

SCHEDULING IN A QUEUING SYSTEM WITH ASYNCHRONOUSLY VARYING SERVICE RATES

MATTHEW ANDREWS, KRISHNAN KUMARAN,
KAVITA RAMANAN, AND ALEXANDER STOLYAR

*Bell Labs
Lucent Technologies
Murray Hill, New Jersey 07974
E-mail: {andrews,kumaran,kavita,stolyar}@research.bell-labs.com*

RAJIV VIJAYAKUMAR

*University of Michigan
Ann Arbor, Michigan
E-mail: rvijayak@engin.umich.edu*

PHIL WHITING

*Bell Labs
Lucent Technologies
Murray Hill, New Jersey 07974
E-mail: pwhiting@research.bell-labs.com*

We consider the following queuing system which arises as a model of a wireless link shared by multiple users. There is a finite number N of input flows served by a server. The system operates in discrete time $t = 0, 1, 2, \dots$. Each input flow can be described as an irreducible countable Markov chain; waiting customers of each flow are placed in a queue. The sequence of server states $m(t), t = 0, 1, 2, \dots$, is a Markov chain with finite number of states M . When the server is in state m , it can serve μ_i^m customers of flow i (in one time slot).

The scheduling discipline is a rule that in each time slot chooses the flow to serve based on the server state and the state of the queues. Our main result is that a simple online scheduling discipline, Modified Largest Weighted Delay First, along with its generalizations, is *throughput optimal*; namely, it ensures that the queues are stable as long as the vector of average arrival rates is within the system maximum stability region.

1. INTRODUCTION

We consider a model motivated by the problem of scheduling transmissions of multiple data users (flows) sharing the same wireless channel (server). The unique “wireless” feature of this problem is the fact that the capacity (service rate) of the channel varies with time randomly and *asynchronously* for different users. The variations of the channel capacity are due to different, random interference levels observed by different users and due to *fast fading* of the signal received by a user. We will refer to this problem as the *variable channel scheduling problem*.

The variable channel problem arises, for example, in the 3G CDMA High Data Rate (HDR) system [6]. (See also [27] for a background on CDMA wireless systems.) In HDR, multiple mobile users in a cell share the same CDMA wireless channel. On the downlink (the link from the cell base station to users), time is divided into fixed-size (1.67-ms) time slots. This slot size is short enough so that (each user’s) channel quality stays approximately constant within one or even a few consecutive time slots. (To be more precise, this is true only for relatively low mobile user velocities; see [27].) In each time slot, data can be transmitted to only one user. Each user constantly reports to the base station its “instantaneous” channel capacity (i.e., the rate at which data can be transmitted if this user is scheduled for transmission in the current time slot).

In the HDR system (and in the generic variable channel model as well), a scheduling algorithm can take advantage of channel variations by giving some form of priority to users with (temporarily) better channels. Since channel capacities of different users vary in time in an asynchronous manner, the quality of service (QoS) of all users can be improved, as compared to scheduling schemes which do not take channel conditions into account. A scheduling rule providing *proportional fairness* in the achieved long-term throughput of different users was proposed and analyzed in [25]. (See also [26].)

The QoS of a data user can be defined in different ways. If data users are *real-time* users, then the packet delays of each flow need to be kept below a certain threshold. This means that the primary goal of a scheduling algorithm is to keep all queues *stable* (i.e., to be able to handle all the offered traffic without queues “blowing up”).

In this article, we consider the generic variable channel scheduling model. Our main result is that a simple online scheduling discipline, *modified largest weighted delay first (M-LWDF)*, is *throughput optimal*; namely it ensures that the queues are stable as long as the vector of average arrival rates is within the system’s *maximum stability region*.

In a time slot t , the M-LWDF discipline serves the flow j for which

$$\gamma_j [W_j(t)]^\beta \mu_j(t) \tag{1}$$

is maximal, where $W_j(t)$ is the head-of-the-line packet delay for flow j , $\mu_j(t)$ is the server capacity for flow j at time t , and β and the γ_j ’s are arbitrary positive constants. (The name M-LWDF is because this discipline is a generalization of the LWDF

discipline [1,22].) Moreover, as we discuss in Section 4, our result actually holds for a quite wide class of disciplines (of which M-LWDF is a member) and a more general class of models. In particular, the throughput optimality holds if instead of maximizing (1), the scheduling rule maximizes

$$\gamma_j [V_j(t)]^\beta \mu_j(t), \tag{2}$$

where $V_j(t) = \eta_j^{(w)} W_j(t) + \eta_j^{(Q)} Q_j(t)$. Here, $\eta_j^{(w)} \geq 0$ and $\eta_j^{(Q)} \geq 0$ are arbitrary parameters for flow j , not equal to zero simultaneously and possibly dependent on j .

Our main stability results are closely related to the series of results on the stability of *MaxWeight-type* scheduling algorithms in queuing networks and in input-buffered crossbar switches. The first results of this type were obtained by Tassiulas and Ephremides [23,24] in the context of wireless systems. For the switch scheduling stability results, see [15,17] and a recent paper [10]. In the context of interactive parallel server systems and systems with randomly varying connectivity, MaxWeight-type stability results were obtained in [3,5]. (See also [4], which is a recent extension of [3].)

The underlying intuition behind the stability of a MaxWeight-type algorithm is the fact that it minimizes the drift of a Lyapunov function of the form $\sum_j [V_j(t)]^{\beta+1}$. Most of the algorithms studied before are for the case $\beta = 1$ and $V_j(t) = Q_j(t)$. As far as we are aware, [17] was the first in which the stability result for a MaxWeight-type rule using flow delays $W_j(t)$ (as opposed to queue lengths $Q_j(t)$) was derived. (A similar result was formulated but not proved in [14].)

The main contribution of this article is that we show that a MaxWeight-type algorithm retains stability properties even if the “weight” of an individual queue j has a form as general as $[V_j(t)]^\beta$. Such a generalization is important because the additional parameters β , $\eta_j^{(w)}$, and $\eta_j^{(Q)}$ allow for a more flexible control of queue lengths and delay distributions, to satisfy a variety of QoS constraints. For example, if we are interested in giving tight delay bounds to a flow j with a low arrival rate, then the “weight” for flow j should be based more on head-of-the-line packet delay than on queue length (i.e., $\eta_j^{(w)}$ should be large relative to $\eta_j^{(Q)}$). Conversely, if flow j has a high arrival rate and we want to bound its buffer space requirements, then $\eta_j^{(Q)}$ should be large relative to $\eta_j^{(w)}$.

To prove our stability results, we use the *fluid limit* technique [7–9,19,20]. (For a MaxWeight-type rule, the technique was also used in [10] in a “switch” model context.) Use of this technique makes the above-described generalization very natural. Roughly speaking, in the “fluid limit” and after some initial period of time, $Q_j(t)$ and $W_j(t)$ stay proportional to each other; thus, MaxWeight algorithms using $Q_j(t)$, $W_j(t)$, or a linear combination $V_j(t)$ are in some sense “indistinguishable” in the fluid limit.

It is shown recently in [21], which analyzes a more general (described in Section 4.2) version of our model, that, in addition to throughput optimality, MaxWeight-type rules have certain asymptotic optimality properties when the system is heavily loaded.

Practical implications of using M-LWDF to provide QoS for real-time data users are addressed in [2]. In particular, we show in [2] that the M-LWDF discipline, with “appropriately” chosen parameters γ_i , provides good QoS defined in terms of the probabilities of packet delays exceeding predefined thresholds.

The rest of the article is organized as follows. In Section 2, we introduce the formal *variable channel scheduling* queuing model. Necessary and sufficient stability conditions are derived and the system stability region is defined in Section 3. In Section 4, we introduce the M-LWDF scheduling rule and formulate our main result—Theorem 3, which states that M-LWDF (along with a wide class of rules generalizing it) is throughput optimal. The proof of Theorem 3 is presented in Section 5.

2. VARIABLE CHANNEL SCHEDULING MODEL

Consider the following queuing system. There is a finite number N of input flows, indexed by $i = 1, 2, \dots, N$, served by a *server*. Each input flow consists of discrete *customers*. (One customer models one byte or bit of data). The system operates in discrete time $t = 0, 1, 2, \dots$. By convention, we will

- (a) identify an (integer) time t , with the unit time interval $[t, t + 1)$, which will sometimes be referred to as the *time slot* t
- (b) assume that all processes we consider are constant within each time slot.

There is a finite set $\{1, \dots, M\}$ of server *states*. This set itself we also denote by M (as well as its cardinality). Associated with each state $m \in M$ is a fixed vector of service rates $(\mu_1^m, \dots, \mu_N^m)$, where all μ_i^m are nonnegative integers. The meaning of μ_i^m is as follows. If in time slot t the server is in state m and the service (in this time slot) is given exclusively to queue i , then μ_i^m type i customers are served from those present at time t (or the entire queue i content at t , whichever is less). We assume that, within each type, customers are served in the order of their arrival in the system.

