

MAXWEIGHT SCHEDULING IN A GENERALIZED SWITCH: STATE SPACE COLLAPSE AND WORKLOAD MINIMIZATION IN HEAVY TRAFFIC

BY ALEXANDER L. STOLYAR

Bell Labs, Lucent Technologies

We consider a *generalized switch* model, which includes as special cases the model of multiuser data scheduling over a wireless medium, the input-queued cross-bar switch model and a discrete time version of a parallel server queueing system. Input flows $n = 1, \dots, N$ are served in discrete time by a switch. The switch *state* follows a finite state, discrete time Markov chain. In each state m , the switch chooses a *scheduling decision* k from a finite set $K(m)$, which has the associated service rate vector $(\mu_1^m(k), \dots, \mu_N^m(k))$.

We consider a heavy traffic regime, and assume a *Resource Pooling* (RP) condition. Associated with this condition is a notion of *workload* $X = \sum_n \zeta_n Q_n$, where $\zeta = (\zeta_1, \dots, \zeta_N)$ is some fixed nonzero vector with nonnegative components, and Q_1, \dots, Q_N are the queue lengths. We study the *MaxWeight* discipline which always chooses a decision k maximizing $\sum_n \gamma_n [Q_n]^\beta \mu_n^m(k)$, that is,

$$k \in \arg \max_i \sum_n \gamma_n [Q_n]^\beta \mu_n^m(i),$$

where $\beta > 0$, $\gamma_1 > 0, \dots, \gamma_N > 0$ are arbitrary parameters. We prove that under MaxWeight scheduling and the RP condition, in the heavy traffic limit, the queue length process has the following properties: (a) The vector $(\gamma_1 Q_1^\beta, \dots, \gamma_N Q_N^\beta)$ is always proportional to ζ (this is “State Space Collapse”), (b) the workload process converges to a Reflected Brownian Motion, (c) MaxWeight minimizes the workload among all disciplines. As a corollary of these properties, MaxWeight asymptotically minimizes the holding cost rate

$$\sum_n \gamma_n Q_n^{\beta+1}$$

at all times, and cumulative cost (with this rate) over finite intervals.

1. Introduction. In this paper, we study the following model in the heavy traffic regime. Multiple input flows, indexed by $n = 1, \dots, N$, each with its own queue, are served by a *generalized switch*. The system operates in discrete time $t = 0, 1, 2, \dots$; in particular, customers can only arrive and depart at integer time points. Switch *states* are random and follow a finite state, discrete time Markov

Received July 2001; revised July 2002, May 2003.

AMS 2000 subject classifications. 60K25, 90B15, 60J70.

Key words and phrases. Generalized switch, MaxWeight scheduling, heavy traffic limit, queueing networks, resource pooling, equivalent workload formulation, state space collapse, asymptotic optimality.

chain. Each state m of the switch has an associated finite set $K(m)$ of *scheduling decisions* (which can also be called *processing modes*). If m is the switch state at time t and a decision $k \in K(m)$ is chosen, then queue *service rates* (at time t) are given by a vector $(\mu_1^m(k), \dots, \mu_N^m(k))$ with nonnegative integer components. This means that $\mu_n^m(k)$ customers of flow n (or the entire queue n content at time t , whichever is less) are served and depart the system by time $t + 1$.

Our primary motivation for considering this model is the problem of scheduling multiple data flow transmissions over a wireless environment (see [1, 18, 27, 30, 31]). In terms of this problem, the N input flows represent data flows which need to be transmitted to N mobile users from a single or multiple radio transmitters. The time is slotted. For each time slot, a (scheduling) decision has to be made as to which transmitters should send data to which mobiles, and at which rates. In the simplest case when there is only one transmitter, only one user can be served in one slot, and transmission rates are fixed, there are exactly N scheduling decisions, namely, “which of the N users to serve.” In general, multiple users can be picked for service in one slot, and the data rates that can be assigned to the transmissions are user dependent (due to differences in radio channel quality) and, moreover, highly interdependent (due to transmitter power constraints and mutual radio signal interference). Our model is not concerned with the exact nature of the transmission rate constraints and dependencies; it simply assumes that there are a finite number of scheduling decisions that can be made, and each decision has an associated vector of transmission (service) rates. Another essential feature of a wireless environment is that it changes randomly with time, because of the random channel quality variations. The “switch state” in our model corresponds to a random state of the wireless environment, and different states typically have different sets of scheduling decisions. Our model assumption that the switch state follows a finite state Markov chain corresponds to the assumption that the radio environment changes with time in a random but “sufficiently stationary” way. We finally note that, in the heavy traffic regime we study in this paper, the time slot duration will be much shorter than typical data packet (customer) delays. This is in fact the case for many modern wireless technologies (see, e.g., [17]).

The generalized switch model also includes as a special case the much studied input-queued cross-bar switch, with L input and output ports (see, e.g., [23, 24]). The $N = L^2$ flows represent input-output port pairs (l_1, l_2) . A scheduling decision k is an input-output “matching,” that is, a subset of L pairs such that each value of l_1 and l_2 appears only once. When a matching k is chosen, flows from this matching are served at a certain (usually constant) rate.

Our model also includes a discrete time version of a parallel server system (see [3, 13, 16, 36]). In this model, N input flows are served by L servers. In each time slot, a server l can choose to serve one of the queues, and if it chooses queue n , it serves it at the (integer) rate $\mu_{nl} > 0$. A “switch” scheduling decision k is then a combination of the decisions of the individual servers, $k = (n_1, \dots, n_L)$, and the rate at which each queue n is served is the sum of its service rates over all servers which pick this queue. We should note that the continuous time parallel

server system studied in [3, 13, 16, 36] is *not* within the framework of a generalized switch. However, the approach of this paper has been subsequently applied in [22] to the analysis of a continuous time parallel server system. In particular, it is shown in [22] that most of the results of this paper (appropriately adjusted) do hold for continuous time systems.

The generalized switch can be viewed as a (discrete time) parallel server system, but with an additional feature that the servers are interdependent in an arbitrarily complex way. (Another additional feature is the randomness of the “environment.”) When such interdependence exists, the notion of “individual server” (or “individual service resource”) loses its significance, and one needs to consider the “server pool” as a single “aggregate service resource,” which we call a “switch.” Such an interpretation, although somewhat superficial, highlights the most important direction in which our model is more general than the parallel server model. It also justifies to some degree our use of the term “resource pooling” in the model where, formally, we do not identify individual resources; in applications, however, what we call a switch may in fact consist of separate resources with some (higher or lower) degree of interdependence.

The issue of *stability* of different versions of the generalized switch model is very well studied ([1, 2, 9, 18, 23, 24, 30, 31]). One of the principal stability results for this type of models is the fact that a simple MaxWeight scheduling discipline attains the maximum *stability region* of the system; results of this type are originally due to Tassiulas and Ephremides [30, 31]. Informally, the (maximum) stability region is the set of mean flow rate vectors $\lambda = (\lambda_1, \dots, \lambda_N)$ such that there exists a scheduling rule making the queue length process stable. The closure of the stability region is a convex polyhedron \bar{V} , which we call (*service rate region*).

In the setting of this paper, the *MaxWeight* discipline is defined as follows. In each time slot choose a scheduling decision

$$k \in \arg \max_{i \in K(m)} \sum_n \gamma_n [Q_n]^\beta \mu_n^m(i),$$

where m is the switch state, $Q = (Q_1, \dots, Q_N)$ is the queue length vector in the time slot, and $\beta > 0$, $\gamma_1 > 0, \dots, \gamma_N > 0$ are an arbitrary set of parameters. (In most of the previous work, MaxWeight with $\beta = 1$ and all $\gamma_n = 1$ was studied.) As demonstrated in the previous work on stability (and can be seen from the above definition), the underlying idea behind the MaxWeight algorithm is that, roughly speaking, it tries to minimize the average drift of the Lyapunov function $\sum_n \gamma_n [Q_n]^{\beta+1}$ at all times. This idea plays an important role in our analysis as well.

In the special case of a (discrete time) parallel server system, the MaxWeight discipline reduces to a particularly simple scheduling rule: each server l serves a queue

$$n \in \arg \max_i \gamma_i \mu_{il} Q_i^\beta.$$

We consider a *heavy traffic regime*, when the vector λ of mean input rates converges to some point v^* on the outer (“north-east”) boundary of the rate region \bar{V} . This roughly means that we study the asymptotic behavior as $r \rightarrow \infty$ of the *diffusion-scaled* queue length process $(1/r)Q(r^2t)$, $t \geq 0$, assuming $\lambda \rightarrow v^*$ fast enough as $r \rightarrow \infty$ so that $r(\lambda - v^*)$ approaches a fixed finite vector.

We assume that the limiting mean input rate vector v^* satisfies the *Resource Pooling* (RP) condition, which we define as uniqueness (up to scaling) of the vector $\zeta = (\zeta_1, \dots, \zeta_N)$, which is the outer normal vector to the rate region \bar{V} at point v^* . The vector ζ has nonnegative components, with at least one being strictly positive. We refer to the quantity $X = \sum_n \zeta_n Q_n$ as the system *workload*, to the vector ζ as the *workload aggregator*, and to its components ζ_n as *workload contributions* of the corresponding flows (or queues).

RP conditions have been introduced in previous work on networks with dynamic routing [20, 21] and parallel server systems [3, 13, 16, 36]. (This work is discussed later in this section.) Also, this notion is closely related to the *equivalent workload formulation* of the Brownian network control problems, developed relatively recently (see [5, 14, 15] and references therein). Although our model is different and we use different (“geometric”) language in the definition of the RP condition, our definition is consistent with the previous work and the general equivalent workload formulation. Very informally, the RP condition implies that, when we consider the process under diffusion scaling, with the appropriate control, the service system (switch in our case) can instantly “transition” a queue length vector Q to any vector $Q + \Delta Q$, such that ΔQ is orthogonal to ζ (and all components of $Q + \Delta Q$ are nonnegative).

Our main result, Theorem 1, can be informally described as follows. Consider the switch in the heavy traffic regime, and assume the RP condition. Then, with probability 1 (“pathwise”), under MaxWeight scheduling, in the heavy traffic (diffusion) limit, the (rescaled) queue length process has the following properties: (a) State Space Collapse—the vector $(\gamma_1 Q_1^\beta, \dots, \gamma_N Q_N^\beta)$ is always proportional to ζ ; (b) the (rescaled) workload process X converges to a one-dimensional Reflected Brownian Motion (RBM); (c) this RBM is a lower bound for the (rescaled) workload process limit under any discipline, that is, MaxWeight asymptotically minimizes (rescaled) workload among all disciplines.

As we will show, it follows from properties (c) and (a) above that MaxWeight asymptotically minimizes the value of $\sum_n \gamma_n Q_n^{\beta+1}$ (with rescaled Q_n) at all times. If $\gamma_n Q_n^{\beta+1}$ is interpreted as the rate at which queue n incurs holding cost, then this implies that MaxWeight asymptotically minimizes cumulative holding costs over finite intervals under diffusion scaling.

Our main result may seem somewhat surprising. Indeed, MaxWeight is an “on-line” rule that only “needs to know” the current queue lengths and the current set of available service rate vector choices. It does *not* need to know the mean input rates λ or stationary probabilities of switch states m , and does not need

any parameters to be “precalculated” before it is applied. And yet it has very strong asymptotic optimality properties—workload minimization and, in addition, holding cost minimization with $\gamma_n Q_n^{\beta+1}$ cost rates. (Again, MaxWeight does not need to know the workload aggregator ζ for these properties to hold.)

In addition, since the parameters β and γ_n can be chosen arbitrarily, the workload can be “distributed” among the queues with positive workload contributions ζ_n arbitrarily. For example, to keep the queue length vector proportional (in the limit) to a fixed vector $b = (b_1, \dots, b_N) \in \mathbb{R}_+^N$ (with the constraint that $b_n > 0$ if and only if $\zeta_n > 0$), it suffices to set $\beta = 1$ and $\gamma_n = \zeta_n/b_n$. This would require knowledge of the workload aggregator ζ , however. Assuming the RP condition holds, the State Space Collapse property of MaxWeight allows one to estimate ζ from queue length measurements. Namely, in applications, γ_n ’s can initially be set arbitrarily (for example all 1’s), then ζ can be estimated by the average observed value of the vector (Q_1, \dots, Q_N) , and then γ_n ’s can be reset to $\gamma_n = \zeta_n/b_n$. We believe that these properties of the MaxWeight discipline make it very attractive for applications.

The basic intuition behind such “nice” system behavior under MaxWeight is roughly as follows. We prove that, under the RP condition, sample paths of the *fluid* process (we call them *fluid sample paths*) corresponding to a critically loaded system (i.e., with input rates on the outer boundary of the rate region), are such that the vector $(\gamma_1 Q_1^\beta, \dots, \gamma_N Q_N^\beta)$ is attracted to the one-dimensional manifold of vectors proportional to ζ . This implies the State Space Collapse property (a) above, and therefore, as long as total queue length is nonzero, the MaxWeight rule “reduces to” the rule choosing

$$k \in \arg \max_{i \in K(m)} \sum_n \zeta_n \mu_n^m(i).$$

This rule is precisely the rule which maximizes the service rate of workload, as long as the total queue is nonzero. More concisely:

Under the RP condition, MaxWeight “induces” State Space Collapse, which in turn is precisely such that MaxWeight maximizes the service rate of the workload.

The general notion of State Space Collapse (SSC), meaning that the limiting process (Q in our case) “lives” on a manifold of lower dimension than the original one (in our case, dimension 1), goes back to the papers of Whitt [33] and Reiman [25, 26]. For *multiclass queueing networks*, a quite general theory of the heavy traffic SSC and convergence of the processes describing network behavior to Semimartingale Reflected Brownian Motions (SRBM) has been recently developed by Bramson [4] and Williams [35]. (See also [6].) Our generalized switch model is *not* within the framework of multiclass networks. The main difference is that each node of a multiclass network has an inherent workload conservation property: even in a pre-limit system, the unfinished work present in the node is served at the maximum possible rate (as long as there is nonzero amount of work). There is *no* such property for the workload of the generalized switch.

(The workload service rate can in principle be “wasted” even when queues are arbitrarily large.) The switch workload conservation property under MaxWeight holds only in the limit, and this is the key part that needs to be proved (and which in turn implies workload minimization among all disciplines). As explained above, given MaxWeight is the scheduling being employed, SSC along a very specific line “needs to occur” to ensure workload conservation in the limit, and MaxWeight does induce precisely such SSC. Despite the difference of the frameworks, the general approach of [4] and [35] can be applied here. More specifically, our technique of demonstrating SSC is similar to that of [4]. The technique involves breaking down an order $O(r^2)$ -long time interval into $O(r)$ -long intervals, and then analyzing behavior of *fluid-scaled* processes [under $(1/r)Q(rt)$ scaling] on those $O(r)$ -long intervals. Since, for large r , sample paths of fluid-scaled processes are close to the fluid sample paths (mentioned above), the attraction property of the latter is used to show that (asymptotically) $(\gamma_1 Q_1^\beta, \dots, \gamma_N Q_N^\beta)$ must stay proportional to ζ .

