l}^v - u(\nu_{2,l}-)]^+. \quad (3.5b)$$

Here, $u(t) \leq \text{WL}(t)$ can be considered to be the control function for the limit system (3.7), (3.8a). It appears only via the discrete values. $u(\nu_{i,l}-)$, $l < \infty$, $i = 1, 2$. The notation in (3.5) can be misleading since the use of the symbol $u(t-)$ suggests either left continuity or that the left hand limit exists at t , or that (some subsequence of) $u^n(\cdot)$ converges weakly. *We do not make these claims for general polling policies*, but the sequence $u^n(\nu_{2,l}^n-)$ will always be tight in n . If the control policy is of the type in (A2.6), then roughly speaking (see Theorem 3.2) the “intervacation sections” of $u^n(\cdot)$ will be tight and have continuous weak sense limits.

The intervacation sections. Let $\bar{\nu}_l^n$ and ν_l^n denote, resp., the (scaled) time of the beginning and end (resp.) of the l -th intervacation interval, irrespective of the source. Thus, by our conventions, $\bar{\nu}_l^n$ is $1/n$ times the real time of the beginning of the l -th vacation. $\bar{\nu}_1^n = 0$. Define the *intervacation sections* as the functions:

$$\text{WL}^n((\bar{\nu}_l^n + t) \wedge \nu_l^n), \quad t \geq 0. \quad (3.6)$$

It is constant for $t \geq \nu_l^n - \bar{\nu}_l^n$.

State space collapse. The physical dimension of the original problem is the number of sources. Mathematically, the dimension is even higher since the set of queue lengths is not Markovian. Theorem 3.1 is an example of what is called state space collapse [7, 24, 25, 26, 32] in the heavy traffic literature. The dimension of the approximating problem is unity and the original $x_i^n(\cdot)$ (which do not necessarily converge weakly) can be asymptotically approximated by a constant (depending on i) times the total workload process $\text{WL}^n(\cdot)$ (which does converge weakly in the sense described in the theorem). Such state space collapse is obviously very helpful in the control problem and for numerical procedures.

Comment on tightness and Theorem 3.1. During a vacation, $\text{WL}^n(\cdot)$ changes in steps of size $O(1/\sqrt{n})$ over an interval of scaled size $O(1/\sqrt{n})$. While the effect of the vacation is (asymptotically) a jump in a well defined sense, because of the way that the jump is realized, $\text{WL}^n(\cdot)$ is not tight in the Skorohod topology. Since the parts of the path between vacations are well behaved, it is convenient to work with the effects of the vacations and the intervacation parts separately. Suppose that the set in (3.6) is tight for each l . It is easy to show this (Theorem 3.1) for $l = 1$. Then the set of its “terminal” conditions $\text{WL}^n(\nu_1^n-)$ is also tight, as is $u^n(\nu_1^n-)$. Thus the set (3.5a) or (3.5b) for the first vacation is also tight (Theorem 3.1). Then, the set of initial conditions $\text{WL}^n(\bar{\nu}_2)$ for the next intervacation interval is tight. Then repeat, as for the first section, etc. In

this way, taking an appropriate subsequence and working section by section, one constructs a “limit” process $WL(\cdot)$. We call this procedure “concatenation.” A primary aim is showing the convergence of the discounted cost functions (4.3b) to (4.3a) for a well defined “limit” process $WL(\cdot)$. One does not need full weak convergence of $WL^n(\cdot)$ for this and the piecewise or concatenation approach is adequate. The various conclusions of the theorem are denoted by (a), (b), etc.

Theorem 3.1. *Assume (A2.0)–(A2.5), and suppose that $(x_i^n(0), i = 1, 2)$ converges weakly to $(x_i(0), i = 1, 2)$. (a) Then $WL^n(0)$ converges weakly to $\sum_i \bar{\Delta}_i^d x_i(0)$. (b) For $i = 1, 2$, the set*

$$\Psi^n = \left(WL^n(\nu_{i,l}^n), u^n(\nu_{i,l}^n), \tau_{i,l}^{v,n}, \tau_{i,l}^{s,n}, \xi_{i,l}^{v,n}, \xi_{i,l}^{s,n}, i = 1, 2, l < \infty \right)$$

is tight in n . (c) The sequence of intervaccation sections of $WL^n(\cdot)$ defined by (3.6) and of $Z^n(\cdot)$ are tight for each i and l , the weak sense limit of any weakly convergent subsequence has continuous paths.

Fix a weakly convergent subsequence of the set Ψ^n and the set of intervaccation sections of $WL^n(\cdot)$ and $Z^n(\cdot)$, and index it by n also (abusing terminology). The weak sense limits are denoted by dropping the n superscript. (d) Then $(\tau_{i,l}^{s,n}, l < \infty)$ converges weakly to $(\tau_{i,l}^s, l < \infty)$, where the $\tau_{i,l}^s$ are exponentially distributed with rate $\bar{\lambda}_i^s$. (e) The differences $WL_i^n(\cdot) - \bar{\Delta}_i^{d,n} x_i^n(\cdot)$ converge weakly to the “zero” process.

(f) Define $WL(\cdot)$ by concatenating the weak sense limits of the successive intervaccation sections of $WL^n(\cdot)$. The weak sense limits of any weakly convergent subsequence are related by

$$WL(t) = WL(0) + bt + w(t) + \sum_i J_i(t) + Z(t), \quad (3.7)$$

where

$$J_i(t) = \sum_{l: \nu_{i,l} \leq t} \xi_{i,l}^v, \quad \nu_{i,l} = \sum_{k=1}^l \tau_{i,k}^s. \quad (3.8a)$$

(g) Also

$$(WL(0), w(\cdot), \tau_{i,l}^v, \tau_{i,l}^s; i = 1, 2, l < \infty) \quad (3.8b)$$

are mutually independent, and $w(\cdot)$ is a Wiener process with variance parameter

$$\sigma^2 = \sum_i [\rho_i [\sigma_{a,i}^2 + \sigma_{d,i}^2]], \quad (3.9)$$

which we assume to be positive. (h) $Z(\cdot)$ is the reflection term. It is continuous, nondecreasing, can increase only at t where $WL(t) = 0$ and satisfies $Z(0) = 0$.

(i) The $\xi_{i,l}^v$ have the representation (3.5). (j) Define the Poisson processes $N_i(\cdot), i = 1, 2$, to be the process with a unit jump at $\nu_{i,l}, l \geq 1$. For each t ,

$$w(t + \cdot) - w(t), N_i(t + \cdot) - N_i(t), i = 1, 2, \quad (3.10)$$

is independent of

$$\begin{aligned} w(s), N_i(s), s \leq t; s \leq t; i = 1, 2, \\ (u(\nu_{i,l}-)I_{\{\nu_{i,l} \leq t\}}, \xi_{i,l}^v I_{\{\nu_{i,l} \leq t\}}, i = 1, 2, l < \infty). \end{aligned} \quad (3.11)$$

(k) The process (3.3) converges weakly to the process with values $\rho_i(t \wedge \tau_{j,l}^v)$.

Comment on the control $u(\nu_{i,l}-)$. Under the general conditions that are used in this theorem to get the weak convergence, we cannot get convergence of the random processes $u^n(\cdot)$, only of the random variables which are the values at selected points. But, in the next section the class of polling policies will be restricted to be in some very reasonable class, and for this class there will be tightness of $u^n(\cdot)$ in an appropriate sense. Then, the weak sense limits will be well defined admissible control functions for the weak sense limit $WL(\cdot)$ process.

Proof. As noted below (3.2), without loss of generality we can suppose that at most one source is on vacation at a time. Given the current real time nt , the real time since the current service started or has to go, or the real time since or until the next arrival are called *residual times*. We define a *residual time error term* to be a random process (to be denoted by $\epsilon^n(\cdot)$) which is bounded by $[\text{constant}/\sqrt{n}]$ times a [finite sum of such residual time terms plus a constant]. Successive uses of $\epsilon^n(\cdot)$ might refer to different residual time error terms. Assumption (A2.1) implies that the $\epsilon^n(\cdot)$ converge weakly to the “zero” process, since the continuity of the limit there implies that the maximum of the first $[nt]$ summands, divided by \sqrt{n} goes to zero in probability as $n \rightarrow \infty$.

The difference between the terms in (2.3) is a residual time error term, thus the process defined by the difference converges weakly to the “zero” process. The proof that $WL^n(0)$ converges as asserted follows from (A2.1) and the representation (2.3), where $t = 0$.

For specificity, we will suppose that the preempt-resume discipline holds for any job which is being served when a vacation of its source starts. Assumption (A2.3) implies that, for any t , the number of vacations on any real time interval $[0, nt]$ is bounded in probability, uniformly in n . Due to this and the fact that (by (A2.1)) the maximum of the first $[nt]$ workloads divided by \sqrt{n} goes to zero in probability as $n \rightarrow \infty$, any of the disciplines cited in Section 2 will yield the same result.

We next prove the weak convergence of $S_i^{a,n}(\cdot)$ to the process with constant values $\bar{\lambda}_i^a t$. Define $\mathcal{T}_i^{a,n}(t) = \sum_{l=1}^{nt} \Delta_{i,l}^{a,n}/n$. By (A2.1), $\mathcal{T}_i^{a,n}(\cdot)$ converges weakly to the process with values $\bar{\Delta}_i^a t = t/\bar{\lambda}_i^a$. Also, possibly modulo a residual time error term,

$$\begin{aligned} S_i^{a,n}(\mathcal{T}_i^{a,n}(t)) &= t \\ \mathcal{T}_i^{a,n}(S_i^{a,n}(t)) &= t. \end{aligned}$$

This and the weak convergence of $\mathcal{T}_i^{a,n}(\cdot)$ imply the asserted weak convergence of $S_i^{a,n}(\cdot)$.

The next step is to show the tightness and asymptotic continuity of $WL^n(\cdot)$ when there are no vacations. In the absence of vacations, we can write

$$WL^n(t) = WL^n(0) + \frac{1}{\sqrt{n}} \sum_i \sum_{l=1}^{nS_i^{a,n}(t)} \Delta_{i,l}^{d,n} - t\sqrt{n} + Z^n(t). \quad (3.12)$$

The term $T^{v,n}(t)$ of (3.2) is not included since, in this part of the proof, we have assumed that there are no vacations. For each i , expand the inner sum in (3.12) as

$$\frac{1}{\sqrt{n}} \sum_{l=1}^{nS_i^{a,n}(t)} [\Delta_{i,l}^{d,n} - \bar{\Delta}_i^{d,n}] + \frac{1}{\sqrt{n}} \sum_{l=1}^{nS_i^{a,n}(t)} \bar{\Delta}_i^{d,n}. \quad (3.13)$$

The first term of (3.13) is $-\bar{\Delta}_i^{d,n} w_i^{d,n}(S_i^{a,n}(t))$. Expand the last term in (3.13) as

$$\frac{1}{\sqrt{n}} \bar{\Delta}_i^{d,n} \sum_{l=1}^{nS_i^{a,n}(t)} \left[1 - \frac{\Delta_{i,l}^{a,n}}{\bar{\Delta}_i^{a,n}} \right] + \frac{1}{\sqrt{n}} \frac{\bar{\Delta}_i^{d,n}}{\bar{\Delta}_i^{a,n}} \sum_{l=1}^{nS_i^{a,n}(t)} \Delta_{i,l}^{a,n}. \quad (3.14)$$

The right hand *sum* in (3.14) equals nt minus the time between nt and the last arrival at or before real time nt . Hence, the right hand term equals $[\bar{\Delta}_i^{d,n}/\bar{\Delta}_i^{a,n}]\sqrt{nt}$ plus a residual time error term.

Summarizing,

$$\begin{aligned} WL^n(t) &= WL^n(0) + \sum_i \bar{\Delta}_i^{d,n} \left[w_i^{a,n}(S_i^{a,n}(t)) - w_i^{d,n}(S_i^{a,n}(t)) \right] \\ &\quad + \sqrt{n} \left[\sum_i \frac{\bar{\Delta}_i^{d,n}}{\bar{\Delta}_i^{a,n}} - 1 \right] t + Z^n(t) + \epsilon^n(t). \end{aligned} \quad (3.15)$$

The hypotheses and the weak convergence of $S_i^{a,n}(\cdot)$ imply the weak convergence of the processes on the right of (3.15) to those of (3.7), except possibly that of $Z^n(\cdot)$, with the given definitions of $w(\cdot)$ and b , but without the jump term. If $\{Z^n(\cdot), n < \infty\}$ were not tight and have continuous weak sense limits, then we would have a contradiction to the facts that $Z^n(\cdot)$ can increase only when $WL^n(t) = 0$ and has jumps of size $1/\sqrt{n}$ only. Thus, by taking a further subsequence if necessary, we can suppose that $Z^n(\cdot)$ converges to the reflection term $Z(\cdot)$ in (3.7). The sequence of processes defined by (3.15) converges weakly and the limit satisfies (3.7) without the jump term.