The random server state process $\mathbf{m} = m(t)$, $t = 0, 1, 2, \dots$ is assumed to be an irreducible (see [12]) discrete-time Markov chain with the (finite) state space M . The (unique) stationary distribution of this Markov chain we denote by $\pi = (\pi_1, \dots, \pi_M)$. Note that, due to irreducibility, $\pi_m > 0$ for all $m \in M$.

We make a nondegeneracy assumption that for each flow i , there is at least one server state $m \in M$ such that $\mu_i^m > 0$. (Otherwise, we would have flows which simply can never be served.)

Denote by $A_i(t)$ the number of type i customers that arrived at time t , and assume by convention that these customers are immediately available for service. We assume that each input process A_i is an irreducible positive recurrent (see [12]) Markov chain with countable state space and that the input processes are mutually independent. (This condition can be relaxed as follows. The aggregate arrival process $A = \{(A_1(t), \dots, A_N(t)), t = 1, 2, \dots\}$ can be described by a finite number of regenerative processes [12] with finite mean regeneration cycles.) Let us denote by λ_i , $i = 1, \dots, N$, the mean arrival rate for flow i (i.e., the mean number of type i

customers arriving in one time slot). The vector of mean arrival rates is denoted by $\lambda \doteq (\lambda_1, \dots, \lambda_N)$.

The random process describing the behavior of the entire system is $(S = S(t), t = 0, 1, 2, \dots)$, where

$$S(t) = \{(U_{i1}(t), \dots, U_{iQ_i(t)}(t)), i = 1, \dots, N; m(t)\},$$

$Q_i(t)$ is the type i queue length at time t , and $U_{ik}(t)$ is the current sojourn time, or *delay*, of the k th type i customer present in the system at time t . (Within each type, the customers are numbered in the order of their arrival.)

A mapping H which takes a system state $S(t)$ in a time slot into a fixed probability distribution $H(S(t))$ on the set of queues N will be called a *scheduling rule*, or a *queueing discipline*. With a fixed discipline H , the queue to serve at time t is chosen randomly according to the distribution $H(S(t))$. So, the number $D_i(t)$ of type i customers served in the time slot t is equal to $\min\{Q_i(t), \mu_i^{m(t)}\}$ if queue i is chosen for service and equal to zero otherwise. According to our conventions, for each time t ,

$$Q_i(t + 1) = Q_i(t) - D_i(t) + A_{i+1}(t), \quad \forall i.$$

Our assumptions imply that with any scheduling rule, S is a discrete-time countable Markov chain. By *stability* of the Markov chain S (and stability of the system) we mean the following property: The set of positive recurrent states is nonempty and it contains a finite subset which is reached with probability one (within finite time) from any initial state. Stability implies the existence of a stationary probability distribution. (If all positive recurrent states are connected, the stationary distribution is unique.)

We conclude this section with some basic notation we use throughout the article. Vector inequalities are understood componentwise; $\lfloor z \rfloor$ and $\lceil z \rceil$ denote the integer part and the “ceiling” of a real number z , respectively. We say that a function $f(t)$ of a real variable t is RCLL if it is right-continuous and has left limit in every point t of its domain. The abbreviation “u.o.c.” in a convergence statement means that the convergence is uniform on any fixed compact subset of the corresponding function domain. We denote by $\mathcal{N} = \{1, 2, \dots\}$ the set of positive natural numbers.

3. NECESSARY AND SUFFICIENT STABILITY CONDITIONS. STABILITY REGION

Suppose a stochastic matrix $\phi = (\phi_{mi}, m \in M, i = 1, \dots, N)$ is fixed, which means that $\phi_{mi} \geq 0$ for all m and i , and $\sum_i \phi_{mi} = 1$ for every m . Consider a *static service split* (SSS) scheduling rule, parameterized by the matrix ϕ . When the server is in state m , the SSS rule chooses for service queue i with probability ϕ_{mi} . (The word *static* in the name of the rule reflects the fact that scheduling decisions depend only on the server state.) Clearly, the vector $v = (v_1, \dots, v_N) = v(\phi)$, where

$$v_i = \sum \pi_m \phi_{mi} \mu_i^m$$

gives the long-term average service rates allocated to different flows. This observation makes the following simple (and quite standard) result very intuitive.

THEOREM 1: *For the existence of a scheduling rule H under which the system is stable, condition (3) is necessary*

$$\lambda \leq v(\phi) \quad \text{for some stochastic matrix } \phi \tag{3}$$

and condition (4) is sufficient

$$\lambda < v(\phi) \quad \text{for some stochastic matrix } \phi. \tag{4}$$

PROOF: The necessity of condition (3) is almost obvious. Consider a rule H under which the system is stable and consider the Markov chain S in a stationary regime. (Such a stationary regime exists, but is not necessarily unique.) We will denote by $H_i(s)$ the probability with which the SSS rule chooses for service the queue i when $S(t) = s$. Then, for any i (and arbitrary fixed time slot t), we can write

$$\begin{aligned} \lambda_i &= EA_i(t) = ED_i(t) = \sum_m \pi_m E(D_i(t) | m(t) = m) \\ &\leq \sum_m \pi_m \sum_s P(S(t) = s | m(t) = m) H_i(s) \mu_i^m \\ &= \sum_m \pi_m \phi_{mi} \mu_i^m, \end{aligned}$$

where

$$\phi_{mi} \doteq \sum_s P(S(t) = s | m(t) = m) H_i(s).$$

Obviously, we have $\sum_i \phi_{mi} = 1$ for each m . The necessity of (3) is proved.

Sufficiency of condition (4) is almost obvious as well. The SSS rule associated with any matrix ϕ satisfying (4) makes the system stable. Indeed, the rates at which service is provided to different flows i is a random process “modulated” by the underlying (ergodic) Markov chain \mathbf{m} , independent of the aggregate arrival process A . Moreover, the average service rate $v_i(\phi)$ available to each flow i is strictly greater than its average arrival rate λ_i . If the Markov chain of interest would be

$$\{(Q_1(t), \dots, Q_N(t)); m(t)\}, \quad t = 0, 1, 2, \dots$$

(viz. its states would track queue lengths only), then, for example, $\max_i Q_i(t)$ can be used as a Lyapunov function to show the stability via standard “drift” criteria, such as those in [18]. However, the states of our Markov chain S include customer sojourn times as well. To accommodate this, the stability proof for the SSS rule (assuming (4)) can be obtained, for example, as a much simplified version of the proof of M-LWDF rule stability (Theorem 3), which is the main result of this article. Since such a proof requires a fair amount of preliminaries, introduced later in the article, we present its details in the Appendix for the interested reader. (We also note that

Theorem 3 itself implies sufficiency of (4). It is, however, more intuitive, simple, and standard to demonstrate this fact via the SSS rule or a similar static rule. That is why we discuss the SSS rule here. ■

The set of all (average arrival rate) vectors λ satisfying condition (4) is usually called the system *maximum stability region*, or just *stability region*.

An SSS rule associated with stochastic matrix ϕ^* will be called *maximal* if the vector $v(\phi^*)$ is not dominated by $v(\phi)$ for any other stochastic matrix ϕ . (We say that vector $v^{(1)}$ is dominated by vector $v^{(2)}$ if $v_i^{(1)} \leq v_i^{(2)}$ for all i and the strict inequality $v_i^{(1)} < v_i^{(2)}$ holds for at least one i .) The following theorem provides a useful characterization of maximal SSS rules.

THEOREM 2: *Consider a maximal SSS rule associated with a stochastic matrix ϕ^* . Suppose, in addition, that all components of $v^* = v(\phi^*)$ are strictly positive. Then, there exists a set of strictly positive constants $\alpha_i, i = 1, 2, \dots, N$, such that for any m and i ,*

$$\phi_{mi}^* > 0 \text{ implies } i \in \arg \max_j \alpha_j \mu_j^m. \tag{5}$$

The theorem says that a maximal SSS rule always chooses for service at any time t a queue i for which $\alpha_i \mu_i^{m(t)}$ is maximal. (It does not say what to do in case of a tie.)

PROOF: Consider the following linear program:

$$\max_{\Lambda, \{\phi_{mi}\}} \Lambda$$

subject to

$$\sum_{m=1}^M \pi_m \mu_i^m \phi_{mi} \geq \Lambda v_i^*, \quad i = 1, \dots, N, \tag{6}$$

$$\sum_{i=1}^N \phi_{mi} = 1, \quad m \in M, \quad 0 \leq \phi_{mi} \leq 1, \quad m \in M, \quad i = 1, \dots, N. \tag{7}$$

From the definition of v^* , we know that $\Lambda = 1$ and $\phi = \phi^*$ solve this linear program, with constraints (6) satisfied as equalities. Then, by the Kuhn–Tucker theorem (see, e.g., [13]), there exists a set of nonnegative Lagrange multipliers $\alpha_0, \alpha_1, \dots, \alpha_N$ such that $\Lambda = 1$ and $\phi = \phi^*$ also solve the following linear program (with the same value of the maximum):

$$\max_{\Lambda, \{\phi_{mi}\}} \alpha_0 \Lambda + \sum_{i=1}^N \alpha_i \left(\sum_{m=1}^M \pi_m \mu_i^m \phi_{mi} - \Lambda v_i^* \right) \tag{8}$$

subject to

$$\sum_{i=1}^N \phi_{mi} = 1, \quad 0 \leq \phi_{mi} \leq 1, \quad \forall m, i. \tag{9}$$

It is easy to verify that all α_i must be strictly positive and $\alpha_0 = 1$. Then, rewriting (8) as

$$\max_{\Lambda, \{\phi_{mi}\}} \Lambda - \Lambda \sum_{i=1}^N \alpha_i v_i^* + \sum_{m=1}^M \pi_m \sum_{i=1}^N \alpha_i \mu_i^m \phi_{mi},$$

we see that condition (5) must hold, because otherwise the maximum would not be achieved by $\phi = \phi^*$. ■

4. THE MODIFIED LARGEST WEIGHTED DELAY FIRST DISCIPLINE

4.1. Main Result

The following natural question arises. Is there a scheduling rule which (unlike SSS) does not use a priori information about the input rates λ_i and the stationary distribution π of the server state, and yet ensures system stability as long as the necessary and sufficient stability condition (4) is satisfied. Theorem 3 shows that the answer is yes.