Several heavy traffic models, related to ours, have been considered in the literature. A continuous time parallel server model under the *Complete Resource Pooling* (CRP) condition was studied by Bell and Williams [3], Harrison [13], Harrison and Lopez [16] and Williams [36]. (The CRP condition requires that, in addition to the RP assumption, workload contributions ζ_n of all flows are strictly positive.) The optimality criterion in this work is the pathwise minimization of both workload and expected discounted cumulative linear holding costs (with cost rates $c_n Q_n$, $c_n > 0$). In [13, 16] *discrete review policies* have been proposed. Their asymptotic optimality was proved in [13] for a two-server model, and conjectured for the general model in [16], based on the derived optimal solution of associated Brownian control problem. *Continuous review threshold rules* for this model were proposed and proved optimal in [3, 36]. Both the discrete review and continuous review threshold rules require *a priori* knowledge of workload contributions ζ_n (in our notation), which in turn depend on the mean rates λ . We note that the MaxWeight rule does *not* solve the linear cost minimization problem for our model. Any solution to this problem requires that, in the limit, all workload is “kept” in the queues n with the smallest ratio c_n/ζ_n . However, the MaxWeight rule can be used to obtain an approximate solution, for example, by setting $\gamma_n = c_n$ for all n and choosing small $\beta > 0$. Another option is to set $\beta = 1$, and set all $\gamma_n = 1$ for all n except for one queue with the smallest c_n/ζ_n ; for this latter queue, γ_n is set to a small positive value. It is easy to see that with both options, almost all workload will be kept in the queues with the smallest c_n/ζ_n .

A queueing network with dynamic routing in the heavy traffic regime has been considered by Laws [21]. (See also [20] for a review of related models.) Although the model in [21] is different from ours, the system behavior is very similar: the paper shows heavy traffic workload minimization and SSC under the RP condition and under a dynamic routing algorithm which seeks to minimize the expected end-to-end delay of each customer. We note that the RP condition in [21] implies

uniqueness of the outer normal vector to the boundary of the rate region, and therefore our RP condition is consistent with that in [21]. The RP condition in [21] also allows zero workload contributions of some flows, and workload contributions do not need to be precomputed.

A model very similar to ours is analyzed by Gans and van Ryzin [11, 12]. Essentially, the model is a generalized switch with only one switch state; that is, there is no random “service environment.” The analysis is quite different, though. The heavy traffic regime is such that the mean rate vector λ (in our terminology) approaches the outer boundary of the rate region “from inside”; in other words, λ always remains within stability region. (Our heavy traffic regime, as well as that in [3, 13, 16, 20, 21, 36], allows λ to approach the rate region boundary from any direction.) Asymptotic optimality is defined in terms of the average total queue length minimization in the stationary regime, and is proved for a certain “batch” processing discipline. The discipline requires *a priori* knowledge of λ .

A multiuser variable channel scheduling model in heavy traffic (motivated by a scheduling problem in wireless systems) has been considered by Buche and Kushner [7]. This model has a constraint that “one user can be served at a time,” although possible generalizations are mentioned. The controls proposed in [7] require that most of the service resources be “preallocated” based on the input rates and channel statistics, and only a small portion of the resources is used for dynamic control. The asymptotic optimality under various criteria is proved.

The rest of the paper is organized as follows. In Section 2 we introduce basic notations, definitions, and conventions used in the paper. We introduce the formal model in Section 3. In Section 4 the conditions defining the system stability region are described, and in Section 5 the MaxWeight scheduling rule is defined. We define the rate region and introduce the RP condition in Section 6. Sections 7 and 8 contain the definition of the heavy traffic regime and scaling, and the statement of our main result, Theorem 1. In Section 9 we discuss main assumptions of Theorem 1, and also conjecture a very plausible result on the asymptotics of stationary distributions. The key technical intuition behind the main results is discussed in Section 10. In Section 11 sample paths of a fluid limit process are defined, and their key attraction property is proved. Finally, the main result, Theorem 1, is proved in Section 12. A discussion and concluding remarks are presented in Section 13.

2. Notation. We will use standard notations \mathbb{R} and \mathbb{R}_+ for the sets of real and real nonnegative numbers, respectively; and a not quite standard notation \mathbb{R}_{++} for the set of strictly positive real numbers. Corresponding N -times product spaces are denoted \mathbb{R}^N , \mathbb{R}_+^N and \mathbb{R}_{++}^N . The space \mathbb{R}^N is viewed as a standard vector-space, with elements $p \in \mathbb{R}^N$ being row-vectors $p = (p_1, \dots, p_N)$. Vector inequalities $p \leq q$ and $p < q$ are understood componentwise. The scalar product (dot-product)

of $p, q \in \mathbb{R}^N$ is

$$p \cdot q \doteq \sum_{n=1}^N p_n q_n;$$

and the norm of q is

$$\|q\| \doteq \sqrt{q \cdot q}.$$

For $\gamma, x \in \mathbb{R}^N$, we will denote

$$\gamma \times q \doteq (\gamma_1 q_1, \dots, \gamma_N q_N).$$

We will slightly abuse notation by applying power operation componentwise:

$$q^\beta \doteq (q_1^\beta, \dots, q_N^\beta) \quad \text{for } q \in \mathbb{R}_+^N, \quad \beta > 0,$$

and often write $q(t)^\beta$ to mean $[q(t)]^\beta$ where it does not cause confusion.

For $\zeta \in \mathbb{R}^N$, $\zeta \neq 0$, and $y \in \mathbb{R}^N$, we denote by

$$L(\zeta) \doteq \{x \in \mathbb{R}^N \mid \zeta \cdot x = 0\}$$

the $(N - 1)$ -dimensional linear subspace orthogonal to ζ , and by

$$L(\zeta, y) \doteq \{x \in \mathbb{R}^N \mid \zeta \cdot x = \zeta \cdot y\}$$

the hyperplane parallel to $L(\zeta)$ and containing y .

The angle between two nonzero vectors $p, q \in \mathbb{R}^N$ is defined in the usual way as

$$\arccos \frac{p \cdot q}{\|p\| \|q\|}.$$

We define the scaling operators Γ^r and $\tilde{\Gamma}^r$, $r > 0$, for a scalar function $h = (h(t), t \in A)$, $A \subseteq \mathbb{R}$, as follows:

$$(\Gamma^r h)(t) \doteq \frac{1}{r} h(rt), \quad t \in A/r \doteq \{\xi/r \mid \xi \in A\}$$

and

$$(\tilde{\Gamma}^r h)(t) \doteq \frac{1}{r} h(r^2 t), \quad t \in A/r^2.$$

For any set of functions the operators Γ^r and $\tilde{\Gamma}^r$ are applied componentwise.

For any scalar function $h = (h(t), t \in \mathbb{R}_+)$, we define the shift operator θ_d , $d \in \mathbb{R}_+$, in the standard way:

$$(\theta_d h)(t) = h(d + t), \quad t \in \mathbb{R}_+;$$

and for sets of functions this operator is applied componentwise. [The standard shift operator θ_d should not be confused with a different “special” shift operator $\Theta(d)$, defined in Section 11.1.] The *oscillation* of function h over a subset $A \subseteq \mathbb{R}_+$ is defined as

$$\text{Osc}(h; A) \doteq \sup_{\xi_1, \xi_2 \in A} |h(\xi_1) - h(\xi_2)|.$$

The $\arg \max_{v \in V} F(v)$ denotes the subset of elements of V maximizing the value of a scalar function $F(v)$. We sometimes write $\arg \max_v F(v)$, if the domain of v is clear from the context.

We denote by $D([0, \infty), \mathbb{R})$ the standard Skorohod space of right-continuous left-limit (RCLL) functions defined on $[0, \infty)$ and taking real values. (See, e.g., [10] for the definition of this space and associated topology and σ -algebra.)

The symbol \xrightarrow{w} denotes convergence in distribution of *random processes* (or other random elements), that is, weak convergence of their *distributions*. Typically, we consider convergence of processes in $D([0, \infty), \mathbb{R})$ or its N -times product space $D^N([0, \infty), \mathbb{R})$ equipped with product topology and σ -algebra.

The symbol $\xrightarrow{\text{u.o.c.}}$ (or the abbreviation u.o.c. after a convergence statement) means *uniform on compact sets* convergence of *elements* of $D([0, \infty), \mathbb{R})$ or its N -times product $D^N([0, \infty), \mathbb{R})$. For functions with a bounded domain $A \subset \mathbb{R}$, the u.o.c. convergence means uniform convergence.

We reserve symbol \Rightarrow for the weak convergence of *elements* of the space $D([0, \infty), \bar{\mathbb{R}})$, which is the space of RCLL functions taking values in the set $\bar{\mathbb{R}}$ of real numbers extended by including two “infinite numbers” $+\infty$ and $-\infty$ (with the natural topology on $\bar{\mathbb{R}}$). If $h, g \in D([0, \infty), \bar{\mathbb{R}})$, then $h \Rightarrow g$ means $h(t) \rightarrow g(t)$ for every $t > 0$ where g is continuous. (Convergence at $t = 0$ is not required.) We will not need any characterization of the topology on $D([0, \infty), \bar{\mathbb{R}})$, beyond the definition of convergence given above.

We will write simply 0 for the zero element of \mathbb{R}^N and for identical zero functions taking values in \mathbb{R} and \mathbb{R}^N .

We denote minimum and maximum of two real numbers ξ_1 and ξ_2 as $\xi_1 \wedge \xi_2$ and $\xi_1 \vee \xi_2$, respectively; and by $\lfloor \xi \rfloor$ and $\lceil \xi \rceil$ the integer part and the ceiling of a real number ξ , respectively.

3. The model. Consider the following queueing system. There is a finite set $\bar{N} = \{1, 2, \dots, N\}$ of input flows served by a *generalized switch*. Each flow n consists of discrete *customers*, which we sometimes call *type n customers*. Customers of each flow (type) waiting for service are queued in a separate queue of infinite capacity. (Customers are never lost.)

The system operates in discrete time $t = 0, 1, 2, \dots$. By convention, we will identify an (integer) time t with the unit time interval $[t, t + 1)$, which will sometimes be referred to as the *time slot t* ; and we will assume that all system variables we consider are constant within each time slot.

The switch has a finite set of *switch states* \bar{M} . In each time slot, the switch is in one of the states $m \in \bar{M}$; and the sequence of states $m(t)$, $t = 0, 1, 2, \dots$, forms an irreducible (finite state) Markov chain with stationary distribution $\{\pi_m, m \in \bar{M}\}$,

$$\pi_m > 0, \forall m \in \bar{M}, \quad \sum_{m \in \bar{M}} \pi_m = 1.$$

Each switch state $m \in \bar{M}$ has an associated finite set $K(m)$ of *scheduling decisions*, which can also be called *processing modes*. If the switch state at time t is m and a decision $k \in K(m)$ is chosen, then the integer number $\mu_n^m(k) \geq 0$ customers of type $n \in \bar{N}$ are served (or the entire queue n content at time t , whichever is less) and depart the system at time $t + 1$. The vector $\mu^m(k) \doteq (\mu_1^m(k), \dots, \mu_N^m(k))$ we will call the *service rate* vector, corresponding to state m and decision $k \in K(m)$. We make a very natural nondegeneracy assumption that

- (1) for any n $\mu_n^m(k) > 0$ for at least one pair of m and $k \in K(m)$.

New customer arrivals occur at times $t = 1, 2, \dots$. Denote by $A_n(t)$, $t \geq 1$, the number of type n customers arrived at time t , and assume by convention that these customers are immediately available for service at time t . We assume that each input process A_n is an ergodic (discrete time) Markov chain with countable state space. We also assume that all input processes and the switch state process $m = (m(t), t = 0, 1, 2, \dots)$ are mutually independent. Let us denote by λ_n , $n \in \bar{N}$, the mean arrival rate of flow n , that is, the mean number of type n customers arriving in one time slot when the Markov chain A_n is in the stationary regime. (The main results of this paper, pertaining to the heavy traffic asymptotic regime, will require additional assumptions on the input flows. We will introduce those assumptions later when we define the heavy traffic regime.)

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

$$S(t) = \{(U_{n,1}(t), \dots, U_{n,Q_n(t)}(t)), A_n(t+1), n \in \bar{N}; m(t)\},$$

$Q_n(t)$ is the type n queue length at time t (including new arrivals at time t), and $U_{nl}(t)$ is the current *delay* (i.e., *age* in the system) of the l th type n customer present in the system at time t . [Within each type, the customers are numbered in the order of their arrivals. If $Q_n(t) = 0$, the vector $(U_{n,1}(t), \dots, U_{n,Q_n(t)}(t))$ is replaced by some special symbol, say “ \emptyset .”] Since the values of $Q_n(t)$, $U_{n,j}$ and $A_n(t+1)$ are nonnegative integers, S is a (discrete time) process with countable state space (defined in the obvious way).

A mapping G which takes a system state $S(t)$ in a time slot into a fixed probability distribution $G(S(t))$ on the set of scheduling decisions $K(m)$ [with $m = m(t)$] will be called a *scheduling rule*, or a *queueing discipline*. With a fixed discipline G , the scheduling decision at time t is chosen randomly according to the distribution $G(S(t))$. If the scheduling decision $k \in K(m(t))$ is chosen at time t , then $D_n(t+1) = \min\{Q_n(t), \mu_n^{m(t)}(k)\}$ of type n customers are served and depart the system at time $t + 1$. According to our conventions, for each flow n ,

$$Q_n(t) = Q_n(t-1) - D_n(t) + A_n(t), \quad t = 1, 2, \dots$$

Our assumptions imply that with any scheduling rule, S is a discrete time, countable state 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 1 (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.)

4. Static Service Split rule and system stability region. Suppose, for each of its states $m \in \bar{M}$, a subprobability measure $\phi_m = (\phi_{mk}, k \in K(m))$ is fixed, which means that $\phi_{mk} \geq 0$ for all $k \in K(m)$, and $\sum_k \phi_{mk} \leq 1$.

Consider a *Static Service Split* (SSS) scheduling rule, parameterized by the set of measures $\phi \doteq (\phi_m, m \in \bar{M})$. When the switch is in state m , the SSS rule chooses one of the scheduling decisions $k \in K(m)$ randomly with probability ϕ_{mk} , and with probability $1 - \sum_k \phi_{mk}$ does not serve any of the queues. Then, clearly, the long-term service rate allocated to flow $n \in \bar{N}$ is equal to

$$v_n = \sum_{m \in \bar{M}} \pi_m \sum_{k \in K(m)} \phi_{mk} \mu_n^m(k).$$

Sometimes, we will call a measure ϕ itself an SSS rule. We will denote by $v(\phi)$ the function which maps an SSS rule ϕ into the corresponding vector $v = (v_1, \dots, v_N)$, as defined above.

The following simple observation is quite standard (and essentially trivial).

PROPOSITION 1. *Let $\lambda \in \mathbb{R}_+^N$ be the vector of mean rates. Then, for the existence of a scheduling rule G under which the system is stable, condition*

$$(2) \quad \lambda \leq v(\phi) \quad \text{for at least one SSS rule } \phi$$

is necessary, and condition

$$(3) \quad \lambda < v(\phi) \quad \text{for at least one SSS rule } \phi$$

is sufficient.

The proof of Proposition 1 is very intuitive, and is outlined as follows. (See, e.g., [1] for more details.) If (3) holds, then the SSS rule ϕ allocates to each flow the average service rate $v_n(\phi) > \lambda_n$, which easily implies stability. If the system is stable under some rule G , then we consider the process in a stationary regime, and denote by ϕ_{mk} the (well-defined) average fraction of time slots when decision k is chosen, among the slots when the switch state is m . These values of ϕ_{mk} form the set ϕ for which (2) must hold; otherwise, one of the queue lengths would run away to infinity with probability 1.

Proposition 1 motivates the following definition. The set

$$V^0 = \{\lambda \in \mathbb{R}_+^N \mid (3) \text{ holds}\}$$

we will call the system *maximum stability region*, or just *stability region*.

