The Annals of Applied Probability
2000, Vol. 10, No. 3, 897–929

DISCRETE-REVIEW POLICIES FOR SCHEDULING STOCHASTIC NETWORKS: TRAJECTORY TRACKING AND FLUID-SCALE ASYMPTOTIC OPTIMALITY

BY CONSTANTINOS MAGLARAS

Columbia University

This paper describes a general approach for dynamic control of stochastic networks based on fluid model analysis, where in broad terms, the stochastic network is approximated by its fluid analog, an associated fluid control problem is solved and, finally, a scheduling rule for the original system is defined by interpreting the fluid control policy.

The main contribution of this paper is to propose a general mechanism for translating the solution of the fluid optimal control problem into an implementable discrete-review policy that achieves asymptotically optimal performance under fluid scaling, and guarantees stability if the traffic intensity is less than one at each station. The proposed policy reviews system status at discrete points in time, and at each such point the controller formulates a processing plan for the next review period, based on the queue length vector observed, using the optimal control policy of the associated fluid optimization problem. Implementation of such a policy involves enforcement of certain safety stock requirements in order to facilitate the execution of the processing plans and to avoid unplanned server idleness.

Finally, putting aside all considerations of system optimality, the following generalization is considered: every initial condition is associated with a feasible fluid trajectory that describes the desired system evolution starting at that point. A discrete-review policy is described that asymptotically tracks this target specification; that is, it achieves the appropriate target trajectory as its fluid limit.

1. Introduction. Today's communication, computer and manufacturing industries offer many examples of technological systems in which "units of work" visit a number of different "servers" in the course of their processing, and in which the workflow is subject to stochastic variability. In this paper such processing systems are modeled as open multiclass queueing networks. These are networks populated by many job classes that may differ in their arrival processes, service requirements, and routes through the network, and there is a many-to-one relation between job classes and servers. The system controller has discretion as to the sequencing of jobs of the various classes at each server, and a rule according to which these decisions are made is called a scheduling policy. For this class of processing networks we study the problem of finding an admissible policy which is optimal or near-optimal under a given performance metric.

Received February 1999; revised October 1999.

AMS 1991 *subject classifications*. Primary 60K25; secondary 90B35, 60F17.

Key words and phrases. Queueing networks, fluid models, discrete-review policies, trajectory tracking, asymptotic optimality.

Dynamic control problems for stochastic processing networks are both analytically and computationally hard. While most often one relies on the use of heuristics that are validated through simulation studies, one approach that has emerged from research over the past ten to fifteen years is based on a hierarchy of approximate models that provide tractable "relaxations" of these problems as a framework for analysis and synthesis. In particular, the analytical theory associated with fluid approximations has produced important insights in understanding how the performance of a multiclass network depends on different design and control parameters. This is our starting point. Specifically, the approach taken here is based on approximating (or replacing) the stochastic network by its fluid analog (this is a model with deterministic and continuous dynamics), solving an associated fluid optimal control problem and then using the derived fluid control policy in order to define an implementable rule for the stochastic network. This procedure is summarized below.

1. Consider a dynamic control problem for the original stochastic network.
2. Form fluid analog of stochastic network and solve the associated fluid optimal control problem.
3. Translate/implement the optimal fluid control in original stochastic network.
4. Consider fluid limit of stochastic network under implemented policy.
5. Verify fluid-scale asymptotic optimality and stability.

Stages 1 to 3 are clear. Stages 4 and 5 describe a criterion for performance analysis under the implemented policy that is consistent with the model approximation adopted at stage 2 in the following sense: the implementation is tested for asymptotic optimality in the limiting regime where the model approximation is valid. This criterion is referred to as *fluid-scale asymptotic optimality* (FSAO). To be more precise, fluid limits are derived through a functional strong law of large numbers (FSLLN) type of scaling, where one observes the system behavior starting from a large initial condition over a proportionally long time horizon, which essentially yields a deterministic transient response model. Thus, the proposed criterion tests whether in the fluid limit regime the system's limiting performance achieves that of the optimal fluid (transient) response that was used in stage 3 in designing the policy under investigation. This is a "minimal" requirement for the implemented policy. In comparison to the original problem at hand, it provides a relaxed notion of (transient) optimality that appears to be simpler and one that hopefully could be achieved even for a general collection of multiclass networks. Finally, apart from FSAO, we also require that the original stochastic network is stable under the implemented policy provided that the traffic intensity parameter at each station is less than one; roughly speaking, stability implies that queue lengths stay finite and that an appropriately defined underlying Markov chain is positive Harris recurrent. Later on, we will show that for reasonable cost criteria (such as convex, increasing cost rate functions over long time horizons), FSAO will imply stability.

The main issue that arises within this policy design framework, and one that has yet to be broached in the existing literature, is in finding a mechanism to translate the solution of the associated fluid optimal control problem into an implementable control policy in the stochastic network in a way that guarantees the criteria described above. The main contribution of this paper is to describe and analyze a family of discrete-review (DR) policies that provide the first general translation mechanism that guarantees fluid-scale asymptotic optimality and stability. In passing, we will give a precise mathematical articulation of the fluid-scale asymptotic optimality criterion. Our second contribution is in generalizing this approach to a family of *trajectory tracking* discrete-review policies. Our primitive fluid control policy is now defined by specifying a feasible fluid trajectory that describes a desired evolution for the queue length process starting at every initial condition. A tracking DR policy will be described that asymptotically under fluid scaling (i.e., as the backlog and the system observation time increase according to a LLN type of scaling) achieves the appropriate target trajectory as its fluid limit. This asymptotic tracking criterion is similar to that of FSAO, and later we will see how the latter can be recovered from the former for a special choice of target trajectories. The generalization to tracking policies has some important practical implications. Target trajectory specifications can now be derived through simpler or alternative methodologies that avoid any explicit optimization in the fluid model which, although vastly simpler than the original problem at hand, can still be prohibitive, especially for on-line applications where the "thinking" time allowed is limited. The idea of trajectory tracking apart from being intuitively appealing, it borrows from a very extensive literature in the area of control theory; there, control policies are designed in order to be able to track a reference signal (the queue length trajectory in our case), which, in turn, is dynamically adjusted as a function of the current state.

The focus on fluid approximations is primarily motivated by recent developments in the area of stability analysis of stochastic networks via fluid model analysis. The important breakthrough in this area was the theory developed by Dai (1995a); see also Chen and Mandelbaum (1991), Rybko and Stolyar (1992), Dai and Meyn (1995), Dai and Weiss (1996), Chen (1995), Stolyar (1995) and Bramson (1998a) for further discussions, refinements and improvements. Simultaneously, there has been a growing interest in using fluid models in a synthesis framework such as the one described here. Early eamples can be found in Chen and Yao (1993), Atkins and Chen (1995), Avram, Bertsimas and Ricard (1995) and Eng, Humphrey and Meyn (1996), where several heuristics based on fluid model optimization were described; in fact, related work can be traced back to Newell (1971).

The papers closest to our work are Meyn (1997b) and Chen and Meyn (1998), where the authors study connections between the optimal policies for the stochastic network and its associated fluid model. Their analytical results are presented in the context of trying to solve directly the stochastic control problem (for the case of Poisson arrivals and exponential service times) by first approximating the optimal value function with the one derived from the solution of the fluid optimization problem, and then using value or policy iteration. In a related paper, Meyn (1997a) first proposed the fluid-scale asymptotic optimality criterion, and showed that the optimal behavior of the stochastic network starting from a large initial condition and over a proportionally long horizon approaches the optimal behavior in the fluid model. In particular, this suggests that the slope of the switching halfspaces of the fluid model equals the slope of the corresponding (nonlinear) switching surfaces of the stochstic system. Following this insight, Meyn (1997b) suggested scheduling policies for the stochastic network based on the switching halfspaces of the fluid model shifted away from the origin by some constant vector—an "affine shift." Finally, recent related results for the case of job shop scheduling problems (with no arrivals and a transient or makespan cost criterion) can be found in Bertsimas, Gamarnik and Sethuraman (1999) and Dai and Weiss (1999).

Apart from the connections related to stability analysis, the use of fluid models is also motivated from the extensive theory on optimal control for deterministic systems with continuous dynamics [see, e.g., the books by Athans and Falb (1966), Bryson and Ho (1975) and Bertsekas (1995)], and from the vast computational simplifications they offer that have been exploited in developing efficient optimization algorithms for their solutions [see Pullan (1993, 1995, 1996)], Weiss (1996, 1997, 1999), Luo and Bertsimas (1996) and Maglaras (1997)]. Moreover, apart from Meyn's structural results, there are other examples that illustrate that the solution of the fluid optimal control problem retains significant information about the solution of the original stochastic network control problem. The $c\mu$ rule for example, has been shown to be optimal for both the underlying stochastic networks and their associated fluid models; see Chen and Yao (1993), Bertsimas, Paschalidis and Tsitsiklis (1995). In the spirit of positive results like these, the premise of the fluid model approach to network control problems is that although we are using a "weak" (FSAO) criterion for what constitutes a "good" control policy, this relaxed notion of optimality will guide us in designing "near-optimal" policies for the original stochastic network control problems.

However, despite the simple structure of fluid models and the apparently modest objective of FSAO, the meaning of the fluid policy in the original network is subtle. This will be demonstrated in the next section through the analysis of the so-called Rybko–Stolyar network, where, although the associated fluid optimization problem is "trivial," each of three "obvious" interpretations in the stochastic network is "wrong." In fact, given the solution of the associated fluid optimal control problem, its translation into an implementable policy for the stochastic network is surprisingly difficult due to the finer structure of the original network model. A similar example was described in Meyn [(1997b), Section 7]. A few simple networks have been analyzed in the papers cited above, but no general mechanism has been constructed that guarantees fluid-scale asymptotic optimality, or even stability, for the policy extracted from the fluid solution. The same translation problem has been observed in the context of the heavy-traffic approach to network control problems. There one follows a similar synthesis procedure based on Brownian approximating models and diffusion scale asymptotic analysis as the traffic intensity at each station approaches 1; see Harrison (1988, 1996a), Harrison and Wein (1989, 1990), Kelly and Laws (1993), Williams (1996), Kushner and Martins (1996).

The first general translation mechanism was proposed by Harrison (1996a) in his BIGSTEP approach to dynamic control for stochastic networks; this was done in the context of Brownian approximations and heavy-traffic limits mentioned above. Harrison (1998) rigorously proved that BIGSTEP is asymptotically optimal (in the heavy-traffic sense) for a simple two-station network. The family of policies described in this paper is an extension, or generalization, of BIGSTEP that hinges on the discrete-review structure proposed by Harrison. Each DR policy in the family to be investigated will be derived from a target trajectory specification, one example being the solution of a fluid optimal control problem. In such a policy, system status is reviewed at discrete points in time, and at each such point the controller formulates a processing plan for the next review period in order to best *track* what the fluid control policy would do starting at that point. Within each period the system is only allowed to process jobs that were present upon the review point, which renders the execution of processing decisions very simple. Implementation of this policy involves enforcement of certain safety stock requirements in order to avoid unplanned server idleness. Review periods and magnitudes of safety stocks increase as queues lengthen, but both grow less-than-linearly as functions of queue length and hence are negligible under fluid scaling. This idea of safety stocks was originally described by Kelly and Laws [(1993), Section 2] in the context of a simple example and then in general form by Harrison (1996). It is also related to the affine shifts proposed by Meyn (1997b).

The rest of the paper is structured as follows. Section 2 analyzes a simple two-station network that should help motivate the general problem being addressed. Section 3 describes open multiclass queueing networks, their associated fluid models and the family of target trajectories to be considered. The family of discrete-review polices under investigation is defined in Section 4, and their associated fluid model is derived in Section 5. Section 6 discusses fluid optimal control problems, defines fluid-scale asymptotic optimality and describes a discrete-review policy that is FSAO and stable. Finally, Section 7 contains some concluding remarks.

2. A motivating example. The simple network shown in Figure 1, studied independently by Kumar and Seidman (1990) and Rybko and Stolyar (1992), will help illustrate some of the relevant issues to be addressed in this paper. In Figure 1 the open-ended rectangles represent buffers in which four distinct job classes reside: classes 1 and 4 are processed by server 1, while classes 2 and 3 are processed by server 2; there is a renewal input flow with average arrival rate λ_1 at buffer 1 and another renewal input flow with arrival rate λ_3 into buffer 3; finally, service times for each job class k are drawn according to some probability distribution with mean $1/\mu_k$. For illustrative

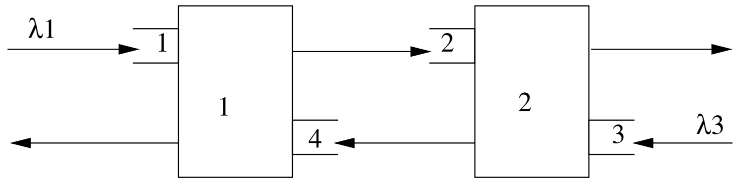

FIG. 1. *The Rybko–Stolyar network.*

purposes we shall consider the following specific numerical data:

$$\lambda_1 = \lambda_3 = 1, \qquad \mu_1 = \mu_3 = 6 \quad \text{and} \quad \mu_2 = \mu_4 = 1.5. \tag{2.1}$$

Control capability in this network is with regard to sequencing decisions between classes 1 and 4 at server 1 and classes 2 and 3 at server 2. Note that the two job classes waiting to be processed at each server differ in their service requirements and routes through the network. Now suppose we wish to find a scheduling policy π that minimizes

$$J_T^\pi(z) = \mathbf{E}_z^\pi \int_0^T \sum_{k=1}^4 Q_k(t)\,dt, \tag{2.2}$$

where $Q_k(t)$ is the class k queue length at time t, and $\mathbf{E}_z^\pi$ denotes the expectation operator with respect to the probability measure $\mathbf{P}_z^\pi$ defined by any admissible policy π and initial condition z.