Let us call the value

$$W_i(t) \doteq U_{i1}(t)$$

(with $W_i(t) = 0$ if $Q_i(t) = 0$ by convention) the *delay* of flow i at time t .

Let a set of positive constants $\gamma_1, \dots, \gamma_N$ and a positive constant $\beta > 0$ be fixed. We define modified largest weighted delay first (M-LWDF) to be the scheduling rule that chooses for service in time slot t a *single* queue

$$i \in \arg \max_j \gamma_j \mu_j^{m(t)} (W_j(t))^\beta.$$

(The “ties” are broken arbitrarily; for example, in favor of the largest index i .)

An analogous rule, which we will call modified largest weighted (unfinished) work first (M-LWWF), chooses a single queue

$$i \in \arg \max_j \gamma_j \mu_j^{m(t)} (Q_j(t))^\beta.$$

THEOREM 3: *Let an arbitrary set of positive constants $\gamma_1, \dots, \gamma_N$ and $\beta > 0$ be fixed. Then, either of the two scheduling rules, M-LWDF or M-LWWF, are throughput optimal; namely, they make the system stable as long as condition (4) holds (i.e., as long as the arrival rate vector λ is within the system stability region).*

As mentioned in Section 1, our proof of Theorem 3 uses the *fluid limit* technique. This technique allows us to “derive” the stability of M-LWDF from the sta-

bility of M-LWDF using the fact that their fluid limits are in a certain sense indistinguishable.

4.2. Generalizations

It will be clear from the proof of Theorem 3 that this result can be significantly generalized. First, the (virtually unchanged) proof allows us to show throughput optimality of the following “mixed” M-LWDF/M-LWDF rule:

Serve queue

$$i \in \arg \max_j \gamma_j \mu_j^{m(t)} (V_j(t))^\beta,$$

where $V_j = \eta_j^{(W)} W_j + \eta_j^{(Q)} Q_j$, and $\eta_j^{(W)}$ and $\eta_j^{(Q)}$ are nonnegative constants that satisfy $\eta_j^{(W)} + \eta_j^{(Q)} > 0$.

In addition, the model assumption that only one queue may be served at a time can be relaxed as follows. For each server state m , there is an associated finite set $K(m)$ of service rate decisions. Associated with each decision $k \in K(m)$ is a service rate vector

$$(\mu_1^m(k), \dots, \mu_1^m(k)).$$

If the decision k is chosen when the server is in state m , then $\mu_j^m(k)$ customers from each queue j (or the entire queue j content $Q_j(t)$ if it is less than $\mu_j^m(k)$) are served within one time slot. Again, a slightly adjusted proof of Theorem 3 allows us to prove that the following MaxWeight-type rule is throughput optimal:

Choose a service rate decision

$$k \in \arg \max_{k \in K(m(t))} \sum_j \gamma_j \mu_j^{m(t)}(k) (V_j(t))^\beta.$$

In the latter general form, our result includes as special cases the throughput optimality results in both the “switch scheduling” model setting [15,17] (and related ones in [3,14]) and the variable channel scheduling setting, which is the main focus of this article.

5. PROOF OF THEOREM 3

Throughout the proof, we consider a system with a fixed set of parameters such that condition (4) holds. It needs to be proved that this system is stable under both M-LWDF and M-LWDF rules.

To simplify notation, the proof will be for the case $\beta = 1$. The generalization of the proof for arbitrary $\beta > 0$ is trivial: The quadratic Lyapunov function in (36) needs to be replaced by the power law function

$$L(y) = \frac{1}{1 + \beta} \sum_1^N \gamma_i y_i^{1+\beta};$$

in the formulations of Lemmas 2 and 6, $q_i(t)$, $q_j(t)$, $w_i(t)$, and $w_j(t)$ need to be replaced by $q_i(t)^\beta$, $q_j(t)^\beta$, $w_i(t)^\beta$, and $w_j(t)^\beta$, respectively; corresponding minor adjustments need to be made throughout the proofs.

5.1. Preliminaries

Let us define the norm of the state $S(t)$ as follows:

$$\|S\| \doteq \sum_i^N Q_i(t) + \sum_i^N W_i(t).$$

Let $S^{(n)}$ denote a process S with an initial condition such that $\|S^{(n)}(0)\| = n$. In the analysis to follow, all variables associated with a process $S^{(n)}$ will be supplied with the upper index (n) .

The following theorem follows from the state-dependent Lyapunov-type stability criteria for countable Markov chains, obtained first by Malyshev and Menshikov [16].

THEOREM 4: *Suppose that there exist $\epsilon > 0$ and an integer $T > 0$ such that for any sequence of processes $\{S^{(n)}, n = 1, 2, \dots\}$, we have*

$$\limsup_{n \rightarrow \infty} E \left[\frac{1}{n} \|S^{(n)}(nT)\| \right] \leq 1 - \epsilon. \tag{10}$$

Then, S is stable.

It was shown by Rybko and Stolyar [19] that a stability condition of the type (10) naturally leads to a fluid-limit approach to the stability problem of queuing systems. This approach was further developed by Dai [8], Chen [7], Stolyar [20], and Dai and Meyn [9]. As the form of (10) suggests, the approach studies a fluid process $s(t)$ obtained as a limit of the sequence of scaled processes $(1/n)S^{(n)}(nt)$, $t \geq 0$. At the heart of the approach in its standard form is a proof that any $s(t)$ starting from any initial state with norm $\|s(0)\| = 1$ reaches zero in finite time T and stays there. It is sufficient, however, to show that for some $\epsilon > 0$, $\|s(T)\| \leq 1 - \epsilon$, which is what we are going to do in this article. (In many cases of interest, a still weaker condition is sufficient: It is enough to verify that any $s(t)$ is such that $\inf_{t \geq 0} \|s(t)\| < 1$, as shown in [20]. This is true in our case as well, as could be shown with a little extra work.) In our setting, we need to define what the scaling $(1/n)S^{(n)}(nt)$ means. In order for this scaling to make sense, we will need an alternative definition of the process.

To this end, let us define the following random functions associated with the process $S^{(n)}(t)$. Let $F_i^{(n)}(t)$ be the total number of type i customers that arrived by time $t \geq 0$, including the customers present at time 0, and let $\hat{F}_i^{(n)}(t)$ be the number

of type i customers that were served by time $t \geq 0$. Obviously, $\hat{F}_i^{(n)}(0) = 0$ for all i . As in [19] and [20], we “encode” the initial state of the system; in particular, we extend the definition of $F_i^{(n)}(t)$ to the negative interval $t \in [-n, 0)$ by assuming that the customers present in the system in its initial state $S^{(n)}(0)$ arrived in the past at some of the time instants $-(n-1), -(n-2), \dots, 0$, according to their delays in the state $S(0)$. By this convention, $F_i^{(n)}(-n) = 0$ for all i and n and $\sum_{i=1}^N F_i^{(n)}(0) = n$. Also, denote by $G_m^{(n)}(t)$ the total number of time slots before time t (i.e., among the slots $0, 1, \dots, t-1$), when the server was in state m , and by $\hat{G}_{mi}^{(n)}(t)$ the number of time slots before time t when the server state was m and the server was allocated to serve queue i . Let us also denote

$$U_i^{(n)}(t) \doteq t - W_i^{(n)}(t), \quad t \geq 0, i = 1, 2, \dots, N.$$

Then, the following relations obviously hold:

$$Q_i^{(n)}(t) \equiv F_i^{(n)}(t) - \hat{F}_i^{(n)}(t), \quad t \geq 0, i = 1, 2, \dots, N, \tag{11}$$

$$U_i^{(n)}(t) \leq t, \quad t \geq 0,$$

$$U_i^{(n)}(t) = \inf\{s \leq t : F_i^{(n)}(s) > \hat{F}_i^{(n)}(t)\}, \quad t \geq 0. \tag{12}$$

It is clear that the process $S^{(n)} = (S^{(n)}(t), t \geq 0)$ is a projection of the process $X^{(n)} = (F^{(n)}, \hat{F}^{(n)}, G^{(n)}, \hat{G}^{(n)}, Q^{(n)}, W^{(n)}, U^{(n)})$, where

$$F^{(n)} = (F_i^{(n)}(t), t \geq -n, i = 1, 2, \dots, N),$$

$$\hat{F}^{(n)} = (\hat{F}_i^{(n)}(t), t \geq 0, i = 1, 2, \dots, N),$$

$$G^{(n)} = (G_m^{(n)}(t), t \geq 0, m \in M),$$

$$\hat{G}^{(n)} = (\hat{G}_{mi}^{(n)}(t), t \geq 0, m \in M, i = 1, 2, \dots, N),$$

$$Q^{(n)} = (Q_i^{(n)}(t), t \geq 0, i = 1, 2, \dots, N),$$

$$U^{(n)} = (U_i^{(n)}(t), t \geq 0, i = 1, 2, \dots, N),$$

$$W^{(n)} = (W_i^{(n)}(t), t \geq 0, i = 1, 2, \dots, N).$$

In other words, a sample path of $X^{(n)}$ uniquely defines the sample path of $S^{(n)}$.