5. The MaxWeight discipline. It has been shown relatively recently that there are scheduling rules which (unlike SSS) do not use *a priori* information about mean rates λ_n and the stationary distribution π of the switch state, and yet ensure system stability as long as condition (3) is satisfied. In particular, the MaxWeight discipline, which we define shortly, has this property. (There are numerous results on MaxWeight stability for different special cases of our model. The first results of this type were obtained probably in [30, 31]. For the model we consider in this paper, the MaxWeight stability was proved in [1].)

Let us call the value

$$W_n(t) \equiv U_{n1}(t)$$

[with $W_n(t) = 0$ if $Q_n(t) = 0$ by convention] the *delay of flow n* at time t .

MaxWeight discipline. Let a set of positive constants β and γ_n , $n \in \bar{N}$, be fixed. When the switch is in state $m \in \bar{M}$ (in a time slot t), a scheduling decision k is chosen from the following subset:

$$(4) \quad k \in \arg \max_{i \in K(m)} \sum_{n \in \bar{N}} \gamma_n [V_n(t)]^\beta \mu_n^m(i),$$

where for each n , $V_n(t) = c_n^Q Q_n(t) + c_n^W W_n(t)$, with fixed constants $c_n^Q \geq 0$ and $c_n^W \geq 0$, $c_n^Q + c_n^W > 0$. (The “ties” are broken arbitrarily, e.g., in favor of the largest index n .)

PROPOSITION 2 ([1]). *Let an arbitrary set of positive constants β and γ_n , $n \in \bar{N}$, be fixed. Then MaxWeight scheduling rule has the maximum stability region; namely, it makes the system stable as long as condition (3) holds.*

The proof of Proposition 2 is a straightforward extension of the proof of Theorem 3 in [1]. In fact, our analysis of the *fluid sample paths* under MaxWeight in Section 11.2 is (in a certain respect) a “superset” of such analysis in [1] which leads to establishing the MaxWeight stability. (See the remark following the statement of Lemma 5 for more detailed comments.)

In the rest of this paper we consider the MaxWeight rule with $V_n(t) = Q_n(t)$, that is, the rule choosing

$$(5) \quad k \in \arg \max_{i \in K(m)} \sum_{n \in \bar{N}} \gamma_n [Q_n(t)]^\beta \mu_n^m(i),$$

although all results (appropriately adjusted) hold for the more general MaxWeight rule defined above, with V_n being a linear combination of queue length Q_n and delay W_n .

6. RP condition. Let us denote by V the set of all service rate vectors $v(\phi)$ corresponding to all possible SSS rules ϕ . In our case, V is a convex polyhedron in \mathbb{R}_+^N (as a linear image of the convex polyhedron of possible values of ϕ). The polyhedron V may turn out to be degenerate (i.e., have dimension less than N), but due to the nondegeneracy assumption (1), it contains at least one vector v with all positive components, $v > 0$.

Let us also consider the set

$$\bar{V} \doteq \{u \in \mathbb{R}_+^N \mid u \leq v \text{ for at least one } v \in V\},$$

which we call system *service rate region*, or just *rate region*. Rate region \bar{V} is also a convex polyhedron, and it is nondegenerate (has dimension N), since V contains at least one vector with all positive components.

Note that rate region \bar{V} is nothing else but the closure of the system stability region V^0 , defined earlier. This of course implies that \bar{V} and V^0 have the same boundary.

Let \bar{V}^* denote the outer (“north-east”) boundary of \bar{V} :

$$\bar{V}^* \doteq \{u \in \bar{V} \mid \text{there is no } v \in \bar{V} \text{ such that } u < v\}.$$

Since \bar{V} is a nondegenerate polyhedron, \bar{V}^* consists of a finite number of $(N - 1)$ -dimensional faces.

DEFINITION. We will say that the RP condition holds for a vector v^* if $v^* \in \bar{V}^*$ and the outer normal vector to \bar{V} at point v^* is unique (up to a scaling). Equivalently, the RP condition holds if v^* lies in the (relative) interior of one of the faces of \bar{V}^* .

Note that the RP condition for v^* implies that $v^* \in \mathbb{R}_{++}^N$, and the corresponding outer normal vector ζ is such that $\zeta \in \mathbb{R}_+^N$. By convention, we assume that if the RP condition holds, then ζ is the vector defined uniquely by the additional condition

$$(6) \quad \|\zeta\| = 1.$$

Note also that the subset of $v \in \bar{V}$ maximizing $\zeta \cdot v$,

$$\bar{V}(\zeta) \doteq \arg \max_{v \in \bar{V}} \zeta \cdot v,$$

is nothing else but the face of \bar{V}^* containing v^* :

$$(7) \quad v^* \in \bar{V}(\zeta).$$

DEFINITION. If vector v^* satisfies the RP condition, and in addition all components of the corresponding normal vector ζ are strictly positive, $\zeta \in \mathbb{R}_{++}^N$, we say that v^* satisfies the CRP condition.

We defer the discussion of the RP and CRP conditions until after the formulation of the main results in the next section.

Our purely geometric definition of the RP condition (as uniqueness of the outer normal vector at point v^*) is consistent with the RP notion introduced in the previous work on networks with dynamic routing [20, 21] and on parallel server systems [3, 13, 16, 36], and with the canonical workload representation for Brownian network control problems [14]. In the work cited above, the requirement that v^* is on the rate region outer boundary, the RP condition and the associated normal ζ are defined in terms of a certain linear program and its dual, with the geometric interpretation (which we use as a definition) being implicit. In our case, a linear programming characterization of the RP condition is also possible, but since it is not required for the MaxWeight definition and is not used in the analysis, we present it in the Appendix.

7. Heavy traffic assumptions. Consider a sequence of systems, indexed by $r \in \mathcal{R} = \{r_1, r_2, \dots\}$, where $r_i > 0$ for all i and $r_i \uparrow \infty$ as $i \rightarrow \infty$. (Hereafter in this paper “ $r \rightarrow \infty$ ” means that r goes to infinity by taking values from the sequence \mathcal{R} , or some subsequence of \mathcal{R} ; the choice of a subsequence will be clear in each case from the context.) From this point on in the paper, the quantities pertaining to the r th system will be supplied with a superscript r .

Assume that, as $r \rightarrow \infty$, the vector of mean rates $\lambda^r = (\lambda_1^r, \dots, \lambda_N^r)$ converges to some fixed vector $\lambda \in \mathbb{R}_{++}^N$ lying on the outer boundary \bar{V}^* of the system rate region \bar{V} :

$$(8) \quad \lambda^r \rightarrow \lambda \in \bar{V}^*,$$

and, moreover,

$$(9) \quad \text{vector } \lambda \text{ satisfies the RP condition.}$$

Let ζ be the (unique) normal vector ζ associated with the RP condition for point λ . Recall that $\zeta \in \mathbb{R}_+^N$ and $\zeta \neq 0$.

Assume, in addition, that the convergence in (8) is such that

$$(10) \quad r(\zeta \cdot \lambda^r - \zeta \cdot \lambda) \rightarrow a,$$

where $a \in \mathbb{R}$.

The quantity

$$X^r(t) \doteq \sum_{n=1}^N \zeta_n Q_n^r(t) = \zeta \cdot Q^r(t),$$

where $Q^r(t) = (Q_1^r(t), \dots, Q_N^r(t))$ is the queue length vector at time t , will be referred to as the *workload* of the switch. The vector ζ will be called the *workload aggregator*, and its components ζ_n will be called *workload contributions* of the corresponding flows (or queues).

We make the following additional assumptions on the input flows. For each $n \in N$ and each r ,

$$(11) \quad A_n^r(t), \quad t = 1, 2, \dots, \text{ are i.i.d.,}$$

which in particular means

$$(12) \quad E[A_n^r(1)] = \lambda_n^r,$$

and also

$$(13) \quad \text{Var}[A_n^r(1)] \rightarrow \sigma_n^2 \geq 0, \quad r \rightarrow \infty$$

and

$$(14) \quad E[A_n^r(1)]^2 I\{A_n^r(1) > z\} \leq \eta(z),$$

where $\eta(\cdot)$ is a fixed function, $\eta(z) \rightarrow 0$ as $z \rightarrow \infty$, and $I\{\cdot\}$ is the indicator function. Assumption (14) is a Lindeberg type condition required to apply a functional central limit theorem (FCLT). This condition will also allow us to use Bramson's weak law estimates ([4], Proposition 4.2) in Section 12.

We assume that the underlying Markov chain for the switch state process

$$(15) \quad m^r(\cdot) \text{ does not change with } r,$$

in the sense that it has the same probability law as the Markov chain $m(\cdot)$ defined earlier.

8. Main results. Let us apply the *diffusion scaling* to the processes Q^r and X^r to define the following scaled processes:

$$\tilde{q}^r(t) \doteq r^{-1} Q^r(r^2 t), \quad t \geq 0, \quad \tilde{x}^r(t) \doteq r^{-1} X^r(r^2 t), \quad t \geq 0.$$

Let v denote the vector such that $\gamma \times v^\beta$ is proportional to ζ and $\zeta \cdot v = 1$:

$$v \doteq \left[\sum_n \gamma_n^{-1/\beta} \zeta_n^{1+1/\beta} \right]^{-1} (\gamma_1^{-1/\beta} \zeta_1^{1/\beta}, \dots, \gamma_N^{-1/\beta} \zeta_N^{1/\beta}).$$

Any vector cv , $c \geq 0$, proportional to v (including the zero vector) we will call an *invariant point*, and the set of all invariant points will sometimes be called the *invariant manifold*.

We assume that initial states of the scaled processes converge to an invariant point:

$$(16) \quad \tilde{q}^r(0) \rightarrow \tilde{q}^\circ(0) = \tilde{x}^\circ(0)v,$$

where $\tilde{x}^\circ(0) \geq 0$ is a fixed constant. Convergence (16) of course implies $\tilde{x}^r(0) \rightarrow \tilde{x}^\circ(0)$.

Let us define the Brownian motion

$$(17) \quad \tilde{w} = (\tilde{x}^\circ(0) + at + \sigma B(t), t \geq 0),$$

where B is a standard (zero drift, unit variance) Brownian motion, a is the parameter in (10), and

$$\sigma^2 \doteq \sigma_s^2 + \sum_n \zeta_n^2 \sigma_n^2,$$

where the parameter σ_s^2 depends on ζ and the Markov chain $m(\cdot)$, and is defined later in (32).

For the Brownian motion process \tilde{w} defined in (17), let us consider the corresponding RBM $\tilde{x}^\circ = (\tilde{x}^\circ(t), t \geq 0)$:

$$(18) \quad \tilde{x}^\circ(t) = \tilde{w}(t) + \tilde{y}^\circ(t),$$

where

$$(19) \quad \tilde{y}^\circ(t) \doteq - \left[0 \wedge \inf_{0 \leq u \leq t} \tilde{w}(u) \right].$$

[This implies that $\tilde{y}^\circ(0) = 0$, and therefore $\tilde{w}(0) = \tilde{x}^\circ(0)$.]

THEOREM 1. *Consider the sequence of systems indexed by $r \in \mathcal{R}$ as described above, and assume that conditions (8)–(16) hold.*

(i) *Suppose that the scheduling rule in the system is MaxWeight. Then, as $r \rightarrow \infty$,*

$$(20) \quad \tilde{x}^r \xrightarrow{w} \tilde{x}^\circ,$$

and, moreover, the following SSC holds:

$$(21) \quad \tilde{q}^r \xrightarrow{w} \tilde{q}^\circ \doteq \tilde{x}^\circ \nu.$$

(ii) *The MaxWeight rule is asymptotically optimal in that it minimizes the workload process. More precisely, the workload process \tilde{x}_G^r corresponding to an arbitrary scheduling discipline G is such that, for any time $t \geq 0$ and any $u \geq 0$,*

$$(22) \quad \liminf_{r \rightarrow \infty} P\{\tilde{x}_G^r(t) > u\} \geq P\{\tilde{x}^\circ(t) > u\}.$$

It will be clear from our proofs that both statements of Theorem 1 in fact hold *pathwise*. That is, the limiting process \tilde{q}° and the sequences of processes under MaxWeight and under arbitrary discipline G (\tilde{q}^r and \tilde{q}_G^r , resp.) can be constructed on a common probability space so that, with probability 1, the following properties hold: *both (20) and (21) hold uniformly on compact sets and*

$$(23) \quad \liminf_{r \rightarrow \infty} \tilde{x}_G^r(t) \geq \tilde{x}^\circ(t) \quad \forall t \geq 0.$$

Moreover, as we show in the proof, $\tilde{q}^\circ(t)$ is the unique vector minimizing the function $\sum_n \gamma_n p_n^{\beta+1}$ among all vectors $p \in \mathbb{R}_+^N$ with workload $\zeta \cdot p = \tilde{x}^\circ(t)$.

This easily implies that, with probability 1, *MaxWeight* minimizes cumulative “power $\beta + 1$ ” costs over any fixed interval $[t_1, t_2]$, that is,

$$(24) \quad \liminf_{r \rightarrow \infty} \int_{t_1}^{t_2} \sum_n \gamma_n [\tilde{q}_{n,G}^r(t)]^{\beta+1} dt \geq \lim_{r \rightarrow \infty} \int_{t_1}^{t_2} \sum_n \gamma_n [\tilde{q}_n^r(t)]^{\beta+1} dt \\ = \int_{t_1}^{t_2} \sum_n \gamma_n [\tilde{q}_n^\circ(t)]^{\beta+1} dt.$$

We see that the workload minimization property (22) [or (23)] is the primary and (as discussed in Section 10) nontrivial “ingredient” implying *MaxWeight* optimality (24) in the sense of cumulative costs. Also, it directly implies that *MaxWeight* produces a “Pareto optimal” vector of queue lengths at any time t . For these reasons, the *MaxWeight* workload minimization property in itself we sometimes call optimality.

We will further discuss Theorem 1, in particular the RP condition and initial condition (16), in the next section.

9. Discussion of main results.

9.1. *A conjecture regarding asymptotics of stationary distributions.* A natural question is that of the asymptotics of stationary distributions of the processes \tilde{q}^r . We do not pursue this question in this paper, but the following short discussion, including Proposition 3 and (very plausible) Conjecture 1, outlines a direction in which this can be done.

PROPOSITION 3. *Consider the sequence of systems indexed by $r \in \mathcal{R}$ as described above, and assume that conditions (8)–(12) and (15) hold, and $a < 0$. Then, under *MaxWeight*, for all sufficiently large r , the system is stable.*

PROOF. Since λ lies in the interior of the face $\bar{V}(\zeta)$ and ζ is an outer normal to that face, conditions $a < 0$, (8) and (10) imply that $\lambda^r \in V^0$ for all large r . Therefore, by Proposition 2, the system is stable for large r . \square

It is well known that, if the drift $a < 0$, then the limiting RBM \tilde{x}° has exponential stationary distribution:

$$P\{\tilde{x}^\circ(\infty) > \xi\} = \exp\{(2a/\sigma^2)\xi\}, \quad \xi \geq 0,$$

where $\tilde{x}^\circ(\infty)$ denotes a random variable distributed as $\tilde{x}^\circ(t)$ is stationary regime. Then it is natural to conjecture that (in the case $a < 0$) the limit of a sequence of stationary distributions is equal to the stationary distribution of the limit.

CONJECTURE 1. Consider the sequence of systems indexed by $r \in \mathcal{R}$ as described above, and assume that conditions (8)–(15) hold, and $a < 0$. Then, under MaxWeight, as $r \rightarrow \infty$,

$$\tilde{q}^r(\infty) \xrightarrow{w} \tilde{x}^\circ(\infty)v,$$

where $\tilde{q}^r(\infty)$ is a random vector with distribution equal to a stationary distribution of \tilde{q}^r .