Following the procedure outlined in the introduction, we proceed by forming the associated fluid model. Fluid models are deterministic and continuous-dynamics approximations of the underlying stochastic networks. Discrete jobs moving stochastically through different queues are replaced by continuous fluids flowing through different buffers, and system evolution is observed starting from any initial state. The deterministic rates at which the different fluids flow through the system are given by the average rates of corresponding stochastic quantities. Specifically, for the Rybko–Stolyar network the fluid model equations are as follows. Denoting by $\dot{\bar{T}}_k(t)$ the instantaneous fraction of effort devoted to serving class k jobs at time t by the associated server, and by $\bar{Q}_k(t)$ the amount of fluid in buffer k at time t, and defining vector functions $\dot{\bar{T}}(t)$ and $\bar{Q}(t)$ in the obvious way, one has

$$\dot{\bar{Q}}(t) = \lambda - R\dot{\bar{T}}(t), \qquad \bar{Q}(0) = z, \tag{2.3}$$

$$\dot{\bar{T}}(t) \ge 0, \qquad \dot{\bar{T}}_1(t) + \dot{\bar{T}}_4(t) \le 1, \qquad \dot{\bar{T}}_2(t) + \dot{\bar{T}}_3(t) \le 1, \qquad \bar{Q}(t) \ge 0, \tag{2.4}$$

where

$$\lambda = \begin{bmatrix} \lambda_1 \\ 0 \\ \lambda_3 \\ 0 \end{bmatrix}, \qquad R = \begin{bmatrix} \mu_1 & 0 & 0 & 0 \\ -\mu_1 & \mu_2 & 0 & 0 \\ 0 & 0 & \mu_3 & 0 \\ 0 & 0 & -\mu_3 & \mu_4 \end{bmatrix}.$$

(More details on the derivation of these equations will be given in Section 3.) The associated fluid optimal control problem will be to choose a control $\dot{\bar{T}}(\cdot)$ for this fluid model that minimizes

$$\bar{J}_T(z) = \int_0^T \sum_{k=1}^{4} \bar{Q}_k(t)\,dt, \tag{2.5}$$

for some fixed $T > 0$. The corresponding value function will be denoted by $\bar{V}_T(z)$.

It can be shown that the optimal control for the fluid model is to give priority at each server to the job class that is closer to exiting the system with server splitting whenever an exiting class is emptied at the other server. That is, each server has responsibility for one incoming buffer and one exit buffer; the exit buffer is given priority unless the other server's exit buffer is empty, in which case server splitting occurs. For an explanation of the latter situation, let us focus on the behavior of server 1 when buffer 2 (the exit buffer for server 2) is empty and buffer 1 is nonempty. In that circumstance, given the data in (2.1), server 1 devotes 25% of its effort to buffer 1 (its own incoming buffer) so that server 2 can remain fully occupied with class 2 jobs, and devotes the other 75% of its effort to draining buffer 4 (its own exit buffer). This policy is myopic in the sense that it removes fluid from the system at the fastest possible instantaneous rate, regardless of future considerations, and it is optimal regardless of the horizon length T.

Given the solution to the fluid optimization problem, which has the structure of a static priority rule together with some "boundary modifications," we now seek to translate the derived policy back into the original stochastic network. The following natural alternatives arise.

LBFS. The simplest candidate policy is to use the static priority rule that emerges from the optimal fluid control law, which gives priority to exiting classes in each server; this is also known as *Last-Buffer-First-Served* (LBFS). As was demonstrated by Lu and Kumar (1991) for the deterministic case and later by Rybko and Stolyar (1992) for the stochastic case, this policy will be unstable for the system parameters specified in (2.1), despite the fact that the nominal utilization rate (or traffic intensity parameter) for each server is equal to 0.833. That is, the static priorities derived from the optimal fluid control policy by neglecting "boundary behavior" have catastrophic performance: they cause instability!

LBFS with priority reversal (LBFS/PR). Here each server uses LBFS as its "default," but switches to the opposite priority when the other server's exit buffer is empty. This policy is stable, but its asymptotic performance is not satisfactory, as will be explained below.

LBFS with server splitting (LBFS/SS). Here we implement exactly the optimal policy derived form the fluid model, splitting server effort in the percentages prescribed by the fluid optimal control policy. Whenever one of the queue lengths is empty, any positive server utilization predicted by the optimal fluid control policy for this class will not be implementable due to the discrete dynamics of the stochastic network. In such cases, this percentage of server utilization will be reallocated to the other class waiting to be processed at that server, if this is nonempty. For example, when $q = (0, 0, +, +)$ the optimal server allocation from the fluid control problem is $(0.25, 1, 0, 0.75)$, yet the implemented server allocation in the stochastic network will be $(0, 0, 1, 1)$; a "+" denotes positive buffer content. Again, this policy is stable, but its asymptotic performance is not satisfactory.

Using any one of the control policies just described, it is natural to consider system behavior under a sequence of initial conditions $\{z^n\}$, such that $|z^n| \to \infty$ as $n \to \infty$, keeping all other system parameters fixed. Consider, for example, the case where $z^n = n[1, 0, 0.5, 1]$. Denoting by $Q^n(\cdot)$ the four-dimensional queue length process with initial state $Q^n(0) = z^n$, we define the fluid scaled version,

$$\bar{Q}^n(t) = \frac{Q^n(nt)}{n}, \qquad 0 \le t \le T, \tag{2.6}$$

and ask whether $\bar{Q}^n$ converges to a limit trajectory as $n \to \infty$ that is optimal in the fluid model. [Due to the LLN type of scaling in (2.6), one expects a deterministic limit to be approached.] Essentially, we are testing whether the system behavior under one of the implemented policies approaches (as n grows and the stochastic problem approaches one of fluid, or transient, optimization) the optimal performance that one started with.

For both LBFS/PR and LBFS/SS, the scaled processes $\bar{Q}^n$ do converge to a deterministic limit as $n \to \infty$, but that limit does not coincide with the optimal fluid trajectory. That is, although these policies may be intended as implementations of the optimal fluid control policy, they do not in fact achieve as their fluid limits a trajectory that is optimal in the fluid model. In detail, both LBFS/PR and LBFS/SS will introduce undesirable idling periods at server 2 while waiting for new class 1 jobs to complete service at station 1. This behavior will not change as n grows, since queue 2 will always have either one or no jobs waiting, and this will lead to the suboptimal behavior claimed above. Both policies fail because the servers are too slow in switching from myopically draining cost out of the system to guarding against idleness that will prevent optimal cost draining in future times. Following this argument one would expect that performance under LBFS/SS will be worse than that under LBFS/PR, which is indeed correct.

The main observation that emerges from the analysis of this example is that undesirable system behavior is observed when one or more of the buffers get depleted. In these instances, the idealizations embodied in the fluid approximating model, especially regarding the continuous versus discrete nature of the dynamics and the deterministic versus stochastic flows, become unrealistic, and the stochastic network can no longer track the target trajectory derived using the fluid model. To guard against such undesirable effects one needs to be cautious about the behavior of the system close to the boundaries. For example, in the Rybko–Stolyar network one just needs to switch priorities when the content of the exit buffer falls below a certain threshold (the magnitude of which will be characterized later) in order to achieve asymptotic optimality; this is in contrast to switching priorities when the buffer content gets depleted. (This is related to the discussion in Meyn [(1997b), Section 7].) However, in more complex networks where intuition can be limited, this policy translation step can be very subtle; successful translation involves relatively fine structure. The remainder of this paper describes a general and easily implementable solution to this problem.

3. Network models.

Open multiclass queueing networks. In the description of a multiclass queueing network we adopt the set-up introduced by Harrison (1988). Consider a queueing network of single server stations indexed by $i = 1, \ldots, S$. (The terms "station" and "server" will be used interchangeably.) The network is populated by job classes indexed by $k = 1, \ldots, K$ and infinite capacity buffers are associated with each class of jobs. Class k jobs are served by a unique station $s(k)$ and their service times are $\{\eta_k(n); n \geq 1\}$. That is, the nth class k job requires $\eta_k(n)$ time units of service from station $s(k)$. Jobs within a class are served on First-In-First-Out (FIFO) basis. Upon completion of service at station $s(k)$, a class k job becomes a job of class m with probability P_{km} and exits the network with probability $1 - \sum_m P_{km}$, independent of all previous history. Assume that the general routing matrix $P = [P_{km}]$ is transient (that is, $I + P + P^2 + \cdots$ is convergent). Let $\{\phi^k(n)\}$ denote the sequence of K-dimensional iid Bernoulli random vectors such that $\phi^k_j(n) = 1$ if upon service completion the nth class k job becomes a class j job and is zero otherwise, and let $\Phi^k(n) = \sum_{j=1}^n \phi^k(j)$. Every job class k can have its own exogenous arrival process with interarrival times $\{\xi_k(n), n \geq 1\}$. The set of classes that have a nonnull exogenous arrival process will be denoted by $\mathscr{E}$ and the notation $E(t)$ will denote the K-dimensional vector of exogenous arrivals in the time interval $[0, t]$. It is assumed that $\mathscr{E} \neq \varnothing$.

The processes (ξ, η) should satisfy the following distributional assumptions: (i) $\xi_1, \ldots, \xi_K$ and $\eta_1, \ldots, \eta_K$ are mutually independent, positive, iid sequences; (ii) $\mathbf{E}[\eta_k(1)] \neq 0$ for $k = 1, \ldots, K$. For some $\theta > 0, \mathbf{E}[\exp \theta \eta_k(1)] < \infty$ for $k = 1, \ldots, K$ and $\mathbf{E}[\exp \theta \xi_k(1)] < \infty$ for $k \in \mathscr{E}$ and (iii) For any $x > 0$, $k \in \mathscr{E}, \mathbf{P}\{\xi_k(1) \geq x\} > 0$. Also, for some positive function $p(x)$ on $\mathbf{R}_+$ with $\int_0^\infty p(x)\,dx > 0$, and some integer j_0, $\mathbf{P}\big\{\sum_{i=1}^{j_0} \xi_k(i) \in dx\big\} \geq p(x)\,dx$. Condition (ii) is stronger than the first moment condition usually imposed [e.g., Dai (1995a)], and it is needed in the derivation of large deviation bounds required in our analysis; (ii) is satisfied by $\{\phi^k(n)\}$. For purposes of proving asymptotic optimality one needs to ensure that the processes (ξ, η, ϕ) satisfy a SLLN, and thus, the iid assumption could be relaxed. However, in order to make use of the general stability theory of Dai (1995a) [see also Dai and Meyn (1995)] we impose (i) as well as the technical conditions in (iii); the latter are never invoked in propositions that are actually proved in this work.

For future reference, let $\lambda_k = 1/\mathbf{E}[\xi_k(1)]$ and $\mu_k = 1/\mathbf{E}[\eta_k(1)] = 1/m_k$ be the arrival and service rates, respectively, for class k jobs, let $\lambda = (\lambda_1, \ldots, \lambda_K)'$, and let $M = \text{diag}\{m_1, \dots, m_K\}$; the $'$ denotes a transpose. The set $\{k\colon s(k) = i\}$, denoted C_i, is called the constituency of the server i, while the $S \times K$ constituency matrix C will be the following incidence matrix:

$$C_{ik} = \begin{cases} 1, & \text{if } s(k) = i, \\ 0, & \text{otherwise.} \end{cases}$$

Given the Markovian routing structure of these networks and the transience of P, one can compute the vector of effective arrival rates, $\alpha = (I - P')^{-1}\lambda$, and the vector of traffic intensities $\rho = CR^{-1}\lambda$, where $R = (I - P')M^{-1}$ and ρ_i denotes the nominal load (or utilization level) for server i. Hereafter, it will be assumed that $\alpha > 0$.

Note. This will help simplify the policy description of Section 4, where class level safety stock requirements will be imposed, and job classes that could start empty and have zero effective arrival rates would have to be treated differently otherwise. This restriction could easily be relaxed. Note that the case where some or all of the classes have no effective arrivals could be of practical interest in the context of transient control. The purely static case, where $\lambda = \alpha = 0$, is addressed by Bertsimas, Gamarnik and Sethuraman (1999) and Dai and Weiss (1999).