Let us also adopt the convention

$$Y(t) = Y(\lfloor t \rfloor) \quad \text{for } Y = S^{(n)}, F_i^{(n)}, \hat{F}_i^{(n)}, G_m^{(n)}, \hat{G}_{mi}^{(n)}, Q_i^{(n)}, U_i^{(n)}, W_i^{(n)},$$

with $t \geq -n$ for $Y = F_i^{(n)}$ and $t \geq 0$ for all other functions. This convention allows us to view the above functions as continuous-time processes defined for all $t \geq 0$ (or $t \geq -n$), but having constant values in each interval $[t, t + 1)$.

Now, consider the scaled process $x^{(n)} = (f^{(n)}, \hat{f}^{(n)}, g^{(n)}, \hat{g}^{(n)}, q^{(n)}, u^{(n)}, w^{(n)})$, where

$$\begin{aligned} f^{(n)} &= (f_i^{(n)}(t), t \geq -1, i = 1, 2, \dots, N), \\ \hat{f}^{(n)} &= (\hat{f}_i^{(n)}(t), t \geq 0, i = 1, 2, \dots, N), \\ g^{(n)} &= (g_m^{(n)}(t), t \geq 0, m \in M), \\ \hat{g}^{(n)} &= (\hat{g}_{mi}^{(n)}(t), t \geq 0, m \in M, i = 1, 2, \dots, N), \\ q^{(n)} &= (q_i^{(n)}(t), t \geq 0, i = 1, 2, \dots, N), \\ u^{(n)} &= (u_i^{(n)}(t), t \geq 0, i = 1, 2, \dots, N), \\ w^{(n)} &= (w_i^{(n)}(t), t \geq 0, i = 1, 2, \dots, N), \end{aligned}$$

and the scaling is defined as

$$z^{(n)}(t) = \frac{1}{n} Z^{(n)}(nt).$$

From (11), we get

$$q_i^{(n)}(t) \equiv f_i^{(n)}(t) - \hat{f}_i^{(n)}(t), \quad t \geq 0, i = 1, 2, \dots, N. \tag{13}$$

The following lemma establishes convergence to a fluid process and is a variant of Theorem 4.1 in [8]. The lemma is a list of basic convergence properties of the scaled sequences $\{x^{(n)}\}$ which we need for future reference. Although the lemma statement is quite long, the properties it describes are rather simple because they follow almost directly from the structure of the model and the strong law of large numbers for the input flow and server state processes.

LEMMA 1: *Consider our system under any scheduling rule such that, within each type i , the customers are served in the order of their arrival in the system. The following statements hold with probability 1. For any sequence of processes $\{X^{(n)}, n \in \mathcal{N}\}$, there exists a subsequence $\{X^{(k)}, k \in \mathcal{K} \subseteq \mathcal{N}\}$ such that as $k \rightarrow \infty$, the scaled subsequence $\{x^{(k)}, k \in \mathcal{K}\}$ has the following convergence properties for each $i \in \{1, \dots, N\}$ and $m \in M$:*

$$(f_i^{(k)}(t), t \geq -1) \Rightarrow (f_i(t), t \geq -1), \tag{14}$$

$$(f_i^{(k)}(t), t \geq 0) \rightarrow (f_i(t), t \geq 0) \text{ u.o.c.}, \tag{15}$$

$$(\hat{f}_i^{(k)}(t), t \geq 0) \rightarrow (\hat{f}_i(t), t \geq 0) \text{ u.o.c.}, \tag{16}$$

$$(q_i^{(k)}(t), t \geq 0) \rightarrow (q_i(t), t \geq 0) \text{ u.o.c.}, \tag{17}$$

$$(g_m^{(k)}(t), t \geq 0) \rightarrow (g_m(t), t \geq 0) \text{ u.o.c.}, \tag{18}$$

$$(\hat{g}_{mi}^{(k)}(t), t \geq 0) \rightarrow (\hat{g}_{mi}(t), t \geq 0) \text{ u.o.c.}, \tag{19}$$

$$(u_i^{(k)}(t), t \geq 0) \Rightarrow (u_i(t), t \geq 0), \tag{20}$$

$$(w_i^{(k)}(t), t \geq 0) \Rightarrow (w_i(t), t \geq 0), \tag{21}$$

where the functions f_i are RCLL nonnegative nondecreasing in $[-1, \infty)$, the functions f_i, \hat{f}_i, g_m and \hat{g}_{mi} are nonnegative nondecreasing Lipschitz-continuous in $[0, \infty)$, functions q_i are continuous in $[0, \infty)$, functions u_i are nondecreasing RCLL in $[0, \infty)$, functions w_i are nonnegative RCLL in $[0, \infty)$, and “ \Rightarrow ” signifies convergence at every continuity point of the corresponding limit function. The limiting set of functions

$$x = (f, \hat{f}, g, \hat{g}, q, u, w)$$

also satisfies the following properties for all $i \in \{1, \dots, N\}$ and $m \in M$:

$$\sum_{i=1}^N f_i(0) \leq 1, \tag{22}$$

$$f_i(t) - f_i(0) = \lambda_i t, \quad t \geq 0, \tag{23}$$

$$\hat{f}_i(0) = 0, \tag{24}$$

$$\hat{f}_i(t) \leq f_i(t), \quad t \geq 0, \tag{25}$$

$$g_m(t) = \pi_m t, \quad t \geq 0, \tag{26}$$

$$q_i(t) = f_i(t) - \hat{f}_i(t), \quad t \geq 0, \tag{27}$$

$$\hat{g}_{mi}(0) = 0, \tag{28}$$

$$\sum_{i=1}^N \hat{g}_{mi}(t) = g_m(t); \tag{29}$$

for any interval $[t_1, t_2] \subset [0, \infty)$,

$$\hat{f}_i(t_2) - \hat{f}_i(t_1) \leq \sum_{m \in M} \mu_i^m (\hat{g}_{mi}(t_2) - \hat{g}_{mi}(t_1)); \tag{30}$$

if $q_i(t) > 0$ for $t \in [t_1, t_2] \subset [0, \infty)$, then

$$\hat{f}_i(t_2) - \hat{f}_i(t_1) = \sum_{m \in M} \mu_i^m (\hat{g}_{mi}(t_2) - \hat{g}_{mi}(t_1)), \tag{31}$$

$$u_i(t) = t - w_i(t); \tag{32}$$

for any fixed $t_1 > 0$, the conditions $u_i(t_1) > 0$ and $\hat{f}_i(t_1) > f_i(0)$ are equivalent and if they hold, then in the interval $[t_1, \infty)$,

$$\lambda_i w_i(t) = q_i(t), \tag{33}$$

which, in particular, implies that w_i and u_i are Lipschitz-continuous in $[t_1, \infty)$.

Remark: The sets of functions x are (“fluid”) limits of the sequences of scaled paths $\{x^{(k)}\}$. As such, its components have the usual natural interpretations. For example, $f_i(t)$ and $\hat{f}_i(t)$ are the amounts of type i “fluid” that arrived into the system and are served by the system by the (scaled) time t , respectively, and $q_i(t) = f_i(t) - \hat{f}_i(t)$ is the amount of unserved type i at time t ; $g_m(t)$ is the total (scaled) time before time t when the server state was m ; $\hat{g}_{mi}(t)$ is the total (scaled) time before time t when the

server state was m and queue i was chosen for service. Property (23) then means that after time 0, the fluid of each type arrives at the constant rate λ_i ; this is generally not true for the interval $[-1, 0]$ because the fluid arrival processes $f_i(t)$ in this interval simply code sojourn times of the customers present at time 0, and these initial sojourn times can be distributed in a “bad” way. Inequality (30) simply means that the amount of fluid served in any interval cannot exceed the “potential” amount which could be served if the server would never incur idleness while serving queue i (the idleness is incurred when queue i is served in a slot at the rate μ_i^m , but there are less than μ_i^m customers in the queue); inequality (31) means that if the amount of unserved fluid $q_i(t)$ in some (scaled) interval is bounded away from zero, then the actual amount of fluid served in this interval is exactly equal to the potential amount of service. The property containing (33) is also simple, but is particularly important for our analysis: It says that if by some fixed (scaled) time t_1 , the amount of type i fluid served is greater than its initial amount (in particular, all of the “initial fluid” is “gone” by time t_1), then for all $t \geq t_1$, the strict linear relation $\lambda_i w_i(t) = q_i(t)$ exists between the amount of fluid $q_i(t)$ and the “head-of-the-line” fluid delay $w_i(t)$. It is this relation which will allow us to, roughly speaking, make a “transition” from the stability of M-LWWF to the stability of M-LWDF by showing that the fluid limit under M-LWDF is in a certain sense indistinguishable from that under M-LWWF, after the system “gets rid” of all the initial fluid.