Now, return to the *original problem*, where there are vacations, and recall that ν_l^n is the scaled time of the start of the first vacation. For the remainder of this proof, Continue using the assumption that the vacations do not overlap. As noted below (3.2), this is accomplished by ignoring all but the first if there are overlaps. The alteration does not change the distribution of the limit processes, since the probability of such a change on any finite interval goes to zero as $n \rightarrow \infty$.

By (A2.3) and (A2.4), the various sequences of times $\tau_{i,l}^{s,n}, \tau_{i,l}^{v,n}, i = 1, 2, l = 1, \dots$, converge weakly and $\bar{\nu}_{l+1}^n - \nu_l^n = \tau_{i,l}^{v,n}/\sqrt{n}$ converges weakly to zero.

By (A2.3) and the weak convergence of the sequence of processes defined by (3.15), $WL^n(\nu_1^n-)$ also converges weakly. Denote the weak sense limits by dropping the superscript n . By (A2.3), the $(\tau_{i,l}^s, l = 1, \dots, i = 1, 2)$ are mutually independent, and exponentially distributed, with rate $\bar{\lambda}_i^s$ for $\tau_{i,l}^s$. By (A2.4), $(\tau_{i,l}^v, l = 1, \dots, i = 1, 2)$ are mutually independent. By (A2.0), (A2.1), (A2.3) and (A2.4), the

$$WL(0), w_i^a(\cdot), w_i^d(\cdot), \tau_{i,l}^v, \tau_{i,l}^s, i = 1, 2, l = 1, \dots \quad (3.16)$$

are mutually independent. For each l , ν_l^n converges weakly and $\nu_{l+1}^n - \nu_l^n$ converges weakly to an exponentially distributed random variable, with rate $\sum_i \bar{\lambda}_i^s$, and the limits are mutually independent, and are independent of the random variables in (3.16) other than $\{\tau_{i,l}^s; i, l\}$.

Suppose for the moment that the jumps $\xi_{i,l}^{v,n}$ are tight for each i, l . Abusing notation, let n index a further subsequence along which all of the $\xi_{i,l}^{v,n}, i = 1, 2, l = 1, \dots$, also converge weakly, and denote the weak sense limits by dropping the superscript n . Then, by repeating the analysis which led to (3.15) on each successive intervaccation interval, we are led to (3.7) with $J_i(\cdot)$ defined by (3.8a) and the independence in (3.8b). Equation (3.7) represents the limit of $WL^n(\cdot)$ in the particular sense that its interjump sections are the weak sense limits of the intervaccation sections (3.6) and its jumps are the limits of the $WL^n(\bar{\nu}_{l+1}^n) - WL^n(\nu_l^n-)$ (for the chosen subsequence).

By taking a further subsequence, if necessary, we can also suppose that, together with the other convergences, $u^n(\nu_{i,l}^n-), i = 1, 2, l = 1, \dots$ converges weakly to random variables which we call $u(\nu_{i,l}-), i = 1, 2, l = 1, \dots$. From the weak convergence of the $u^n(\nu_{i,l}^n-), \nu_{i,l}^n, i = 1, 2, l = 1, \dots$, we have the weak convergence of the $u^n(\nu_l^n-), l = 1, \dots$

We will next show that $\xi_{i,l}^{v,n}$ is tight for each i and l and that (3.5) characterizes the weak sense limits. To simplify the notation, we will start with the first vacation and let the first vacation be that of source 1. With this simplifying assumption, we can write $\tau_{1,1}^{s,n} = \nu_{1,1}^n = \nu_1^n$, and we will use these variables (and their weak sense limits) interchangeably. By the weak convergence of $WL^n(\cdot)$ (with the weak sense limit of $WL^n(\cdot)$ being continuous) when there are no vacations and of $\tau_{1,1}^{s,n}$, $WL_1^n(\tau_{1,1}^{s,n}-) = WL_1^n(\nu_1^n-) = u_1^n(\nu_1^n-)$ converges weakly to the random variable which we denote by $u(\nu_{1,1}^s-)$. We will show that, under (3.4a),

$$WL^n(\nu_{1,1}^n + \tau_{1,1}^{v,n}/\sqrt{n}) - WL^n(\nu_{1,1}^n) \Rightarrow 0, \quad (3.17)$$

and under (3.4b),

$$WL_2^n(\nu_{1,1}^n + \tau_{1,1}^{v,n}/\sqrt{n}) \Rightarrow 0, \quad (3.18)$$

$$WL_1^n(\nu_{1,1}^n + \tau_{1,1}^{v,n}/\sqrt{n}) - WL_1^n(\nu_{1,1}^n) \text{ is tight}, \quad (3.19)$$

and the weak sense limit (along the selected weakly convergent subsequence) of (3.19) is defined by (3.5). Since the (scaled) vacation interval in question is $[\nu_{1,1}^n, \nu_{1,1}^n + \tau_{1,1}^{v,n}/\sqrt{n}]$, strictly speaking, the arguments in the functions in (3.19) should be $(\nu_{1,1}^n + \tau_{1,1}^{v,n}/\sqrt{n})-$, and analogously for (3.17) and (3.18). However,

since the processes defined by $WL^n(t) - WL^n(t-)$ and $u^n(t) - u^n(t-)$ converge weakly to the “zero” process, one can always replace $t-$ by t without changing any of the weak sense limits. We will do this to simplify the notation.

Since source 2 is being polled during this vacation, for $t \leq \tau_{1,1}^{v,n}$ we can write

$$WL_2^n(\nu_{1,1}^n + t/\sqrt{n}) = [WL^n(\nu_{1,1}^n) - u^n(\nu_{1,1}^n)] - t + T^{v,n}(\nu_{1,1}^n + t/\sqrt{n}) + \bar{A}_{2,1}^{1,n}(t). \quad (3.20)$$

We will next show that the process defined by

$$\bar{A}_{i,1}^{1,n}(t \wedge \tau_{1,1}^{v,n}) - \rho_i^n(t \wedge \tau_{1,1}^{v,n})$$

converges weakly to the “zero” process. This is what was meant by the statement below (3.4a) to the effect that work can be assumed to arrive continuously during a vacation and it is the last assertion (k) of the theorem. Note that the local fluid time scale defined below (3.3) is used in (3.20), so that t denotes an interval of length $\sqrt{n}t$ in real time or t/\sqrt{n} in scaled time.

Use the representation (3.3) (and dropping the $-$ in the indices of summation without changing the end result) to write $\bar{A}_{i,1}^{1,n}(t)$ as

$$\begin{aligned} & \frac{1}{\sqrt{n}} \sum_{l=nS_i^{a,n}(\nu_{1,1}^n)+1}^{nS_i^{a,n}(\nu_{1,1}^n+(t \wedge \tau_{1,1}^{v,n})/\sqrt{n})} [\Delta_{i,l}^{d,n} - \bar{\Delta}_i^{d,n}] \\ & + \frac{1}{\sqrt{n}} \sum_{l=nS_i^{a,n}(\nu_{1,1}^n)+1}^{nS_i^{a,n}(\nu_{1,1}^n+(t \wedge \tau_{1,1}^{v,n})/\sqrt{n})} \bar{\Delta}_i^{d,n}. \end{aligned} \quad (3.21)$$

The first term in (3.21) goes weakly to zero by the tightness of $\nu_{i,l}^n, \tau_{i,l}^{v,n}$ in n , the weak convergence of $S_i^{a,n}(\cdot)$ and condition (A2.1). We need to characterize the right hand term of (3.21). Write it as:

$$\begin{aligned} & \frac{\bar{\Delta}_i^{d,n}}{\sqrt{n}} \sum_{l=nS_i^{a,n}(\nu_{1,1}^n)+1}^{nS_i^{a,n}(\nu_{1,1}^n+(t \wedge \tau_{1,1}^{v,n})/\sqrt{n})} \left[1 - \frac{\Delta_{i,l}^{a,n}}{\bar{\Delta}_i^{a,n}} \right] \\ & + \frac{\bar{\Delta}_i^{d,n}}{\bar{\Delta}_i^{a,n} \sqrt{n}} \sum_{l=nS_i^{a,n}(\nu_{1,1}^n)+1}^{nS_i^{a,n}(\nu_{1,1}^n+(t \wedge \tau_{1,1}^{v,n})/\sqrt{n})} \Delta_{i,l}^{a,n}. \end{aligned}$$

Just as for the first term in (3.21), the first term in the above expression goes weakly to the “zero” process as $n \rightarrow \infty$. The real time difference between the arguments in the upper and lower indices in the last expression is $\sqrt{n}[t \wedge \tau_{1,1}^{v,n}]$. Hence, the sum in the second term times $1/\sqrt{n}$ is $t \wedge \tau_{1,1}^{v,n}$, modulo a residual time error term. Thus, the difference between the second term and

$$\frac{\bar{\Delta}_i^d}{\bar{\Delta}_i^a}(t \wedge \tau_{1,1}^{v,n}) = \rho_i(t \wedge \tau_{1,1}^{v,n}) \quad (3.22)$$

converges weakly to the “zero” process.

The above computations concerning the arriving scaled work during the vacation show that the net change $\xi_{1,1}^{v,n}$ in the total workload is tight, and that we can suppose, asymptotically and in the local fluid time scale defined below (3.3), that scaled work arrives at the queues continuously (i.e., as a fluid process) at the mean rate ρ_i during the vacation.

Consider the case (3.4a). By what has just been proved, the scaled work process arriving to queue 1 during the the first vacation is arbitrarily well approximated (in the local fluid time scale) by (3.22) for $i = 1$. Similarly, the scaled work that departs queue 2 during that time (local fluid time scale) is (asymptotically) equal to $\tau_{1,1}^{v,n}$ minus the idle time in the local fluid time scale. But, due to the fluid approximation and the condition (3.4a), this idle time is zero, asymptotically. Thus, by the above computation, the net increase in the workload (input minus output) of queue 2 during the vacation is (asymptotically) equal to $[\rho_2 - 1]\tau_{1,1}^{v,n}$. Thus, by the heavy traffic condition (A2.2), adding the changes in the two queues, we see that the net change in the total workload during the vacation converges weakly to zero as $n \rightarrow \infty$.

Now, consider the condition (3.4b), continue to use the approximation (3.22), and recall the definition of $\hat{\tau}_{i,l}^{v,n}$ from (3.4b). It is then clear that $\tau_{i,l}^{v,n} - \hat{\tau}_{i,l}^{v,n}$ is the net contribution to $T^{v,n}(\cdot)$ during this vacation interval, and (3.5) follows from this and the results of the last paragraph. Let $\xi_{1,1}^v$ denote the weak sense limit of $\xi_{1,1}^{v,n}$. Thus, we have obtained (3.7) and (3.8) up to and including the time of the first vacation.

Now, with $WL_i^n(\nu_{1,1}^n + \tau_{1,1}^{v,n}/\sqrt{n})$ well defined and tight, restart the $WL_i^n(\cdot)$ at scaled time $\nu_{1,1}^n + \tau_{1,1}^{v,n}/\sqrt{n}$ and repeat the above approximation and limit procedure. Then an induction argument yields the asserted limit relations for all of the intervacaion sections and jumps. Take a weakly convergent subsequence of Ψ^n in the theorem statement, and obtain (3.7) and (3.8) by concatenating the intervacaion sections.

Next, we will prove the (asymptotic) linear relationship (e) between $WL_i^n(\cdot)$ and $x_i^n(\cdot)$. We say that a sequence of real-valued processes $q^n(\cdot)$ is *bounded in probability* if, for each $T > 0$,

$$\lim_{N \rightarrow \infty} \limsup_n P \left\{ \sup_{t \leq T} |q^n(s)| \geq N \right\} = 0. \quad (3.23)$$

It will be seen that the tightness of $WL^n(\cdot)$ implies that $x_i^n(\cdot)$ satisfies (3.23). Let us assume this for the moment. We will use the representation (2.3). Write $\Delta_{i,l}^{d,n} = \bar{\Delta}_{i,l}^{d,n} + [\Delta_{i,l}^{d,n} - \bar{\Delta}_{i,l}^{d,n}]$. With this representation, expand each of the two terms in the brackets in (2.3) into two components, analogous to what was done to get the expression below (3.21). The first component of the expansion of (say) the first term inside the brackets in (2.3) is

$$\bar{\Delta}_i^{d,n} x_i^n(t). \quad (3.24)$$

The second component of (again) the first term inside the brackets of (2.3) is

$$\frac{1}{\sqrt{n}} \sum_{l=L_i^n(t)}^{L_i^n(t)+\sqrt{n}x_i^n(t)-1} \left[\Delta_{i,l}^{d,n} - \bar{\Delta}_{i,l}^{d,n} \right], \quad (3.25)$$

and it converges weakly to the “zero” process by the weak convergence assumption (A2.1) on $w_i^{d,n}(\cdot)$ since $x_i^n(\cdot)$ is assumed to satisfy (3.23).