Denote by $Q_k(t)$ the total number of class k jobs in the system at time t, and by $Q(t)$ the corresponding K-vector of "queue lengths." A generic value of $Q(t)$ will be denoted by q, and the size of this vector is defined as $|q| = \sum_k q_k$. A scheduling policy is a rule for making server allocation decisions over time. It takes the form of a K-dimensional cumulative allocation process $\{T^y(t),\ t \geq 0;\ T^y(0) = 0\}$, where $T_k^y(t)$ denotes the time allocated by server $s(k)$ into serving class k jobs up to time t, and the superscript "y" denotes the dependence on the initial condition that may include additional information apart from the initial queue length configuration in the system. The cumulative allocation process should be nondecreasing, $T^y(0) = 0$, $CT^y(t) \leq t\mathbf{1}$ (where $\mathbf{1}$ denotes the vector of ones of appropriate dimension), and finally it must be nonanticipating; the last restriction implies that current allocations only depend on information available up to time t. Each server can only process one job at a time.

Given any admissible scheduling policy, a Markovian state descriptor can be constructed and an underlying Markov chain can be identified for the controlled network. The Markovian state at time t will be denoted by $Y(t)$ and the corresponding normed state space will be $(\mathbf{Y}, \|\cdot\|)$; see the comments by Dai and Meyn [(1995), Section IIb] or Bramson [(1998a), Section 3] regarding the choice of $\|\cdot\|$. In the next section, a Markov chain will be constructed for the family of DR policies under investigation. Other examples can be found in Dai (1995a).

Fluid models. The fluid models associated with multiclass queueing networks are deterministic and have continuous dynamics. They are formally derived through a FSLLN type of scaling, where the network processes are studied starting from large initial conditions and observed over a proportionally long time horizon. In more detail, one considers a sequence of initial conditions $\{y^n\} \subset \mathbf{Y}$ such that $\|y^n\| \to \infty$ as $n \to \infty$. For any real valued process $\{f^y(t),\ t \geq 0\}$ that is right continuous with left limits, its fluid scaled counterpart is defined by

$$\bar{f}^n(t) = \frac{1}{\|y^n\|} f^{y^n}(\|y^n\|t). \tag{3.1}$$

Applying this scaling on the queue length and cumulative allocation processes, Dai [(1995a), Theorem 4.1] showed that, provided that ω does not belong to an exceptional set of measure zero, there exists a subsequence $\{y^{n_j}(\omega)\}$ such that $(\bar{Q}^{y^{n_j}}(\cdot,\omega), \bar{T}^{y^{n_j}}(\cdot,\omega)) \to (\bar{Q}(\cdot,\omega),\ \bar{T}(\cdot,\omega))$ u.o.c. as $j \to \infty$, and the pair $(\bar{Q}(\cdot,\omega), \bar{T}(\cdot,\omega))$ satisfies equations

$$\bar{Q}(t) = \bar{Q}(0) + \lambda t - (I - P')M^{-1}\bar{T}(t), \tag{3.2}$$

$$\bar{Q}(t) \geq 0, \qquad C\dot{\bar{T}}(t) \leq \mathbf{1}, \qquad \dot{\bar{T}}(t) \geq 0 \quad \text{for } t \geq 0, \tag{3.3}$$

together with some additional conditions specific to the scheduling policy employed. Fluid limits depend on ω through the converging subsequence $\{y^{n_j}\}$ and the limiting initial condition. They are neither deterministic nor unique, but their dynamics are captured by the deterministic and continuous equations of evolution (3.2) and (3.3). Whenever possible, the dependence on ω will be suppressed from the notation. In the sequel, the overbar notation will signify fluid scaled quantities and a superscript n will be used to signify the scaled processes corresponding to the initial condition y^n. The use of the overbar notation without any superscript will denote the fluid limit of the appropriate variable; for example, $\bar{T}(\cdot)$ as the limit of $\bar{T}^n(\cdot)$.

Let $R_a(0)$ be the $|\mathscr{E}|$-vector and $R_s(0)$ be the K-vector of residual times until the first exogenous arrivals or service completions respectively. Equations (3.2) and (3.3) are referred to as the *undelayed fluid model*, and they implicitely assume that $(\bar{R}_a^{n_j}(0),\ \bar{R}_s^{n_j}(0)) \to (0,0)$ almost surely. The limit processes $(\bar{Q}, \bar{T})$ are Lipschitz continuous, and therefore, time derivatives $(\dot{\bar{Q}}(t), \dot{\bar{T}}(t))$ will exist a.e. [see Dai and Weiss (1996), Lemma 2.1]. Hence, for almost all times $t \geq 0$ the system dynamics can also be expressed in the differential form: $\dot{\bar{Q}}(t) = \lambda - R\dot{\bar{T}}(t)$; this representation will be useful later on. We will say that $(\bar{Q}, \bar{T}) \in FM$, or equivalently that it is a *fluid solution*, if this pair of state and input trajectories satisfy equations (3.2) and (3.3). Excellent expositions of fluid models, their derivation, their properties, as well as a discussion of delayed versus undelayed limits can be found in Dai (1995a, b) and Bramson [(1998), Section 4].

Target trajectory specifications. Denote by $\mathbf{C}_{R_+^K}[0,\infty)$ the space of continuous functions of a parameter $t \in [0,\infty)$ taking values on $\mathbf{R}_+^K$ and by $\mathbf{AC}_{R_+^K}[0,\infty)$ the corresponding space of absolutely continuous functions. A fluid trajectory $\{\bar{Q}(t), t \geq 0\}$ is feasible if there exists an (fluid) allocation process $\{\bar{T}(t);\, t \geq 0\}$ such that $(\bar{Q}, \bar{T})$ satisfies (3.2) and (3.3), for all $t \geq 0$. In general, the set of feasible fluid trajectories starting from an initial condition z, denoted by $\mathscr{Q}_z$, is defined by

$$\mathscr{Q}_z = \{\bar{Q}(\cdot) \in \mathbf{AC}_{R_+^K}[0, \infty)\colon\ \bar{Q}(0) = z, \exists \bar{T}(t),\ t \geq 0 \text{ such that } (\bar{Q}, \bar{T}) \in FM\}. \tag{3.4}$$

Define a fluid trajectory map $\Psi\colon \mathbf{R}_+^K \to \mathbf{AC}_{R_+^K}[0, \infty)$, that assigns to every initial condition z a unique target fluid trajectory $\{\Psi(t; z),\ t \geq 0\}$, which is an element in $\mathscr{Q}_z$. That is,

$$\Psi(\cdot; z) \in \mathscr{Q}_z \quad \text{for all } z \geq 0. \tag{3.5}$$

$\Psi(\cdot; z)$ denotes the target trajectory starting at z, and $\Psi(t; z)$ denotes the state vector at time t along this trajectory. The mapping Ψ is measurable in the following sense: for any measurable set A, where $y \in A$ implies that $y \in \mathscr{Q}_z$ for some $z \geq 0$, we have that $(\Psi)^{-1}(A) = \{z\colon z = y(0),\ \forall\, y \in A\}$, which is measurable with respect to the Borel measure on $\mathbf{AC}_{R_+^K}[0, \infty)$.

We will require that Ψ satisfies two conditions. The first is the semigroup property that

$$\Psi(t + s; z) = \Psi(s; \Psi(t; z)) \quad \text{for all } z \geq 0,\ t \geq 0,\ s \geq 0. \tag{S}$$

This is a memoryless property: the target trajectory starting at z is independent of the path followed by the system until it reached that point. The second is a smoothness condition for the vector field associated with Ψ. First, note that from the Lipschitz continuity of the fluid trajectories it follows that $\dot{\Psi}(t; z) \triangleq d\Psi(t; z)/dt = \dot{\Psi}(0; \Psi(t; z))$ exists almost everywhere. We will also require that

$$\begin{gathered} \Psi(\cdot; z) \text{ is continuous in } z \text{ and that} \\ \dot{\Psi}(0; z) \text{ is continuous for almost all } z \geq 0. \end{gathered} \tag{C}$$

The main restriction here is that of continuity of Ψ with z, which roughly says that for two initial conditions that are close to each other the corresponding target trajectories will also be close. The other requirement in (C) appears to be a mild one.

Given the trajectory $\Psi(\cdot; z)$ the cumulative allocation process up to any time t is uniquely defined by $\bar{T}_\Psi(t; z) = R^{-1}(z + \lambda t - \Psi(t; z))$. Finally, in order to uniquely define derivatives, the following convention will be followed. At points where $\dot{\Psi}(0; z)$ either does not exist or is not continuous we assign $\dot{\Psi}(0; z) \equiv \dot{\Psi}(0^+; z) = \lim_{h \downarrow 0} \dot{\Psi}(h; z)$; this limit exists since there are only countably many points of discontinuity for Ψ or $\dot{\Psi}$. This convention makes no distinction at points where the derivative exists and is continuous, and, of course, the trajectories $\Psi(\cdot; z)$ and cumulative allocation processes $\bar{T}_\Psi(\cdot; z)$ remain unaffected.

Under the standing assumptions, one could describe the target trajectory mapping either by specifying $\Psi(\cdot, z)$ for all $z \geq 0$, or, more compactly, by specifying $\dot{\Psi}(0; z)$ or $\dot{T}_\Psi(0; z)$ for all $z \geq 0$. The last two specifications are in state feedback form.

A target trajectory map Ψ describes the desired system behavior that the controller will try to track. One such choice will be the solution that emerges from the associated fluid optimal control problem. According to the general procedure described in the introduction, one needs to formulate a performance criterion, like the one of stages 4 and 5, that is consistent with this trajectory tracking formulation. This criterion will be referred to as *asymptotic tracking* and is defined as follows.

DEFINITION 3.1. Consider a multiclass open queueing network under an admissible policy π and let $Q_\pi(\cdot)$ be the corresponding queue length process. For any sequence of initial conditions $\{y^n\} \subset \mathbf{Y}$ such that $\|y^n\| \to \infty$ as $n \to \infty$, assume that for some random variable Z_0 and each converging subsequence $\{y^{n_j}\}$, $\bar{Q}^{n_j}(0) \to Z_0$ and $(\bar{R}_a^{n_j}(0),\ \bar{R}_s^{n_j}(0)) \to (0, 0)$ almost surely. The policy π is said to be asymptotically tracking with respect to the specification Ψ if for almost all ω,

$$\bar{Q}_\pi^{n_j}(\cdot, \omega) \to q_\pi(\cdot, \omega) = \Psi(\cdot; Z_0(\omega)). \tag{3.6}$$

This definition assumes that the limiting initial condition converges almost surely to some well-defined random variable [i.e., $\bar{Q}^n(0) \to Z_0$]. This appears to be a mild assumption, and it will be put into use later in the context of the fluid-scale asymptotic optimality criterion. The restriction to undelayed fluid limits could be relaxed by extending the mapping Ψ to deal with delayed initial conditions; however, such an extension appears to have limited practical significance and will not be addressed here.

An example: minimum time control. A natural candidate for a target trajectory specification is to try to optimize transient performance by minimizing the time to drain the system starting from any initial condition. As it was shown by Weiss (1996) for the case of reentrant lines, this problem has a remarkably simple solution, which is to go to the origin along a straight line. Assume that $\rho < \mathbf{1}$ and let t be the draining time under some control $\bar{T}(\cdot)$. Then,

$$\bar{Q}(t) = z + \lambda t - R\bar{T}(t) = 0,$$

which implies that $\bar{T}(t) = R^{-1}z + R^{-1}\lambda t$. The capacity constraints imply that

$$\begin{aligned} C\bar{T}(t) \ \leq\ \mathbf{1}t \ &\Rightarrow\ CR^{-1}z \ \leq\ (\mathbf{1} - \rho)t \\ &\Rightarrow\ t \geq\ t^*(z) \triangleq \max_{1 \leq i \leq S} \frac{(CR^{-1}z)_i}{1 - \rho_i}. \end{aligned} \tag{3.7}$$

We now show how to achieve the lower bound $t^*(z)$. Define the instantaneous control $\dot{\bar{T}}(t) = R^{-1}\lambda + R^{-1}z/t^*(z)$, for $t \leq t^*(z)$, and $\dot{\bar{T}}(t) = R^{-1}\lambda$, for $t > t^*(z)$. To verify that this is a feasible control that achieves the bound in (3.7), note that (a) since $R^{-1} = M^{-1}(I - P')^{-1} = M^{-1}(I + P' + P^2 + \cdots)$ which is elementwise nonnegative, $\dot{\bar{T}}(t) \geq 0$ and (b) that (3.7) implies that $C\dot{\bar{T}}(t) = \rho + CR^{-1}z/t^*(z) \leq \mathbf{1}$.

The corresponding trajectory mapping, denoted $\Psi^{\min}$, is defined by $\Psi^{\min}(\cdot; z) = \bar{Q}_z^{\min}(\cdot)$, where

$$\bar{Q}_z^{\min}(t) = z\left(1 - \frac{t}{t^*(z)}\right) \quad \text{for} \quad t \leq t^*(z) \quad \text{and} \tag{3.8}$$
$$\bar{Q}_z^{\min}(t) = 0 \quad \text{for} \quad t > t^*(z).$$

It is easy to verify that $\Psi^{\min}$ is continuous in z, that $\dot{\Psi}^{\min}(\cdot; z)$ exists and is continuous a.e. [the only points of discontinuity being at $t = t^*(z)$ when $\Psi^{\min}(t^*(z); z) = 0$], and finally, that $\Psi^{\min}(t; \bar{Q}_z^{\min}(s)) = \Psi^{\min}(t + s; z)$. Hence, (S) and (C) are satisfied.