PROOF OF LEMMA 1: It follows from the strong law of large numbers that, with probability 1 for every i ,

$$(f_i^{(n)}(t) - f_i^{(n)}(0), t \geq 0) \rightarrow (\lambda_i t, t \geq 0) \quad \text{u.o.c.}$$

To prove (15), (22), and (23), it suffices to choose a subsequence $\{x^{(k)}\}$ such that for every i , $\lim_{k \rightarrow \infty} f_i^{(k)}(0)$ exists, and denote the limit by $f_i(0)$. Since all $f_i^{(k)}$ and $u_i^{(k)}$ are nondecreasing, we can always choose a further subsequence such that (14) and (20) hold. Then, (21) follows from (20).

Properties (18) and (26) follow from the ergodicity of the server state process.

Also, for any fixed $0 \leq t_1 \leq t_2$, for every i, m , and any n , we have (using the notation $\mu^* \doteq \max_{m,j} \mu_j^m$)

$$\hat{f}_i^{(n)}(t_2) - \hat{f}_i^{(n)}(t_1) \leq \sum_{m \in M} \mu_i^m \left(\hat{g}_{mi}^{(n)}(t_2) - \hat{g}_{mi}^{(n)}(t_1) + \frac{1}{n} \right) \leq \mu^* \left(t_2 - t_1 + \frac{1}{n} \right).$$

From this inequality, we deduce the existence of a subsequence (of the subsequence already chosen) such that the convergences (16) and (19) take place and (30) holds.

Relations (24), (25), (28), (29), and (32) follow from the corresponding relations which trivially hold for the prelimit functions (for any index $n \in \mathcal{N}$). The convergence (17) and identity (27) trivially follow from identity (13).

Suppose that $q_i(t) > 0$ for $t \in [t_1, t_2] \subset [0, \infty)$. Let us fix $\delta \in (0, \min_{t \in [t_1, t_2]} q_i(t))$. The Lipschitz continuity of $q_i(\cdot)$, along with u.o.c. con-

vergence of $q_i^{(k)}$ to q_i , implies that (with probability 1) the sequence $\{X^{(k)}\}$ is such that for all sufficiently large k , the following inequalities hold:

$$\min_{t \in [t_1 k, t_2 k + 1]} Q_i^{(k)}(t) > \delta k > \max_m \mu_i^m.$$

The latter property implies that if the queue i was chosen for service anywhere in the interval $[[t_1 k], t_2 k + 1]$ when the server state was m , then exactly μ_i^m type i customers were served. So, we must have

$$|\hat{F}_i^{(k)}(kt_2) - \hat{F}_i^{(k)}(kt_1) - \sum_{m \in M} \mu_i^m (\hat{G}_{mi}^{(k)}(kt_2) - \hat{G}_{mi}^{(k)}(kt_1))| \leq 2 \max_m \mu_i^m.$$

Multiplying the last inequality by $1/k$ and taking the limit $k \rightarrow \infty$, we obtain (31).

Property (33) easily follows from the fact that in the interval $[0, \infty)$, the scaled input flow function $f_i^{(k)}(\cdot)$ converges u.o.c. to the strictly increasing linear function $f_i(0) + \lambda_i t$. We omit details. ■

Since some of the component functions included in x (viz. $f_i(\cdot)$, $\hat{f}_i(\cdot)$, $g_m(\cdot)$, $\hat{g}_{mi}(\cdot)$, and $q_i(\cdot)$) are Lipschitz in $[0, \infty)$, they are absolutely continuous. Therefore, at almost all points $t \in [0, \infty)$ (with respect to Lebesgue measure), the derivatives of all those functions exist. We will call such points *regular*.

In the rest of this article, when we consider a fixed limiting set of functions x , as defined in Lemma 1, we always assume that a sequence of prelimit paths $\{x^{(k)}\}$, which “defines it” (viz. the convergence properties of Lemma 1 hold), is fixed as well, along with the corresponding sequence of unscaled paths $\{X^{(k)}\}$.

5.2. Proof of Theorem 3 for the M-LWWF Discipline

The meaning of the following auxiliary lemma is that if relation (34) holds at some (scaled) time t , then by virtue of the M-LWWF scheduling rule, in some neighborhood of point t , flow i cannot be served.

LEMMA 2: Consider the system with the M-LWWF discipline. With probability 1, a limiting set of functions x , as defined in Lemma 1, satisfies the following additional property. If

$$\gamma_i \mu_i^m q_i(t) < \max_j \gamma_j \mu_j^m q_j(t) \tag{34}$$

for some regular point $t \geq 0$, for some i and m , then

$$\hat{g}_{mi}^t(t) = 0. \tag{35}$$

PROOF: Let us pick a j at which the maximum in inequality (34) is attained. In a similar manner to the proof of property (31) (in Lemma 1), we can fix a small positive $\delta_1 > 0$ such that, for all sufficiently large k , for the unscaled path $X^{(k)}$ we must have

$$\max_{\zeta \in [(t-\delta_1)k, (t+\delta_1)k]} \gamma_i \mu_i^m Q_i^{(k)}(\zeta) < \min_{\zeta \in [(t-\delta_1)k, (t+\delta_1)k]} \gamma_j \mu_j^m Q_j^{(k)}(\zeta).$$

(If $t = 0$, then the time interval should be $[0, \delta_1 k]$.) This means that in the interval $[(t - \delta_1)k + 1, (t + \delta_1)k - 1]$, queue i cannot be served in any time slot when the server is in state m because it would contradict the M-LWWF scheduling rule. Thus, for all sufficiently large k , we must have

$$\hat{g}_i^{(k)}\left(t + \frac{\delta_1}{2}\right) - \hat{g}_i^{(k)}\left(t - \frac{\delta_1}{2}\right) = 0,$$

which implies $\hat{g}_i(t + \delta_1/2) - \hat{g}_i(t - \delta_1/2) = 0$, and we are done. ■

Let us introduce a quadratic Lyapunov function

$$L(y) = \frac{1}{2} \sum_1^N \gamma_i y_i^2 \tag{36}$$

for a vector $y = (y_1, \dots, y_N)$.

The following lemma embodies the key idea behind MaxWeight-type scheduling rules: They try to maximize the rate of decrease of the Lyapunov function $L(q(t))$. So, roughly speaking, since there exists at least one scheduling rule (e.g., an SSS rule with ϕ such that $\lambda < v(\phi)$) under which $L(q(t))$ has a negative drift (when $L(q(t)) > 0$), the drift of $L(q(t))$ under M-LWWF has to be negative as well.

LEMMA 3: Consider a system with the M-LWWF discipline. For any $\delta_1 > 0$, there exists $\delta_2 > 0$ such that the following holds. With probability 1, a limiting set of functions x , as defined in Lemma 1, satisfies the following additional properties:

$L(q(t)), t \geq 0$, is an absolutely continuous function,

$$L(q(0)) \leq \frac{1}{2} \sum_1^N \gamma_i, \tag{37}$$

and at any regular point t ,

$$L(q(t)) \geq \delta_1 \text{ implies } \frac{d}{dt} L(q(t)) \leq -\delta_2. \tag{38}$$

PROOF: Let us pick a fixed stochastic matrix ϕ such that $\lambda_i < v_i(\phi)$ for all i . (The existence of such a matrix is condition (4).)