Conjecture 1 along with Theorem 1 implies that the stationary distribution of \tilde{x}° is a stochastic lower bound of any weak limit of a sequence of stationary distributions of \tilde{x}_G^r under any scheduling rule G .

9.2. *RP condition.* Our main result on the optimality of the MaxWeight rule assumes the RP condition. However, the form of the MaxWeight rule obviously does *not* involve the RP condition in any way, and in particular it does *not* require *a priori* knowledge of the workload aggregator ζ . In addition, the definition of the RP condition shows that almost all vectors λ within the outer (north-east) boundary \bar{V}^* (with respect to the natural Lebesgue measure on \bar{V}^*) *do* satisfy the RP condition. (Equivalently, a ray starting from the origin and having a random direction p distributed uniformly on $\{\|p\| = 1\} \cap \mathbb{R}_+^N$ almost surely hits \bar{V}^* at a point λ satisfying the RP condition.) This explains why the RP condition is common in applications (although it is not the only situation of interest, of course), and highlights the importance of relaxing the CRP to the RP condition. (The subset of points of \bar{V}^* , where the CRP condition is not satisfied, typically has nonzero Lebesgue measure.) Thus, one of the important points of this paper is that, in many applications, one can use the MaxWeight rule without *a priori* verifying the RP condition and/or computing the workload aggregator ζ (or doing any other “precalculation”), and yet achieve optimality properties, as described in the previous section.

The following terminology and notation will be useful later. A flow (queue) n with positive workload contribution $\zeta_n > 0$ we will call a *critical* flow (queue); otherwise (if $\zeta_n = 0$), the flow (queue) will be called *noncritical*. By $\bar{N}^{(C)}$ and $\bar{N}^{(NC)} = \bar{N} \setminus \bar{N}^{(C)}$ we will denote the subsets of critical and noncritical flows, respectively.

This terminology is natural (and consistent with that in [14, 21]). Indeed, a small increase of the (limiting) input rates λ_n of noncritical flows will not move the vector λ out of the rate region boundary, and, moreover, this will neither violate the RP condition nor change the workload contributions of different flows: this only moves λ within the interior of the same face of \bar{V}^* . Thus, for any noncritical flow n , the switch has “spare capacity” to serve this flow without sacrificing service rates of other flows. And, as our main result above shows, the MaxWeight scheduling rule is able to exploit this circumstance and “automatically identify and

isolate” the critical subset of flows, so that neither the (limiting) behavior of the process nor the workload minimization property is affected by a small increase of noncritical flow rates. For a rule as parsimonious as MaxWeight, which does not know in advance which flows are critical, such an “automatic critical subsystem isolation” property is nontrivial.

On the other hand, an increase of the (limiting) input rate of any critical flow will move vector λ out of the set \bar{V} . In order for λ to stay within the same face of \bar{V}^* , the input rates of critical flows can only be “traded off” in a way such that $\zeta \cdot \lambda$ (i.e., the limiting mean workload arrival rate) remains constant.

It is also easy to observe that the RP condition implies the CRP condition for the reduced “critical subsystem,” obtained from the original one by removing noncritical flows and excluding noncritical components from the service rate vectors $\mu^m(k)$.

In the rest of the paper, for a given vector $p \in \mathbb{R}^N$ (typically it will be a scaled or unscaled vector of queue lengths q), $p^{(C)}$ and $p^{(NC)}$ will denote projections of p on the subspaces of critical and noncritical components, respectively. [In other words, $p^{(C)}$ is obtained from p by replacing p_n with 0 for $n \in \bar{N}^{(NC)}$, and similarly $p^{(NC)}$ is obtained by replacing critical components with 0.] Trivially, $p = p^{(C)} + p^{(NC)}$.

9.3. *Initial condition* (16). Condition (16) in Theorem 1 requires that initial states $\tilde{q}^r(0)$ converge to an invariant point. (Convergence to zero is included because the zero vector is an invariant point.) However, Theorem 1 can be generalized for the case when initial condition (16) does not hold, that is,

$$\tilde{q}^r(0) \rightarrow p \neq \tilde{x}^\circ(0)v,$$

where $\tilde{x}^\circ(0) \doteq \zeta \cdot p > 0$. Informally, in this case, similarly to the situation described by Theorem 3 in [4], a “weak limit” process \tilde{q} under MaxWeight may experience a “jump” at time 0 such that $\tilde{q}(0) = \tilde{x}(0)v$, where the initial workload (of the limit process) $\tilde{x}(0) \geq \tilde{x}^\circ(0)$. (We put “weak limit” and “jump” in quotation marks because in this case weak convergence on the interval $[0, \infty)$ does not hold, and needs to be replaced by convergence in the open interval $(0, \infty)$; and the proofs would also be pathwise, using Skorohod representation. So, by $\tilde{q}(0)$ and $\tilde{x}(0)$ we actually mean $\tilde{q}(0+)$ and $\tilde{x}(0+)$, resp.) It can be shown, using our Lemmas 7 and 8, that the ratio $\tilde{x}(0)/\|p\|$ (and $\|\tilde{q}(0)\|/\|p\|$) is bounded above by a fixed constant. Moreover, if p is close enough to the invariant manifold in the sense that the angle between p and v is small enough, then there is *no* initial jump of workload, $\tilde{x}(0) = \tilde{x}^\circ(0)$. [But the queue length vector always jumps, i.e., $\tilde{q}(0) \neq p$, as long as p is not an invariant point.]

After a possible initial jump, the limiting process \tilde{q} behaves the same way as \tilde{q}° : \tilde{x} is a one-dimensional RBM with drift a and variance σ^2 , and $\tilde{q} = \tilde{x}v$.

This possible workload jump at initial time 0, in the case of “bad” initial state, is

due to the parsimonious nature of MaxWeight. Since MaxWeight “knows” neither the mean rate vector λ nor the exact geometry of the rate region \bar{V} , it takes some transient period (which becomes infinitely short on the diffusion time scale) for the queue lengths vector to “adapt,” namely, get close enough to the invariant manifold. Within this initial transience period (which, again, shrinks to zero in the diffusion limit), the workload service rate could be wasted. A different rule which *a priori* knows rates λ and the geometry of \bar{V} , could possibly prevent such wastage and avoid the initial workload jump.

Note, however, that MaxWeight behavior with a “bad” initial state has little implications for the MaxWeight optimality in practical situations when $a < 0$. First, MaxWeight minimizes workload starting at time t_* when $\tilde{q}(t)$ hits 0 for the first time, and this occurs with probability 1 in the case $a < 0$. Second, if $a < 0$, MaxWeight stochastically minimizes workload in the stationary regime (assuming Conjecture 1 holds).

10. Intuition behind the main results. In this section we define additional random functions associated with the system for each value of the scaling parameter r . This will allow us to provide intuition behind our main results, and introduce notation used in the proofs.

We denote by

$$F_n^r(t) \doteq \sum_{l=1}^{\lfloor t \rfloor} A_n^r(l)$$

the cumulative number of type n customers arrived by time t (i.e., in the interval $[0, t]$, excluding customers present at time 0). Let

$$(25) \quad \hat{F}_n^r(t) \doteq \sum_{l=1}^{\lfloor t \rfloor} D_n^r(l)$$

denote the number of type n customers that were served and have departed by time $t \geq 0$. Also, denote by $G_m^r(t)$ the total number of time slots by (and including) time $t - 1$, when the server was in state m ; and by $\hat{G}_{mk}^r(t)$ the number of time slots by (and including) time $t - 1$ when the server state was m and the scheduling decision $k \in K(m)$ was chosen.

Recall that the probability law of the Markov chain $m^r(\cdot)$ describing the switch state process is the same for each r . Let us introduce the following function of a switch state:

$$\bar{\mu}^m \doteq \max_{k \in K(m)} \zeta \cdot \mu^m(k), \quad m \in \bar{M},$$

which is the maximum possible amount of workload that could potentially be

served in one time slot when the switch is in state m . Then

$$(26) \quad \bar{\mu} \doteq \sum_{m \in \bar{M}} \pi_m \bar{\mu}^m = \zeta \cdot \lambda$$

is the maximum possible average rate at which switch can serve workload. [The second equality in (26) is established, e.g., as follows. Consider an SSS rule ϕ such that, in each switch state m , a fixed decision $k \in K(m)$ maximizing $\zeta \cdot \mu^m(k)$ is chosen. Then, we see that $\bar{\mu} = \zeta \cdot v(\phi)$; this $v(\phi)$ maximizes $\zeta \cdot v$ over all $v \in \bar{V}$, and so does λ .]

For each $t \geq 0$, we denote by

$$(27) \quad H^r(t) \doteq \sum_{l=1}^{\lfloor t \rfloor} \bar{\mu}^{m^r(l-1)} = \sum_{m \in \bar{M}} \bar{\mu}^m G_m^r(t)$$

the potential amount of workload that could be served by time t , and by

$$Y^r(t) \doteq H^r(t) - \zeta \cdot \hat{F}^r(t) \equiv \sum_{l=1}^{\lfloor t \rfloor} [\bar{\mu}^{m^r(l-1)} - \zeta \cdot D^r(l)],$$

the amount of workload service “wasted” by time t . The following process

$$W^r(t) \doteq X^r(0) + \zeta \cdot F^r(t) - H^r(t), \quad t \geq 0,$$

depends only on the initial workload and model primitives and is invariant w.r.t. a scheduling discipline.

The following relations obviously hold for all $t \geq 0$ and any $n \in \bar{N}$:

$$(28) \quad F_n^r(0) = \hat{F}_n^r(0) = 0,$$

$$Q_n^r(t) = Q_n^r(0) + F_n^r(t) - \hat{F}_n^r(t),$$

$$(29) \quad X^r(t) = W^r(t) + Y^r(t).$$

Assumptions (11)–(14) imply a FCLT for each input flow:

$$(30) \quad \{r^{-1}(F_n^r(r^2t) - \lambda_n^r r^2t), t \geq 0\} \xrightarrow{w} \{\sigma_n B(t), t \geq 0\},$$

where B is a standard (zero drift, unit variance) Brownian motion.

From the FCLT for Markov chains, we have the following FCLT for the potential workload service process H^r . For any initial states of the (switch state) Markov chains $m^r(\cdot)$, as $r \rightarrow \infty$:

$$(31) \quad \{r^{-1}(H^r(r^2t) - \bar{\mu} r^2t), t \geq 0\} \xrightarrow{w} \{\sigma_s B(t), t \geq 0\},$$

where

$$(32) \quad \sigma_s^2 = \lim_{n \rightarrow \infty} n^{-1} E \left[\sum_{t=1}^n \bar{\mu}^{m^r(t-1)} - \bar{\mu} n \right]^2.$$

(The parameter σ_s is of course independent of r .)

For future reference, we define the process Z^r , describing system evolution:

$$Z^r = (Q^r, X^r, W^r, Y^r, F^r, \hat{F}^r, G^r, H^r, \hat{G}^r),$$

where

$$Q^r = (Q^r(t) = (Q_1^r(t), \dots, Q_N^r(t)), t \geq 0),$$

$$X^r = (X^r(t), t \geq 0),$$

$$W^r = (W^r(t), t \geq 0),$$

$$Y^r = (Y^r(t), t \geq 0),$$

$$F^r = (F^r(t) = (F_1^r(t), \dots, F_N^r(t)), t \geq 0),$$

$$\hat{F}^r = (\hat{F}^r(t) = (\hat{F}_1^r(t), \dots, \hat{F}_N^r(t)), t \geq 0),$$

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

$$H^r = (H^r(t), t \geq 0),$$

$$\hat{G}^r = ((\hat{G}_{mk}^r(t), m \in \bar{M}, k \in K(m)), t \geq 0).$$

Recall our convention that all component functions, as functions of t , are defined for $t \in \mathbb{R}_+$ and are constant within each time slot $[t, t+1)$, $t = 0, 1, 2, \dots$

Now, let us apply diffusion scaling $\tilde{\Gamma}^r$ to the processes W^r and Y^r . Thus, in addition to $\tilde{q}^r(\cdot)$ and $\tilde{x}^r(\cdot)$, we consider similarly defined

$$\tilde{w}^r(t) \doteq r^{-1} W^r(r^2 t), \quad t \geq 0,$$

and $\tilde{y}^r(\cdot)$.

From (29) we have

$$(33) \quad \tilde{x}^r(t) = \tilde{w}^r(t) + \tilde{y}^r(t), \quad t \geq 0,$$

and it follows from the definition of W^r , (30), (31) and (10), that \tilde{w}^r weakly converges to a Brownian motion \tilde{w} defined in (17):

$$(34) \quad \tilde{w}^r \xrightarrow{w} \tilde{w}.$$

Relation (33) is of course key for our analysis. The process \tilde{w}^r is a “driving” process, depending only on the system primitives, and converging to a Brownian motion. The process \tilde{x}^r is nonnegative, and so the nondecreasing process \tilde{y}^r appears to be a “regulation” (or “pushing”) process, which keeps \tilde{x}^r from going below zero. However, the unusual feature of relation (33) in our case is that, unlike a “conventional” regulation process, \tilde{y}^r can increase (i.e., workload service can be wasted) even when \tilde{x}^r is arbitrarily large. Thus, the nontrivial part of our analysis is to show that, as Theorem 1(i) claims, under MaxWeight discipline, \tilde{y}^r in fact converges to the conventional regulation process \tilde{y}° [defined in (19)] which does not increase when $\tilde{x}^\circ(t) = \tilde{w}(t) + \tilde{y}^\circ(t) > 0$.

11. Fluid sample paths for MaxWeight discipline.

11.1. *Definition and basic properties.* In this section we study sequences of processes Z^r under *fluid scaling* and under MaxWeight discipline. In fact, the definition of a fluid sample path below only involves sample paths of the processes Z^r under fluid scaling, and their limits. Within this section, we consider sequences $\{Z^r, r \in \mathcal{R}_f\}$, where \mathcal{R}_f can be an arbitrary (possibly completely unrelated to \mathcal{R}) nondecreasing sequence of positive numbers tending to infinity.

Recall our convention that all component functions of Z^r , as functions of t , are defined for $t \in \mathbb{R}_+$ and are constant within each time slot $[t, t + 1), t = 0, 1, 2, \dots$. Since we are going to consider the process Z^r restarted at different, not necessarily integer, times, from this point on it will be convenient to generalize the definition of Z^r . We will allow Z^r to be either a process defined in Section 10 or its version restarted at an arbitrary fixed time $d \geq 0$, namely, a process $\Theta(d)Z^r \doteq (\theta_d Q^r, \theta_d X^r, \theta_d W^r - W^r(d) + X^r(d), \theta_d Y^r - Y^r(d), \theta_d F^r - F^r(d), \theta_d \hat{F}^r - \hat{F}^r(d), \theta_d G^r - G^r(d), \theta_d H^r - H^r(d), \theta_d \hat{G}^r - \hat{G}^r(d))$. (Such a generalization is nothing more than a convention that, in the definition of Z^r , the very first time slot can be shorter than 1.)

For each r , consider the scaled process

$$\Gamma^r Z^r \doteq z^r = (q^r, x^r, w^r, y^r, f^r, \hat{f}^r, g^r, h^r, \hat{g}^r).$$

Thus, the component functions of z^r are piecewise constant, but a “time slot” has the length $1/r$, except for the first slot which may be shorter than $1/r$. The “special” shift operator $\Theta(d)z^r$ (for $d \geq 0$) acts on z^r analogously to the way it acts on Z^r .