Hereafter, it will be assumed that the specification Ψ is known a priori. The enormous theoretical and algorithmic simplifications offered by fluid models renders the computation of Ψ simple. The reader is referred to some of the references mentioned in the introduction for more details.

4. Discrete-review policies. The family of policies we propose is based on the recent idea of a discrete-review structure introduced by Harrison (1996a) in his BIGSTEP approach to dynamic flow management in multiclass networks. Discrete-review policies, and specifically, policies that step through time in large intervals within which a deterministic planning logic is employed, have also been proposed by other researchers in the areas of applied probability and network control. Some examples that are closer to our work can be found in Bertsimas and van Ryzin (1991), Bambos and Warland (1993), Tassiulas and Papavassiliou (1995), Gans and van Ryzin (1997), but other related papers can be found as well. The main idea in all these papers is that the scheduling and execution steps in the corresponding systems become more efficient as the planning horizons become longer or, equivalently, as the amount of work to be processed within each period increases.

A discrete-review policy is defined by or is derived from the trajectory mapping Ψ: $\mathbf{R}_+^K \to \mathbf{AC}_{R_+^K}[0, \infty)$, a function l: $\mathbf{R}_+ \to \mathbf{R}_+$, plus a K-dimensional vector β that satisfy the following restrictions. First, Ψ satisfies assumptions (S) and (C). Second, $l(\cdot)$ is real valued, strictly positive, concave, and further satisfies

$$\frac{l(x)}{\log(x)} > c_0 \quad \text{and} \quad \frac{l(x)}{\log(x)} \to \infty \quad \text{as } x \to \infty \tag{4.1}$$

and

$$\frac{l(x)}{x} \to 0 \quad \text{as } x \to \infty. \tag{4.2}$$

And third, β is a vector in $\mathbf{R}_+^K$ that satisfies

$$\beta > \mu. \tag{4.3}$$

Initialization. Given the initial condition $Q(0)$, we set the length of the review periods, l, and the vector of safety stock levels, θ, to be

$$l = l(|Q(0)|) \quad \text{and} \quad \theta = \beta l. \tag{4.4}$$

Planning. Under any of the policies to be considered, system status will be observed at a sequence of times $0 = t_0 < t_1 < t_2 < \cdots$; we call t_j the jth review point and the time interval between t_j and t_{j+1} the jth planning period. Given the queue length vector $q = Q(t_j)$ observed at t_j, server activities over the next planning period are determined by a two-step procedure.

STEP 1 (Target setting and enforcement of safety stock requirements). Given the observed queue length vector q, we choose a target z that represents the system's nominal ending state upon completion of this review period as follows. Let $\bar{l} = l/|Q(0)|$, $(\bar{q} - \bar{\theta})^+ = (q - \theta)^+/|Q(0)|$ where $x^+ = \max(x, 0)$, and rewrite q as $q = q \wedge \theta + (q - \theta)^+$, where $x \wedge y = \min(x, y)$. The first term corresponds to the safety stock requirement we are striving to maintain. The second term is our effective queue length vector that we want to drain using the specification Ψ. Upon completion of the review period the nominal ending state maps $q \wedge \theta$ to θ and $\theta + (q - \theta)^+$ to $|Q(0)|\Psi(\bar{l}; (\bar{q} - \bar{\theta})^+)$. The latter involves a step of rescaling time and space by a factor of $|Q(0)|$ that reduces the tracking problem to a normalized scale consistent with the fluid model specification Ψ. The target state is

$$z = \theta + |Q(0)|\Psi(\bar{l}; (\bar{q} - \bar{\theta})^+). \tag{4.5}$$

STEP 2 (Tracking). We now compute a K-vector of time allocations, denoted x, that will nominally steer the state from q to its target z over the ensuing period. Following the decomposition of q and (4.5), we express x in the form $x_\theta + x_\Psi$, where x_θ denotes the allocation vector over an interval of length t_θ that will steer $q \wedge \theta$ to θ, and x_Ψ denotes the corresponding vector that will steer the state from $(q - \theta)^+$ to $|Q(0)|\Psi(\bar{l}; (\bar{q} - \bar{\theta})^+)$. From the definition of Ψ we have that

$$x_\Psi = |(Q(0)| T_\Psi \left(\bar{l}; (\bar{q} - \bar{\theta})^+\right). \tag{4.6}$$

Next, we compute x_θ as the minimum time control from $q \wedge \theta$ to θ. Suppose that under some feasible allocation $T(t)$ we have that at time t, $\lambda t - RT(t) \geq (\theta - q)^+ \Rightarrow RT(t) - \left[\lambda t - (\theta - q)^+\right] \leq 0$. Multiplying by R^{-1}, which is elementwise nonnegative, and since $T(t) \geq 0$, we get that

$$t \geq \max_k \frac{(R^{-1}(\theta - q)^+)_k}{(R^{-1}\lambda)_k} \triangleq t_\theta. \tag{4.7}$$

That is, t_θ is a lower bound on the time required to effect a change of $(\theta - q)^+$, and the corresponding vector of nominal time allocations x_θ is given by

$$x_\theta = R^{-1}\lambda t_\theta - R^{-1}(\theta - q)^+. \tag{4.8}$$

(Similarly to the example of Section 3, linear translation will achieve that bound.)

In total, the vector of time allocations over the ensuing period, denoted by x, is given by

$$x = x_\theta + x_\Psi. \tag{4.9}$$

Note that if $q \geq \theta$, then $x = x_\Psi$.

Execution. Given this vector of nominal time allocations, a plan expressed in units of jobs of each class to be processed over the ensuing period, and a nominal idleness plan expressed in units of time for each server to remain idle over the same period are formed as follows:

$$\begin{aligned} p_k &= \left\lfloor \frac{x_k}{m_k} \right\rfloor \quad &&\text{for } k = 1, \dots, K \quad \text{and} \\ u_i &= (l + t_\theta - (Cx)_i)^+ \quad &&\text{for } i = 1, \dots, S. \end{aligned} \tag{4.10}$$

The execution of these decisions is as follows. First, the plan p is implemented in open-loop fashion; that is, each server i processes sequentially p_k jobs for each class $k \in C_i$. Condition (4.3) implies that if q is close to $\theta = \beta l$, then $q > \mu l$ and thus, any feasible processing plan over the ensuing period will be implementable from jobs present upon the review point. The ordering of the execution sequence is not important for the asymptotic analysis considered here, although presumably it could affect performance for the stochastic problem. If $q \not\geq \mu l$, then the execution of p is more subtle, since not all jobs needing to be processed are present at their respective buffers at time t_j. A detailed solution for this case is described in Lemma 4.2 of Maglaras (1999). Later on, we will show that as $|Q(0)|$ increases, the probability that $q \not\geq \mu l$ vanishes. Let d_i denote the time taken to complete processing of all jobs at server i. Next, each server i will idle for $u_i \wedge (l + t_\theta - d_i)^+$ time units. The completion of this idling period signals the beginning of the $(j+1)$st review period. That is, $t_{j+1} = t_j + \max(l + t_\theta, d_1, \dots, d_S)$.

Hereafter, the notation $\mathbf{DR}(\Psi, l, \beta)$ will denote the discrete-review policy derived from the trajectory map Ψ, the function $l(\cdot)$, and the vector β that satisfy (S) and (C), (4.1), and (4.2) and (4.3), respectively. Also, we shall differentiate between review periods by writing $q(j)$, $x(j)$ and so on.

In order to make use of Dai's stability theory, an underlying Markov chain for the family of DR policies just described will be constructed as follows. Assume that $t_j \leq t < t_{j+1}$. Let $p(t)$ be a K-vector, where $p_k(t)$ is the number of class k jobs that remain to be processed at time t according to the processing plan $p(j)$. Let $d(t)$ be an S-vector, where $d_i(t)$ is the time spent so far in the jth execution period by server i in the processing of jobs. (Then $u(t) = u \wedge (l + t_\theta - d(t))^+$ is the S-vector of remaining idling times for the ensuing period.)

Finally, let $R_a(t)$ be the $|\mathscr{E}|$-vector and $R_s(t)$ be the K-vector of residual times until the next exogenous arrivals or service completions respectively.

The Markovian state descriptor is

$$Y(t) = [Q(t);\ Q(t_j);\ p(t);\ d(t);\ R_a(t);\ R_s(t);\ |Q(0)|], \tag{4.11}$$

and $\mathbf{Y}$ will represent the underlying state space. Imitating Dai (1995a) and using the strong Markov property for piecewise deterministic processes of Davis (1984), it is easy to show that the process $\{Y(t),\ t \geq 0\}$ is a strong Markov process with state space $\mathbf{Y}$. The associated norm will be

$$\|Y(t)\| = |Q(t)| + |p(t)| + |d(t)| + |R_s(t)| + |R_a(t)|.$$

We conclude this section with some general remarks. To avoid undesirable boundary behavior, we have "shifted" the target trajectories by θ; that is, the point where all queues are empty is mapped to $Q(t) = \theta$ and these policies will track desired target trajectories that will tend to steer the state toward that shifted origin; this is achieved through x_Ψ. Steps 1 and 2 extend the trajectory mapping to initial states where $q \ngeq \theta$ by providing a target path to be followed until the queue length vector is again above θ. (The minimum time control used in x_θ is not essential, but what is needed is an adequate positive drift to be prescribed that will steer the queue length process toward θ.)

There are three relevant time scales in the planning and execution of DR policies: first, the system is evolving in a time horizon proportional to $|Q(t)|$; second, the DR policy is planning and executing in time intervals of length $l(|Q(0)|)$ and third, individual events, such as service completions or external arrivals, occur in a time scale proportional to μ and λ. Note that for purposes of transient (or fluid-scale) asymptotic analysis, $|Q(0)|$ is the appropriate measure of the scale of the system. This motivates the choices in (4.4). For steady-state control under DR policies, where the significance of the initial condition is limited, $|Q(0)|$ should be replaced by the average size of the queue length vector, which is $\mathscr{O}(1/(1-\rho))$; this is consistent with the heavy traffic analysis and the BIGSTEP approach described in Harrison (1996). Alternatively, one could adjust the length of the review periods and the magnitude of safety stocks dynamically, by setting $l(j) = l(|q(j)|)$ and $\theta(j) = \beta l(j)$; this was proposed and analyzed in Maglaras (1999). For purposes of this paper, all three of these choices yield identical results; we are using the simplest of the three.

Finally, to gain intuition on the design of these policies, it is instructive to consider the asymptotic behavior of the system as $|Q(0)|$ increases. One would expect the following to be observed: (i) under (4.1) and (4.2), the three time scales will separate; (ii) tracking within each period will become very accurate [this follows from (4.1) and elementary large deviations analysis]; (iii) accurate tracking implies that $Q(t) \gtrapprox \theta$, which in turn implies that no capacity is lost in enforcing the safety stock constraints, that is, $x(j) \approx x_\Psi(j)$ and (iv) both θ and the error between the "shifted" and true trajectories become negligible. Properties (i)–(iv) will yield asymptotic tracking of Ψ.

The reader is referred to Harrison (1996) and Maglaras (1999) for more comments on DR policies.

5. The main result: asymptotic tracking under $\mathbf{DR}(\Psi, l, \beta)$. Broadly speaking, the operation of a DR policy is as follows: the controller reviews status at a discrete point in time, uses Ψ to select the target trajectory starting at that point, chooses a target state along this trajectory l time units into the future, computes a vector of time allocations to move to this target and implements the corresponding plan. Upon completion, status is reviewed again. The state of the system is potentially different from the target specified in the previous review period, so a new trajectory is selected through Ψ and the planning and execution procedure is repeated.

In this section, the asymptotic behavior of these policies is analyzed, and in particular, it will be shown that the policy $\mathbf{DR}(\Psi, l, \beta)$ is *asymptotically tracking* with respect to the specification Ψ. For example, for the sequence of initial conditions $Q^n(0) = nz$, asymptotic tracking implies that the fluid limit under the policy $\mathbf{DR}(\Psi, l, \beta)$ achieves the desired trajectory $\Psi(\cdot; z)$. In the stochastic network, this translates into accurate tracking of the target fluid trajectory when the system is operating at large population levels. Note that this result does not depend on the traffic intensity vector ρ of the system; however, depending on ρ, the feasible specifications Ψ will differ, and the behavior of the controlled system will vary accordingly.