Note that the difference between the two terms inside the brackets of (2.3) is a residual time error term $\epsilon^n(t)$, where $\epsilon^n(\cdot)$ converges weakly to the “zero” process, and this is true irrespective of whether or not (3.23) holds for $x_i^n(\cdot)$.

By the tightness of the sections of $WL^n(\cdot)$ between the $(l-1)$ st and l -th vacations (for each l) and of the associated set of jumps, (3.23) holds for $WL_i^n(\cdot)$, hence it holds for the process defined by the first term in the brackets in (2.3). Suppose that (3.23) does not necessarily hold for $x_i^n(\cdot)$. The expansion of the first term in (2.3) into (3.24) and (3.25) still holds. By the assumption (A2.1) on the $w_i^{d,n}(\cdot)$, the fact that the upper index of summation in (3.25) is no bigger than $nS_i^{a,n}(\cdot) + \sqrt{n}x_i^n(0)$ and the weak convergence of $S_i^{a,n}(\cdot)$ and of $x_i^n(0)$, the process defined by (3.25) satisfies (3.23). Then, since the sum of (3.24) and (3.25) is the first term in (2.3), which satisfies (3.23), so must the process defined by (3.24). Hence $x_i^n(\cdot)$ satisfies (3.23). The fact that (3.23) holds for $x_i^n(\cdot)$ and the assumption (A2.1) imply that (3.25) converges weakly to the “zero” process. Hence, $WL_i^n(\cdot)$ is asymptotically equivalent to $\bar{\Delta}_i^{d,n} x_i^n(\cdot)$.

Only the nonanticipativity (j) needs to be proved. But this is a consequence of the independence in (3.16) ■

More than 2 sources. Suppose that there are an arbitrary number of sources, with the natural extensions of the assumptions (A2.0)–(A2.5) and the notation holding. Then, on any (scaled) interval $[0, T]$, with a probability that goes to one as $n \rightarrow \infty$, there are still only a finite number of vacations and at most one source can be on vacation at a time. Because of this, the analogs of (3.4) and (3.5) can easily be written.

Consider the l -th vacation of source i . If

$$\tau_{i,l}^{v,n} < [WL^n(\nu_{i,l}^n -) - WL_i^n(\nu_{i,l}^n -)] + \sum_{j \neq i} \bar{A}_{j,l}^{i,n},$$

then the method of Theorem 3.1 can be used to show that (asymptotically) the vacation at source i ends before the other queues are emptied, and the vacation has no immediate effect on the total workload. On the other hand, if

$$\tau_{i,l}^{v,n} > [WL^n(\nu_{i,l}^n -) - WL_i^n(\nu_{i,l}^n -)] + \sum_{j \neq i} \bar{A}_{j,l}^{i,n} \equiv \hat{\tau}_{i,l}^{v,n},$$

then there is (asymptotically) a forced idle time during the vacation. Dropping the n -superscripts, the increase in scaled work during the l -th vacation of

source i has the asymptotic form

$$\xi_{i,l}^v = \left[\tau_{i,l}^v - [WL(\nu_{i,l}-) - WL_i(\nu_{i,l}-)] - \tau_{i,l}^v \sum_{j \neq i} \rho_j \right]^+$$

which equals

$$[\rho_i \tau_{i,l}^v - [WL(\nu_{i,l}-) - WL_i(\nu_{i,l}-)]]^+,$$

where $WL_i(\nu_{i,l}-)$ is the weak sense limit of (a suitable weakly convergent subsequence of) $WL_i^n(\nu_{i,l}^n-)$. This is the only change in Theorem 3.1.

The control form (A2.6a). The control form specified by (A2.6a) seeks to force the relationship (asymptotic) $x_2^n(t) \sim \phi(x_1^n(t))$, at least when possible between vacations, when both sources are available. The control, in practice, might be based on either the queue sizes or on the workloads, depending on what information is available to the controller at the server. Theorem 3.1 implies that we can do either, due to the asymptotic equivalence $WL_i^n(\cdot) \sim \bar{\Delta}_i^{d,n} x_i^n(\cdot)$. To a control represented by the function $\phi(\cdot)$ in (A2.6a), there is one in terms of the workload in the sense of asymptotic equivalence. We will now see that (A2.6a) is asymptotically equivalent to the existence of a continuous and nondecreasing function $\theta(\cdot)$ such that we poll source 1 if $WL_1^n(t) \geq \theta(WL^n(t))$ and poll source 2 otherwise, provided that the source is available.

To get $\theta(\cdot)$, we use the asymptotic equivalence $x_2^n(t) \sim \phi(WL_1^n(t)/\bar{\Delta}_1^d)$ and

$$WL_2^n(t) \sim \bar{\Delta}_2^d \phi(WL_1^n(t)/\bar{\Delta}_1^d).$$

Since, asymptotically, using the policy $\phi(\cdot)$ between vacations,

$$WL^n(t) \sim \sum_i WL_i^n(t) = WL_1^n(t) + \bar{\Delta}_2^d \phi(WL_1^n(t)/\bar{\Delta}_1^d),$$

we can define an ‘‘asymptotic inverse’’ $\theta(\cdot)$ to the function $\phi(\cdot)$ in that (between vacations) we poll source 1 if $WL_1^n(t) \geq \theta(WL^n(t))$ and poll source 2 otherwise. The inverse is obtained from:

$$WL - WL_1 = \bar{\Delta}_2^d \phi(WL_1/\bar{\Delta}_1^d).$$

Note that we could have started with $\theta(\cdot)$ and derived $\phi(\cdot)$ from it: I.e., suppose that we are a given nonnegative, continuous and nondecreasing function $\theta(\cdot)$ satisfying $\theta(WL) \leq WL$, and use the rule: between vacations, poll source 1 if $WL_1^n(t) \geq \theta(WL^n(t))$ and poll source 2 otherwise. This can be turned into a (asymptotic) rule based on the $x_i^n(\cdot)$ by polling source 1 if

$$x_1^n(t) \bar{\Delta}_1^d \geq \theta(x_1^n(t) \bar{\Delta}_1^d + x_2^n(t) \bar{\Delta}_2^d)$$

and polling source 2 otherwise.

The function $\phi(\cdot)$ in terms of the numbers queued is often (but certainly not always) the more pertinent in applications. But the dynamic programming

equation will be in terms of the total workload and the system (3.7), since the basic weak convergence result is in terms of the total workload. The total workload formulation is also much more convenient from the computational point of view due to the “state space collapse” for which, no matter how many sources there are, the problem is one dimensional. So, it is important to be able to travel back and forth between the queued number and total workload forms.

Note on realizing the relationship $WL_1^n(t) \sim \theta(WL^n(t))$, or its equivalent in terms of the number queued, under (A2.6). Suppose that both sources are available at scaled time t , and that we wish to change $WL_1^n(\cdot)$ to the value $WL_1^{*,n} > WL_1^n(t)$ as quickly as possible. In heavy traffic, by polling queue 2, the scaled queue of source 1 increases at a mean rate of $\bar{\lambda}_1^a \bar{\Delta}_1^d \sqrt{n}$ in scaled time. Thus, it takes approximately $[WL_1^{*,n} - WL_1^n(t)] / [\bar{\lambda}_1^a \bar{\Delta}_1^d \sqrt{n}]$ units of scaled time for the transition. Thus, in the heavy traffic limit, with neither source on vacation, any desired change can be realized instantaneously.

The relationship $WL_1^n(t) \sim \theta(WL^n(t))$ cannot be realized arbitrarily well, uniformly (for large n) on the entire interval between vacations. This is because the uncontrollable changes in the $WL_i^n(\cdot)$ during a vacation will cause it to be violated for a short interval just after the vacation ends, while we “catch up.” But there are $\epsilon_n \rightarrow 0$ as $n \rightarrow \infty$ such that the sections of the differences

$$x_1^n(\cdot) - \frac{\theta(WL^n(\cdot))}{\bar{\Delta}_1^d} \quad (3.26a)$$

and

$$x_2^n(\cdot) - \frac{WL^n(\cdot) - \theta(WL^n(\cdot))}{\bar{\Delta}_2^d} \quad (3.26b)$$

starting (scaled time) ϵ_n after a vacation *begins* and stopping at the start of the next vacation converge to the zero process as $n \rightarrow \infty$. This will be sufficient for our purposes.

With the control $\theta(\cdot)$ in (A2.6b) used, we can write the $\xi_{i,l}^v$ of (3.5) as

$$\xi_{1,l}^v \equiv [(1 - \rho_2) \tau_{1,l}^v - [WL(\nu_{1,l}) - \theta(WL(\nu_{1,l}))]]^+, \quad (3.27a)$$

and

$$\xi_{2,l}^v \equiv [(1 - \rho_1) \tau_{2,l}^v - \theta(WL(\nu_{2,l}))]^+. \quad (3.27b)$$

The following theorem codifies the last part of the above discussion and the proof follows from the computations done in Theorem 3.1. The last sentence of the theorem holds because of the weak convergence of the intervacaion sections and the fact that, between vacations, $WL(\cdot)$ behaves like a Wiener process, provided that $\sigma_{\alpha,i}^2 > 0$ for some α, i .

Theorem 3.2. *Assume the conditions of Theorem 3.1 and (A2.6a) as well. Then, there are positive real numbers $\epsilon_n \rightarrow 0$ such that $x_2^n(\cdot) - \phi(x_1^n(\cdot))$ converges weakly to the “zero” process on each interval $[\nu_{i,l}^n + \epsilon_n, \nu_{i,l+1}^n]$. So does $WL_1^n(\cdot) -$*

$\theta(WL^n(\cdot))$ and the processes defined in (3.26), where $\theta(\cdot)$ is defined from $\phi(\cdot)$ as above the theorem.

Now assume (A2.6b) in lieu of (A2.6a). Then, excluding an arbitrarily small neighborhood of the times where $WL^n(t)$ is a point of discontinuity of $\theta(\cdot)$, the last sentence of the previous paragraph holds for $\theta(\cdot)$. Assume that at least one of the $\sigma_{\alpha,i}^2, \alpha = a, d, i = 1, 2$, is positive. Given $\epsilon > 0$ and $t_1 > 0$, let $T_\epsilon^n(t_1)$ denote the Lebesgue measure of the closure set of time points on $[0, t_1]$ at which $WL^n(t)$ is within ϵ of a point of discontinuity of $\theta(\cdot)$. Then, for each $\delta > 0$ and $t_1 > 0$,

$$\lim_{\epsilon \rightarrow 0} \limsup_n P \{T_\epsilon^n(t_1) \geq \delta\} = 0. \quad (3.28)$$

Correlated vacations of the sources. Up to now, we have supposed that the vacation processes of the two sources are independent of each other. This would be the case if they were due to movement in independent environments or to extraneous interference if the sources were far apart. If the vacations were due to extraneous interference which affected the sources in a similar manner, then the vacation intervals would be correlated. The main problem in introducing such correlation is algebraic, in that it complicates the expressions.

Let us first suppose that, in addition to the mutually independent vacations specified by (A2.3) and (A2.4), there are also simultaneous vacations of the two sources, as defined by the following condition.

A3.1. For each n , the intervals between the end of a simultaneous vacation and the start of the next one are denoted by $n\tau_l^{m,n}, l = 1, \dots$. They are mutually independent, exponentially distributed, independent of all the other “driving” random variables and have rate $\bar{\lambda}^{m,n}/n$, where $\bar{\lambda}^{m,n}$ converges to $\bar{\lambda}^m > 0$ as $n \rightarrow \infty$.

A3.2. For each n , there are mutually independent and identically distributed random variables $\tau_l^{m,v,n}, l = 1, \dots$, such that the duration of the l -th simultaneous vacation interval is $\sqrt{n}\tau_{i,l}^{m,v,n}$. Also, $\tau_l^{m,v,n}$ converges weakly to a random variable $\tau_l^{m,v}$ as $n \rightarrow \infty$. The $\tau_l^{m,v,n}, l = 1, \dots$, are independent of all other “driving” random variables.

If (A3.1) and (A3.2) are added to the conditions of Theorem 3.1 or Theorem 3.2, then the results would be the same, except for the addition of another (independent) jump process $J^m(\cdot)$. Let $\nu_l^{m,n}$ denote the (scaled) starting times of the successive mutual vacations, and ν_l^m the weak sense limits, The weak sense limit (in the sense used in Theorem 3.1) equation is

$$WL(t) = WL(0) + bt + w(t) + \sum_i J_i(t) + J^m(t) + Z(t), \quad (3.29)$$

where

$$J^m(t) = \sum_{l: \nu_l^m \leq t} \xi_l^m, \quad (3.30)$$

where ξ_l^m is the weak sense limit of (see (3.3))

$$\bar{A}_l^{m,n} = \frac{1}{\sqrt{n}} \sum_{l=nS_1^{a,n}(\nu_l^{m,n} + \tau_l^{m,n}/\sqrt{n})-1}^{nS_1^{a,n}(\nu_l^{m,n} + \tau_l^{m,n}/\sqrt{n})} \Delta_{1,l}^{d,n} + \frac{1}{\sqrt{n}} \sum_{l=nS_2^{a,n}(\nu_l^{m,n} + \tau_l^{m,n}/\sqrt{n})-1}^{nS_2^{a,n}(\nu_l^{m,n} + \tau_l^{m,n}/\sqrt{n})} \Delta_{2,l}^{d,n}, \quad (3.31)$$

and the limit is just τ_l^m , owing to the analysis done for the $\bar{A}_{i,l}^{j,n}$ in Theorem 3.1 and (A2.2).