From (28) we get

$$(35) \quad q_n^r(t) \equiv q_n^r(0) + f_n^r(t) - \hat{f}_n^r(t), \quad t \geq 0, n \in \bar{N}.$$

DEFINITION. A fixed set of functions $z = (q, x, w, y, f, \hat{f}, g, h, \hat{g})$ we will call a *fluid sample path* (FSP) if there exists a sequence \mathcal{R}_f of values of r , and a sequence of sample paths (of the corresponding processes) $\{z^r\}$ such that, as $r \rightarrow \infty$ along sequence \mathcal{R}_f ,

$$(36) \quad z^r \rightarrow z \quad \text{u.o.c.,}$$

and in addition

$$(37) \quad \|q(0)\| < \infty,$$

$$(38) \quad (f_n^r(t), t \geq 0) \rightarrow (\lambda_n t, t \geq 0) \quad \text{u.o.c., } \forall n \in \bar{N},$$

$$(39) \quad (g_m^r(t), t \geq 0) \rightarrow (\pi_m t, t \geq 0) \quad \text{u.o.c., } \forall m \in \bar{M}.$$

We emphasize that a sequence \mathcal{R}_f whose existence is required in the above definition may be completely unrelated to the sequence \mathcal{R} we introduced earlier in the definition of the heavy traffic regime. For an FSP z , any sequence of scaled paths $\{z^r, r \in \mathcal{R}_f\}$ which satisfies the conditions of the above definition will be called a sequence *defining* z .

REMARK. The definition of an FSP does *not* require any assumptions on the vector of mean rates λ besides $\lambda \in \mathbb{R}_+^N$. (To be precise, this is true if we exclude components x^r, w^r, y^r, h^r from z^r , and corresponding components x, w, y, h from z . Or, alternatively, we can assume that workload aggregator $\zeta \in \mathbb{R}_+^N$ is just some fixed vector, possibly unrelated to λ .) Many of the FSP properties established in this paper (as can be seen from their proofs) hold for any $\lambda \in \mathbb{R}_+^N$, not necessarily satisfying RP condition. In particular, this applies to properties in Lemma 1 [excluding (43)–(45)] and Lemmas 2, 3 and 5.

The following lemma establishes some basic properties of fluid sample paths, easily implied by their definition.

LEMMA 1. *For any fluid sample path z , all its component functions are Lipschitz continuous and, in addition,*

$$(40) \quad f(t) = \lambda t, \quad t \geq 0,$$

$$(41) \quad g_m(t) = \pi_m t, \quad t \geq 0, \quad m \in \bar{M},$$

$$(42) \quad q(t) = q(0) + f(t) - \hat{f}(t), \quad t \geq 0,$$

$$(43) \quad h(t) = \sum_{m \in \bar{M}} \bar{\mu}^m g_m(t) = \bar{\mu} t = \zeta \cdot \lambda t, \quad t \geq 0,$$

$$(44) \quad w(t) = x(0) + \zeta \cdot f(t) - h(t) = x(0) = w(0), \quad t \geq 0,$$

$$(45) \quad x(t) = \zeta \cdot q(t) = w(t) + y(t) = x(0) + y(t), \quad t \geq 0,$$

$$(46) \quad \begin{aligned} & \hat{f}(t_2) - \hat{f}(t_1) \\ & \leq \sum_{m \in \bar{M}} \sum_{k \in K(m)} [\hat{g}_{mk}(t_2) - \hat{g}_{mk}(t_1)] \mu^m(k), \quad t_2 \geq t_1 \geq 0, \end{aligned}$$

$$(47) \quad \sum_{k \in K(m)} [\hat{g}_{mk}(t_2) - \hat{g}_{mk}(t_1)] = g_m(t_2) - g_m(t_1), \quad t_2 \geq t_1 \geq 0, \quad m \in \bar{M}.$$

PROOF. Properties (40) and (41) follow directly from the definition of an FSP, which of course means that each function f_n and g_m is Lipschitz. Lipschitz continuity of limiting functions \hat{f}_n and \hat{g}_{mk} follows from the fact that the increments of the corresponding prelimit (unscaled) functions \hat{F}_n^r and \hat{G}_{mk}^r within

one time slot are uniformly bounded due to our model assumptions. [See (25) and the paragraph following (25) for the definitions of \hat{F}_n^r and \hat{G}_{mk}^r .] The Lipschitz continuity of all components of FSP z easily follows. Due to (35), property (42) also follows from the FSP definition. In (43), the first equality follows from the definitions of H^r [in (27)] and FSP, and the second and third ones from (41) and (26). The definitions of W^r and FSP imply the first equality in (44), which in turn implies the remaining equalities in (44) in view of (40) and (43). The first two equalities in (45) follow from the definition of X^r and (29) (along with FSP definition), and the last one follows from (44). Inequality (46) is implied by the inequality

$$\hat{f}^r(t_2) - \hat{f}^r(t_1) \leq \sum_{m \in \bar{M}} \sum_{k \in K(m)} [\hat{g}_{mk}^r(t_2) - \hat{g}_{mk}^r(t_1)] \mu^m(k), \quad t_2 \geq t_1 \geq 0,$$

which is a trivial consequence of the fact that if the switch state in a time slot is m and a scheduling decision $k \in K(m)$ is chosen, then the number of type n customers served in this slot cannot exceed $\mu_n^m(k)$. Similarly, (47) follows from

$$\sum_{k \in K(m)} [\hat{g}_{mk}^r(t_2) - \hat{g}_{mk}^r(t_1)] = g_m^r(t_2) - g_m^r(t_1), \quad t_2 \geq t_1 \geq 0, \quad m \in \bar{M},$$

which in turn is a consequence of the fact that when the switch is in state m , at least one of the decisions $k \in K(m)$ is chosen. \square

Lemma 2 is closely related to Lemma 1 and is also a simple corollary of the FSP definition. The lemma will be used in the proofs of our main (heavy traffic limit) results.

LEMMA 2. *Suppose a sequence of sample paths $\{z^r\}$, with $r \rightarrow \infty$ along some sequence \mathcal{R}_{f1} , is such that the conditions (38) and (39) hold and, for some nonnegative constants $c_1 \leq c_2$, $\|q^r(0)\| \in [c_1, c_2]$ for all r . Then, there exists a subsequence $\mathcal{R}_f \subseteq \mathcal{R}_{f1}$ along which z^r converges (u.o.c.) to an FSP z with $\|q(0)\| \in [c_1, c_2]$.*

PROOF. We have $(f^r, g^r) \rightarrow (f, g)$ u.o.c. by assumption, which also implies that $h^r \rightarrow h$ u.o.c., with f, g, h defined by (40), (41) and (43). We can always choose a subsequence $\mathcal{R}_{f2} \subseteq \mathcal{R}_{f1}$ along which we have $(\hat{f}^r, \hat{g}^r) \Rightarrow (\hat{f}, \hat{g})$ (since the functions in the left-hand side are nondecreasing). The limit functions \hat{f}, \hat{g} must be Lipschitz continuous (because the increments of prelimit functions \hat{F}_n^r and \hat{G}_{mk}^r within one time slot are uniformly bounded). This implies that, in fact, $(\hat{f}^r, \hat{g}^r) \rightarrow (\hat{f}, \hat{g})$ u.o.c. It remains to choose a further subsequence $\mathcal{R}_f \subseteq \mathcal{R}_{f2}$ along which $q^r(0) \rightarrow q(0)$ with $\|q(0)\| \in [c_1, c_2]$. Along this subsequence we must have $(q^r, x^r, w^r, y^r) \rightarrow (q, x, w, y)$ u.o.c., with q, x, w, y defined by the other components and relations (42), (44) and (45). Thus, z^r converges (u.o.c.) to $z = (q, x, w, y, f, \hat{f}, g, h, \hat{g})$, which is an FSP. \square

The scaling and shift properties, described in the following Lemma 3, and their proof are analogous to those of Lemmas 6.1 and 6.2 in [28]. Since FSPs are defined as limits, they basically just inherit these properties from the pre-limit paths.

LEMMA 3. *For any fluid sample path z , the following properties hold.*

- (i) Scaling (“Similarity”). *For any $c > 0$, $\Gamma^c z$ is also a fluid sample path.*
- (ii) Shift. *For any $d \geq 0$, $\Theta(d)z$ is also a fluid sample path.*

PROOF. Consider a fixed FSP z and a sequence of scaled paths $\{z^r, r \in \mathcal{R}_f\}$ which defines it. Let $\{Z^r, r \in \mathcal{R}_f\}$ be the corresponding sequence of unscaled paths, that is, $z^r = \Gamma^r Z^r$ for all r . For any fixed $c > 0$, we have $\Gamma^c z^r = \Gamma^{cr} Z^r$ for all r . Therefore, $\{\Gamma^c z^r\}$ is a valid sequence of scaled paths, obtained from paths Z^r (which can be relabeled as Z^{cr}) by operators Γ^{cr} . We can now verify directly that the sequence $\{\Gamma^c z^r\}$ defines FSP $\Gamma^c z$, which proves (i). Similarly, for any $d \geq 0$, $\{\Theta(d)z^r\}$ is a valid sequence of scaled paths, which defines FSP $\Theta(d)z$. This proves (ii). \square

11.2. *Uniform attraction of fluid sample paths.* In this section we prove that the family of fluid sample paths is such that, as $t \rightarrow \infty$, $q(t)$ converges uniformly [up to scaling by the initial state norm $\|q(0)\|$] to an invariant point. Recall that an invariant point is any vector $q^* \in \mathbb{R}_+^N$ such that $q^* = cv$ for some $c \geq 0$ (or, equivalently, such that $\gamma \times [q^*]^\beta = c\zeta$ for some $c \geq 0$). Note that an invariant point q^* is the unique point where function

$$(48) \quad \Psi(q) \doteq \frac{1}{\beta + 1} \sum_n \gamma_n q_n^{\beta+1},$$

restricted to the hyperplane $\zeta \cdot q = \zeta \cdot q^*$ [i.e., the hyperplane $L(\zeta, q^*)$], attains its minimum. Recall that the set of all invariant points (a ray in our case) is called the invariant manifold.

The underlying intuition for the FSP attraction property is that, as we will see shortly in Lemmas 5 and 6, there are two “Lyapunov functions” $\Psi(q(t))$ and $x(t) = \zeta \cdot q(t)$, the former nonincreasing and the latter nondecreasing, which “sandwich” the values of $q(t)$.

For vectors $q \in \mathbb{R}_+^N$, consider the following (Lyapunov) function:

$$(49) \quad G(q) \doteq \frac{\Psi(q)}{\Psi(q^*)} - 1,$$

where $q^* = (\zeta \cdot q)v$ is the unique invariant point lying in the hyperplane $L(\zeta, q)$, and we use the following conventions: if $q \neq 0$ and $q^* = 0$ [meaning $\Psi(q^*) = 0$], we put $G(q) = \infty$; if $q = 0$, we put $G(q) = 0$.

It follows from the definition of an invariant point that $G(q) \geq 0$ and, moreover, $G(q) > 0$ unless q is an invariant point. We also note that $G(q)$ is *invariant with respect to scaling* of q ; that is, $G(q) = G(cq)$ for any $c > 0$.

For vectors $q \in \mathbb{R}_+^N$, let us also introduce the function $\alpha(q)$ which is the angle between vectors $\gamma \times q^\beta$ and vector ζ , namely,

$$\alpha(q) = \arccos \frac{(\gamma \times q^\beta) \cdot \zeta}{\|\gamma \times q^\beta\| \|\zeta\|}$$

if $q \neq 0$, and $\alpha(q) = 0$ by convention if $q = 0$. Trivially, $\alpha(q)$ is also invariant with respect to scaling of q .

It follows directly from the above definitions that all three following conditions are equivalent: $G(q) = 0$, $q = q^*$ and $\alpha(q) = 0$. Moreover, using the fact that $\Psi(q)$ is a continuous strictly convex function, it is easy to observe that, for nonzero vectors q , all three following convergence properties are equivalent: $G(q) \rightarrow 0$, $\|q - q^*\|/\|q^*\| \rightarrow 0$, $\alpha(q) \rightarrow 0$. This in particular implies that

$$(50) \quad \sup\{G(q) | q \neq 0, \alpha(q) \leq \Delta\} \downarrow 0 \quad \text{as } \Delta \downarrow 0,$$

$$(51) \quad \sup\{\alpha(q) | q \neq 0, G(q) \leq \Delta\} \downarrow 0 \quad \text{as } \Delta \downarrow 0$$

and

$$(52) \quad \sup\left\{\frac{\|q - (\zeta \cdot q)v\|}{\|(\zeta \cdot q)v\|} \mid q \neq 0, G(q) \leq \Delta\right\} \downarrow 0 \quad \text{as } \Delta \downarrow 0.$$

For a fixed switch state m and a nonzero vector $p \in \mathbb{R}_+^N$, let us define $K^*(m, p) \doteq \arg \max_{k \in K(m)} p \cdot \mu^m(k)$. Since $K(m)$ is a finite set for every m and function $p \cdot \mu^m(k)$ is continuous on $p \in \mathbb{R}_+^N$, we obtain the following simple fact.

LEMMA 4. *Suppose a nonzero vector $p \in \mathbb{R}_+^N$ is fixed. Then there exists small $\Delta > 0$ (depending on p) such that, for all nonzero vectors q forming an angle smaller than Δ with p , we have*

$$K^*(m, q) \subseteq K^*(m, p) \quad \forall m \in \bar{M}.$$

We now introduce some definitions and conventions regarding derivatives of FSP components. Since all component functions of an FSP are Lipschitz, they are absolutely continuous, and therefore almost all points $t \in \mathbb{R}_+$ (with respect to Lebesgue measure) are such that all component functions of z have derivatives; we will call such points *regular*. For any regular point $t \geq 0$, we have

$$(53) \quad \frac{d}{dt} q(t) = \lambda - \hat{v}(t),$$

where

$$(54) \quad \hat{v}(t) \doteq \hat{f}'(t) \leq v(t) \doteq \sum_{m \in \bar{M}} \sum_{k \in K(m)} \hat{g}'_{mk}(t) \mu^m(k)$$

with the inequality following from (46). Necessarily, $v(t) \in \bar{V}$ [and therefore $\hat{v}(t) \in \bar{V}$]. Indeed, by (47) and (41), for any m , we have $\sum_{k \in K(m)} \hat{g}'_{mk}(t) = \pi_m$, and thus

$$v(t) = \sum_{m \in \bar{M}} \pi_m \sum_{k \in K(m)} (\hat{g}'_{mk}(t)/\pi_m) \mu^m(k) \in \bar{V}.$$

In the rest of the paper we use the following convention: when we write an expression containing derivatives of FSP components (or derivatives of functions of FSP components) at time $t \geq 0$, we always mean that it holds under the additional condition that t is regular, even if we do not state this condition explicitly.

The following Lemma 5 describes key properties of FSPs leading eventually to the uniform attraction property. In particular, (57) is the key differential inclusion an FSP (for MaxWeight rule) must satisfy. Property (58) [which is a corollary of (57)] shows that the derivative of the (Lyapunov) function $\Psi(q(t))$ is minimized at all times. This property of the FSPs is a manifestation of the underlying “principle” behind the MaxWeight rule pointed out in the Introduction.