THEOREM 5.1. *Consider a multiclass open queueing network under the policy* $\mathbf{DR}(\Psi, l, \beta)$. *For any sequence of initial conditions* $\{y^n\} \subset \mathbf{Y}$, *such that* $\|y^n\| \to \infty$ *as* $n \to \infty$, *assume that for some random variable* Z_0 *and each converging subsequence* $\{y^{n_j}\}$, $\bar{Q}^{n_j}(0) \to Z_0$ *and* $(\bar{R}_a^{n_j}(0),\ \bar{R}_s^{n_j}(0)) \to (0,0)$ *almost surely. Then for almost all sample paths* ω,

$$(\bar{Q}^{n_j}(\cdot,\omega),\ \bar{T}^{n_j}(\cdot,\omega),\ \bar{p}^{n_j}(\cdot,\omega),\ \bar{d}^{n_j}(\cdot,\omega)) \to (\bar{Q}(\cdot,\omega),\ \bar{T}(\cdot,\omega),\ 0,\ 0) \quad \text{u.o.c.}, \tag{5.1}$$

and $(\bar{Q}, \bar{T})$ *further satisfies* (3.2) *and* (3.3) *together with the policy specific equations*

$$\bar{Q}(\cdot,\omega) = \Psi(\cdot; \bar{Q}(0,\omega)) \quad \text{and} \quad \bar{T}(\cdot,\omega) = \bar{T}_\Psi(\cdot; \bar{Q}(0,\omega)). \tag{5.2}$$

In particular, the policy $\mathbf{DR}(\Psi, l, \beta)$ *is asymptotically tracking with respect to the specification* Ψ.

PROOF. Without loss of generality, in the sequel we work directly with the converging subsequence, thus avoiding the use of double subscripts, and we assume that $\|y^n\| = n$ for all $n \geq 1$.

Existence of a converging subsequence and the convergence of the queue length and allocation processes follow from Theorem 4.1 in Dai (1995a). Recall the construction of $Y(t)$. For all $t \geq 0$, $p(t) < M^{-1}l\mathbf{1}$ and $d(t) < \hat{t}_\theta + l\mathbf{1}$, where $\hat{t}_\theta = \max_k (R^{-1}\theta)_k/(R^{-1}\lambda)_k$. From (4.2) we have that $l(n)/n \to 0$, which implies that $(\bar{p}^n(\cdot), \bar{d}^n(\cdot)) = (p^n(n\cdot)/n, d^n(n\cdot)/n) \to (0,0)$ almost surely. This establishes (5.1).

We now turn into analyzing the fluid limit of $\bar{T}^n(\cdot)$. First, we introduce some useful notation. Let $j_{\max}(nt) = \min\{j\colon t_j \geq nt\}$ be the number of review periods up to time nt. Let $T^n(j)$ be the actual vector of time allocations over the execution of the jth review period; that is, $T^n_k(j)$ is the sum of the service times for all class k jobs processed during the jth period. Also, $x^n_\theta(j)$, $x^n_\Psi(j)$ will denote the vector of nominal time allocations for steps 1 and 2 of the DR planning logic, respectively, and $x^n(j) = x^n_\theta(j) + x^n_\Psi(j)$ will denote the total nominal allocation over the jth review period. The queue length vector observed at time t_j will be denoted by $q^n(j)$ and $\bar{q}^n(j) = q^n(j)/n$. Finally, for two sequences a^n, b^n of right continuous functions with left limits, we will say that $a^n(\cdot) \approx b^n(\cdot)$, if $|a^n(\cdot) - b^n(\cdot)| \to 0$, uniformly on compact sets, as $n \to \infty$. Then,

$$\begin{aligned}
\bar{T}^n(t) &= \frac{1}{n} T^n(nt) \\
&= \frac{1}{n} \sum_{j=1}^{j_{\max}(nt)} T^n(j) \\
&= \frac{1}{n} \sum_{j=1}^{j_{\max}(nt)} x^n(j) + \frac{1}{n} \sum_{j=1}^{j_{\max}(nt)} (T^n(j) - x^n(j)) \qquad (5.3) \\
&\overset{\text{(a)}}{\approx} \frac{1}{n} \sum_{j=1}^{j_{\max}(nt)} x^n(j) \\
&\overset{\text{(b)}}{\approx} \frac{1}{n} \sum_{j=1}^{j_{\max}(nt)} x^n_\Psi(j) \\
&= \frac{1}{n} \sum_{j=1}^{j_{\max}(nt)} n T_\Psi(\bar{l}^n; \bar{q}^n(j)) \\
&= \sum_{j=1}^{j_{\max}(nt)} \frac{T_\Psi(\bar{l}^n; \bar{q}^n(j))}{\bar{l}^n} \bar{l}^n \qquad (5.4) \\
&\overset{\text{(c)}}{\longrightarrow} \bar{T}(t) \\
&\overset{\text{(d)}}{=} \int_0^t \dot{\bar{T}}_\Psi(s; \bar{Q}(0, \omega))\, ds. \qquad (5.5)
\end{aligned}$$

The remainder of the proof will justify steps (a)–(d). It is divided in three parts. The first is the derivation of some large deviations estimates that describe the operation of the DR structure, and the second is an application of FSLLN. These parts are used in proving statements (a) and (b). The third part is used in establishing (c) and (d), and it involves the actual properties of the trajectory specification Ψ.

The following result was proved in Lemmas 4.1 and 4.3 of Maglaras (1999). *If $q(j) \geq \mu l$, then $p(j) \leq q(j)$ (this implies that the processing plan will be implemented from jobs present at time t_j) and moreover, for any $\varepsilon > 0$ and for*

l (or $|Q(0)|$) sufficiently large,

$$\mathbf{P}(q(j+1) \ngeq (1-\varepsilon)\theta \mid q(j) \geq \mu l) \leq \exp(-f(\varepsilon)l), \tag{5.6}$$

where $f(\varepsilon)$ is a positive and convex rate function that depends on the distributional characteristics of the interarrival and service time processes.

If $q(j) \ngeq \mu l$, then it may be that the processing plan cannot be implemented from work present at the corresponding buffers at time t_j, and as a result some undesirable idleness may be incurred while shifting work through the network. A specific execution methodology was described in Lemma 4.2 of Maglaras (1999) for which the following result was proved.

For any $\varepsilon > 0$ and $|Q(0)|$ sufficiently large,

$$\mathbf{P}(q(j+1) \ngeq (1-\varepsilon)\theta \mid q(j) \ngeq \mu l) \leq \exp(-h(\varepsilon)l), \tag{5.7}$$

where $h(\varepsilon)$ is a positive and convex rate function. Moreover, for some constant $L > 0$, independent of ε and $|Q(0)|$, and some positive and convex function $g(L)$,

$$\mathbf{P}(t_{j+1} - t_j \geq Ll) \leq \exp(-g(L)l). \tag{5.8}$$

Now, $|Q^n(0)| = n$ and $l^n = l(|Q^n(0)|)$. From condition (4.1), it follows that for any positive constant κ, we can choose n sufficiently large such that $(f(\varepsilon) \wedge h(\varepsilon))l^n \geq \kappa \log(n)$. This will make the upper bounds in (5.6)–(5.8) decay at the polynomial rate $n^{-\kappa}$.

Define the sequence of events $\{A_n\}$, where $A_n = \{\omega \colon q^n(j) \ngeq (1-\varepsilon)\theta^n$ for some $1 \leq j \leq j_{\max}(nt)\}$. Suppose that $|Q(0)| > N(\varepsilon, \kappa)$, where $N(\varepsilon, \kappa)$ is a sufficiently large constant. Then,

$$\begin{aligned}
\mathbf{P}(A_n) &= \mathbf{P}(q^n(1) \ngeq (1-\varepsilon)\theta^n) \\
&\quad + \sum_{j=1}^{j_{\max}} \mathbf{P}(q^n(j+1) \ngeq (1-\varepsilon)\theta^n,\ q^n(i) \geq (1-\varepsilon)\theta^n,\ i \leq j) \\
&\leq e^{-\kappa \log(n)} + \sum_{j=1}^{j_{\max}(nt)} \mathbf{P}(q^n(j+1) \ngeq (1-\varepsilon)\theta^n \mid q^n(i) \geq (1-\varepsilon)\theta^n,\ i \leq j) \\
&\qquad \times \mathbf{P}(q^n(i) \geq (1-\varepsilon)\theta^n,\ i \leq j) \\
&\leq \frac{1}{n^\kappa} + \sum_{j=1}^{j_{\max}(nt)} \mathbf{P}(q^n(j+1) \ngeq (1-\varepsilon)\theta^n \mid q^n(i) \geq (1-\varepsilon)\theta^n,\ i \leq j) \\
&\leq \frac{j_{\max}(nt)}{n^\kappa}.
\end{aligned}$$

For $\kappa \geq 3$,

$$\sum_n \mathbf{P}(A_n) \leq \sum_{n \leq N(\varepsilon,\kappa)} \mathbf{P}(A_n) + \sum_{n > N(\varepsilon,\kappa)} \frac{j_{\max}(nt)}{n^\kappa}$$
$$\leq N(\varepsilon,\kappa) + c \sum_{n > N(\varepsilon,\kappa)} \frac{1}{n^{\kappa-1}} < \infty.$$

Applying the Borel–Cantelli lemma we get that $\mathbf{P}(\limsup_n A_n) = 0$ for any $\varepsilon > 0$. That is, as $n \to \infty$, $Q^n(t_j) \geq \theta^n$ for all $j \geq 1$ a.s. To correct for the fact that potentially some of the components of the initial condition at time t_0 were not above θ^n, one simply incurs a maximum penalty of Ll^n time units that, in turn, is negligible under fluid scaling.

Following the last remark and applying the FSLLN for the interarrival, service time and switching processes we establish step (a); see Proposition 4.2 in Maglaras (1999) for a proof. Furthermore, it follows that for some constant a that depends on μ, λ and P,

$$\sum_{j=1}^{j_{\max}(nt)} \bar{x}^n_\theta(j) \leq \frac{1}{n}(Ll^n + na\varepsilon).$$

Since ε can be made arbitrarily small, it follows that $\sum_j \bar{x}^n_\theta(j) \to 0$ a.s. This establishes step (b).

Existence of the limit in step (c) follows from Theorem 4.1 of Dai (1995a). It remains to prove (d). We know that the Reimann sum in (5.4) converges to a definite integral of the form $\int_0^t \dot{T}_*(s)\,ds$; note that the control process $\dot{T}_*(\cdot)$ has not been specified yet. Using the Lipschitz continuity of $\bar{T}(\cdot)$ we also know that $\dot{\bar{T}}(t)$ exists almost everywhere, and that $\bar{T}(t)$ can also be expressed in the form $\bar{T}(t) = \int_0^t \dot{\bar{T}}(s)\,ds$. Hence, it must be that $\dot{T}_*(t) = \dot{\bar{T}}(t)$ almost everywhere on the real line, and thus, for our purposes, it suffices to characterize the limit $\dot{T}_*(t)$. Let

$$h^n(\bar{q}^n(j)) = \frac{T_\Psi(\bar{l}^n; \bar{q}^n(j))}{\bar{l}^n} = R^{-1}\lambda + R^{-1}\frac{\bar{q}^n(j) - \Psi(\bar{l}^n; \bar{q}^n(j))}{\bar{l}^n}.$$

Fix n, and recall that from (C) we have that $\Psi(\cdot; z)$ is continuous in z. It follows that $h^n(z)$ is continuous in z for every n, and thus

$$\lim_{n\to\infty} h^n(z) = R^{-1}\lambda - R^{-1}\dot{\Psi}(0^+; z) \equiv R^{-1}\lambda + R^{-1}\dot{\Psi}(0; z) = \dot{\bar{T}}_\Psi(0; z). \tag{5.9}$$

To examine the limit of $h^n(\bar{q}^n(j))$, we fix $t \geq 0$ and let j vary with n such that $t_j < nt \leq t_{j+1}$. In this case, we have that $\bar{q}^n(j) = \bar{Q}^n(t - \varepsilon(t)/n)$ for some $\varepsilon(t) < l^n$. From Theorem 4.1 of Dai (1995a), we have that for almost all ω, $\bar{Q}^n(t - \varepsilon(t)/n) \to \bar{Q}(t)$. Using the continuity of h^n and (5.9), we have that almost surely,

$$h^n(\bar{q}^n(j)) = h^n(\bar{Q}^n(t - \varepsilon(t)/n)) \to \dot{\bar{T}}_\Psi(0; \bar{Q}(t)).$$

This completes the proof of (d). [This is the only part where the specific properties of Ψ, and in particular, assumption (C), are invoked.]

Since (a)–(d) are all true almost surely, we have proved that for almost all ω and any time $t \geq 0$,

$$\bar{T}^n(t, \omega) \ \rightarrow \ \bar{T}_\Psi(t; \bar{Q}(0, \omega)).$$

By the Lipschitz continuity of the allocation processes it follows that all of the above conditions are true for any s such that $0 \leq s \leq t$, which implies that the convergence is uniform on compact sets for all $t \geq 0$. This completes the proof of (5.2). Finally, asymptotic tracking follows directly from the definition of this property together with (5.1) and (5.2). □

REMARK. The Chernoff bounds in (5.6)–(5.8) follow from the existence of exponential moments for the interarrival and service time processes; this was postulated as distributional assumption (ii) in Section 3. Also, from the derivation of these large deviations estimates, it appears that $\mathscr{O}(\log(n))$ is a lower bound on the amount of safety stocks that will still guarantee the desired asymptotic performance; this is certainly true in the context of DR policies, but it has also been suggested by other researchers in the recent past in the context of their work. This justifies the choice of review period lengths according to (4.1). To avoid the use of safety stocks altogether one would have to incorporate a "smarter" execution methodology. However, this appears to come with a severe penalty in complexity that renders the implementation of these policies unrealistic.