For any regular $t \geq 0$ such that $L(q(t)) > 0$, the derivative of $L(q(t))$ can be written

$$\frac{d}{dt} L(q(t)) = \sum_{i=1}^N \gamma_i q_i(t) (\lambda_i - \hat{f}_i'(t)) \tag{39}$$

$$= \sum_{i=1}^N \gamma_i q_i(t) (\lambda_i - v_i(\phi)) + \sum_{i=1}^N \gamma_i q_i(t) v_i(\phi) - \sum_{i=1}^N \gamma_i q_i(t) v_i(\hat{\phi}), \tag{40}$$

where

$$\hat{\phi}_{mi}(t) \doteq \frac{\hat{g}'_{mi}(t)}{\pi_m},$$

and we use the fact (following from property (31)) that

$$\hat{f}'_i(t) = \sum_m \mu_i^m \hat{g}'_{mi}(t) \quad \text{if } q_i(t) > 0.$$

Let us choose $\delta_3 > 0$ such that $L(y) \geq \delta_1$ implies $\max_i y_i \geq \delta_3$. Then, the first sum in (40) is bounded as follows:

$$\sum_{i=1}^N \gamma_i q_i(t) (\lambda_i - v_i(\phi)) \leq -\left(\min_i \gamma_i\right) \delta_3 \min_i (v_i(\phi) - \lambda_i) \doteq -\delta_2.$$

It remains to show that

$$K(\hat{\phi}(t), q(t)) \geq K(\phi, q(t)), \tag{41}$$

where $K(\xi, y)$ denotes the function of a stochastic $M \times N$ matrix ξ and a nonnegative N -dimensional vector y , defined as

$$K(\xi, y) \doteq \sum_{i=1}^N \gamma_i y_i v_i(\xi) = \sum_i \gamma_i y_i \sum_m \pi_m \xi_{mi} \mu_i^m = \sum_m \pi_m \sum_i \xi_{mi} \gamma_i \mu_i^m y_i.$$

It is easy to see that for any nonnegative vector y , a stochastic matrix ξ maximizes $K(\xi, y)$ if and only if the following condition holds for every i and m : If $\gamma_i \mu_i^m y_i < \max_j \gamma_j \mu_j^m y_j$, then

$$\xi_{mi} = 0. \tag{42}$$

However, property (35) shows that (42) is satisfied for $y = q(t)$ and $\xi = \hat{\phi}(t)$. This proves (41) and the lemma. ■

LEMMA 4: Consider a system with the M -LWFF discipline. For any $\delta > 0$, there exists $T > 0$ such that with probability 1, a limiting set of functions x , as defined in Lemma 1, satisfies the following additional property:

$$L(q(t)) \leq \delta, \quad t \geq T. \tag{43}$$

The proof follows from Lemma 3.

PROOF OF THEOREM 3 FOR M -LWFF: According to Lemmas 1–4, for any fixed $\epsilon_1 > 0$ we can always choose a large enough integer $T > 0$ such that for any sequence of random processes $\{X^{(n)}\}$, there exists a subsequence $\{X^{(k)}\}$ such that with probability 1, the convergence to a limiting set of functions x takes place and, moreover,

$$\sum_i q_i(T) \leq \epsilon_1. \tag{44}$$

If we recall that T is large, then it follows from (44) that

$$\hat{f}_i(T) = f_i(T) - q_i(T) > f_i(0), \quad \forall i, \tag{45}$$

implying (by (33)) that

$$w_i(T) = \frac{q_i(T)}{\lambda_i}, \quad \forall i. \tag{46}$$

This, in turn, implies (since ϵ_1 is small) that

$$\sum_i q_i(T) + \sum_i w_i(T) \leq \left(1 + \frac{1}{\left(\min_i \lambda_i\right)}\right) \epsilon_1 \doteq 1 - \epsilon < 1.$$

Therefore, with probability 1,

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \|S^{(n)}(nT)\| \leq 1 - \epsilon. \tag{47}$$

Since

$$\|S^{(n)}(nT)\| \leq n + \sum_i [F_i^{(n)}(nT) - F_i^{(n)}(0)] + N[n + nT],$$

our input process assumptions easily imply that the sequence $\{(1/n)\|S^{(n)}(nT)\|\}$ is uniformly integrable. This, along with (47), verifies condition (10). The proof is complete. ■

The following supplemental statement about the M-LWWF discipline will play an important role in the stability proof for the M-LWDF discipline.

Consider a *generalized* system with a given discipline H . The generalization is to assume that some time slots are unavailable for service of any queue. In each available for service time slot, the scheduling rule is H . In a generalized system, let $G_m^{(n)}(t)$ denote the number of *available for service* time slots (by time t) when the server is in state m . (Such a generalized system arises later, when we want to study the service dynamics of a *subset* of queues. To do that, we will view the time slots allocated to any other queue as unavailable for service of the subset of queues on which we focus.)

LEMMA 5: *Let positive constants K_0 and K_1 be fixed. Consider a sequence of fixed sample paths $\{X^{(k)}\}$ of the generalized system under M-LWWF such that as $k \rightarrow \infty$, all properties described in Lemmas 1 and 2 hold with the following modifications:*

Property (22) is replaced by

$$\sum_{i=1}^N f_i(0) \leq K_0 < \infty, \tag{48}$$

property (26) is replaced by

$$g_m(t) = \pi_m t - h_m(t), \quad t \geq 0, \tag{49}$$

where each function h_m is nondecreasing Lipschitz-continuous, $h_m(0) = 0$, and

$$\sum_m \lim_{t \rightarrow \infty} h_m(t) \leq K_1.$$

Then, the function $L(q(t))$ has the upper bound $C < \infty$, which depends only on K_0 and K_1 :

$$L(q(t)) \leq C, \quad t \geq 0. \tag{50}$$

PROOF: The idea of the proof is simple: the total “amount” of (scaled) time when service is unavailable to the queues is finite, bounded above by K_1 . During the “rest of the time,” when the service is available, the Lyapunov function $L(q(t))$ cannot increase, due to the “reasons” presented in the proof of Lemma 3. However, we need to apply this idea in a continuous time setting, which requires some care with the estimates. We now proceed with the details.

We will use the notation $\bar{L}(t) \doteq L(q(t))$. Let us choose $\delta > 0$ small enough so that the following holds for regular points t . If $g'_m(t) \geq \pi_m - \delta$ for each m , then $(d/dt)\bar{L}(t) < 0$. (The existence of such a δ is easily obtained using the argument and the estimates used in the proof of Lemma 3.) Note that $\sum_m h'_m(t) \leq \delta$ implies $g'_m(t) \geq \pi_m - \delta$ for each m .

Let us denote by Λ the Lebesgue measure and by \mathcal{L} the σ -algebra of Lebesgue-measurable subsets of $[0, \infty)$. Consider the subset

$$B \doteq \left\{ t \in [0, \infty) : t \text{ is regular, } \sum_m h'_m(t) > \delta \right\}.$$

It is easy to check that $B \in \mathcal{L}$ and

$$\Lambda(B) \leq \frac{K_1}{\delta}.$$

Define the measure ν on \mathcal{L} as follows:

$$\nu(A) \doteq \Lambda(A \cap B).$$

Notice that $\nu([0, \infty)) = \Lambda(B)$.

For future reference, we note that for some fixed positive c_1 and c_2 and all regular t ,

$$\bar{L}'(t) \leq c_1 + c_2 \bar{L}(t), \tag{51}$$

which follows from the estimate

$$\bar{L}'(t) = \sum_i \gamma_i q_i(t) q'_i(t) \leq (\max \lambda_i) \sum_i \gamma_i q_i(t) \leq (\max \lambda_i) \sum_i \gamma_i [1 + (q_i(t))^2].$$

We see that the derivative $\bar{L}'(t)$ is bounded above as in (51) at regular points $t \in B$, and it is negative at regular points $t \in [0, \infty) \setminus B$. We can write

$$\begin{aligned} \bar{L}(t) &\leq \bar{L}(0) + \int_{[0, t] \cap B} \bar{L}'(y) \Lambda(dy) = \bar{L}(0) + \int_0^t \bar{L}'(y) \nu(dy) \\ &\leq \bar{L}(0) + c_1 \nu([0, t]) + c_2 \int_0^t \bar{L}(y) \nu(dy) \\ &\leq \bar{L}(0) + c_1 \nu([0, \infty)) + c_2 \int_0^t \bar{L}(y) \nu(dy). \end{aligned}$$

Applying Gronwall’s inequality [11, p. 498], we obtain

$$\bar{L}(t) \leq [\bar{L}(0) + c_1 \nu([0, \infty))] \exp\{c_2 \nu([0, \infty))\}$$

and, finally,

$$\bar{L}(t) \leq \left[K_0 + \frac{c_1 K_1}{\delta} \right] \exp\left\{ \frac{c_2 K_1}{\delta} \right\}, \quad t \geq 0,$$

which proves the lemma. ■

5.3. Proof of Theorem 3 for the M-LWDF Discipline

The following lemma describes the key property of the M-LWDF discipline which is analogous to the M-LWWF property described in Lemma 2.