LEMMA 5. *For any FSP, the following hold. [Recall the definitions of $v(t)$ and $\hat{v}(t)$ in (54).]*

(i) *For any $n \in \bar{N}$,*

$$(55) \quad q_n(t) > 0 \quad \text{implies} \quad \hat{v}_n(t) = v_n(t), \quad t \geq 0.$$

(ii) *We have*

$$(56) \quad v(t) \in \arg \max_{v \in \bar{V}} (\gamma \times q(t)^\beta) \cdot v, \quad t \geq 0.$$

(iii) *We have*

$$(57) \quad \frac{d}{dt} q(t) = \lambda - \hat{v}(t), \quad \hat{v}(t) \in \arg \max_{v \in \bar{V}} (\gamma \times q(t)^\beta) \cdot v, \quad t \geq 0,$$

$$(58) \quad \begin{aligned} [\Psi(q(t))]' &= (\gamma \times q(t)^\beta) \cdot (\lambda - \hat{v}(t)) \\ &= \min_{v \in \bar{V}} (\gamma \times q(t)^\beta) \cdot (\lambda - v), \quad t \geq 0, \end{aligned}$$

REMARK. As we mentioned earlier in the remark following the FSP definition, Lemma 5 holds for any $\lambda \in \mathbb{R}_+^N$, not only λ satisfying the RP condition. If λ is within the stability region V^0 , then (by the definition of V^0) $\lambda < v$ for some fixed $v \in \bar{V}$, and therefore [by (58)] for any FSP the derivative $[\Psi(q(t))]' \leq (\gamma \times q(t)^\beta) \cdot (\lambda - v)$ is strictly negative and separated from zero as long as $\|q(t)\|$ is separated from zero. Thus, $\lambda \in V^0$ implies that, uniformly on $q(0)$ with $\|q(0)\| = 1$, for any $\varepsilon > 0$, the norm $\|q(t)\|$ must reach level ε at or before some fixed finite time, depending only on ε . This property of FSPs im-

plies (using the fluid limit technique) stability of our system with the vector of mean rates $\lambda \in V^0$, and this is exactly how the MaxWeight stability is proved in [1]. (Establishing the MaxWeight stability in the case of more general weights, as in (4), requires an additional step, which can also be found in [1].)

PROOF OF LEMMA 5. Throughout this proof we consider a fixed FSP z , a sequence of scaled paths $\{z^r, r \in \mathcal{R}_f\}$, which defines it, and the corresponding sequence of unscaled paths $\{Z^r, r \in \mathcal{R}_f\}$, that is, $z^r = \Gamma^r Z^r$ for all r .

(i) Suppose $q_n(t) > 0$. We know that $q_n(\cdot)$ is continuous. Then the following observation is true:

There exist small fixed $\Delta_1 > 0$ and $\Delta_2 > 0$ [both depending on $q_n(t)$] such that, for all sufficiently large r , the unscaled paths Z^r are such that $Q_n^r(\xi) > r\Delta_2$ for all $\xi \in [rt, rt + r\Delta_1]$.

This means that, for any $\xi \in [t, t + \Delta_1]$ and all large r ,

$$(59) \quad \hat{f}_n^r(\xi) - \hat{f}_n^r(t) = \sum_{m \in \bar{M}} \sum_{k \in K(m)} [\hat{g}_{mk}^r(\xi) - \hat{g}_{mk}^r(t)] \mu_n^m(k),$$

where (we remind) $\hat{f}_n = \Gamma^r \hat{F}_n$, $\hat{g}_{mk} = \Gamma^r \hat{G}_{mk}$, with \hat{F}_n and \hat{G}_{mk} defined in (25) and the paragraph following (25). Taking the $r \rightarrow \infty$ limit in (59), we obtain analogous equality for the corresponding FSP components, which by taking derivative on ξ at t implies $\hat{v}_n(t) = v_n(t)$.

(ii) Suppose $q(t) \neq 0$. Consider the sets $K^*(m, \gamma \times q(t)^\beta)$ for each m . From Lemma 4, the fact that $q(\cdot)$ is continuous, and the form of the MaxWeight rule, we easily obtain the following property:

There exists a small fixed $\Delta_3 > 0$ such that, for all sufficiently large r , the unscaled paths Z^r are such that in the time interval $[rt, rt + r\Delta_3]$ for any switch state m , only the decisions from the subset $K^*(m, \gamma \times q(t)^\beta)$ can be chosen.

By the definition of $K^*(m, \gamma \times q(t)^\beta)$, the value of $(\gamma \times q(t)^\beta) \cdot \mu^m(k)$ is same for all $k \in K^*(m, \gamma \times q(t)^\beta)$. Therefore, for any $\xi \in [t, t + \Delta_3]$, we have

$$\begin{aligned} & [\gamma \times q(t)^\beta] \cdot \sum_{m \in \bar{M}} \sum_{k \in K(m)} (\hat{g}_{mk}^r(\xi) - \hat{g}_{mk}^r(t)) \mu^m(k) \\ &= \sum_{m \in \bar{M}} \sum_{k \in K(m)} (\hat{g}_{mk}^r(\xi) - \hat{g}_{mk}^r(t)) \left\{ \max_{i \in K(m)} [\gamma \times q(t)^\beta] \cdot \mu^m(i) \right\}. \end{aligned}$$

Taking the $r \rightarrow \infty$ limit in the last equality and using (47) and (39), we obtain

$$\begin{aligned} & [\gamma \times q(t)^\beta] \cdot \sum_{m \in \bar{M}} \sum_{k \in K(m)} (\hat{g}_{mk}(\xi) - \hat{g}_{mk}(t)) \mu^m(k) \\ &= \sum_{m \in \bar{M}} \pi_m(\xi - t) \left\{ \max_{i \in K(m)} [\gamma \times q(t)^\beta] \cdot \mu^m(i) \right\} \\ &= (\xi - t) \left\{ \max_{v \in \bar{V}} [\gamma \times q(t)^\beta] \cdot v \right\}. \end{aligned}$$

Finally, taking the derivative on ξ at t , we see that

$$[\gamma \times q(t)^\beta] \cdot v(t) = \max_{v \in \bar{V}} [\gamma \times q(t)^\beta] \cdot v,$$

which proves (56) [for $q(t) \neq 0$]. If $q(t) = 0$, then (56) holds trivially, because the $\arg \max$ in the right-hand side is equal to the entire set \bar{V} , and $v(t)$ always belongs to \bar{V} . The proof of (56) is complete.

(iii) The property (57) follows from (53), (56) and (55). The left-hand side equality in (58) is obtained by differentiating $\Psi(q(t))$ on t and applying the equality in (57). The right-hand side equality in (58) follows from the inclusion in (57). \square

The following lemma is basically just a corollary of Lemma 5.

LEMMA 6. *For any FSP, the following hold:*

(i) *We have $x'(t) \geq 0$ and $[\Psi(q(t))]' \leq 0$ for all $t \geq 0$. Therefore, $x(t)$ is nondecreasing, $\Psi(q(t))$ is nonincreasing, and $q(0) = 0$ implies $q(t) \equiv 0$ for all $t \geq 0$.*

(ii) *There exists $\varepsilon > 0$ (depending only on the model parameters) such that*

$$[\Psi(q(t))]' \leq -\varepsilon \rho(t),$$

where

$$\rho(t) \doteq \|\gamma \times q(t)^\beta\| \sin \alpha(q(t))$$

is the distance from point $\gamma \times q(t)^\beta$ to the invariant manifold.

(iii) *For any $\varepsilon_1 > 0$, there exists $T_1 = T_1(\varepsilon_1)$ depending on ε_1 such that, for any FSP with $q(0) \neq 0$, $\min\{t \geq 0 \mid \rho(t) \leq \varepsilon_1 \|q(0)\|\} \leq \|q(0)\| T_1$.*

PROOF. (i) We have $x'(t) = \zeta \cdot (\lambda - \hat{v}(t)) \geq 0$, because λ maximizes $\zeta \cdot v$ over \bar{V} and $\hat{v}(t) \in \bar{V}$. Inequality $[\Psi(q(t))]' \leq 0$ follows from (58) and the fact that $\lambda \in \bar{V}$.

(ii) For an $\varepsilon > 0$, consider the subset $U(\varepsilon)$ consisting of all vectors $\lambda + \xi$ such that $\zeta \cdot \xi = 0$ and $\|\xi\| = \varepsilon > 0$. Let us fix $\varepsilon > 0$ small enough so that $U(\varepsilon) \subset \bar{V}(\zeta)$. [Recall that $\bar{V}(\zeta)$ is the face of the boundary of \bar{V} , orthogonal to ζ and containing λ in its interior.] It follows from (58) that

$$[\Psi(q(t))]' \leq \min_{v \in U(\varepsilon)} (\gamma \times q(t)^\beta) \cdot (\lambda - v) = -\varepsilon \|\gamma \times q(t)^\beta\| \sin \alpha(q(t)).$$

(iii) Due to the scaling property of FSPs, without loss of generality we can assume $\|q(0)\| = 1$. We know that $\Psi(q(t)) \geq 0$ and $\Psi(q(t)) \leq \Psi(q(0)) \leq \sum_n \gamma_n$ for all $t \geq 0$. But, by (ii), $[\Psi(q(t))]' \leq -\varepsilon \varepsilon_1 < 0$ as long as $\rho(t) > \varepsilon_1$. Therefore, if $\rho(0) > \varepsilon_1$, $\rho(t)$ must hit the value ε_1 by the time $T_1 = \sum_n \gamma_n / (\varepsilon \varepsilon_1)$. \square

We note for future reference that, for any vector $q \in \mathbb{R}_+^N$ with workload $x = \zeta \cdot q > 0$ [or, equivalently, with $q^{(C)} \neq 0$], we have

$$(60) \quad \frac{\|q^{(C)}\|}{x} < \kappa, \quad \frac{x}{\|q^{(C)}\|} < \kappa, \quad \frac{x}{\|q\|} < \kappa, \quad \frac{x}{\max_n q_n} < \kappa,$$

where $\kappa > 1$ is a universal constant, depending only on the number of flows N and workload aggregator ζ . In addition, for any two nonzero vectors p and q such that $\Psi(p) = \Psi(q)$, the ratio $\|p\|/\|q\|$ is bounded above by a universal constant (depending only on N, β, γ); without loss of generality assume this bound to be strictly less than κ :

$$(61) \quad \max\{\|p\|/\|q\| \mid q \in \mathbb{R}_+^N, q \neq 0, \Psi(p) = \Psi(q)\} < \kappa.$$

The following result describes the uniform attraction property of the FSPs, along with other properties used later in the paper.

THEOREM 2. *For any FSP, the following hold.*

(i) *As functions of $t \geq 0$, $\Psi(q(t))$ is continuous nonincreasing, $x(t) = \zeta \cdot q(t)$ is continuous nondecreasing, and $G(q(t))$ is nonincreasing and it is continuous in every point t where it is finite [i.e., where $x(t) > 0$].*

(ii) *If $q(0) = 0$, then $q(t) \equiv 0$ for all $t \geq 0$. If $q(0) \neq 0$, then $\sup_{t \geq 0} \|q(t)\| < \kappa \|q(0)\|$ and $\sup_{t \geq 0} x(t) < \kappa^2 \|q(0)\|$.*

(iii) *Uniformly on FSPs with $\|q(0)\| = 1$,*

$$G(q(t)) \downarrow 0 \quad \text{as } t \rightarrow \infty.$$

PROOF. (i) This follows from Lemma 1, Lemma 6(i), and the definition of $G(q(t))$.

(ii) These properties follow from the fact that $\Psi(q(t))$ can only decrease and “universal” bounds (61) and (60).

(iii) Let $\|q(0)\| = 1$. Let us fix arbitrary small $\delta_1 > 0$. We will prove that the time for $\alpha(q(t))$ to reach $\delta_1 > 0$ is uniformly bounded by some fixed constant depending only on δ_1 . This will imply the desired uniform convergence $G(q(t)) \downarrow 0$ because $G(q(t))$ is nonincreasing and [by (50)], for any $\delta_2 > 0$, we can always choose $\delta_1 > 0$ small enough so that $\alpha(q(t)) \leq \delta_1$ implies $G(q(t)) \leq \delta_2$.

Given $\delta_1 > 0$ fixed above, let us choose a sufficiently small $\varepsilon_1 > 0$ such that, for any vector $p \in \mathbb{R}_+^N$,

$$(62) \quad \|\gamma \times p^\beta\| \leq \varepsilon_1 / \sin \delta_1 \quad \text{implies} \quad \|p\| \leq 1/2.$$

(The meaning of this choice will become clear later in the proof.) For this ε_1 we choose the corresponding T_1 as in Lemma 6(iii).

Consider the (possibly finite) sequence $\tau_0 \leq \tau_1 \leq \tau_2 \leq \dots$ constructed as follows. Let $\tau_0 = 0$. Let τ_1 be the smallest time $t \geq \tau_0$ such that either

$$(63) \quad \alpha(q(t)) \leq \delta_1$$

or $\|q(t)\| \leq \|q(\tau_0)\|/2$. If (63) holds for $t = \tau_1$, the sequence stops. If not, we define τ_2 as the smallest time after τ_1 such that either (63) holds or $\|q(t)\| \leq \|q(\tau_1)\|/2$. In the former case we stop, and in the latter case we define τ_3 analogously, and so on. We claim that, for any $i \geq 1$ (for which τ_i is well defined), we have

$$(64) \quad \tau_i - \tau_{i-1} \leq \|q(\tau_{i-1})\|T_1.$$

Due to scaling property of the FSPs, it suffices to prove (64) for $i = 1$, that is, to show that $\tau_1 \leq T_1$. Indeed, by Lemma 6(iii), at some time $t_1 \leq T_1$ we have $\rho(t_1) = \|\gamma \times q(t_1)^\beta\| \sin \alpha(q(t_1)) \leq \varepsilon_1$, which means that either $\alpha(q(t_1)) \leq \delta_1$ or [by (62)] $\|q(t_1)\| \leq 1/2 = \|q(0)\|/2$. This means $\tau_1 \leq t_1$, which proves claim (64).

Whether the sequence $\{\tau_i\}$ is finite or not, we have

$$\tau \doteq \sup_i \tau_i \leq \|q(0)\|T_1/(1 - 1/2) = 2T_1.$$

If sequence $\{\tau_i\}$ is finite, then $\alpha(q(\tau)) \leq \delta_1$ by construction. If sequence $\{\tau_i\}$ is infinite, then [by continuity of $q(t)$] $q(\tau) = \lim_i q(\tau_i) = 0$, which means $\alpha(q(\tau)) = 0$. Thus, the time for $\alpha(q(t))$ to reach δ_1 is uniformly bounded by $2T_1$. The proof is complete. \square

REMARK. Suppose $\lambda \in \bar{V}^*$ and $\lambda \in \mathbb{R}_{++}^N$, but the RP condition for λ does *not* necessarily hold. Our analysis of the FSPs (under MaxWeight) can be easily generalized to show that, in this case, the FSP uniform attraction property still holds, with the invariant manifold defined more generally as follows. Let $\bar{\zeta}$ be the convex cone of outward-pointing normal vectors to \bar{V} at point λ . (This cone is simply a ray, if RP condition holds.) Then, the set of invariant points (or the invariant manifold) is defined as the set of vectors $p \in \mathbb{R}_+^N$ such that $\gamma \times p^\beta \in \bar{\zeta}$. This invariant manifold is also a cone, although not necessarily convex. As mentioned earlier, the differential inclusion (57) and property (58) hold for any $\lambda \in \mathbb{R}_+^N$. Moreover, when $\lambda \in \bar{V}^*$ and $\lambda \in \mathbb{R}_{++}^N$, Lemma 6(ii) and (iii) holds as well, if $\alpha(q(t))$ is understood more generally as the angle between $\gamma \times q^\beta$ and the cone $\bar{\zeta}$. The fact that the “workload is nondecreasing” also holds if it is understood as the property that $\zeta \cdot q(t)$ is nondecreasing in t for each $\zeta \in \bar{\zeta}$. We do not provide details on the convergence proof modification, which we believe can be easily “recovered” by an interested reader. It should be emphasized that—if the RP condition fails—the uniform attraction of the FSPs to the invariant manifold in itself does *not* imply any heavy traffic (diffusion limit) optimality properties, such as those described in our main Theorem 1.