Finally, we include some remarks regarding the stability properties of these tracking policies.

DEFINITION 5.1 [Dai (1995a)].

(a) An open multiclass queueing network is stable under a specified scheduling policy if the underlying Markov chain is positive Harris recurrent.

(b) The fluid model associated with a scheduling policy is *stable* if there exists a time $\delta > 0$ such that for any solution $(\bar{Q}(\cdot), \bar{T}(\cdot))$ with $|\bar{Q}(0)| = 1$, $\bar{Q}(t) = 0$, for $t \geq \delta$.

The main result we use is due to Dai [(1995a), Theorem 4.2].

THEOREM 5.2. *A multiclass open queueing network is stable under a scheduling policy if the associated fluid model is stable.*

Hence, stability of $\mathbf{DR}(\Psi, l, \beta)$ can be deduced by examining the properties of the fluid specification Ψ. In the context of tracking policies this is natural: "a DR tracking policy is stable, if it is tracking a stable fluid trajectory specification Ψ." The result follows directly from these definitions.

Note. To prove stability one needs to be able to consider initial conditions where $\bar{R}_a = \bar{R}_s \neq 0$, for which Ψ is not defined yet. To fix this problem, we pick a time t_0 such that $t_0 \geq \bar{R}_a$ and $\bar{T}(t) \geq \bar{R}_s$ for all $t \geq t_0$, and then strengthen Proposition 5.1 to hold for all $|z| < B$, for some B that depends on t_0.

PROPOSITION 5.1. *An open multiclass queueing network with* $\rho < \mathbf{1}$ *is stable under* $\mathbf{DR}(\Psi, l, \beta)$, *if there exists a* $\delta > 0$, *such that for all initial conditions* z *with* $|z| < 1$, $\Psi(t; z) = 0$, *for all* $t \geq \delta$.

This result is important from a practical viewpoint, since practitioners can freely design fluid trajectory specifications Ψ without any stringent stability constraint, apart from that target trajectories should drain at some point in the future; the latter is largely irrelevant when one is concerned with the current actions of a tracking policy.

6. Fluid-scale asymptotic optimality. We now return to the problem of optimal control of stochastic processing networks and the policy design procedure outlined in the introduction. We will show that by choosing the target trajectory specification Ψ as the solution of the associated fluid optimal control problem, the results of the previous section can be extended to establish fluid-scale asymptotic optimality of DR policies.

6.1. *Fluid optimal control specifications.* The family of network control problems addressed in stage 1 of the policy design flowchart are defined as follows. Let $g: \mathbf{R}_+^K \to \mathbf{R}_+$ be a $\mathbf{C}^2$ convex cost rate function such that, for some constants $\underline{b}, \underline{c}, \bar{b}, \bar{c} > 0$, where $\underline{b} \leq \bar{b}$ and $\underline{c} \leq \bar{c}$,

$$\underline{b}|x|^{\underline{c}} \leq g(x) \leq \bar{b}|x|^{\bar{c}}. \tag{6.1}$$

Given the cost rate function g, the following stochastic network control problem is considered: choose an allocation process $T(t)$, or equivalently, an admissible policy π, in order to minimize

$$J_T^{\pi}(z) = \mathbf{E}_z^{\pi} \int_0^T g(Q(t))\, dt, \tag{6.2}$$

where $\mathbf{E}_z^{\pi}$ denotes the expectation operator with respect to the probability measure $\mathbf{P}_z^{\pi}$ defined by any admissible policy π and initial condition z. The use of T with no time argument denotes a time horizon and should not be confused with the cumulative allocation $T(t)$. It is natural to think of T as long but finite. Note that if $T < \infty$, then the problem in (6.2) remains meaningful even when $\rho \not< \mathbf{1}$, where, for example, long run averages will not exist. This will allow for an easy extension of our results to the *heavy-traffic* regime, where $\rho \to \mathbf{1}$.

In stage 2 of the flowchart, one proceeds to analyze an associated fluid optimization problem which is defined by

$$\bar{V}_T^g(z) = \min_{\bar{T}(\cdot)} \left\{ \int_0^T g(\bar{Q}(t))\, dt \colon \bar{Q}(0) = z \text{ and } (\bar{Q}, \bar{T}) \in FM \right\}. \tag{6.3}$$

$\bar{V}_T^g(z)$ denotes the value function of the fluid optimization problem starting from the initial condition z and over a control horizon T, which is the same as the one in (6.2), and the superscript g denotes the dependence on the cost rate function.

Consider the set $\mathscr{Q}_z$ of feasible fluid trajectories starting at z. It is easy to show that $\mathscr{Q}_z$ is compact. Consider the set $J(\mathscr{Q}_z) = \{\int_0^T g(\bar{Q}(t))\,dt\colon \bar{Q} \in \mathscr{Q}_z\}$. From the properties of g and the Lipschitz dynamics of fluid models it follows that $J(\mathscr{Q}_z)$ will be a compact set of Lipschitz continuous and convex functionals in z. Existence of $\bar{V}_T^g(z)$ that is defined as the pointwise minimum over the compact set $J(\mathscr{Q}_z)$ follows from the Weierstrass theorem. Moreover, $\bar{V}_T^g(z)$ will itself be Lipschitz continuous and convex in z, and thus almost everywhere differentiable. Its gradient will be denoted by $\nabla\bar{V}_T^g$. The optimal instantaneous allocation, denoted $\dot{\bar{T}}^g(\bar{Q}(t), t)$, can be characterized by a direct application of the dynamic programming principle; see, for example, Bertsekas [(1995), Section 3.2] for background information and a derivation based on the Hamilton–Jacobi–Bellman equation and the maximum principle. It is computed by the following linear program:

$$\begin{aligned} \dot{\bar{T}}^g(t) &= \operatorname{argmin}_{v\in\mathscr{V}(\bar{Q}(t))}\ \nabla\bar{V}_T^g(\bar{Q}(t), t)'(\lambda - Rv) \\ &= \operatorname{argmax}_{v\in\mathscr{V}(\bar{Q}(t))}\ \nabla\bar{V}_T^g(\bar{Q}(t), t)'Rv, \end{aligned} \tag{6.4}$$

where $\mathscr{V}(\bar{Q}(t)) = \{v\colon v \geq 0,\ Cv \leq \mathbf{1},\ (Rv)_k \leq \lambda_k \text{ for all } k \text{ such that } \bar{Q}_k(t) = 0\}$ is the set of admissible controls when the state is $\bar{Q}(t)$.

Next, we analyze the fluid model under the policy (6.4), where we show that optimality implies stability, provided that the control horizon of the fluid optimization problem is sufficiently long. A related result for a linear holding cost infinite horizon criterion is due to Meyn (1997a). Our result holds for convex cost rate functions. The nature of the result is explained by the fact that when the control horizon is small, the optimal trajectory can never reach the origin, and as a result cost minimizing actions need not produce a stable fluid trajectory. In passing, we show that if T is increased above a certain threshold, then the solution of (6.3) coincides with that of the infinite horizon problem. This is very useful in numerical computation of optimal controls, where only finite horizon problems can be addressed.

PROPOSITION 6.1. *Assuming that* $\rho < \mathbf{1}$, *there exists a constant* Δ_g *that depends on the cost rate function* g, *such that if the control horizon* T *in* (6.3) *is such that* $T > \Delta_g$, *then the fluid model* (3.2), (3.3) *and* (6.4) *is stable*.

PROOF. Given an initial condition z, we use the minimum time control described in Section 3 to move back to the origin. The corresponding cost accrued will be denoted by $\hat{V}^g(z)$ and it is computed by

$$\hat{V}^g(z) = \int_0^{t^*(z)} g\left(z(1 - t/t^*(z))\right) dt.$$

This is an upper bound on the value function $\bar{V}^g_T(z)$. [This upper bound is valid even if $t^*(z) \geq T$.]

Let $\bar{Q}^*(\cdot)$ be any fluid solution of (3.2), (3.3) and (6.4). By construction, $\bar{Q}^*(\cdot)$ achieves the minimum draining cost $\bar{V}^g_T(z)$. Given the control horizon T in (6.3), we have that $\min_{t\leq T} g(\bar{Q}^*(t)) \leq \hat{V}_g(z)/T$. Let τ be the time that this minimum is attained. Given the properties of g, we have that $g(\bar{Q}^*(\tau)) \geq \underline{b}|\bar{Q}^*(\tau)|^{\underline{c}}$, which implies that $|\bar{Q}^*(\tau)|^{\underline{c}} \leq \hat{V}^g(z)/(T\underline{b})$. Next, we choose a sufficiently long control horizon T, such that for any $0 < \gamma < 1$, we have that $|\bar{Q}^*(\tau)| \leq \gamma$, independent of the initial condition z. Let

$$\zeta = \max_{|z|=1} \max_i \frac{(CR^{-1}z)_i}{1-\rho_i} \quad \text{and} \quad \delta = \max_{|z|\leq 1} g(z) = \bar{b}. \tag{6.5}$$

For any initial condition such that $|z| = 1$, $\hat{V}^g(z) \leq \delta t^*(z) \leq \bar{b}\zeta$. Let $\Delta_g = \bar{b}\zeta/(\gamma^{\underline{c}}\underline{b})$. Then, for any $T > \Delta_g$, there exists a time $\tau \in (0, \Delta_g]$ such that $|\bar{Q}^*(\Delta)| \leq \gamma$.

The remainder of this proof imitates the arguments in Theorem 6.1 of Stolyar (1995). For $m = 1, 2, \ldots$, let $\tau_m = \min\{t > 0\colon |\bar{Q}^*(t)| \leq \gamma^m,\ |\bar{Q}^*(0)| = \gamma^{m-1}\}$. Modifying (6.5), we can define $\zeta^m = \zeta\gamma^{m-1}$ and $\delta^m = \bar{b}(\gamma^{m-1})^{\bar{c}}$. It follows that

$$\tau_m \leq \Delta_g(\gamma^{\bar{c}+1-\underline{c}})^{m-1}.$$

Clearly, $\sum_m \tau_m \leq \Delta_g/(1-\gamma^{\bar{c}+1-\underline{c}}) \triangleq T_0$. Continuity of $|\bar{Q}^*(t)|$ in t implies that $\lim_m |\bar{Q}^*(\sum_m \tau_m)| = 0$, and therefore, that $\sup\{t > 0\colon |\bar{Q}^*(t)| = 0,\ |\bar{Q}^*(0)| = 1\} \leq \sum_m \tau_m \leq \Delta_g/(1-\gamma^{\bar{c}+1-\underline{c}})$. By observing that the fluid trajectories under the optimal fluid control policy will remain empty once they drain for the first time we complete the proof. □

COROLLARY 6.1. *Consider* (6.3) *with an infinite control horizon* $T = \infty$ *and let* $\bar{Q}^\infty(\cdot)$ *denote the corresponding optimal trajectory. For any initial condition* $|z| \leq 1$ *and any* $\gamma \in (0,1)$; $\bar{Q}^\infty(t) = 0$, *for* $t \geq T_0(\gamma)$, *where*

$$T_0(\gamma) \triangleq \frac{\bar{b}\ \zeta}{\gamma^{\underline{c}}\ \underline{b}} \ \frac{1}{1-\gamma^{\bar{c}+1-\underline{c}}}. \tag{6.6}$$

Minimizing over γ, one gets the smallest upper bound $T_0 \triangleq T_0(\gamma^*)$ at

$$\gamma^* = \left(\frac{\underline{c}}{\bar{c}+1}\right)^{1/(\bar{c}+1-\underline{c})}.$$

When $\rho = \mathbf{1}$, the optimal fluid trajectory will be the cost minimizing path to the state of least achievable cost, which is no longer the origin. In this case, a similar bound can be derived for T_0 by considering an augmented cost rate function, where one has subtracted the cost of the terminal state. Hereafter, it will be assumed that $T \geq T_0$, and thus, that the optimal solution of the finite horizon problem is equal to $\bar{Q}^\infty(\cdot)$.

An alternative characterization of the optimal control policy is based on the optimal fluid trajectories starting at every initial condition. Let $(\bar{Q}^*, \bar{T}^*)$ be a pair of optimal state and control trajectories for (6.3). Then, for all $t \in [0, T]$,

$$\bar{Q}^*(t) = z + \lambda t - R\bar{T}^*(t) \quad \text{and} \quad \int_0^T g(\bar{Q}^*(t))\,dt = \bar{V}_T^g(z). \tag{6.7}$$

Let $\mathscr{Q}_z^g$ be the set of all such minimizers defined by

$$\mathscr{Q}_z^g = \Big\{\bar{Q}^* \in \mathbf{AC}_{R_+^K}[0, T]\colon\ \bar{Q}^*(0) = z,\ \exists\ \bar{T}^*(t),$$
$$t \in [0, T] \text{ s.t. } (\bar{Q}^*, \bar{T}^*) \text{ satisfies (3.3), (6.7)}\Big\}.$$

It has already been established that $\mathscr{Q}_z^g$ is nonempty. Each of the (possibly many) elements in this set is a pair of trajectories that is optimal for (6.3). Define the fluid trajectory map $\Psi^g\colon \mathbf{R}_+^K \to \mathbf{AC}_{R_+^K}[0, \infty)$ that maps an initial condition z to a single target fluid trajectory $\{\Psi^g(t; z),\ t \geq 0\}$, by

$$\begin{aligned} &\Psi^g([0, T]; z) \in \mathscr{Q}_z^g \quad \text{for all} \quad z \geq 0 \quad \text{and} \\ &\Psi^g(t; z) = \Psi^g(T; z) \quad \text{for all } t \geq T. \end{aligned} \tag{6.8}$$

The extension of the optimal fluid trajectory to times beyond the control horizon T of the fluid control problem is justified by the fact that since $T \geq T_0$, $\Psi^g(t; \Psi^g(T; z)) = \Psi^g(T; z)$, for all $t \geq 0$.