4 The Limit Control Problem

The limit dynamical model. Theorem 3.1 enables us to write the correct limit control problem. As usual in heavy traffic modeling, the aim is to use the limit control problem to get good controls for the physical problem and approximations to its optimal costs, under heavy traffic. The limit dynamics are defined by (3.7) and (3.8) where the jumps are defined by (3.5), where an admissible control $u(\cdot)$ satisfies $u(t) \leq WL(t)$ and is nonanticipative in the sense that it is a measurable process and, for each t , $u(t)$ is independent of $w(t + \cdot) - w(t)$, $N_i(t + \cdot) - N_i(t-)$, $i = 1, 2$. Since $WL(\cdot)$ is continuous at all t where there are no jumps and has a left hand limit at t if there is a jump there, $WL(t-)$ is well defined for all t . But, $u(t-)$ is not necessarily defined. If the control for (3.7), (3.8) is defined via a function such as the $\theta(\cdot)$ in (A2.6b), then we would have $u(t) = \theta(WL(\cdot))$ and $u(t-)$ is well defined for almost all t , which assures that $u(\nu_{i,l}-)$ is well defined with probability one. This is one of the main reasons for our interest in control functions such as $\theta(\cdot)$. Alternatively, we could write the jumps as

$$\xi_{1,l}^v = [\rho_1 \tau_{1,l}^v - [WL(\nu_{1,l}-) - u(\nu_{1,l})]]^+, \quad (4.1a)$$

$$\xi_{2,l}^v = [\rho_2 \tau_{2,l}^v - u(\nu_{2,l})]^+, \quad (4.1b)$$

where $u(\cdot)$ is a “predictable” process [15, 23]. All that is important is the nonanticipativity as defined above, so that, for each t , $u(t)$ is independent of any jump that might occur at t .

The cost function. Let $c_i(\cdot)$ be strictly increasing and continuous real-valued function on $[0, \infty)$ with $c_i(0) = 0$, and satisfying $c_i(x) \leq Kx + K$ for some $K < \infty$. We will work with a discounted cost function. The cost rate will depend on whether we are penalizing queued jobs or queued workload. In the latter case, we penalize the workloads individually and simply use the cost rate

$$\sum_i c_i(WL_i^n(\cdot)) = c_1(u^n(\cdot)) + c_2(WL^n(\cdot) - u^n(\cdot)) \equiv c(WL^n(\cdot), u^n(\cdot)).$$

In the former case, we would like to penalize the queue sizes individually. I.e, with a cost rate $\sum_i c_i(x_i^n(t))$. Since, in general, we do not have a weak convergence result for the $x_i^n(\cdot)$ for an arbitrary admissible control policy, we are

still forced to work with the workload formulation. Then we asymptotically approximate in terms of workload as

$$\sum_i c_i(x_i^n(t)) \approx c_1 \left(\frac{u^n(t)}{\Delta_1^{d,n}} \right) + c_2 \left(\frac{WL^n(t) - u^n(t)}{\Delta_2^{d,n}} \right) \equiv c(WL^n(t), u^n(t)). \quad (4.2)$$

Let $\beta > 0$ be the discount factor. The cost function for the limit system will be

$$W_\beta(WL(0), u(\cdot)) = E \int_0^\infty e^{-\beta t} c(WL(t), u(t)) dt, \quad (4.3a)$$

and for the physical system it will be

$$W_\beta^n(WL^n(0), u^n(\cdot)) = E \int_0^\infty e^{-\beta t} c(WL^n(t), u^n(t)) dt. \quad (4.3b)$$

Define $V_\beta(WL(0)) = \inf_u W_\beta(WL(0), u(\cdot))$ and $V_\beta^n(WL^n(0)) = \inf_{u^n} W_\beta^n(WL^n(0), u^n(\cdot))$, where the $u^n(\cdot)$ and $u(\cdot)$ are admissible. Under a feedback control $\theta^n(\cdot)$ and associated polling policy satisfying (A2.6b), we can write

$$W_\beta^n(WL^n(0), \theta^n(\cdot)) = E \int_0^\infty e^{-\beta t} c(WL^n(t), \theta^n(WL^n(t))) dt. \quad (4.4)$$

A restriction of the class of controls and a redefinition of the inf. As noted, we would like to show that a nearly optimal control for the limit problem is nearly optimal for the physical problem for large n and that

$$V_\beta^n(WL^n(0)) \rightarrow V_\beta(WL(0)) \quad (4.5)$$

if $WL^n(0)$ converges weakly to $WL(0)$, analogously to what was done in [2, 20, 21]. This is hard to do, since the control appears in the dynamics (3.7) only via the magnitude of the jumps.

The usual method [2, 20, 21] for showing (4.5) involves writing the control in some form such that the sequence of optimal (or ϵ -optimal) controls for the physical problem is tight, and any weak sense limit of the (state process, control process) is an admissible limit control problem. For example, suppose that we have a problem where the control is a vector-valued function which takes values in a compact set. Then, we would write the controls in relaxed control form [19] with the weak topology on them. Since, any sequence of such relaxed controls is tight (in the weak topology which is normally used), there is always a weakly convergent subsequence. One could attempt the same thing here. The sequence of relaxed control representations of the control will be tight. The problem is that what is our $u(\cdot)$ is the derivative of the relaxed control. This derivative is defined only almost everywhere, and in particular, it is not guaranteed that the weak sense limit of $u^n(\nu_{i,l}^n-)$ would be $u(\nu_{i,l}-)$, where this $u(\cdot)$ is the derivative of the weak sense limit of the relaxed control representations. The problem is that, while the cost rate can be written as a linear function of the relaxed control, the jump distribution depends only on specific values of the $u^n(\cdot)$.

One can probably circumvent these difficulties. However, in order to apply any control which is nearly optimal for the limit system to the physical system, the form in which $u(\cdot)$ appears in (3.5) essentially implies that it should be a feedback control which is continuous “most of the time.” Because of these facts, we will be concerned with the more restricted class of controls of (A2.6b) defined by the following.

A4.1. *Given an integer M , let Θ denote a class of functions, each member of which satisfies (A2.6b), has at most M points of discontinuity, and on each finite interval $[0, WL]$, the functions in Θ are equicontinuous between discontinuities. The controls will be restricted to such a class, with the polling policy being as defined in (A2.6b).*

Fix the class Θ for some M and modulus of equicontinuity. Redefine $V_\beta^n(WL(0))$ and $V_\beta(WL(0))$ to be the infima over controls in the class Θ .

A slight alteration of the proof of Theorem 3.1 yields the following. Assume the conditions of Theorem 3.1 and let $\theta^n(\cdot) \in \Theta$. Choose the weakly convergent subsequence such that $\theta^n(\cdot)$ also converges (in the Skorohod topology) to, say $\theta(\cdot)$. Then the conclusions of Theorems 3.1 and 3.2 hold. Suppose, in addition, that the function whose expectation is being taken in (4.4) is uniformly (in n) integrable. Then the weak convergence in Theorems 3.1 and 3.2 implies that the expected value in (4.4) converges to the expected value for the controlled limit system. We state this in the following more restrictive way, since that is the way that it will be verified. The proof is simpler than those in [2, 20, 21] for other control problems under heavy traffic, owing to the more restricted class of allowed controls. The proof implies that a good control for the limit problem will be good for the physical problem under heavy traffic.

Theorem 4.1. *Assume (A4.1) and the conditions of Theorem 3.1. Let $c_i(\cdot)$ satisfy the conditions imposed at the beginning of this section. Suppose that there is a real K_1 such that*

$$\sup_{\theta(\cdot) \in \Theta} E |WL^n(t)|^2 \leq K_1 t + K_1. \quad (4.6)$$

Then the function whose expectation is being taken in (4.4) is uniformly integrable and

$$V_\beta^n(WL^n(0)) \rightarrow V_\beta(WL(0)).$$

Comments on the proof. Let $\epsilon > 0$ be small and arbitrary. Let $\theta^{\epsilon,n}(\cdot)$ and $\theta^\epsilon(\cdot)$ be ϵ -optimal controls in Θ for the processes $WL^n(\cdot)$ and $WL(\cdot)$, with the initial conditions $WL^n(0)$ and $WL(0)$, resp. Condition (A4.1) implies that, by choosing a subsequence if necessary, we can suppose that $\theta^{\epsilon,n}(\cdot)$ converges to some $\bar{\theta}(\cdot) \in \Theta$ in the Skorohod topology. Then

$$\epsilon + V_\beta^n(WL^n(0)) \geq W_\beta^n(WL^n(0), \theta^{\epsilon,n}(\cdot)) \rightarrow W_\beta(WL(0), \bar{\theta}(\cdot)) \geq V_\beta(WL(0)).$$

Thus,

$$\liminf_n V^n(WL^n(0)) \geq V_\beta(WL(0)).$$

Now, apply $\theta^\epsilon(\cdot)$ to $WL^n(\cdot)$ to get

$$V_\beta^n(WL^n(0)) \leq W_\beta^n(WL^n(0), \theta^\epsilon(\cdot)) \rightarrow W_\beta(WL(0), \theta^\epsilon(\cdot)) \leq V_\beta(WL(0)) + \epsilon.$$

These inequalities yield the theorem. ■

On the condition (4.6). Condition (4.6) is not a consequence of the other conditions. Write (3.15) for the general case where the vacations are included as

$$WL^n(t) = h^n(t) + Z^n(t).$$

It follows from the estimates given for the general Skorohod problem in [10, Theorem 2.2.] that there is a real K such that

$$WL^n(t) \leq K \sup_{s \leq t} |h^n(s)|, \quad \text{all } t.$$

Thus, a sufficient condition for (4.6) is that the following inequalities hold.

$$\sup_n E |WL^n(0)|^2 < \infty, \quad (4.7)$$

$$E \sup_{s \leq t} |w_i^{\alpha,n}(S_i^{\alpha,n}(s))|^2 = O(t), \quad \alpha = a, d, i = 1, 2, \quad (4.8)$$

$$\bar{J}^n(t) = \sum_j E \left| \sum_{k: \nu_{j,k}^n \leq t} \xi_{j,k}^{v,n} \right|^2 = O(t^p), \quad \text{some } p > 0. \quad (4.9)$$

Dealing in detail with (4.8) and (4.9) will take us far afield, but they do hold under quite broad conditions. To illustrate one of the possibilities, we will give some of the details under the following condition.

A4.2. For each n , the random variables $(\Delta_{i,l}^{\alpha,n}, l < \infty)$ are mutually independent and identically distributed for each $i = 1, 2, \alpha = a, d$, and the absolute third moments are uniformly bounded. There are $\bar{\Delta}_i^{\alpha,n}$ and $\bar{\Delta}_i^\alpha$ such that $E\Delta_{i,l}^{\alpha,n} = \bar{\Delta}_i^{\alpha,n} \rightarrow \bar{\Delta}_i^\alpha$, $\alpha = a, d$. Also, the second moments of $\tau_{i,l}^{v,n}$ are uniformly bounded.

Theorem 4.2. Assume (A2.3), (A2.4) and (A4.2). Then (4.8) and (4.9) hold.

Proof. First, we consider (4.8). Fix α and i and define $\psi_{i,l}^{\alpha,n} = (1 - \Delta_{i,l}^{\alpha,n} / \bar{\Delta}_i^{\alpha,n})$. Let $\mathcal{F}_{i,l}^{\alpha,n}$ denote the minimal ν -algebra which measures $\{\psi_{i,j}^{\alpha,n}, j \leq l\}$, and write $E_{i,l}^{\alpha,n}$ for the associated conditional expectation. The $\psi_{i,l}^{\alpha,n}$ are martingale differences in that $E_{i,l}^{\alpha,n} \psi_{i,l+1}^{\alpha,n} = 0$ with probability one for all l . There is $K_2 < \infty$ such that

$$E_{i,l}^{\alpha,n} \left| \psi_{i,l+1}^{\alpha,n} \right|^2 \leq K_2.$$

Define

$$N_i^{\alpha,n}(t) = \frac{1}{n} \times \min \left\{ n : \sum_{l=1}^n \Delta_{i,l}^{\alpha,n} \geq nt \right\}.$$

The $S_i^{\alpha,n}(t)$ and $N_i^{\alpha,n}(t)$ will differ by at most $1/n$. The $N_i^{\alpha,n}(t)$ have the advantage that they are stopping times with respect to the filtrations $\mathcal{F}_{i,l}^{\alpha,n}$. In particular, $\{\omega : N_i^{\alpha,n}(t) \geq l\} \in \mathcal{F}_{i,l-1}^{\alpha,n}$. We have

$$E \max_{s \leq t} |w_i^{\alpha,n}(S_i^{\alpha,n}(s))|^2 \leq E \max_{m \leq nN_i^{\alpha,n}(t)} \frac{1}{n} \left| \sum_{l=1}^m \psi_{i,l}^{\alpha,n} \right|^2. \quad (4.10)$$

Owing to the martingale properties, the right hand side of (4.10) is bounded by $K_2 EN_i^{\alpha,n}(t)$. Thus, we need to bound $EN_i^{\alpha,n}(t)$.