12. Proof of Theorem 1. As we mentioned in the Introduction, our proof of the heavy traffic SSC property, which relies on the attraction property of FSPs, follows the general approach developed in [4] and [35].

For each $r \in \mathcal{R}$, consider the following process, obtained by a diffusion scaling:

$$\tilde{\Gamma}^r(Q^r, X^r, W^r, Y^r, F^r, G^r, H^r) \doteq (\tilde{q}^r, \tilde{x}^r, \tilde{w}^r, \tilde{y}^r, \tilde{f}^r, \tilde{g}^r, \tilde{h}^r).$$

To prove properties (20)–(22), it will suffice to show that, for any fixed subsequence $\mathcal{R}_1 \subseteq \mathcal{R}$, there exists another subsequence $\mathcal{R}_2 \subseteq \mathcal{R}_1$ such that these properties hold when $r \rightarrow \infty$ along \mathcal{R}_2 . To do that, it in turn suffices to show that we can (using Skorohod representation) construct all processes (for all $r \in \mathcal{R}$) on the same probability space and choose subsequence \mathcal{R}_2 in a way such that the desired properties hold with probability 1 (or are implied by certain probability 1 properties) as $r \rightarrow \infty$ along \mathcal{R}_2 . In this section we do just that to prove (20) and (21) [which is Theorem 1(i)]. The proof of (22) [which is Theorem 1(ii)] will be even simpler—the common probability space will be such that the desired probability 1 property holds along the sequence \mathcal{R} itself.

We construct the underlying probability space as follows. According to the Skorohod representation theorem (see, e.g., [10]), for each n , the sequence of the input processes $\{F_n^r\}$ and a standard Brownian motion B_n can be constructed on a probability space such that, as $r \rightarrow \infty$ along \mathcal{R} , the convergence in (30) holds u.o.c. with probability 1 (w.p.1):

$$(65) \quad (\tilde{f}_n^r(t) - \lambda_n^r r t, t \geq 0) \xrightarrow{\text{u.o.c.}} \{\sigma_n B_n(t), t \geq 0\}.$$

Similarly, the sequence of processes $\{(H^r, G^r)\}$ [with distributions defined by Markov chains $m^r(\cdot)$] and a standard Brownian motion B_s can be constructed on a probability space such that the convergence in (31) holds u.o.c. w.p.1, which can be written as

$$(66) \quad (\tilde{h}^r(t) - \bar{\mu} r t, t \geq 0) \xrightarrow{\text{u.o.c.}} (\sigma_s^2 B_s(t), t \geq 0).$$

This is done in two steps. First, given distributions of the processes H^r and the weak convergence (31), we use the Skorohod representation to construct the sequence of H^r on a probability space such that (66) holds. Then, using the existence of a regular conditional distribution of G^r on H^r (see, e.g., Theorem 8.1 in the Appendix of [10]), we can augment the probability space so that processes G^r are defined on this space as well, and the pair (H^r, G^r) for each r has the correct joint distribution.

We can and do assume that the underlying probability space (Ω, \mathcal{F}, P) is the direct product of the $N + 1$ probability spaces specified above, and (without loss of generality) assume that this probability space is complete. By ω we will denote elements of Ω .

Then we have (33):

$$(67) \quad \tilde{x}^r(t) = \tilde{w}^r(t) + \tilde{y}^r(t), \quad t \geq 0,$$

and the probability 1 convergence version of (34):

$$(68) \quad (\tilde{w}^r(t), t \geq 0) \xrightarrow{\text{u.o.c.}} (\tilde{w}(t), t \geq 0),$$

where

$$\tilde{w}^r(t) = \tilde{x}^r(0) + \zeta \cdot \tilde{f}^r(t) - \tilde{h}^r(t)$$

and

$$\tilde{w}(t) \doteq \tilde{x}^\circ(0) + at + \sigma B(t), \quad t \geq 0,$$

is a Brownian motion, defined as in (17). By the definition of a Brownian motion, the sample paths of \tilde{w} are continuous.

Using Large Deviations estimates for Markov chains and the Borel–Cantelli lemma, it is easy to show (as, e.g., in [27]) that, as $r \rightarrow \infty$ along \mathcal{R} , the following properties hold with probability 1, for any fixed $T_3 > 0$:

$$(69) \quad \max_{0 \leq l \leq T_3 r^{3/2}} \left| \sqrt{r} g_m^r \left(\frac{l+1}{\sqrt{r}} \right) - \sqrt{r} g_m^r \left(\frac{l}{\sqrt{r}} \right) - \pi_m \right| \rightarrow 0, \quad m \in \bar{M},$$

$$(70) \quad \max_{0 \leq l \leq T_3 r^{3/2}} \left| \sqrt{r} h^r \left(\frac{l+1}{\sqrt{r}} \right) - \sqrt{r} h^r \left(\frac{l}{\sqrt{r}} \right) - \bar{\mu} \right| \rightarrow 0,$$

where (70) follows from (69).

To simplify notation, without loss of generality we assume that properties (65)–(70) hold for all $\omega \in \Omega$ (not just for almost all ω).

Now suppose an arbitrary subsequence $\mathcal{R}_1 \subseteq \mathcal{R}$ is fixed.

Since our arrival processes are i.i.d., we can use Bramson’s weak law estimate ([4], Proposition 4.2) to choose a subsequence $\mathcal{R}_2 \subseteq \mathcal{R}_1$ such that the following property holds. (We relegate its derivation to the Appendix.) With probability 1, as $r \rightarrow \infty$ along \mathcal{R}_2 , for any $T_3 > 0$ and any $n \in \bar{N}$, we have

$$(71) \quad \max_{0 \leq l \leq T_3 r} \sup_{0 \leq \xi \leq 1} |f_n^r(l + \xi) - f_n^r(l) - \lambda_n \xi| \rightarrow 0.$$

For the rest of this section, we define $\Omega_2 \subseteq \Omega$ as the (measurable, probability 1) subset of $\omega \in \Omega$ such that property (71) holds [in addition to (65)–(70)] along the subsequence \mathcal{R}_2 .

For any $r \in \mathcal{R}$, \tilde{y}^r is a nondecreasing RCLL function. Therefore, for any fixed $\omega \in \Omega$, from any subsequence $\mathcal{R}_3(\omega) \subseteq \mathcal{R}$ (which may depend on ω !) we can choose a further subsequence $\mathcal{R}_4(\omega) \subseteq \mathcal{R}_3(\omega)$ along which

$$(72) \quad \tilde{y}^r \Rightarrow \tilde{y},$$

where \tilde{y} is some nonnegative nondecreasing function in $D([0, \infty), \bar{\mathbb{R}})$. [This means that $\tilde{y}(t)$ may take infinite values $+\infty$. Also, recall that “ \Rightarrow ” means convergence in every point of continuity of the limit function except maybe point 0.] It is possible that $\tilde{y}(0) > 0$.

We note that (72) implies that

$$\tilde{x}^r \Rightarrow \tilde{x} \doteq \tilde{w} + \tilde{y},$$

and therefore $\tilde{x}(t) < \infty$ if and only if $\tilde{y}(t) < \infty$.

The key part of the proof of Theorem 1(i) is proving that if $\omega \in \Omega_2$, $\mathcal{R}_3(\omega) \subseteq \mathcal{R}_2$ [implying $\mathcal{R}_4(\omega) \subseteq \mathcal{R}_2$] and the scheduling discipline is MaxWeight, then $\tilde{y} = \tilde{y}^\circ$, where \tilde{y}° is the regulation function defined in (19). The proof of Theorem 1(ii) is much simpler: we will show that, for any $\omega \in \Omega$, $\mathcal{R}_3(\omega) \subseteq \mathcal{R}$ and arbitrary scheduling discipline, $\tilde{y}(t) \geq \tilde{y}^\circ(t)$ for all $t \geq 0$.

To divide the proof of Theorem 1(i) into manageable parts, two auxiliary Lemmas 7 and 8, pertaining to MaxWeight discipline, are formulated below. Lemma 7 contains the key fact that if workload $\tilde{x}^r(t)$ at some time t stays bounded above and separated from zero as $r \rightarrow \infty$, and the ratio $\|\tilde{q}^r(t)\|/\tilde{x}^r(t)$ stays bounded (which is automatically true if the CRP holds, but not necessarily true under the RP condition), then the limiting path $\tilde{y}(\cdot)$ cannot increase in a small interval to the right of t . This fact will be the main tool used in the proof of Theorem 1(i). The proof of Lemma 7 is presented in Section 12.1. Lemma 8 shows that in fact $\|\tilde{q}^{r,(NC)}(\cdot)\|$ is uniformly small (on compact sets) for large r , which allows us to apply Lemma 7, because we always have $\|\tilde{q}^{r,(C)}(t)\|/\tilde{x}^r(t) < \kappa$ for nonzero $\tilde{x}^r(t)$. The proof of Lemma 8 uses an argument similar to that in the proof of Lemma 7. Also, the proof of Theorem 1 in the case when the CRP condition holds does not require Lemma 8, so the reader may choose to skip Lemma 8 and its proof at first reading. For these reasons we relegate the proof of Lemma 8 to the Appendix.

LEMMA 7. (i) *Suppose the scheduling discipline is MaxWeight. Suppose $\omega \in \Omega_2$ and a subsequence $\mathcal{R}_4(\omega) \subseteq \mathcal{R}_2$ are fixed such that, along this subsequence, (72) holds. Suppose a sequence $\{\tilde{t}^r, r \in \mathcal{R}_4(\omega)\}$ is fixed such that*

$$(73) \quad \tilde{t}^r \rightarrow t' \geq 0,$$

$$(74) \quad \tilde{x}^r(\tilde{t}^r) \rightarrow C > 0$$

and

$$(75) \quad \limsup_{r \rightarrow \infty} \|\tilde{q}^r(\tilde{t}^r)\| < \kappa_1 C,$$

where $\kappa_1 > 1$ is a fixed constant. Let $\delta > 0$ be such that

$$\varepsilon = \text{Osc}(\tilde{w}; [t' - 3\delta, t' + 3\delta] \cap \mathbb{R}_+) < C/2.$$

Then,

- (a) \tilde{y} (and \tilde{x}) is finite in $[0, t' + \delta]$;
- (b) \tilde{y} does not increase in $(t', t' + \delta]$; that is, $\tilde{y}(t' + \delta) - \tilde{y}(t') = 0$;
- (c) the following bounds hold, with $C_1 \doteq \kappa^2 \kappa_1 C + 2\varepsilon$:

$$(76) \quad C - 2\varepsilon \leq \tilde{x}(t) \leq C_1 \quad \forall t \in [t', t' + \delta],$$

$$(77) \quad \tilde{x}^r(t) < C_1 \quad \text{and} \quad \|\tilde{q}^r(t)\| < 2\kappa C_1 \quad \forall t \in [\tilde{t}^r, \tilde{t}^r + \delta]$$

for all large $r \in \mathcal{R}_4(\omega)$;

(d) for any $\delta' > 0$, $(\tilde{q}^r(t) - \tilde{x}^r(t)v, t \in [t' + \delta', t' + \delta]) \xrightarrow{u.o.c.} 0$.

(ii) Suppose conditions of (i) hold and, in addition, $\tilde{t}^r = t'$ for all r and $\tilde{q}^r(t') \rightarrow Cv$. Then,

(c') $\tilde{x}(t') = C [= \lim \tilde{x}^r(t')]$;

(d') $(\tilde{q}^r(t) - \tilde{x}^r(t)v, t \in [t', t' + \delta]) \xrightarrow{u.o.c.} 0$.

REMARK. If the CRP condition holds, then condition (74) “automatically” implies (75), with any fixed $\kappa_1 > \kappa$.

LEMMA 8. (i) Suppose the scheduling discipline is MaxWeight. Suppose $\omega \in \Omega_2$ and a subsequence $\mathcal{R}_4(\omega) \subseteq \mathcal{R}_2$ are fixed such that, along this subsequence, (72) holds. Suppose a sequence $\{\tilde{t}^r, r \in \mathcal{R}_4(\omega)\}$ is fixed such that

$$\tilde{t}^r \rightarrow t' \geq 0$$

and

$$\|\tilde{q}^r(\tilde{t}^r)\| \rightarrow C_2 \geq 0.$$

Let C and $\delta > 0$ be fixed such that $C > \kappa^2 C_2$ and

$$(78) \quad \varepsilon = \text{Osc}(\tilde{w}; [t' - 9\delta, t' + 9\delta] \cap \mathbb{R}_+) < C/2.$$

Then,

(a) the following bounds hold:

$$\tilde{x}(t) \leq C_1 \quad \forall t \in [t', t' + \delta]$$

and

$$\tilde{x}^r(t) < C_1 \quad \text{and} \quad \|\tilde{q}^r(t)\| < 2\kappa C_1 \quad \forall t \in [\tilde{t}^r, \tilde{t}^r + \delta]$$

for all large $r \in \mathcal{R}_4(\omega)$,

with $C_1 = \kappa^2 \kappa_1 C + 2\varepsilon$, $\kappa_1 = 2\kappa$;

(b) for any $\delta' > 0$, $(\tilde{q}^{r,(NC)}(t), t \in [t' + \delta', t' + \delta]) \xrightarrow{u.o.c.} 0$.

(ii) Suppose the conditions of (i) hold and, in addition, $\tilde{t}^r = t'$ for all r , and $\tilde{q}^r(t')$ converges to an invariant point, that is, $\tilde{q}^r(t') \rightarrow (C_2/\|v\|)v$. Then we have

(b') $(\tilde{q}^{r,(NC)}(t), t \in [t', t' + \delta]) \xrightarrow{u.o.c.} 0$.

The proof is presented in the Appendix.

REMARK. Since for any $\delta > 0$ there exists a sufficiently large C satisfying (78), we see that Lemma 8(ii) holds for any $\delta > 0$.

12.1. *Proof of Lemma 7.* To establish the asymptotic properties claimed by Lemma 7 for the diffusion-scaled paths in the interval $[\tilde{t}^r, \tilde{t}^r + \delta]$, $\delta > 0$, in this section we will first study the fluid-scaled paths z^r in the corresponding interval $[r\tilde{t}^r, r\tilde{t}^r + r\delta]$. More precisely, we consider the following *family* of fluid-scaled paths restarted at times which are a constant $T > 0$ apart from each other. For each r and for each integer $l \in [0, 2\delta r/T - 1]$, consider the path

$$(79) \quad \bar{z}^{r,l} \doteq \Theta(r\tilde{t}^r + Tl)z^r$$

and let $\bar{x}^{r,l}, \bar{w}^{r,l}, \bar{y}^{r,l}, \bar{q}^{r,l}$ denote the corresponding components of $\bar{z}^{r,l}$. (The choice of constant T will be given later in Lemma 10.) We will focus on the behavior of each path $\bar{z}^{r,l}(\cdot)$ within the interval $[0, T]$. Obviously, if $\tilde{t}^r \rightarrow t'$, for all large r and any integer $l \in [0, 2\delta r/T - 1]$, a time $u \in [0, T]$ for the path $\bar{z}^{r,l}$ corresponds to the time

$$\tilde{t}^r + lT/r + u/r \in [t' - 3\delta, t' + 3\delta] \cap \mathbb{R}_+$$

on the diffusion time scale, that is, that of $\bar{x}^r, \bar{w}^r, \bar{y}^r, \bar{q}^r$. This in particular means that the time T for $\bar{z}^{r,l}$ and time 0 for $\bar{z}^{r,l+1}$ correspond to the same time on the diffusion time scale.