The semigroup property (S) follows from the principle of optimality and the fact that $T > T_0$. Using the continuity of $\bar{V}_T^g(z)$ and the properties of g we deduce that $\Psi^g(\cdot; z)$ is continuous in z. From the Lipschitz continuity of the fluid trajectories it follows that $\dot{\Psi}^g(0; z)$ exists a.e. Smoothness of g implies that $\dot{\Psi}^g(0; z)$ will also be a.e. continuous. Hence, assumption (C) holds.

These properties were established for linear cost functions in Pullan (1995), where the optimal controls were shown to be piecewise constant with finitely many discontinuities. For general cost rate functions, these properties follow from classical results in optimal control theory.

6.2. *The fluid-scale asymptotic optimality criterion.* Fluid-scale asymptotic optimality is a relaxed notion of optimality in comparison to the original criterion in (6.2), consistent with the policy design procedure described in the introduction. The following definition is adapted from Meyn (1997b).

DEFINITION 6.1. Consider any sequence of initial conditions $\{y^n\} \subset \mathbf{Y}$ such that $\|y^n\| \to \infty$ as $n \to \infty$ and assume that for every converging subsequence $\{y^{n_j}\}$ and some random variable Z_0, $\bar{Q}^{n_j}(0) \to Z_0$ and $(\bar{R}_a^{n_j}(0),\ \bar{R}_s^{n_j}(0)) \to (0, 0)$ almost surely. Then a policy π^* is said to be asymptotically optimal under fluid scaling if for all admissible scheduling policies π,

$$\liminf_{n\to\infty} \left[\mathbf{E}_{y^n}^{\pi^*} \int_0^T g(\bar{Q}^n(t))\,dt - \mathbf{E}_{y^n}^{\pi} \int_0^T g(\bar{Q}^n(t))\,dt\right] \leq 0. \tag{6.9}$$

Meyn (1997b) stated this definition for the case of linear costs, where he considered the limit as $T \to \infty$, and restricted attention to stable scheduling policies. By focusing on a finite horizon cost, one need not impose this stability restriction, which is difficult to check. Furthermore, the finite horizon criterion remains meaningful even when the traffic intensity at every station is not restricted to be strictly less than one. Moreover, for $T > T_0$, we recover Meyn's criterion in this finite horizon setting. Finally, the assumption regarding the a.s. convergence of the initial fluid scaled queue length vector appears to be a mild one and it will be motivated shortly.

One would like to establish a criterion of fluid-scale asymptotic optimality that depends on the fluid limit trajectories and not the prelimit of fluid scaled sequences as in (6.9). Given that g is nonnegative and all processes $\bar{Q}^n$ are defined in the same probability space, $(\mathbf{Y}, \mathscr{B}_Y)$ equipped with the probability measure P^π [see Section 2.2 in Dai (1995a) for a formal discussion], we have, by Fatou's lemma [see Royden (1963), page 86],

$$\liminf_n \mathbf{E}^\pi_{y^n}\left[\int_0^T g(\bar{Q}^n(t))\,dt\right] \geq \mathbf{E}^\pi\left[\liminf_n \int_0^T g(\bar{Q}^n(t))\,dt\right]. \tag{6.10}$$

For almost every ω by the definition of the liminf, there is a subsequence $\{n_j(\omega)\}$ of $\{n\}$ such that

$$\liminf_n \int_0^T g(\bar{Q}^n(t))\,dt = \lim_j \int_0^T g(\bar{Q}^{n_j}(t))\,dt.$$

From Theorem 4.1 in Dai (1995a), for a.e. ω we have that there is a subsequence $\{n_{j_k}(\omega)\}$ of $\{n_j(\omega)\}$ and a pair of solutions $(\bar{Q}(\cdot,\omega), \bar{T}(\cdot,\omega))$ of the fluid equations such that $(\bar{Q}^{n_{j_k}}(\cdot,\omega), \bar{T}^{n_{j_k}}(\cdot,\omega)) \to (\bar{Q}(\cdot,\omega), \bar{T}(\cdot,\omega))$ u.o.c. as $k \to \infty$. Since g is continuous, it follows that

$$\lim_j \int_0^T g(\bar{Q}^{n_j}(t,\omega))\,dt = \lim_k \int_0^T g(\bar{Q}^{n_{j_k}}(t,\omega))\,dt = \int_0^T g(\bar{Q}(t,\omega))\,dt.$$

Now, by the definition of (6.3), the last integral is bounded below by $\bar{V}^g_T(\bar{Q}(0,\omega))$ and so combining the above we obtain

$$\liminf_n \int_0^T g(\bar{Q}^n(t))\,dt \geq \bar{V}^g_T(\bar{Q}(0,\omega)).$$

In order to take expectations, one needs that $\bar{V}^g_T(\bar{Q}(0,\omega))$ be a random variable. This is postulated as an assumption in this set-up by assuming that $\lim_n \bar{Q}^n(0)$ converges a.s. to some random variable Z_0. Alternatively, one could assume that $\bar{Q}^n(0)$ converges to a random variable in distribution and proceed using some tightness arguments; see Theorems 3 and 4 in Puhalskii and Reiman (1998). Given the assumption regarding the convergence to Z_0, one has that $\bar{Q}(0,\omega) = Z_0(\omega)$ a.s. Second, $\bar{V}^g_T(\cdot)$ is a continuous and convex function of its argument and thus it is measurable with respect to the Borel measure on $\mathbf{R}_+^K$. Hence, $\bar{V}_T^g(\bar{Q}(0,\omega))$ is a properly defined random variable and

$$\liminf_n \mathbf{E}_{y^n}^{\pi}\left[\int_0^T g(\bar{Q}^n(t))\,dt\right] \geq \mathbf{E}^{\pi}\left[\bar{V}_T^g(Z_0)\right], \tag{6.11}$$

where the right member of (6.11) does not depend on π since Z_0 does not depend on π.

Given this set-up, the goal in establishing FSAO will be to exhibit a policy π^* such that for each sequence of initial conditions $\{y^n\}$ for which $\bar{Q}^n(0) \to Z_0$ a.s. as $n \to \infty$, we have that $\lim_n \mathbf{E}_{y^n}^{\pi^*}\left[\int_0^T g(\bar{Q}^n(t))\,dt\right]$ exists and

$$\lim_n \mathbf{E}_{y^n}^{\pi^*}\left[\int_0^T g(\bar{Q}^n(t))\,dt\right] = \mathbf{E}\left[\bar{V}_T^g(Z_0)\right]. \tag{6.12}$$

From (6.11) and (6.12) we have that for any admissible policy π,

$$\lim_n \mathbf{E}_{y^n}^{\pi^*}\left[\int_0^T g(\bar{Q}^n(t))\,dt\right] \leq \liminf_n \mathbf{E}_{y^n}^{\pi}\left[\int_0^T g(\bar{Q}^n(t))\,dt\right],$$

which is (6.9) from Definition 6.1. It remains, of course, to establish that such a policy π^* exists, since otherwise condition (6.12) would be overly stringent.

In conclusion, FSAO can be checked by analyzing the fluid limits under a candidate policy π. For the sequence of initial conditions $\bar{Q}^n(0) = nz$, where $\lim_n \bar{Q}^n(0) = z$ a.s., FSAO is reduced to checking whether starting from an arbitrary initial condition z, the fluid limit under a candidate policy achieves the optimal cost $\bar{V}_T^g(z)$; this was the check performed in Section 2.

6.3. *The policy* $\mathbf{DR}(\Psi^g, l, \beta)$. We now use the trajectory mapping Ψ^g, derived from the solution of the fluid optimization problem (6.3), to define the discrete-review policy $\mathbf{DR}(\Psi^g, l, \beta)$ in the usual way. Proceeding as in Section 5, we get that the fluid model associated with $\mathbf{DR}(\Psi^g, l, \beta)$ is the set of equations (3.2) and (3.3) together with the policy specific equation

$$\bar{Q}(\cdot,\omega) = \Psi^g(\cdot; Z_0(\omega)). \tag{6.13}$$

Next, we follow the discussion above to show that this policy is FSAO.

PROPOSITION 6.2. *An open multiclass queueing network under the policy* $\mathbf{DR}(\Psi^g, l, \beta)$ *is asymptotically optimal under fluid scaling.*

PROOF. We have to show that condition (6.12) is satisfied, where the policy π is $\mathbf{DR}(\Psi^g, l, \beta)$. By the definition of the liminf and from Theorem 4.1 in Dai (1995a) for a.e. ω we have that there is a subsequence $\{n_{j_k}(\omega)\}$ of $\{n_j(\omega)\}$ and a pair of solutions $(\bar{Q}(\cdot,\omega), \bar{T}(\cdot,\omega))$ of the fluid equations (3.2), (3.3) and (6.13), such that $(\bar{Q}^{n_{j_k}}(\cdot,\omega), \bar{T}^{n_{j_k}}(\cdot,\omega)) \to (\bar{Q}(\cdot,\omega), \bar{T}(\cdot,\omega))$ u.o.c. as $k \to \infty$. Since g is continuous, it follows that

$$\lim_k \int_0^T g(\bar{Q}^{n_{j_k}(\omega)}(t,\omega))\,dt = \int_0^T g(\bar{Q}(t,\omega))\,dt.$$

Equation (6.13) and the definition of the set $\mathscr{D}^g_{Z_0(\omega)}$ imply that

$$\int_0^T g(\bar{Q}(t,\omega))\,dt = \bar{V}^g_T(Z_0(\omega)).$$

By the continuity of g it follows that

$$\lim_n \int_0^T g(\bar{Q}^n(t,\omega))\,dt = \bar{V}^g_T(Z_0(\omega)). \tag{6.14}$$

It remains to show that $\lim_n \ \mathbf{E}^{\pi}_{y^n}[\int_0^T g(\bar{Q}^n(t))\,dt]$ exists and in particular, that the limit and the expectation operator can be interchanged, for which a uniform integrability condition is needed. The following estimate on the queue length process was given in Lemma 4.5 in Dai (1995a),

$$\frac{1}{\|y\|} Q^y(\|y\|\delta) \le \sum_{k\in\mathscr{E}} \frac{1}{\|y\|} E^y_k(\|y\|\delta) + 1. \tag{6.15}$$

Given that the cost rate function g is of polynomial growth [this follows from (6.1)], and that the distributional assumption (ii) of Section 3 ensures the integrability of any polynomial moment of $E(t)$, it follows from (6.15) that $g(\bar{Q}^n(\cdot))$ is uniformly integrable. Hence, passing to the limit when taking expectations in (6.14) one gets that

$$\lim_n \mathbf{E}^{\pi}_{y^n}\left[\int_0^T g(\bar{Q}^n(t))\,dt\right] = \mathbf{E}^{\pi}\left[\int_0^T g(\bar{Q}(t))\,dt\right] = \mathbf{E}\left[\bar{V}^g_T(Z_0)\right], \tag{6.16}$$

which establishes the desired result. □

Propositions 5.1 and 6.1 establish that the DR translation, as well as any other mechanism, that leads to a FSAO policy will also be stable; this addresses a question raised in Meyn (1997b).

COROLLARY 6.2. *An open multiclass queueing network with* $\rho < \mathbf{1}$ *is stable under* $\mathbf{DR}(\Psi^g, l, \beta)$.