For an integer $m > 0$, write

$$\frac{m \wedge nN_i^{\alpha,n}(t)}{n} = \frac{1}{n} \sum_{l=1}^{m \wedge (nN_i^{\alpha,n}(t))} 1 = \frac{1}{n} \sum_{l=1}^{m \wedge (nN_i^{\alpha,n}(t))} \psi_{i,l}^{\alpha,n} + \frac{1}{n\bar{\Delta}_i^{\alpha,n}} \sum_{l=1}^{m \wedge (nN_i^{\alpha,n}(t))} \Delta_{i,l}^{\alpha,n}. \quad (4.11)$$

The expectation of the first term on the right is zero. Dropping that term and letting $m \rightarrow \infty$ yields

$$EN_i^{\alpha,n}(t) = \frac{t}{\bar{\Delta}_i^{\alpha,n}} + \frac{1}{n} E[(\text{first time of arrival to source } i \text{ at or after } nt) - nt]. \quad (4.12)$$

Thus

$$EN_i^{\alpha,n}(t) = \frac{t}{\bar{\Delta}_i^{\alpha,n}} + \delta_n,$$

where $\lim_n \delta_n = 0$, and (4.8) holds. [The proof of the renewal theorem for the ‘‘excess life’’ in [13, pp 192–193] implies that $\delta_n \rightarrow 0$.]

Now turn to the proof of (4.9). To bound (4.9), we can use the expression

$$E \left| \xi_{j,k}^{v,n} \right|^2 \leq \sum_i \frac{1}{n} E \left| \sum_{l=nS_i^{\alpha,n}(\nu_{j,k}^n) + 1}^{nS_i^{\alpha,n}(\nu_{j,k}^n + \tau_{j,k}^{v,n}/\sqrt{n})} \Delta_{i,l}^{d,n} \right|^2. \quad (4.13)$$

Writing $\Delta_{i,l}^{d,n} = [\Delta_{i,l}^{d,n} - \bar{\Delta}_i^{d,n}] + \bar{\Delta}_i^{d,n}$ in (4.13) and splitting the upper bound in (4.13) into the two corresponding parts yields a bound on $E|\xi_{j,k}^{v,n}|^2$ as (twice) the sum of:

$$\sum_i \left[\bar{\Delta}_i^{d,n} \right]^2 E \left| w_i^{d,n}(\nu_{j,k}^n + \tau_{j,k}^{v,n}/\sqrt{n}) - w_i^{d,n}(\nu_{j,k}^n) \right|^2 \leq K_2 \sum_i \left[\bar{\Delta}_i^{d,n} \right]^2 E \tau_{j,k}^{v,n}/\sqrt{n}$$

and

$$\sum_i \frac{[\bar{\Delta}_i^{d,n}]^2}{n} E \left[\# \text{arrivals at queue } i \text{ in real time } \left[n\nu_{j,k}^n, n\nu_{j,k}^n + \sqrt{n}\tau_{j,k}^{v,n} \right] \right]^2.$$

The first expression is $O(1/\sqrt{n})$. Following the idea in the expansion (4.11), for the second expression we get the bound, for some real K_3 :

$$K_3 E \left[\tau_{j,k}^{v,n} \right]^2 + K_3 E [\text{a residual time term}]^2 / n.$$

But, by the cited proof of the renewal theorem for the “excess life” in [13, pp 192–193], and the third moment condition in (A4.2), the mean square value of the residual time term is bounded, uniformly in all indices.

We have obtained a bound for the mean square value of each of the jumps. To complete the proof, we need to average over the number of vacations on real time $[0, nt]$. We do this by ignoring the vacation durations in computing the distribution of the number of vacations on any real time interval $[0, nt]$, which gives an upper bound. Then the number has a Poisson distribution for each n , with the rate parameter being bounded in n , and the number is independent of the jump sizes. If there are L vacations on $[0, nt]$, then (4.9) is bounded by L^2 times the bound on the mean square value of each jump. Finally, using the Poisson distribution, average over L to get (4.9) for $p = 2$. ■

The Bellman equation for the limit system. Let \mathcal{L} denote the differential generator of the pure diffusion part of (3.7): I.e., for smooth real-valued $f(\cdot)$, $\mathcal{L}f(WL) = \sigma^2 f_{ww}(WL)/2 + bf_w(WL)$. Write the control in feedback form $u(t) = \theta(WL(t)) \leq WL(t)$ for some measurable function $\theta(\cdot)$. The jump part of the differential operator, acting on a measurable real-valued function $f(\cdot)$, is

$$\sum_i \bar{\lambda}_i^s E [f(WL + \xi_i^v) - f(WL)] \quad (4.14)$$

where ξ_i^v is the jump due to a vacation of source i , and E denotes the expectation of the jump given the WL and the control just before the start of the jump. The boundary condition is $f_w(0) = 0$.

Define $\bar{V}_\beta(WL)$ to be the inf of the cost over all admissible controls, not only those of the form in (A4.1). Define the function

$$H(\bar{V}_\beta, WL) = \min_{\theta(WL) \leq WL} \left\{ c(WL, \theta(WL)) + \sum_i \bar{\lambda}_i^s E [\bar{V}_\beta(WL + \xi_i^v) - \bar{V}_\beta(WL)] \right\}. \quad (4.15)$$

The formal Bellman equation is the partial differential integral equation

$$\mathcal{L}\bar{V}_\beta(WL) - \beta\bar{V}_\beta(WL) + H(\bar{V}_\beta, WL) = 0, \quad (4.16)$$

with the boundary condition $\bar{V}_{\beta, WL}(0) = 0$. The subscript WL denotes the derivative.

Conjecture and assumption. We have not been able to find anything in the literature concerning the PDE (4.16), where the jump magnitudes are controlled. To fully justify the restriction (A4.1), it is necessary to show both that (4.16)

has a unique (either classical or viscosity sense) solution which is the minimal cost and that the minimizing $\theta(\cdot)$ in (4.15) is of the type in (A2.6b). This seems to be a very reasonable expectation, although we have not been able to demonstrate it. As noted in the next section, it is essentially obvious in certain special cases; e.g., where $c_i(x) = x_i$, and we expect that it holds under broad conditions on $c(\cdot)$. Note that this is not an impulse control problem. The jump times are those of a Poisson process.

Thus, we assume that our conjecture is true: namely, that the minimum cost satisfies (4.16) and that the optimal control, given by the minimizer in (4.15), satisfies (A2.6b). Under this assumption, an optimal control for the limit problem is nearly optimal for the physical problem under heavy traffic, if the controls for the physical problem are restricted to a large enough class of the type in (A4.1).

5 Extensions and Comments

In special cases, the weak convergence results and the form of the limit problem suggest nearly optimal strategies for the physical problem, without much additional analysis. A case of current interest will be discussed.

Minimizing the total expected workload. Suppose that the cost rates $c_i(\cdot)$, written in terms of workload, satisfy $c_i(WL_i) = WL_i$. Then $c(WL, \theta(WL)) = WL$ and the control problem is the minimization of the expectation of the integral of the discounted total workload. The mean total workload $EWL(t)$ for the limit problem is minimized, uniformly in t , by using the policy $\theta(\cdot)$ that minimizes the mean jump, namely,

$$Q := \bar{\lambda}_1^s E [\rho_1 \tau_1^v - [WL - \theta(WL)]]^+ + \bar{\lambda}_2^s E [\rho_2 \tau_2^v - \theta(WL)]^+. \quad (5.1)$$

Example: the case of exponentially distributed vacation intervals. As an example, assume that τ_i^v is exponentially distributed with parameter v_i . Note that for any real number y and any random variable τ , exponentially distributed with parameter w , we have

$$E(\tau - y)^+ = w \int_y^\infty e^{-wx}(x - y)dx = w \int_0^\infty e^{-w(y+z)}zdz = \frac{e^{-wy}}{w}.$$

Denote for $i = 1, 2$, $j = 1, 2$, $j \neq i$,

$$w_i = \frac{v_i}{1 - \bar{\lambda}_j^a / \bar{\lambda}_j^d}.$$

Then we obtain

$$Q = \frac{\bar{\lambda}_1^s e^{-w_1(WL - \theta)}}{w_1} + \frac{\bar{\lambda}_2^s e^{-w_2 \theta}}{w_2}.$$

Thus, for each WL , Q is convex with respect to θ , and its minimum is obtained at θ for which $dQ(\theta)/d\theta = 0$, provided that this solution satisfies $\theta \in (0, WL)$.

If it does not, then the minimum over $\theta \in [0, WL]$ is obtained on one of the boundaries. Differentiating with respect to θ yields

$$\bar{\lambda}_1^s e^{-w_1(WL-\theta)} - \bar{\lambda}_2^s e^{-w_2\theta} = 0.$$

Solving this equation yields:

$$\theta(WL) = \frac{\log(\bar{\lambda}_2^s/\bar{\lambda}_1^s)}{w_1 + w_2} + \frac{w_1}{w_1 + w_2}WL.$$

The optimal policy $\theta(\cdot)$ for the limit problem is then given by

$$\theta^*(WL) = \left(\min \left(WL, \frac{\log(\bar{\lambda}_2^s/\bar{\lambda}_1^s)}{w_1 + w_2} + \frac{w_1}{w_1 + w_2}WL \right) \right)^+.$$

Example: the symmetrical case. In the special case where the two sources have the same rates, it is obvious that $\theta(WL) = WL/2$. Thus, under the conditions of Theorem 3.1 and the uniform integrability conditions of Theorem 4.1, the minimization of (5.1) yields a nearly optimal strategy for large n .

No vacations. The asymptotic optimality of the $c\mu$ -rule. Suppose that there are no vacations, and the basic desired cost rate is $\bar{c}_1 x_1^n + \bar{c}_2 x_2^n$, where $\bar{c}_i > 0$. Write the limit form of the cost rate in terms of the workload as

$$\bar{\lambda}_1^d \bar{c}_1 \theta(WL) + \bar{\lambda}_2^d \bar{c}_2 [WL - \theta(WL)]. \quad (5.2)$$

The minimizer of (5.2) is just the $c\mu$ -rule. Namely, poll source 1 if $\bar{\lambda}_1^d \bar{c}_1 > \bar{\lambda}_2^d \bar{c}_2$ and there are jobs there, and conversely for source 2. Under the conditions of Theorem 3.1 and the uniform integrability conditions, such a rule would be asymptotically optimal for the physical system. In this case, the limit workload does not depend on the polling policy, only the cost rate does. This is an asymptotic form of the well known $c\mu$ rule [31]. The asymptotic optimality of this rule under heavy traffic was given in [30].

The $c\mu$ rule gives priority to one of the queues, and this might lead to unacceptably long waits in the nonpriority queue. This can be alleviated with a nonlinear weighting. For example, queue 1 might have a smaller cost rate than queue 2 for moderate queue lengths. But to discourage the complete priority of queue 1, we might use a nonlinear cost rate for queue 2.

Remark. Note that the optimal policies in all the examples in this section indeed satisfy assumption A4.1, which is required in Theorem 4.1. Thus the class of policies described by A4.1 is rich enough to contain an optimal policy within it for these asymptotic problems.

6 Stability

Definition: Stability, uniformly in n for large n . Suppose that there are real n_0 and \bar{W} such that

$$E[\text{time for } WL^n(t), t \geq t_0, \text{ to return to the value } WL^n(t) \leq \bar{W} \mid \text{data to real time } nt_0, WL^n(t_0) = q] \leq K(q),$$

for all $n \geq n_0, t_0, q$, where $K(q)$ is bounded on each bounded q -set. Then we say that $WL^n(\cdot)$ is *stable, uniformly in n* .

Definition: Stability for fixed n . Fix n , and suppose that the above conditional mean return time property holds for all q and t_0 . Then, for that value of n , the queue is said to be *stable*.

We will use the following assumption, an modification of (A2.2).