LEMMA 6: Consider a system with the M-LWDF discipline. With probability 1, a limiting set of functions x , as defined in Lemma 1, satisfies the following additional property. If in some interval $[t_1, t_2]$, $0 \leq t_1 < t_2 < \infty$, for some fixed m and fixed i and j we have

$$\sup_{t_1 \leq t \leq t_2} \gamma_i \mu_i^m w_i(t) < \inf_{t_1 \leq t \leq t_2} \gamma_j \mu_j^m w_j(t), \tag{52}$$

then

$$\hat{g}_{mi}(t_2) - \hat{g}_{mi}(t_1) = 0. \tag{53}$$

PROOF: The proof is analogous to the proof of Lemma 2. (The only additional difficulty is the fact that the functions $w_i(\cdot)$ may not be continuous.) Note that condition (52) implies that $\mu_j^m > 0$. We will consider only the nontrivial case when $\mu_i^m > 0$. (The case $\mu_i^m = 0$ is treated analogously to and simpler than this case.) Let us fix positive constants α and δ such that

$$\sup_{t_1 \leq t \leq t_2} \gamma_i \mu_i^m w_i(t) < \alpha - \delta < \alpha + \delta < \inf_{t_1 \leq t \leq t_2} \gamma_j \mu_j^m w_j(t). \tag{54}$$

Then, for all $t \in [t_1, t_2]$, we have

$$u_i(t) > t - \frac{\alpha - \delta}{\gamma_i \mu_i^m}$$

and

$$u_j(t) < t - \frac{\alpha + \delta}{\gamma_j \mu_j^m}.$$

Since for each i , $u_i(\cdot)$ and all $u_i^{(k)}(\cdot)$ are nondecreasing and we have the convergence $u_i^{(k)}(t) \rightarrow u_i(t)$ for every t where u_i is continuous, we see that for all sufficiently large k and for all $t \in [t_1, t_2]$,

$$u_i^{(k)}(t) > t - \frac{\alpha}{\gamma_i \mu_i^m}$$

and

$$u_j^{(k)}(t) < t - \frac{\alpha}{\gamma_j \mu_j^m}.$$

From the latter two inequalities, we see that

$$\gamma_i \mu_i^m w_i^{(k)}(t) < \alpha < \gamma_j \mu_j^m w_j^{(k)}(t), \quad t \in [t_1, t_2].$$

Just as in the proof of Lemma 2, we observe that the latter property implies that for all large k ,

$$\hat{g}_{mi}^{(k)}\left(t_2 - \frac{1}{k}\right) - \hat{g}_{mi}^{(k)}\left(t_1 + \frac{1}{k}\right) = 0$$

because the corresponding unscaled path $X^{(k)}$ is such that queue i may not be served in any time slot in the interval $[kt_1 + 1, kt_2 - 1]$ when the server is in state m . (Otherwise, we would get a violation of the M-LWDF scheduling rule.) Taking the limit $k \rightarrow \infty$ completes the proof. ■

The following lemma shows that under M-LWDF, all fluid limits x are such that after some fixed time T_N , all of the “initial fluid” is served (and, therefore, the linear relation $q_i(t) = \lambda_i w_i(t)$ holds) for all $t \geq T_N$ and all queues i .

LEMMA 7: *Consider a system with the M-LWDF discipline. There exists $T_N > 0$ such that with probability 1, a limiting set of functions x , as defined in Lemma 1, satisfies the following additional property:*

$$\hat{f}_i(T_N) > f_i(0), \quad i = 1, \dots, N.$$

To illustrate the intuition behind the formal proof, we present the following informal discussion. Suppose we consider the system with two flows $i = 1, 2$ and assume that by some fixed time $T_1 \geq 0$, we have $\hat{f}_1(T_1) > f_1(0)$ (i.e., all of the initial fluid of type 1 has been served). Consider a fixed sufficiently large time T_2 . Let us show why the assumption that the initial type 2 fluid is not served by time T_2 , namely

$$\hat{f}_2(T_2) \leq f_2(0), \tag{55}$$

leads to a contradiction. We observe that, first, the flow 2 delay $w_2(t) \geq t$ for all $t \in [T_1, T_2]$. Second, the amount of time unavailable to flow 1 in $[T_1, T_2]$ is bounded above: $f_2(0) \leq 1$. Then, according to Lemma 5, $q_1(t)$ —and therefore $w_1(t) = q_1(t)/\lambda_1$ —is bounded above in $[T_1, T_2]$ by a constant independent of T_2 . Therefore, during most of the interval $[T_1, T_2]$, the ratio of the waiting times $w_2(t)/w_1(t)$ is very large. This means that (during most of the interval $[T_1, T_2]$) as long as the server state m is such that flow 2 can be served at strictly positive rate μ_2^m , the M-LWDF rule must choose for service queue 2 over queue 1. This means that the amount of time when queue 2 is served is of the order of T_2 , which is large. However, then all of initial type 2 fluid, the amount of which is upper bounded by 1, must be served by time T_2 —a contradiction to assumption (55).

PROOF OF LEMMA 7: Let us fix an arbitrary $\epsilon_2 > 0$. We have

$$f_i(\epsilon_2) = f_i(0) + \lambda_i \epsilon_2 > f_i(0), \quad \forall i,$$

and

$$\sum_i q_i(\epsilon_2) \leq \sum_i f_i(\epsilon_2) \leq K_1 \doteq 1 + \left(\sum_i \lambda_i \right) \epsilon_2.$$

We will show the existence of T_N such that

$$\hat{f}_i(T_N) \geq f_i(\epsilon_2), \quad i = 1, \dots, N. \tag{56}$$

The proof of (56) is by induction.

Induction Base. There exists $T_1 > 0$ such that for at least one i ,

$$\hat{f}_i(T_1) \geq f_i(\epsilon_2).$$

Let us set $T_1 \doteq \epsilon_2 + K_1/\pi^*$, where π^* is the sum of the stationary probabilities π_m over server states m such that $\mu_j^m > 0$ for at least one j . Suppose the statement of the induction base, with this T_1 , does not hold. Then, for all sufficiently large k , we must have

$$\sum_i [\hat{f}_i^{(k)}(T_1) - \hat{f}_i^{(k)}(\epsilon_2)] \geq \pi^*(T_1 - \epsilon_2) + o(1) = K_1 + o(1),$$

where $o(1)$ is a term vanishing as $k \rightarrow \infty$. Taking the $k \rightarrow \infty$ limit, we obtain

$$\sum_i [\hat{f}_i(T_1) - \hat{f}_i(\epsilon_2)] \geq K_1,$$

which means (see the definition of K_1) that $\sum_i \hat{f}_i(T_1) \geq \sum_i f_i(\epsilon_2)$ and, therefore, $\hat{f}_i(T_1) \geq f_i(\epsilon_2)$ for at least one i . This contradiction proves the induction base.

Induction Step. Suppose that there exists $T_l > 0$, $1 \leq l < N$ such that for at least one subset $N_l \subset \{1, \dots, N\}$ of cardinality l , we have

$$\hat{f}_j(T_l) \geq f_j(\epsilon_2) \tag{57}$$

for all $j \in N_l$. Then, there exists $T_{l+1} \geq T_l$ such that (57) holds for all j within at least one subset N_{l+1} of cardinality $l + 1$.

We will prove the induction step for $l = 1$. (The generalization for arbitrary l is straightforward.) Thus, we need to prove the existence of $T_2 \geq T_1$ such that for at least two different flows i and r , (57) holds for $j = i, r$, with T_1 being the constant from the induction base statement.

Let us fix i for which

$$\hat{f}_i(t) \geq f_i(\epsilon_2), \quad t \geq T_1,$$

according to the induction base. Suppose

$$\hat{f}_j(T_1) < f_j(\epsilon_2) \quad \text{for all } j \neq i. \tag{58}$$

We observe that

$$\sum_{j \neq i} (f_j(\epsilon_2) - \hat{f}_j(T_1)) \leq K_1,$$

where K_1 is as defined earlier, and

$$q_i(T_1) \leq K_0 \doteq 1 + \lambda_i T_1.$$

Suppose that a constant $T_2 > T_1$ is fixed such that

$$\hat{f}_r(T_2) < f_r(\epsilon_2) \quad \text{for all } r \neq i. \tag{59}$$

(Below, we provide a choice of T_2 such that assumption (59) leads to a contradiction.)

Let us view each unscaled path $X^{(k)}$ after time kT_1 as a generalized system (described just above Lemma 5) with the single input flow of type i and with time slots allocated to any other flow being unavailable to flow i . (By convention, only the slots in which at least one customer of at least one flow $r \neq i$ was actually served are considered unavailable to flow i .) Then, for the scaled generalized system, starting at time T_1 , we have

$$h_m(t) = \sum_{r \neq i} [\hat{g}_{mr}(t) - \hat{g}_{mr}(T_1)] \leq K_1, \quad T_1 \leq t \leq T_2, m \in M. \tag{60}$$

Since x is such that the simple linear relation $\lambda_i w_i(t) = q_i(t)$ holds for flow i for all $t \geq T_1$, the generalized system with the M-LWDF discipline satisfies all of the prop-

erties of the generalized system with the M-LWWF discipline (including Lemma 5), with each γ_i replaced by γ_i/λ_i . Thus, from Lemma 5, we have

$$\frac{1}{2} \left(\frac{\gamma_i}{\lambda_i} \right) q_i^2(t) \leq C,$$

where the left-hand side is the “ $L(q(t))$ ” for the generalized system and $C \geq 0$ is the constant defined in Lemma 5, depending only on the constants K_0 and K_1 specified in this proof. From the last display we have the estimate

$$q_i(t) \leq C_1 \doteq \sqrt{\frac{2C\lambda_i}{\gamma_i}}, \quad t \in [T_1, T_2]. \tag{61}$$

Note that C_1 does not depend on the choice of T_2 .