Returning to the Rybko–Stolyar example, a fluid-scale asymptotically optimal discrete-review policy can be defined using the optimal control description of Section 2. The choices of threshold values and planning horizon lengths can be defined from any vector β and function $l(\cdot)$ that satisfy the appropriate conditions. This policy can be further simplified by exploiting the structure of the Rybko–Stolyar network in order to form a threshold or continuous-review policy that achieves the same asymptotic performance. The desired threshold policy will be one that gives priority to the exiting class at each server, unless the exit class at the other server is below the associated threshold requirement, in which case the incoming class gets higher priority. This is LBFS with priority reversal below the desired threshold and is denoted by $\mathbf{CR}(\Psi^g, \theta)$. Figure 2 illustrates the tracking behavior of these policies by overlaying simulated

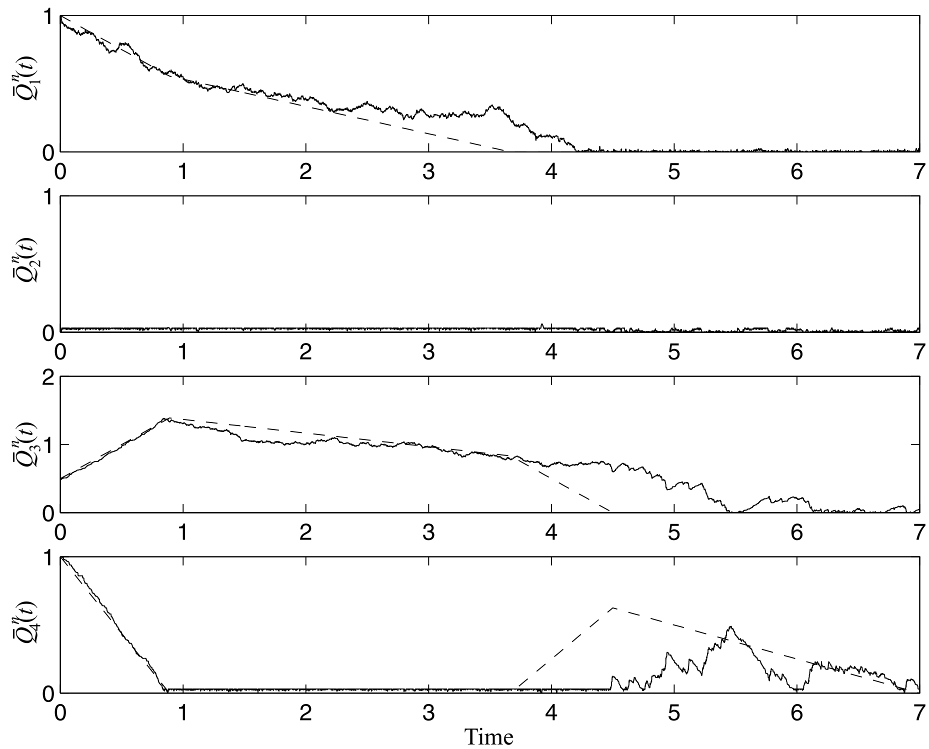

FIG. 2. *Trajectory tracking: optimal fluid trajectories versus state trajectories under* $\mathbf{CR}(\Psi^g, \theta)$; $Q(0) = 200[1, 0, 0.5, 1]$, $(n = 200)$.

trajectories for the continuous-review implementation with the optimal fluid trajectories. [A general discussion of continuous-review policies can be found in Maglaras (1998).]

7. Concluding remarks. We have described a general methodology for dynamic control of stochastic processing networks based on fluid model approximations. While their deterministic and continuous dynamics make them simple to analyze, optimize and control, they still retain much of the structure of the original systems under investigation, and as a result they provide a very promising framework for synthesis of stochastic network controllers.

The main result of this paper was to describe the first general mechanism that can translate an almost arbitrary fluid control policy into a simple, implementable discrete-review tracking rule for the stochastic system that is (a) asymptotically tracking under fluid-scaling, and (b) stable, provided that the fluid control policy one started with was itself stable. Each policy in this family steps through time in *discrete* intervals, and during each of these intervals, system status is *reviewed* and scheduling decisions are made in order to *track* what the fluid control policy would do starting at that point. Successful implementation involves the use of long review periods for accurate tracking, and the enforcement of safety stocks that prevent the system from approaching the boundaries, where the discrete and stochastic dynamics of the network can be significantly different from their fluid idealization.

Many interesting future directions of research arise. The results on fluid scale asymptotic analysis need to be refined in order to capture the speed of convergence of the system dynamics to their fluid limit; this appears to be a tractable undertaking due to the good structure of the proposed policies. Such an extension will provide bounds on how close to the true optimal performance these discrete-review policies are and give new insights on the inherent tradeoff of the proposed framework, where one has restricted attention to deterministic fluid model analysis. On a separate front, one should contrast our results to those obtained using diffusion approximations and heavy-traffic analysis and also explore the applicability of the techniques developed in this paper to the translation of Brownian control policies.

Acknowledgments. I am deeply grateful to Michael Harrison for his continuous support throughout the course of this work and to Sunil Kumar for his numerous comments, suggestions, and overall help. I am also grateful to the anonymous referees and Associate Editors for their detailed comments and suggestions that have greatly improved the paper. In particular, the general articulation of the fluid-scale asymptotic optimality criterion follows one of their earlier suggestions, and the simplified version of DR policies that does not involve the use of an LP was motivated by one of their comments.

REFERENCES

ATHANS, M. and FALB, P. (1966). *Optimal Control*. McGraw-Hill, New York.

ATKINS, D. and CHEN, H. (1995). Performance evaluation of scheduling control of queueing networks: fluid model heuristics. *Queueing Sys.* **21** 391–413.

AVRAM, F., BERTSIMAS, D. and RICARD, M. (1995). Fluid models of sequencing problems in open queueing networks; an optimal control approach. In *Stochastic Networks*, *Proceedings of the IMA*(F. Kelly and R. Williams, eds.) **71** 199–234. Springer, New York.

BAMBOS, N. and WALRAND, J. (1993). Scheduling and stability aspects of a general class of parallel processing systems. *Adv. in Appl. Probab.* **25** 176–202.

BERTSEKAS, D. P. (1995). *Dynamic Programming and Optimal Control* **1**. Athena Scientific, Belmont, MA.

BERTSIMAS, D., GAMARNIK D. and SETHURAMAN, J. (1999). Asymptotically optimal solutions for job shop scheduling problems. Technical report, Operations Research Center, MIT.

BERTSIMAS, D., PASCHALIDIS, I. C. and TSITSIKLIS, J. N. (1995). Branching bandits and Klimov's problem: achievable region and side constraints. *IEEE Trans. Automat. Control* **40** 2063–2075.

BERTSIMAS, D. and VAN RYZIN, G. (1991). A stochastic and dynamic vehicle routing problem in the Euclidean plane. *Oper. Res.* **39** 601–615.

BRAMSON, M. (1998). Stability of two families of queueing networks and a discussion of fluid limits. *Queueing Sys.* **28** 7–31.

BRYSON, A. E. and HO, Y. C. (1975). *Applied Optimal Control*. Hemisphere Publishing, New York.

CHEN, H. (1995). Fluid approximations and stability of multiclass queueing networks: work-conserving policies. *Ann. Appl. Probab.* **5** 637–655.

CHEN, H. and MANDELBAUM, A. (1991). Discrete flow networks: bottleneck analysis and fluid approximations. *Math. Oper. Res.* **16** 408–446.

CHEN, H. and YAO, D. (1993). Dynamic scheduling of a multiclass fluid network. *Oper. Res.* **41** 1104–1115.

CHEN, R.-R. and MEYN, S. P. (1999). Value iteration and optimization of multiclass queueing networks. *Queueing Sys.* **32** 65–97.

DAI, J. G. (1995a). On positive Harris recurrence of multiclass queueing networks: a unified approach via fluid limit models. *Ann. Appl. Probab.* **5** 49–77.

DAI, J. G. (1995b). Stability of open multiclass queueing networks via fluid models. In *Stochastic Networks Proceedings of the IMA* (F. Kelly and R. Williams, eds.) Springer, New York. **71** 71–90.

DAI, J. G. and MEYN, S. (1995). Stability and convergence of moments for multiclass queueing networks via fluid limit models. *IEEE Trans. Autmat. Control* **40** 1889–1904.

DAI, J. G. and WEISS, G. (1996). Stability and instability of fluid models for certain re-entrant lines. *Math. Oper. Res.* **21** 115–134.

DAI, J. G. and WEISS, G. (1999). A fluid heuristic for job shop scheduling. Technical report, Dept. Statistics, Univ. Haifa.

DAVIS, M. H. A. (1984). Piecewise-deterministic Markov processes: a general class of nondiffusion stochastic models. *J. Roy Statist. Soc. Ser. B* **46** 353–388.

ENG, D., HUMPHREY, J. and MEYN, S. P. (1996). Fluid network models: linear programs for control and performance bounds. In *13th World Congress of International Federation of Automatic Control* **2** 19–24 North-Holland, Amsterdam.

GANS, N. and VAN RYZIN, G. (1997). Optimal control of a multiclass, flexible queueing system. *Oper. Res.* **45** 677–693.

HARRISON, J. M. (1988). Brownian models of queueing networks with heterogeneous customer populations. In *Stochastic Differential Systems, Stochastic Control Theory and Applications Proceedings of the IMA* (W. Fleming and P. L. Lions, eds.) **10** 147–186. Springer, New York.

HARRISON, J. M. (1996). The BIGSTEP approach to flow management in stochastic processing networks. In *Stochastic Networks: Theory and Applications* (F. Kelly, S. Zachary and I. Ziedins, eds.) 57–90. Oxford Univ. Press.

HARRISON, J. M. (1998). Heavy traffic analysis of a system with parallel servers: Asymptotic optimality of discrete-review policies. *Ann. Appl. Probab.* **8** 1–24.

HARRISON, J. M. and WEIN, L. M. (1989). Scheduling network of queues: heavy traffic analysis of a simple open network. *Queueing Sys.* **5** 265–280.

HARRISON, J. M. and WEIN, L. M. (1990). Scheduling networks of queues: heavy traffic analysis of a two-station closed network. *Oper. Res.* **38** 1052–1064.

KELLY, F. and LAWS, N. (1993). Dynamic routing in open queueing networks: Brownian models, cut constraints and resource pooling. *Queueing Sys.* **13** 47–86.

KUMAR, P. R. and SEIDMAN, T. I. (1990). Dynamic instabilities and stabilization methods in distributed real-time scheduling of manufacturing systems. *IEEE Trans. Autmat. Control* **35** 289–298.

KUSHNER, H. J. and MARTINS, L. F. (1996). Heavy traffic analysis of a controlled multiclass queueing network via weak convergence methods. *SIAM J. Control Optim.* **34** 1781–1797.

LU, S. H. and KUMAR, P. R. (1991). Distributed scheduling based on due dates and buffer priorities. *IEEE Trans. Autmat. Control* **36** 1406–1416.

LUO, X. and BERTSIMAS, D. (1996). A new algorithm for state constrained separated linear programs. Technical report, MIT.

MAGLARAS, C. (1997). A methodology for dynamic control policy design for stochastic processing networks via fluid models. In *Proceedings IEEE Conference on Decision and Control* 1208–1214. IEEE, New York.

MAGLARAS, C. (1999). Dynamic scheduling in multiclass queueing networks: stability under discrete-review policies. *Queueing Sys.* **31** 171–206.

MAGLARAS, C. (2000). Tracking Policies for control of stochastic networks: Brownian models, heavy traffic and asymptotic optimality. Technical report, Columbia University.

MEYN, S. P. (1997a). The policy improvement algorithm for Markov decision processes with general state space. *IEEE Trans. Autmat. Control* **42** 1663–1680.

MEYN, S. P. (1997b). Stability and optimization of queueing networks and their fluid models. In *Mathematics of Stochastic Manufacturing Systems*. *Lectures in Applied Math.* 175–200. Amer. Math. Soc. Providence, RI.

NEWELL, G. F. (1971). *Applications of Queueing Theory*. Chapman and Hall, London.

PUHALSKII, A. and REIMAN, M. (1998). A critically loaded multirate link with trunk reservation. *Queueing Sys.* **28** 157–190.

PULLAN, M. C. (1993). An algorithm for a class of continuous linear programs. *SIAM J. Control Optim.* **31** 1558–1577.

PULLAN, M. C. (1995). Form of optimal solutions for separated continuous linear programs. *SIAM J. Control Optim.* **33** 1952–1977.

PULLAN, M. C. (1996). A duality theory for separated continuous linear programs. *SIAM J. Control Optim.* **34** 931–965.

ROYDEN, L. H. (1963). *Real Analysis*. Macmillan, New York.

RYBKO, A. N. and STOLYAR, A. L. (1992). Ergodicity of stochastic processes describing the operations of open queueing networks. *Problems Inform. Transmission* **28** 199–220.

STOLYAR, A. L. (1995). On the stability of multiclass queueing networks: a relaxed sufficient condition via limiting fluid processes. *Markov Processes Related Fields* **1** 491–512.

TASSIULAS, L. and PAPAVASSILIOU, S. (1995). Optimal anticipative scheduling with asynchronous transmission opprtunities. *IEEE Trans. Autmat. Control* **40** 2052–2062.

WEISS, G. (1996). Optimal draining of fluid re-entrant lines: some solved examples. In *Stochastic Networks: Theory and Applications* (F. Kelly, S. Zachary and I. Ziedins, eds.) 19–34. Oxford Univ. Press.

WEISS, G. (1997). Algorithm for minimum wait draining of two station fluid re-entrant line. Technical report, Dept. Statistics, Univ. Haifa.

WEISS, G. (1999). Scheduling and control of manufacturing systems via optimal control of multiclass fluid networks. In *Proceedings of 37th Allerton Conference on Communications, Control and Computing*.

WILLIAMS, R. J. (1996). On the approximation of queueing networks in heavy traffic. In *Stochastic Networks: Theory and Applications* (F. Kelly, S. Zachary, and I. Ziedins, eds.) 35–56. Oxford Univ. Press.

COLUMBIA BUSINESS SCHOOL
409 URIS HALL
3022 BROADWAY
NEW YORK, NEW YORK 10027
E-MAIL: c.maglaras@columbia.edu