A6.1. *There is a real $b_0 < 0$: $\sqrt{n} \left[\sum_i \rho_i^n - 1 \right] \leq b_0$ for all n .*

Comments on stability. Under (A6.1), it is trivial to prove the stability of the weak sense limit system (3.7) using classical stochastic Liapunov function methods, as in [14, 16, 17]. Stability is one of the most important properties of physical systems, and should be proved under broad conditions. It is essentially an assertion on the robustness of the system, and should hold under reasonable perturbations of the basic data. Stability of the physical system is not automatically guaranteed by stability of the weak sense limit. The technique to be employed is versatile and gets the desired stability property, uniformly in reasonable perturbations of the basic data, in the sense that the function $K(q)$ will not depend on the exact form of the data, under a reasonable mixing-type condition. The first definition above concerns large n . It will be seen that if there are no vacations then we can set $n_0 = 1$ in that definition, under broad conditions on the data.

Under the conditions of Theorem 3.1, the ratio of time on vacation to total time goes to zero as $n \rightarrow \infty$. If n is fixed and small, then it is conceivable that this ratio would be large enough so that the accumulation of data during the vacations will not be offset by the processing between vacations, as is necessary for stability. But, from the point of view of stability with vacations, there is an equivalence between large n and small $\bar{\lambda}_i^s$. This explains the last assertion of Theorem 6.2.

First, we will provide the motivation for the perturbed Liapunov function approach. Then it is used for the problem without vacations and stability uniformly in n (not just in large n) is proved. Then, vacations are added. We will simplify the algebra by supposing that arrivals to the queues can occur only

at multiples of (real time) $\delta > 0$, which can be as small as desired. Otherwise, we would use integrals in lieu of sums, but the results would be the same. Also, again for notational simplicity and with little loss of generality, we also suppose that vacations start and stop only at integral multiples of δ , and modify the assumptions (A2.3) and (A2.4) appropriately. We will also use (A2.5) plus other (weak) conditions to be imposed below. Let $E_{k\delta}^n$ denote the expectation, given all of the system data up to and including *real time* $k\delta$. Let $I_{i,k\delta}^{a,n}$ be the indicator function of an arrival at real time $k\delta$ from source i , and let $\Delta_{i,k\delta}^{d,n}$ be the associated work, if there is an arrival.

Motivation and background: Perturbed Liapunov functions. A perturbed Liapunov function method will be used [6, 17, 18, 22]. The classical Liapunov function method is quite limited, for problems such as ours, since there is not usually a “contraction” at each step to yield the local supermartingale property of a classical Liapunov function. The perturbed Liapunov function method is a powerful extension of the classical method. In the perturbed Liapunov function method, one adds a small perturbation to the original Liapunov function. As will be seen, when this perturbation can be well defined it provides an “averaging” which is needed to get the local supermartingale property.

The primary Liapunov function will be simply $WL(\cdot)$. The final Liapunov function will be of the form $W^n(\cdot) = WL^n(\cdot) + \delta W^n(\cdot)$, where $\delta W^n(\cdot)$ is bounded. Suppose that there is no vacation at real time $k\delta$. Then, for $WL^n(k\delta/n) \geq \delta$, we can write

$$E_{k\delta}^n WL^n(k\delta/n + \delta/n) - WL^n(k\delta/n) = -\frac{\delta}{\sqrt{n}} + \frac{1}{\sqrt{n}} \sum_i E_{k\delta}^n I_{i,k\delta+\delta}^{a,n} \Delta_{i,k\delta+\delta}^{d,n}. \quad (6.1)$$

The right hand term needs to be “averaged,” and this is done with the use of a perturbation function $\delta W^n(\cdot)$.

Motivation using a simpler problem. Before defining the actual perturbation which will be used, for motivation we will discuss the general principle with a simpler form when there are no vacations. Even for this problem, stability of the physical queues is not guaranteed by stability of the limit system. Let $\bar{\Delta}_i^{a,n} = 1/\bar{\lambda}_i^{a,n}$ and $\bar{\Delta}_i^{d,n}$ be centering constants such that the corresponding ρ_i^n satisfy (A6.1) for some $b_0 < 0$. More will be said about them later.

Proceeding formally until further notice, define the first suggested perturbation:

$$\delta \tilde{W}^n(k\delta/n) = \frac{1}{\sqrt{n}} \sum_i \sum_{j=k+1}^{\infty} E_{k\delta}^n \left[I_{i,j\delta}^{a,n} \Delta_{i,j\delta}^{d,n} - \delta \bar{\lambda}_i^{a,n} \bar{\Delta}_i^{d,n} \right]. \quad (6.2)$$

Clearly, the centering constants must be such that the sum is well defined, and we return to this point below. If $WL^n(k\delta/n) \geq \delta$, then we get

$$\begin{aligned} E_{k\delta}^n \delta \tilde{W}^n(k\delta/n + \delta/n) - \delta \tilde{W}^n(k\delta/n) \\ = -\frac{1}{\sqrt{n}} \sum_i E_{k\delta}^n \left[I_{i,k\delta+\delta}^{a,n} \Delta_{i,k\delta+\delta}^{d,n} - \delta \bar{\lambda}_i^{a,n} \bar{\Delta}_i^{d,n} \right]. \end{aligned}$$

Define $\tilde{W}^n(k\delta/n) = WL^n(k\delta/n) + \delta W^n(k\delta/n)$. Then (6.1) and the last expression yield

$$E_{k\delta}^n \tilde{W}^n(k\delta/n + \delta/n) - \tilde{W}^n(k\delta/n) = -\frac{\delta}{\sqrt{n}} + \frac{1}{\sqrt{n}} \sum_i \left[\delta \bar{\lambda}_i^{a,n} \bar{\Delta}_i^{d,n} \right].$$

By the condition (A6.1), the right side of the last expression is asymptotically $\leq b_0\delta/n$. Thus, it is less than the negative constant b_0 times the scaled time interval δ/n . Hence, when $WL^n(t) \geq \delta$, $\tilde{W}^n(k\delta/n)$ has the supermartingale property and we can use this to get the desired (uniform in n) stability if $\tilde{W}^n(\cdot)$ is “well defined and bounded.”

Let us examine the sum in (6.2) more closely to see why it is well defined and bounded under broad mixing conditions. Since $\bar{\Delta}_i^{d,n}$ and $\bar{\lambda}_i^{a,n}$ are merely *centering* constants for the *entire* sequence, the actual mean values or rates can vary with time (say, being periodic, etc.). Fix k and let $\mu_{i,1}\delta$ and $\mu_{i,2}\delta$ be the real times of the first two arrivals to queue i after real time $k\delta$. Formally, consider the part of the inner sum in (6.2) given by

$$E_{k\delta}^n \sum_{j=\mu_{i,1}+1}^{\mu_{i,2}} \left[I_{i,j\delta}^{a,n} \Delta_{i,j\delta}^{d,n} - \delta \bar{\lambda}_i^{a,n} \bar{\Delta}_i^{d,n} \right].$$

This equals

$$E_{k\delta}^n \left[\Delta_{i,\mu_{i,2}\delta}^{d,n} - (\mu_{i,2} - \mu_{i,1}) \delta \bar{\lambda}_i^{a,n} \bar{\Delta}_i^{d,n} \right]. \quad (6.3)$$

Next, suppose that the interarrival times and workloads are mutually independent, with the members of each set being mutually independent and identically distributed, with finite second moments, and means $\bar{\Delta}_i^{a,n}, \bar{\Delta}_i^{d,n}$. Then (6.3) equals zero, since $E_{k\delta}^n(\mu_{i,2} - \mu_{i,1})\delta = \bar{\Delta}_i^{a,n}$. Obviously, for any integer m , $\mu_{i,1}, \mu_{i,2}$ can be the m -th and $(m+1)$ -st arrival times with the same result. Thus, under the independence assumptions, (6.2) is just

$$E_{k\delta}^n \left[\Delta_{i,\mu_{i,1}\delta}^{d,n} - \delta(\mu_{i,1} - k) \bar{\lambda}_i^{a,n} \bar{\Delta}_i^{d,n} \right] = \bar{\Delta}_i^{d,n} E_{k\delta}^n \left[1 - \frac{\delta(\mu_{i,1} - k)}{\bar{\Delta}_i^{a,n}} \right],$$

where $E_{k\delta}^n(\mu_{i,1} - k)\delta$ is just the conditional expectation of the mean time to the next arrival after $k\delta$, given the data to time $k\delta$. For use below, keep in mind that this quantity is bounded uniformly in k , under the above assumptions on the independence and the moments. Hence, formally, $\delta \tilde{W}^n(t)$ is of the order of $1/\sqrt{n}$, uniformly in all variables.

Now, suppose that the interarrival times are as in the last paragraph, but the service times are correlated, still with centering constant $\bar{\Delta}_i^{d,n}$. Let $\mu_{i,j}, j = 1, \dots$, denote the sequence of arrival times after $k\delta$. Then (6.3) equals

$$E_{k\delta}^n \left[\Delta_{i,\mu_{i,2}\delta}^{d,n} - \bar{\Delta}_i^{d,n} \right]. \quad (6.4)$$

Then, grouping terms and formally speaking, we see that the sum (6.2) is just (6.3)/ \sqrt{n} , plus a series

$$\frac{1}{\sqrt{n}} \sum_i \sum_{l=\mu_{i,2}}^{\infty} E_{k\delta}^n \left[\Delta_{i,l}^{d,n} - \bar{\Delta}_{i,l}^{d,n} \right].$$

Clearly, the inner sum is bounded under quite broad mixing conditions. All that is needed is that $E_{k\delta}^n \left[\Delta_{i,l}^{d,n} - \bar{\Delta}_{i,l}^{d,n} \right] \rightarrow 0$ is a summable way as $l - k \rightarrow \infty$. A similar computation can be done if the $\Delta_{i,l}^{a,n}$ are correlated.

If the inner sum in (6.2) is well defined and bounded (uniformly in n, k, ω), then Theorem 6.1, which summarizes the above discussion, proves stability, uniformly in (all) n and the discounting which is used there is not needed. While the inner sums of (6.2) are well defined under broad conditions, there are interesting examples where they are not. For example, consider the case where $\Delta_{i,l}^{a,n} = H\delta = 1/\lambda_i^{a,n}$, where H is an integer and the work in all jobs is just the constant $\bar{\Delta}_i^{d,n}$. Then the inner sum, taken from $k+1$ to m , oscillates between zero and $-(H-1)\delta\bar{\lambda}_i^{a,n}\bar{\Delta}_i^{d,n}/\sqrt{n}$, as $m \rightarrow \infty$. The most convenient way of circumventing this problem is to suitably discount the defining sums [22, 27]. Thus, for some small $\beta > 0$, consider the alternative ‘‘discounted’’ perturbation:

$$\delta W_{\beta}^n(k\delta/n) = \frac{1}{\sqrt{n}} \sum_i \sum_{j=k+1}^{\infty} E_{k\delta}^n e^{-(j-k-1)\beta\delta/n} \left[I_{i,j\delta}^{a,n} \Delta_{i,j\delta}^{d,n} - \delta\bar{\lambda}_i^{a,n} \bar{\Delta}_i^{d,n} \right]. \quad (6.5)$$

The sum (6.5) is well defined if $E|\Delta_{i,l}^{d,n}|$ is uniformly bounded, and then the conditional expectation can be taken either inside or outside of the summation.

Stability without vacations. We now proceed to prove the stability results. It is simpler to start with the assumption that there are no vacations. The following additional assumption will be used. The above discussion shows that the assumption covers many cases of interest.