From this point, we “switch back” to interpreting $X^{(k)}$ as a path of the original system. Let us denote by $M(i)$ the subset of elements $m \in M$ such that $\mu_j^m > 0$ for at least one flow $j \neq i$ and denote $\pi^*(i) \doteq \sum_{m \in M(i)} \pi_m$. Let us choose $T'_2 > T_1$ large enough so that for any pair of $j \neq i$ and $m \in M$ such that $\mu_j^m > 0$, we have

$$\frac{\gamma_i \mu_i^m C_1}{\lambda_i} < \gamma_j \mu_j^m (T'_2 - \epsilon_2). \tag{62}$$

Finally, let us choose $T_2 > T'_2$ large enough so that

$$\pi^*(i)(T_2 - T'_2) > K_1.$$

Our choice of T'_2 in (62) guarantees that for all sufficiently large k , the unscaled path $X^{(k)}$ must be (according to the M-LWDF rule) such that in the interval $[kT'_2, kT_2]$, in every time slot in which the state of the server belongs to the set $M(i)$, one of the flows $r \neq i$ is chosen for service. This observation implies that in the $k \rightarrow \infty$ limit for the corresponding scaled paths, we must have

$$\sum_{r \neq i} [\hat{g}_{mr}(T_2) - \hat{g}_{mr}(T_1)] \geq \pi^*(i)(T_2 - T'_2) > K_1.$$

This is a contradiction to (60), which shows that, for the T_2 chosen above, (59) cannot hold, and, therefore,

$$\hat{f}_r(T_2) \geq f_r(\epsilon_2) \tag{63}$$

for at least one $r \neq i$.

We have proved claim (63), assuming condition (58). However, the opposite of condition (58) means that, trivially, (63) holds for some $r \neq i$ and any $T_2 \geq T_1$. Thus, (63) holds for the chosen T_2 regardless of condition (58).

Our choice of T_2 depended on i . However, since there is only a finite number of possible values of i , we can choose T_2 so that (63) holds for some $r \neq i$ no matter what i is. The proof of the induction step is complete. ■

PROOF OF THEOREM 3 FOR M-LWDF: We proved the existence of $T_N > 0$ such that for any sequence of random processes $\{X^{(n)}\}$, there exists a subsequence $\{X^{(k)}\}$ such that with probability 1, the convergence to a limiting set of functions x takes place, and, moreover, x is such that the linear relation exists for all i :

$$\lambda_i w_i(t) = q_i(t), \quad t \geq T_N.$$

This fact, along with Lemma 6, means that with probability 1 in the interval $[T_N, \infty)$ the set x also satisfies all the properties described in Lemmas 2–4 if only in their formulations we replace γ_i by γ_i/λ_i , replace (37) by the condition

$$L(q(T_N)) \leq \frac{1}{2} \sum_1^N \gamma_i (1 + \lambda_i T_N)^2,$$

and move the time origin to T_N . Therefore, for any $\epsilon_1 > 0$, there exists $T \geq T_N$ such that with probability 1, x satisfies the condition

$$\sum_i q_i(T) \leq \epsilon_1.$$

The rest is exactly as in the proof of the theorem for M-LWWF. The only difference is that we obtain (46) directly from the property (33) and Lemma 7, not from (45).

6. CONCLUSIONS

We consider the variable channel scheduling queuing model which naturally arises in wireless communications. We show that a wide class of online scheduling rules, including the M-LWDF and M-LWWF rules (and their generalizations), are throughput optimal (i.e., they make all queues stable as long as the flow arrival rates are within the system stability region). One of the main contributions of this work is that we show that the throughput optimality of MaxWeight-type scheduling rules is preserved when flow waiting times are used as queue state variables in place of (or in conjunction with) the queue lengths.

We believe that the class of scheduling algorithms we study in this article can be efficiently used in applications to provide flexible control of quality of service to multiple data flows—in particular flows sharing a time-varying wireless link.

Acknowledgment

We would like to thank Sem Borst for numerous useful discussions.

References

1. Andrews, M., Kumaran, K., Ramanan, K., Stolyar, A., & Whiting, P. (1999). Data Rate Scheduling Algorithms and Capacity Estimates for the CDMA Forward Link. Bell Labs Technical Memorandum.
2. Andrews, M., Kumaran, K., Ramanan, K., Stolyar, A., Vijayakumar, R., & Whiting, P. (2000). CDMA Data QoS Scheduling on the Forward Link with Variable Channel Conditions. Bell Labs Technical Memorandum.

3. Armony, M. & Bambos, N. (1999). Queueing networks with interacting service resources. In *Proceedings of the 40th Annual Allerton Conference on Communication, Control, and Computing*. Monticello, pp. 42–51.
4. Armony, M. & Bambos, N. (2003). Queueing dynamics and maximal throughput scheduling in switched processing systems. *Queueing Systems: Theory and Applications* 44: 209–252.
5. Bambos, N. & Michalidis, G. (2002). On parallel queueing with random server connectivity and routing constraints. *Probability in the Engineering and Informational Sciences* 16: 185–203.
6. Bender, P., Black, P., Grob, M., Padovani, R., Sindhushayana, N., & Viterbi, A. (2000). CDMA/HDR: A bandwidth efficient high speed wireless data service for nomadic users. *IEEE Communications Magazine* 38: 70–77.
7. Chen, H. (1995). Fluid approximations and stability of multiclass queueing networks: Work-conserving disciplines. *Annals of Applied Probability* 5: 637–665.
8. Dai, J.G. (1995). On the positive Harris recurrence for open multiclass queueing networks: A unified approach via fluid limit models. *Annals of Applied Probability* 5: 49–77.
9. Dai, J.G. & Meyn, S.P. (1995). Stability and convergence of moments for open multiclass queueing networks via fluid limit models. *IEEE Transactions on Automatic Control* 40: 1889–1904.
10. Dai, J.G. & Prabhakar, B. (2000). The throughput of data switches with and without speedup. In *Proceedings of the INFOCOM'2000*.
11. Ethier, S.N. & Kurtz, T.G. (1986). *Markov process: Characterization and convergence*. New York: Wiley.
12. Feller, W. (1950). *An introduction to probability theory and its applications*. New York: Wiley.
13. Gill, P.E. & Murray, W. (1974). *Numerical methods for constrained optimization*. London: Academic Press.
14. Kahale, N. & Wright, P.E. (1997). Dynamic global packet routing in wireless networks. In *Proceedings of the INFOCOM'97*, pp. 1414–1421.
15. McKeown, N., Anantharam, V., & Walrand, J. (1996). Achieving 100% throughput in an input-queued switch. In *Proceedings of the INFOCOM'96*, pp. 296–302.
16. Malyshev, V.A. & Menshikov, M.V. (1979). Ergodicity, continuity, and analyticity of countable Markov chains. *Transactions of Moscow Mathematical Society* 39: 3–48.
17. Mekkitikul, A. & McKeown, N. (1996). A starvation free algorithm for achieving 100% throughput in an input-queued switch. In *Proceedings of the ICCCN'96*, pp. 226–231.
18. Moustafa, M.D. (1957). Input-output Markov processes. *Proc. Koninklijke Nederlandse Academie der Wetenschappen* 60: 112–118.
19. Rybko, A.N. & Stolyar, A.L. (1992). Ergodicity of stochastic processes describing the operation of open queueing networks. *Problems of Information Transmission* 28: 199–220.
20. Stolyar, A.L. (1995). On the stability of multiclass queueing networks: A relaxed sufficient condition via limiting fluid processes. *Markov Processes and Related Fields* 1(4): 491–512.
21. Stolyar, A.L. (2004). MaxWeight scheduling in a generalized switch: State space collapse and workload minimization in heavy traffic. *Annals of Probability*, to appear.
22. Stolyar, A.L. & Ramanan, K. (2001). Largest weighted delay first scheduling: Large deviations and optimality. *Annals of Applied Probability* 11: 1–48.
23. Tassiulas, L. & Ephremides, A. (1992). Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks. *IEEE Transactions on Automatic Control* 37: 1936–1948.
24. Tassiulas, L. & Ephremides, A. (1993). Dynamic server allocation to parallel queues with randomly varying connectivity. *IEEE Transactions on Information Theory* 39: 466–478.
25. Tse, D. (1999). Forward Link Multiuser Diversity Through Proportional Fair Scheduling. Presentation at Bell Labs.
26. Viswanath, P., Tse, D., & Laroia, R. (2002). Opportunistic beamforming using dumb antennas. *IEEE Transactions on Information Theory* 48(6): 1277–1294.
27. Viterbi, A.J. (1995). *CDMA. Principles of spread spectrum communication*. Reading, MA: Addison-Wesley.

APPENDIX

Details of the Proof of Sufficiency in Theorem 1

Lemma 1 holds for any scheduling rule, including the SSS rule associated with the matrix ϕ . For this rule, with probability 1, a limiting set of functions x is such that

$$\hat{g}_{mi}(t) = \phi_{mi} g_m(t) = \phi_{mi} \pi_m t, \quad t \geq 0.$$

From this and the argument analogous to that used in (39) and (40), we see that at any regular point $t \geq 0$, condition $q_i(t) > 0$ implies

$$q_i'(t) = \lambda_i - \hat{f}_i'(t) = \lambda_i - v_i(\phi) < 0.$$

Therefore, $q(t) \equiv 0$ for all $t \geq \max_i 1/(v_i(\phi) - \lambda_i)$. The rest of the proof is the same as in the proof of Theorem 3 for the M-LWWF rule, which follows Lemma 4 in Section 5.2. ■