A6.2. *There is real B such that w.p.1 $|\delta W_{\beta}^n(k\delta/n)| \leq B/\sqrt{n}$, for all $\beta > 0$ and all n, k , where $\delta W_{\beta}^n(\cdot)$ is defined by (6.5).*

Define the final perturbed Liapunov function

$$W_{\beta}^n(k\delta/n) = WL^n(k\delta/n) + \delta W_{\beta}^n(k\delta/n). \quad (6.6)$$

Theorem 6.1. *Let $WL^n(0)$ be tight, suppose that there are no vacations, and assume (A2.5), (A6.1) and (A6.2). Then the process $WL^n(\cdot)$ is stable, uniformly in n .*

Proof. We have

$$\begin{aligned} & E_{k\delta}^n \delta W_{\beta}^n(k\delta/n + \delta/n) - \delta W_{\beta}^n(k\delta/n) \\ &= -\frac{1}{\sqrt{n}} \sum_i E_{k\delta}^n \left[I_{i,k\delta+\delta}^{a,n} \Delta_{i,k\delta+\delta}^{d,n} - \delta\bar{\lambda}_i^{a,n} \bar{\Delta}_i^{d,n} \right] + \epsilon_k^n. \end{aligned} \quad (6.7)$$

where

$$\epsilon_k^n = E_{k\delta}^n \left[1 - e^{-\beta\delta/n} \right] \delta W_\beta^n(k\delta/n + \delta/n). \quad (6.8)$$

Thus, adding (6.1) and (6.7),

$$E_{k\delta}^n W_\beta^n(k\delta/n + \delta/n) - W_\beta^n(k\delta/n) = -\frac{\delta}{\sqrt{n}} + \frac{1}{\sqrt{n}} \sum_i \left[\delta \bar{\lambda}_i^{a,n} \bar{\Delta}_i^{d,n} \right] + \epsilon_k^n. \quad (6.9)$$

By the condition (A6.1), the right side of (6.9) is asymptotically no greater than

$$\frac{b_0\delta}{n} + \epsilon_k^n. \quad (6.10)$$

Assumption (A6.2) implies that $|\epsilon_k^n| = B[\beta\delta/n]\gamma^n$, where $\gamma^n \rightarrow 0$. Thus, for small β

$$E_{k\delta}^n W_\beta^n(k\delta/n + \delta/n) - W_\beta^n(k\delta/n) \leq \frac{b_0\delta}{2n}, \quad \text{for } WL^n(k\delta/n) \geq \delta. \quad (6.11)$$

Inequality (6.11) implies that $W_\beta^n(k\delta/n)$ has the supermartingale property when $WL^n(k\delta/n) \geq \delta$. Suppose that $W_\beta^n(k\delta/n) = B_2 > B + \delta$ and let $B_2 > B_1 > B + \delta$. Then, by standard stability arguments [16, 17], the mean number of steps (of real time length δ and conditioned on the data to real time $k\delta$) for $W_\beta^n(j\delta/n)$, $j \geq k$, to return to the set where $W_\beta^n(j\delta/n) \leq B_1$ is bounded by

$$\frac{W_\beta^n(k\delta/n)}{[-b_0\delta/2n]} \leq \frac{B + WL^n(k\delta/n)}{[-b_0\delta/2n]}.$$

Since $|\delta W_\beta^n(t)| \leq B/\sqrt{n}$, the return time estimate also holds for $WL^n(\cdot)$. Thus, in the time scale which is used to define $WL^n(\cdot)$, where time is compressed by a factor of n , the conditional mean return time is asymptotically bounded by $2[B + WL^n(t)]/[-b_0]$. Hence, we have stability, uniformly in n . ■

Stability, with vacations. Now, we add the vacations. Again, to simplify the notation, suppose a preempt-resume discipline, so that we do not have to concern ourselves with redoing all of an interrupted job. The analysis for the latter case follows similar lines.

The polling policy is subject only to the unrestrictive condition (A6.3). The condition is motivated by the fact that the $\xi_{i,l}^v$ defined by (3.5) go to zero as the individual workloads go to infinity, since the larger the work remaining in the nonvacationing sources, the less likely it is that the server will have idle time during a vacation. If (A6.3) does not hold, then there might not be stability for *each* $b < 0$, uniformly in large n . For example, suppose that the polling policy is to give source 1 priority. Then $WL_1^n(t)$ will be arbitrarily close to zero, except possibly during and for a short interval just after a vacation of that source. Consequently, the mean or conditional mean jump in the total workload during a vacation of source 2 which starts at scaled time t will not go to zero as $WL^n(t)$ goes to infinity. The condition (A6.3) excludes exhaustive polling (but

only when the workload is very large), where a source is polled until its queue is empty, unless a vacation of that source intervenes. But, when the total workload is large, we might not want to use exhaustive polling anyway. While we work with two sources for notational simplicity, the idea is the same no matter what the number of sources. By our convention, this real time vacation interval is the half open interval $[k\delta + \delta, k\delta + \delta + \sqrt{n}\tau_{i,k\delta}^{v,n})$.

A6.3. *The polling policy is unrestricted, except for the following. There are constants $\bar{W}_a \ll \bar{W}_b$, which will be as large as we wish. If $WL^n(t) \leq \bar{W}_b$, then there are no restrictions. If $WL^n(t) > \bar{W}_b$, then the only restriction is that if*

$$WL_i^n(t) \geq \bar{W}_a \tag{6.12}$$

is not satisfied for some i and the other source is not on vacation, then we poll the other source.

A6.4. *For each $\epsilon > 0$, there is $\bar{W} < \infty$ such that for $i, j : i \neq j$,*

$$E \left[\xi_{i,l}^{v,n} \mid \text{data to scaled time } \nu_{i,l}^n, WL_j^n(\nu_{i,l}^n -) \geq \bar{W} \right] \leq \epsilon.$$

The assumption (A6.4) holds under the conditions of Theorem 4.2. The condition (6.12) cannot be guaranteed for all time. It will sometimes not hold during a vacation or for a vanishingly short (in scaled time) interval after. The real time interval between vacations is $O(n)$. Let $\alpha(\cdot)$ be a real valued function on $[0, \infty)$ such that $\alpha(n)/\sqrt{n} \rightarrow \infty$ and $\alpha(n)/n \rightarrow 0$. Suppose that a vacation ends at real time t_0 and the next one begins at real time t_1 , with $t_1 - t_0 = O(n)$. Then with a probability (conditioned on the data up to t_0) that goes to unity as $n \rightarrow \infty$, one can poll such that (6.12) is guaranteed on $[t_0 + \alpha(n), t_1)$. The excluded interval is just $\alpha(n)/n$ in scaled time. Condition (A6.3) works since the probability that two successive vacations will be within $\alpha(n)$ in real time is $O(\alpha(n)/n)$.

Theorem 6.2. *Let $\{WL^n(0)\}$ be tight and assume (A2.3)–(A2.5) and (A6.1)–(A6.4). Then the process $WL^n(\cdot)$ is stable, uniformly in n . Fix n . Then for small enough $\bar{\lambda}_i^s, i = 1, 2$, $WL^n(\cdot)$ is stable.*

Note on the stability of the limit problem (3.7). Let \mathcal{L} denote the differential generator of (3.7) when $WL > 0$. Then

$$\mathcal{L}WL(t) = b + \sum_i \bar{\lambda}_i^s E_{WL(t-), u(t-)} \xi_i^v.$$

Since, for the limit problem, the condition (6.12) can always be guaranteed for $WL(t) > \bar{W}_a$ (except at the jump instants) and $\bar{W}_b \gg \bar{W}_a$ are as large as we wish, it can always be assured that the sum in the above expression is

arbitrarily small for large $WL^n(t)$. Then, since $b < 0$, $WL(\cdot)$ is stable. The proof below attempts to duplicate this idea.

Proof. In this proof, it is more convenient to work in scaled time. Thus, let $E_t^{s,n}$ denote the expectation conditioned on all data to scaled time t . All scaled times are integral multiples of δ/n . Suppose that no source is on vacation at scaled time t and a vacation of some source starts at scaled time $t + \delta/n$. Then, let $\tau_t^{v,n}/\sqrt{n}$ denote the scaled time which passes until *no* source is on vacation, and let $\xi_t^{v,n}$ denote the total jump in the workload due to all vacations which start at scaled time $t + \delta/n$ and end at scaled time $t + \delta/n + \tau_t^{v,n}/\sqrt{n}$. Thus, it might cover a single vacation, or several overlapping or abutting vacations.

Define $\sigma_k^n, k \geq 0$, recursively as follows. Start with $\sigma_0^n = 0$. Given σ_k^n , if no vacation starts at scaled time $\sigma_k^n + \delta/n$, then set $\sigma_{k+1}^n = \sigma_k^n + \delta/n$. If a vacation starts at scaled time $\sigma_k^n + \delta/n$, then set $\sigma_{k+1}^n = \sigma_k^n + \delta/n + \tau_{\sigma_k^n}^{v,n}/\sqrt{n}$. Thus, the σ_k^n are the sequence of scaled times $k\delta/n$, but with the instants where some source is on vacation skipped. To prove the stability it is sufficient to work with $WL^n(\sigma_k^n)$ and $WL^n(\sigma_k^n) \geq \bar{W}_a$ only.

Until further notice, suppose that the condition (A6.3) holds at the start of each vacation. The event that this is not the case is very rare, for large n and will be accounted for later. Recall the definition of $W_\beta^n(\cdot)$ in (6.6). Let $E_t^{v,n}$ denote the expectation, conditioned on all data to scaled time t and the event that a vacation starts at scaled time $t + \delta/n$. By the computations in Theorem 6.1, we have (whether or not (6.12) holds)

$$\begin{aligned} & E_{\sigma_k^n}^{s,n} W_\beta^n(\sigma_{k+1}^n) - W_\beta^n(\sigma_k^n) \\ & \leq \prod_i \left(1 - \frac{\bar{\lambda}_i^{s,n}}{n} + o\left(\frac{\bar{\lambda}_i^{s,n}}{n}\right) \right) \frac{b_0 \delta}{2n} + \left[\sum_i \frac{\delta}{n} \bar{\lambda}_i^{s,n} + o\left(\frac{\delta \sum_i \bar{\lambda}_i^{s,n}}{n}\right) \right] E_{\sigma_k^n}^{v,n} \xi_{\sigma_k^n}^{v,n}. \end{aligned} \quad (6.13)$$

The $o(\cdot)$ will be ignored henceforth. Now, by (A6.4) the term $E_{\sigma_k^n}^{v,n} \xi_{\sigma_k^n}^{v,n}$ in (6.13) goes to zero as $WL^n(\sigma_k^n)$ goes to infinity, which yields the stability, uniformly in n for large n .

Next let us consider the possibility that we might not always have $WL_i^n(t) \geq \bar{W}_a, i = 1, 2$, at the start of a vacation, when $WL^n(t) \geq \bar{W}_b$, for $\bar{W}_b \gg \bar{W}_a$, both being sufficiently large. Let $I_{t+\delta/n}^{v,n}$ denote the event that a vacation starts at scaled time $t + \delta/n$, with $WL_i^n(t) \leq \bar{W}_a$ for some i and $WL^n(t) \geq \bar{W}_b$. Let μ_l^n denote the scaled time of the l -th occurrence of this event. We exploit the fact that this event is ‘‘rare’’ for large n , by introducing another perturbation to the Liapunov function. We need to add

$$E_{\sigma_k^n}^{s,n} I_{\sigma_k^n + \delta/n}^{v,n} \tau_{\sigma_k^n}^{v,n} \quad (6.14)$$

to the right side of (6.13) (and multiply the right hand term there by $(1 - I_{\sigma_k^n + \delta/n}^{v,n})$), which leaves the estimates for that term unchanged. By (A2.4), (6.14) is bounded by $K_1 E_{\sigma_k^n}^{s,n} I_{\sigma_k^n + \delta/n}^{v,n}$ for some constant K_1 . Introduce the additional

perturbation to the Liapunov function:

$$\delta\bar{W}_\beta^n(k\delta/n) = K_1 \sum_{l=k+1}^{\infty} e^{-(l-k-1)\beta\delta/n} E_{k\delta/n}^{s,n} I_{l\delta/n}^{v,n}. \quad (6.15)$$

This equals (dropping $o(\delta/n)$ terms for simplicity)

$$K_1 \frac{\delta \sum_i \bar{\lambda}_i^{s,n}}{n} \sum_{l: \mu_l^n > k\delta/n} E_{k\delta/n}^{s,n} e^{-\beta(\mu_l^n - k\delta/n - \delta/n)}. \quad (6.16)$$

Write the sum as $K_\beta^n(k)$. Then, for each $\beta > 0$, there is $n(\beta) < \infty$ such that $K_\beta^n(k) \leq 2$ for $n \geq n(\beta)$ and all k .

Note the difference of the conditional expectations:

$$E_{\sigma_k^n}^{s,n} \delta\bar{W}_\beta^n(\sigma_{k+1}^n) - \delta\bar{W}_\beta^n(\sigma_k^n) = -K_1 E_{\sigma_k^n}^{s,n} I_{\sigma_k^n + \delta/n}^{v,n} + \epsilon_k^{v,n}, \quad (6.17)$$

where

$$\epsilon_k^{v,n} = K_1 \left[1 - e^{-\beta\delta/n}\right] \sum_{l=k+1}^{\infty} e^{-(l-k-2)\beta\delta/n} E_{k\delta/n}^{s,n} I_{l\delta/n}^{v,n}. \quad (6.18)$$

For $n \geq n(\beta)$,

$$\epsilon_k^{v,n} \leq 2K_1\beta \sum_i \frac{\bar{\lambda}_i^{s,n}}{n}.$$

Now, use the new perturbed Liapunov function defined by $W_\beta^n(k\delta/n) + \delta\bar{W}_\beta^n(k\delta/n)$. The conditional difference (6.17) cancels (6.14), modulo the error $\epsilon_k^{v,n}$, which is $O(\delta\beta/n)$ and β can be made as small as desired for large enough n . The proof is then completed as in Theorem 6.1 ■

7 Unreliable Channels

Up to now, it was supposed that any data sent from any source to the server arrived without error. In this section, we suppose that an error during transmission (as distinct from a vacation) is possible. Suppose that the server time is divided into “slots,” of duration $\delta > 0$. I.e., the work in each arrival is an integral multiple of δ , and each δ -interval is devoted to either a job from one of the sources, or to idling if there is no work present at the beginning of the slot. If a vacation starts during a slot, it is assumed that that the data in that slot (if any) is retransmitted later. But, this has no effect on the heavy traffic limit. In case of an error in transmission, the data in the slot must be retransmitted. A variety of such situations can be readily incorporated into our general model.

If the sequence of errors are mutually independent in time, or if they do not depend on the source, then the modeling and analytical problem is relatively simple. The errors would not depend on the source if the corrupting noise were at the server/base station, or was due to, say, a general atmospheric condition

which affects all sources in the same way. On the other hand, if the errors are correlated (say, channels with bursty noise) and are *source dependent* as well, then the modeling problem is complicated by the fact that the server is allowed to poll the sources in a rather arbitrary way. For example, suppose that the noise is bursty for the channel connecting to one of the sources but that the channel connecting to the other is noise free. Then, depending on how the sources are sequenced, the correlation between the errors can take many forms. Thus, it is hard to know the relation between the sequencing of the polling of the sources and the channel noise. One could try to poll taking into account the correlation. But, even if this were feasible, it is beyond our goals. For these reasons, we will assume that the disturbing noise does not depend on the source, despite the importance of the general problem.

The error model. Let $I_l^{e,n}$ denote the indicator function of the event that the data transmitted in the l -th time slot was not acceptable and needed to be retransmitted. Let $S_i^{e,n}(t)$ (resp., $S^{e,n}(t)$) denote $1/n$ times the number of slots transmitted (successfully or not) from source i (resp., from both sources) by real time nt . Let $I_l^{a,n}$ denote the indicator function of the event that there is available data to be transmitted in time slot l from any source not on vacation. The (scaled) work that must be retransmitted by real time nt is

$$L^n(t) = \frac{\delta}{\sqrt{n}} \sum_{l=1}^{nt/\delta} I_l^{e,n} I_l^{a,n}. \quad (7.1)$$

For some centering constant $p^{e,n}$, write $L^n(t)$ as

$$L^n(t) = \frac{\delta}{\sqrt{n}} \sum_{l=1}^{nt/\delta} [I_l^{e,n} - p^{e,n}] I_l^{a,n} + \sqrt{n} p^{e,n} S^{e,n}(t) \delta \quad (7.2)$$

The last term on the right of (7.2) is (see (3.2))

$$p^{e,n} [\sqrt{nt} - T^{v,n}(t) - Z^n(t)]. \quad (7.3)$$

Define

$$w^{e,n}(t) = \frac{\delta}{\sqrt{n}} \sum_{l=1}^{nt/\delta} [I_l^{e,n} - p^{e,n}] I_l^{a,n}. \quad (7.4)$$

We will use the following assumptions.

A7.1. *The error process is independent of all of the other driving processes. Also, $p^{e,n}$ converges to the constant p^e as $n \rightarrow \infty$ and $w^{e,n}(\cdot)$ converges weakly to a Wiener process, which will be denoted by $w^e(\cdot)$, and whose variance is $\delta\sigma_e^2$.*

A7.2. [The new heavy traffic condition.] *There is a constant b such that*

$$\lim_n \sqrt{n} \left[\sum_i \rho_i^n + p^{e,n} - 1 \right] = b.$$

(A7.1) is an assumption on the channel and will be returned to below. By adding the work to be retransmitted, (3.15) becomes

$$\begin{aligned} WL^n(t) &= WL^n(0) + \sum_i \bar{\Delta}_i^{d,n} \left[w_i^{a,n}(S_i^{a,n}(t)) - w_i^{d,n}(S_i^{a,n}(t)) \right] + w^{\epsilon,n}(t) \\ &+ \sqrt{n} \left[\sum_i \rho_i^n + p^{\epsilon,n} - 1 \right] t + (1 - p^{\epsilon,n}) [Z^n(t) + T^{v,n}(t)] + \epsilon^n(t), \end{aligned} \quad (7.5)$$

where $\epsilon^n(\cdot)$ is a residual time error process.

Under the conditions of Theorem 3.1, with (A7.1) added and the new heavy traffic condition (A7.2) used, Theorem 3.1 continues to hold, with the following changes. The process $w^\epsilon(\cdot)$ is added to $w(\cdot)$. The jumps are computed by first showing that (in the local fluid time scale) the processes of completed work *during a vacation* can be asymptotically approximated by a fluid process with slope $1 - p^\epsilon$, and they are

$$\begin{aligned} \xi_{1,l}^v &= [((1 - p^\epsilon) - \rho_2) \tau_{1,l}^v - [WL(\nu_{1,l-}) - u(\nu_{1,l-})]]^+ \\ &= [\rho_1 - [WL(\nu_{1,l-}) - u(\nu_{1,l-})]]^+. \end{aligned} \quad (7.6a)$$

$$\xi_{2,l}^v = [((1 - p^\epsilon) - \rho_1) \tau_{2,l}^v - u(\nu_{2,l-})]^+ = [\rho_2 \tau_{2,l}^v - u(\nu_{2,l-})]^+. \quad (7.6b)$$

$w^\epsilon(\cdot)$ is independent of $w(\cdot)$. With these changes, Theorem 3.2 also holds. Theorem 4.2 will continue to hold with these changes, provided that $E|w^{\epsilon,n}(t)|^2 = O(t)$. Similarly, the analogs of the stability results hold.

Remark. It is not possible to account for the retransmissions by simply enlarging the work in each job by an amount that has the same distribution as the retransmitted work does. This is because the controls are based on either the current queued work or queued numbers, and not what might be expected due to future errors and retransmissions.

Comments concerning $w^{\epsilon,n}(\cdot)$. First, suppose that the errors are independent from slot to slot with $P\{I_l^{\epsilon,n} = 1\} = p^{\epsilon,n}$. Then Donsker' Theorem [12] implies that $w^{\epsilon,n}(\cdot)$ is tight and converges weakly to a Wiener process process with variance $\delta p^\epsilon(1 - p^\epsilon)$.

Now, turn to the correlated error problem. The error process concerns the channel, and is defined whether or not there is something to be transmitted. Suppose that the error process is Markov and doesn't depend on n , for notational simplicity. In particular, assume that

$$P\{I_{l+1}^{\epsilon,n} = 1 | I_l^{\epsilon,n} = 0\} = p, \quad P\{I_{l+1}^{\epsilon,n} = 0 | I_l^{\epsilon,n} = 1\} = q,$$

where p and q are in $(0, 1)$. Then $p^\epsilon = p/(p + q)$. Let I_l^ϵ denote the stationary error process. Again, it is not hard to verify that $w^{\epsilon,n}(\cdot)$ converges weakly to a Wiener process with variance

$$\delta E [I_l^\epsilon - p^\epsilon]^2 + 2\delta E \sum_{l=1}^{\infty} [I_l^\epsilon - p^\epsilon] [I_0^\epsilon - p^\epsilon]$$

[12, 17].

We have $I_l^{a,n} = 0$ if both sources are on vacation, both queues are empty or one source is on vacation and the other queue is empty at real time $l\delta$. These possibilities have negligible effect asymptotically.

Lévy processes. Many other models are possible for the error process (7.1) and a couple of other possibilities will be outlined. One approach, which does not require the addition of (A7.1) and uses (3.5) for the jumps, is to simply suppose that $L^n(\cdot)$ converges weakly to a general Lévy jump process. For example, suppose that the noise occurs in occasional bursts, where the rate at which the bursts occur (in real time) is $\bar{\lambda}^{e,n}/n$ and the duration (in real time) is $\sqrt{n}\tau_l^{e,n}$, where the durations (and the process of starting) are mutually independent and independent of the other “driving” random variables in the system. In this model the bursts are rare, and occur at a rate which is of the order of that of the vacations. But $\bar{\lambda}^{e,n}$ might be much larger than $\bar{\lambda}^{s,n}$, and the $\tau_l^{e,n}$ much smaller than $\tau_{i,i}^{v,n}$.

The scheme in the last paragraph supposed a finite rate $\bar{\lambda}^{e,n}$ for the bursts. The rate could depend on the duration, so that shorter durations have higher rates, with the rate going to infinity as the duration goes to zero, but in such a way that there is a limit Lévy process.

References

- [1] E. Altman, D. Kofman, and U. Yechiali. Discrete time queues with delayed information. *Queueing Systems*, 19:361–376, 1995.
- [2] E. Altman and H.J. Kushner. Admission control for combined guaranteed performance and best effort communications systems under heavy traffic. *SIAM J. Control and Optimization*, 37:1780–1807, 1999.
- [3] E. Altman and A. Shwartz. Optimal priority assignment: a time sharing approach. *IEEE Trans. on Automatic Control*, 34:1098–1102, 1989.
- [4] J.S. Baras, D.J. Ma, and A.M. Makowski. K competing queues with geometric service requirements and linear costs: the μc rule is always optimal. *Systems and Control Letters*, 6:173–180, 1985.
- [5] P. Billingsley. *Convergence of Probability Measures*. Wiley, New York, 1968.
- [6] G. Blankenship and G.C. Papanicolaou. Stability and control of systems with wide band noise disturbances. *SIAM J. Appl. Math.*, 34:437–476, 1978.
- [7] M. Bramson. State space collapse with application to heavy traffic limits for multiclass queueing networks. *Queueing Systems*, 30:89–148, 1998.

- [8] M. Carr and B. Hajek. Scheduling with asynchronous service opportunities with applications to multiple satellite systems. *IEEE Trans. on Automatic Control*, 38:1820–1833, 1993.
- [9] D. R. Cox and W. L. Smith. *Queues*. Methuen, London, 1961.
- [10] P. Dupuis and H. Ishii. On Lipschitz continuity of the solution mapping to the Skorokhod problem, with applications. *Stochastics Stochastics Rep.*, 35:31–62, 1991.
- [11] B.R. Elbert. *The Satellite Communication Applications Handbook*. Artech House, Boston, London, 1997.
- [12] S.N. Ethier and T.G. Kurtz. *Markov Processes: Characterization and Convergence*. Wiley, New York, 1986.
- [13] S. Karlin and H.M. Taylor. *A First Course in Stochastic Processes*. Academic Press, New York, second edition, 1975.
- [14] R.Z. Khasminskii. *Stochastic Stability of Differential Equations*. Sijthoff, Noordhoff, Alphen aan den Rijn, Amsterdam, 1982.
- [15] H. Kunita and S. Watanabe. On square integrable martingales. *Nagoya Math. J.*, 30:209–245, 1967.
- [16] H.J. Kushner. *Stochastic Stability and Control*. Academic Press, New York, 1967.
- [17] H.J. Kushner. *Approximation and Weak Convergence Methods for Random Processes with Applications to Stochastic Systems Theory*. MIT Press, Cambridge, Mass., 1984.
- [18] H.J. Kushner. *Weak Convergence Methods and Singularly Perturbed Stochastic Control and Filtering Problems*, volume 3 of *Systems and Control*. Birkhäuser, Boston, 1990.
- [19] H.J. Kushner and P. Dupuis. *Numerical Methods for Stochastic Control Problems in Continuous Time*. Springer-Verlag, Berlin and New York, 1992.
- [20] H.J. Kushner, D. Jarvis, and J. Yang. Controlled and optimally controlled multiplexing systems: A numerical exploration. *Queueing Systems*, 20:255–291, 1995.
- [21] H.J. Kushner and K.M. Ramachandran. Optimal and approximately optimal control policies for queues in heavy traffic. *SIAM J. Control Optim.*, 27:1293–1318, 1989.
- [22] H.J. Kushner and G. Yin. *Stochastic Approximation Algorithms and Applications*. Springer-Verlag, Berlin and New York, 1997.

- [23] R. Liptser and A.N. Shiryaev. *Statistics of Random Processes*. Springer-Verlag, Berlin and New York, 1977.
- [24] W.P. Peterson. Diffusion approximations for networks of queues with multiple customer types. *Math. Oper. Res.*, 9:90–118, 1951.
- [25] M.I. Reiman. Some diffusion approximations with state space collapse. In F. Baccelli and G. Fayolle, editors, *Proc. Int. Seminar on Modelling and Performance Evaluation Methodology*, pages 209–240. Springer Verlag, 1983.
- [26] M.I. Reiman. A multiclass feedback queue in heavy traffic. *Advances in Appl. Prob.*, 20:179–207, 1998.
- [27] V. Solo and X. Kong. *Adaptive Signal Processing Algorithms*. Prentice-Hall, Englewood Cliffs, NJ, 1995.
- [28] L. Tassiulas and A. Ephremides. Dynamic server allocation to parallel queues with randomly varying connectivity. *IEEE Trans. on Automatic Control*, 39:466–478, 1993.
- [29] L. Tassiulas and S. Papavassiliou. Optimal anticipative scheduling with asynchronous transmission opportunities. *IEEE Trans. on Automatic Control*, 40:2052–2062, 1995.
- [30] J.A. van Mieghem. Dynamic scheduling with convex delay costs: The generalized $c\mu$ rule. *Ann. Appl. Prob.*, 5:809–833, 1995.
- [31] J. Walrand. *An Introduction to Queuing Networks*. Prentice-Hall, Englewood Cliffs, 1988.
- [32] R.J. Williams. Diffusion approximations for open multiclass queueing networks: sufficient conditions involving state space collapse. *Queueing Systems*, 30:27–88, 1998.