We start with the following simple lemma which shows that sequences (on r) of paths $\bar{z}^{r,l}$, with l possibly dependent on r , have FSPs as their limits.

LEMMA 9. *Suppose $\omega \in \Omega_2$, a subsequence $\mathcal{R}_4(\omega) \subseteq \mathcal{R}_2$ and a corresponding bounded nonnegative sequence $\{\tilde{t}^r, r \in \mathcal{R}_4(\omega)\}$ are fixed. Consider the family of paths $\bar{z}^{r,l}$ defined by (79), associated with sequence $\{\tilde{t}^r\}$ and some constants $T > 0$ and $\delta > 0$. Assume that $c_1 \leq \|\bar{q}^{r,l(r)}(0)\| \leq c_2$ for all $r \in \mathcal{R}_5 \subseteq \mathcal{R}_4(\omega)$, where $l(r) \in [0, 2\delta r/T - 1]$ are integers and $0 \leq c_1 \leq c_2 < \infty$ are constants. Then, there exists a subsequence $\mathcal{R}_f \subseteq \mathcal{R}_5$ along which $\bar{z}^{r,l(r)}$ converges (u.o.c.) to an FSP z with $\|q(0)\| \in [c_1, c_2]$.*

PROOF. Using properties (71) and (69), it is easy to verify that conditions (38) and (39) hold for the sequences $\{\bar{f}^r = \Theta(r\tilde{t}^r + l(r)T)f^r, r \in \mathcal{R}_5\}$ and $\{\bar{g}^r = \Theta(r\tilde{t}^r + l(r)T)g^r, r \in \mathcal{R}_5\}$. Thus, we can apply Lemma 2. \square

The following lemma describes key properties of the family of paths $\{\bar{z}^{r,l}\}$, which will imply Lemma 7 almost directly.

LEMMA 10. (i) *Suppose the conditions of Lemma 7(i) hold. Let $C_1 = \kappa^2\kappa_1C + 2\varepsilon$. Then, for any sufficiently small $\varepsilon_2 > 0$, there exists $T > 0$ such that, for all sufficiently large r , the following properties (80)–(82) hold for integer $l \in [1, 2\delta r/T - 1]$ and the properties (83)–(84) hold for integer*

$l \in [0, 2\delta r/T - 1]$:

$$(80) \quad G(\bar{q}^{r,l}(u)) < 2\varepsilon_2 \quad \text{for } u = 0 \text{ and } u = T,$$

$$(81) \quad G(\bar{q}^{r,l}(u)) < 3\varepsilon_2 \quad \text{for all } u \in [0, T],$$

$$(82) \quad \bar{y}^{r,l}(u) \equiv \bar{y}^{r,l}(u) - \bar{y}^{r,l}(0) = 0 \quad \text{for all } u \in [0, T],$$

$$(83) \quad C - 2\varepsilon < \bar{x}^{r,l}(u) < C_1 \quad \text{for all } u \in [0, T],$$

$$(84) \quad (C - 2\varepsilon)/\kappa < \|\bar{q}^{r,l}(u)\| < 2\kappa C_1 \quad \text{for all } u \in [0, T].$$

(ii) Suppose the conditions of Lemma 7(ii) hold. Let $C_1 = C + 2\varepsilon$. Then, for any sufficiently small $\varepsilon_2 > 0$, there exists $T > 0$ such that, for all sufficiently large r , properties (80)–(84) hold for all integer $l \in [0, 2\delta r/T - 1]$ (including 0).

PROOF. (i) Suppose $\varepsilon_2 > 0$ is sufficiently small so that conditions $q \neq 0$ and $G(q) \leq 3\varepsilon_2$ imply [by (51), (52) and Lemma 4] that $\|q\| < 2\|q^{(C)}\| < 2\kappa(\zeta \cdot q)$, $q_n > \varepsilon_3(\zeta \cdot q)$ for all $n \in \bar{N}^{(C)}$ and some fixed $\varepsilon_3 > 0$, and $K^*(m, \gamma \times q^\beta) \subseteq K^*(m, \zeta)$ for all m .

According to Theorem 2, we can choose a constant $T_2 \geq 0$ (depending only on ε_2) such that, for any FSP with $q(0) \neq 0$,

$$G(q(t)) \leq \varepsilon_2 \quad \forall t \geq \|q(0)\|T_2.$$

Then we choose arbitrary

$$T > 2\kappa C_1 T_2.$$

(As the formulation of the lemma suggests, we will prove that C_1 is the upper bound on the workload $\bar{x}^{r,l}(u)$ for all $l \in [0, 2\delta r/T - 1]$ and all $u \in [0, T]$. Then, κC_1 and $2\kappa C_1$ will be the upper bounds on $\|\bar{q}^{r,(C),l}(u)\|$ and $\|\bar{q}^{r,l}(u)\|$, resp.) Our choice of T ensures that any FSP with the initial state norm $\|q(0)\| \leq 2\kappa C_1$ is such that $G(q(T)) \leq \varepsilon_2$.

The basic idea of the proof is as follows. First, we show that in $[0, T]$ both $\bar{x}^{r,0}(u)$ and $\|\bar{q}^{r,0}(u)\|$ remain bounded above and away from 0, and $\bar{q}^{r,0}(T)$ must be close to the invariant manifold. Second, we show that, for all integer $1 \leq l \leq 2\delta r/T - 1$, queue length vector $\bar{q}^{r,l}(u)$ stays close to the invariant manifold in $[0, T]$ (in addition to its norm being bounded above and away from 0), which implies that $\bar{y}^{r,l}(u)$ cannot increase (i.e., workload service rate cannot be “wasted”) in $[0, T]$.

Let $l = 0$. For all large r , we have the following upper bound:

$$(85) \quad \limsup_{r \rightarrow \infty} \sup_{[0, T]} \|\bar{q}^{r,0}(u)\| < \kappa \limsup_{r \rightarrow \infty} \|\bar{q}^{r,0}(0)\| < \kappa \kappa_1 C.$$

Indeed, if the left-hand side inequality would not hold, then using Lemma 9 we would be able to choose a subsequence of paths $\bar{z}^{r,0}$ converging to an FSP z with $q(0) \neq 0$ and the norm $\|q(u)\|$ increasing to at least $\kappa\|q(0)\|$ at some

time $u \in [0, T]$, which is not possible according to Theorem 2. Bound (85) automatically implies

$$(86) \quad \limsup_{r \rightarrow \infty} \sup_{[0, T]} \bar{x}^{r,0}(u) < \kappa^2 \kappa_1 C.$$

Similarly [using Lemma 9 and the fact that, for any FSP, the workload $x(\cdot)$ is nondecreasing] we obtain the lower bound:

$$(87) \quad \liminf_{r \rightarrow \infty} \inf_{[0, T]} \bar{x}^{r,0}(u) \geq C.$$

Finally, for all large r , we must have

$$(88) \quad G(\bar{q}^{r,0}(T)) = G(\bar{q}^{r,1}(0)) < 2\varepsilon_2.$$

If the inequality above would not hold, then using Lemma 9 and the facts that

$$\begin{aligned} \limsup \|\bar{q}^{r,0}(0)\| &< \kappa \kappa_1 C < 2\kappa C_1, \\ \liminf \|\bar{q}^{r,0}(0)\| &> C/\kappa, \end{aligned}$$

we would be able to construct an FSP z such that $C/\kappa \leq \|q(0)\| \leq 2\kappa C_1$ and $G(q(T)) \geq 2\varepsilon_2$, which is impossible due to our choice of T .

Now consider the behavior of $\bar{z}^{r,l}$ for $l \geq 1$. Suppose properties (80)–(84) do not hold. Then, there exists a subsequence $\mathcal{R}_5 \subseteq \mathcal{R}_4(\omega)$ and a corresponding sequence of integers $l' = l'(r) \in [1, 2\delta r/T - 1]$ such that (80)–(84) hold for all $1 \leq l \leq l' - 1$, but at least one of the properties (80)–(84) does not hold for $l = l'$. (The case $l' = 1$ is possible.) This construction and (88) imply that (80) must hold for $l = l'$ and $u = 0$, namely, $G(\bar{q}^{r,l'}(0)) < 2\varepsilon_2$. Also by construction [and (87)] we have

$$\bar{x}^{r,l'}(0) > C - 2\varepsilon$$

and

$$(C - 2\varepsilon)/\kappa < \|\bar{q}^{r,l'}(0)\| < 2\kappa C_1.$$

This implies that, for all large r (along \mathcal{R}_5), the properties (80) and (81) hold for $l = l'$ and, in addition, we have

$$(89) \quad \liminf_r \inf_{u \in [0, T]} \bar{x}^{r,l'}(u) \geq C - 2\varepsilon > 0.$$

[Otherwise, using Lemma 9, we would be able to construct an FSP z with $x(0) \geq C - 2\varepsilon$ and $\|q(0)\| \leq 2\kappa C_1$, violating either the property that $x(t)$ is nondecreasing, or the property that $G(q(t))$ is nonincreasing, or the property that $G(q(t)) \leq \varepsilon_2$ for all $t \geq T$.]

Properties (81) and (89) imply that (82) must hold for $l = l'$ for all large r . Indeed, they (and our choice of ε_2) imply that workload $\bar{x}^{r,l'}(u)$ and all critical queue lengths $\bar{q}_n^{r,l'}(u)$ stay bounded away from 0 in $[0, T]$. Then, (81) (along with our choice of ε_2) implies that, for any $u \in (0, T]$, when switch is in a state m , only

decisions from the corresponding subset $K^*(m, \zeta)$ can be chosen; this [and the fact that critical queue lengths $\bar{q}_n^{r,l'}(u)$ are bounded away from 0] means that $\bar{y}^{r,l'}(\cdot)$ cannot increase in $(0, T]$, which is equivalent to (82) with $l = l'$.

Now, let us show that (83) and (84) must hold for $l = l'$ for all large r (along \mathcal{R}_5). Given that (82) holds for all $1 \leq l \leq l'$, we can write

$$\begin{aligned}
 (90) \quad \bar{x}^{r,l'}(u) &= \bar{x}^{r,0}(T) + \sum_{l=1}^{l'-1} [\bar{x}^{r,l}(T) - \bar{x}^{r,l}(0)] + \bar{x}^{r,l'}(u) - \bar{x}^{r,l'}(0) \\
 &= \bar{x}^{r,0}(T) + \sum_{l=1}^{l'-1} [\bar{w}^{r,l}(T) - \bar{w}^{r,l}(0)] + \bar{w}^{r,l'}(u) - \bar{w}^{r,l'}(0) \\
 &= \bar{x}^{r,0}(T) + \bar{w}^r(\tilde{t}^r + l'T/r + u/r) - \bar{w}^r(\tilde{t}^r + T/r),
 \end{aligned}$$

where

$$|\bar{w}^r(\tilde{t}^r + l'T/r + u/r) - \bar{w}^r(\tilde{t}^r + T/r)| < 2\varepsilon$$

for all large r due to our choice of ε and u.o.c. convergence $\bar{w}^r \rightarrow \bar{w}$. Therefore, if we recall (86) and (87), we obtain (83) for $l = l'$. The upper bound in property (84) (for $l = l'$ and large r) follows from (83) and (81) (and our choice of ε_2), and the lower bound follows from that in (83) automatically (by our choice of κ).

Thus, we have proved that, for all large r along \mathcal{R}_5 , all properties (80)–(84) must hold with $l = l'$. This is a contradiction with the definition of the sequence $\{l'(r), r \in \mathcal{R}_5\}$. The proof of statement (i) is complete.

(ii) The proof of this statement is a simplified version of the proof of (i) in that we do not need a “special treatment” of the case $l = 0$. Namely, the convergence $\lim_r \bar{x}^{r,0}(0) = C$ is employed in place of the estimates (86) and (87), and $G(\bar{q}^{r,0}(0)) < 2\varepsilon_2$ in place of (88). Then, (80)–(84) are proved by contradiction, by constructing $\mathcal{R}_5 \subseteq \mathcal{R}_4$ and the sequence $\{l'(r)\}$ the same way as in the proof of (i) except we allow $l'(r) \geq 0$ [as opposed to $l'(r) \geq 1$]. Finally, we use

$$\bar{x}^{r,l'}(u) = \bar{x}^{r,0}(0) + \sum_{l=0}^{l'-1} [\bar{x}^{r,l}(T) - \bar{x}^{r,l}(0)] + \bar{x}^{r,l'}(u) - \bar{x}^{r,l'}(0)$$

in place of (90). We omit further details. \square

PROOF OF LEMMA 7. (i) Let us choose a small $\varepsilon_2 > 0$ and a corresponding T as in Lemma 10(i). If we recall that time $u \in [0, T]$ for $\bar{z}^{r,l}$ corresponds to the time $\tilde{t}^r + lT/r + u/r$ on the diffusion time scale, we see that Lemma 10 implies that, for all large r , we have

$$(91) \quad C - 2\varepsilon < \bar{x}^r(t) < C_1 \quad \text{for all } t \in [\tilde{t}^r, \tilde{t}^r + (3/2)\delta],$$

$$(92) \quad \|\tilde{q}^r(t)\| < 2\kappa C_1 \quad \text{for all } t \in [\tilde{t}^r, \tilde{t}^r + (3/2)\delta],$$

$$(93) \quad \tilde{y}^r(t) - \tilde{y}^r(\tilde{t}^r + T/r) = 0 \quad \text{for all } t \in [\tilde{t}^r + T/r, \tilde{t}^r + (3/2)\delta],$$

$$(94) \quad G(\tilde{q}^r(t)) < 3\varepsilon_2 \quad \text{for all } t \in [\tilde{t}^r + T/r, \tilde{t}^r + (3/2)\delta].$$

Since $\tilde{y}^r \Rightarrow \tilde{y}$, $\tilde{x}^r \Rightarrow \tilde{x}$, and both \tilde{y} and \tilde{x} are RCLL, the properties (91)–(93) easily imply statements (a)–(c). We know from (52) that $\|p - (\zeta \cdot p)v\| / \|(\zeta \cdot p)v\|$ is uniformly small for nonzero vectors $p \in \mathbb{R}_+^N$ with small $G(p)$. Therefore, (94) and (91), along with the fact that ε_2 can be chosen arbitrarily small, imply statement (d).

(ii) We choose a small $\varepsilon_2 > 0$ and a corresponding T as in Lemma 10(ii). Then, (c') follows from right-continuity of \tilde{x} , estimate (83) (with $C_1 = C + 2\varepsilon$ and $\tilde{t}^r = t'$), and the fact that ε in (83) can be made arbitrarily small by choosing a sufficiently small $\delta > 0$. To prove (d'), we notice that Lemma 10(ii) implies that both (94) and (91) hold for all $t \in [t', t' + (3/2)\delta]$. Then, analogously to (d), property (d') follows from (94) and (91). \square

12.2. *Proof of Theorem 1(i).* To prove the convergences (20) and (21), it will suffice to prove the following statement:

As $r \rightarrow \infty$ along \mathcal{R}_2 , for any $\omega \in \Omega_2$ (and, therefore, with probability 1), we have the following convergences:

$$(95) \quad (\tilde{y}^r(t), t \geq 0) \xrightarrow{u.o.c.} (\tilde{y}^\circ(t), t \geq 0),$$

$$(96) \quad (\tilde{q}^r(t), t \geq 0) \xrightarrow{u.o.c.} (\tilde{q}^\circ(t), t \geq 0),$$

where \tilde{y}° is defined by (19), $\tilde{q}^\circ = \tilde{x}^\circ v$, and $\tilde{x}^\circ = \tilde{w} + \tilde{y}^\circ$.

Let us prove (95) and (96). Consider arbitrary fixed $\omega \in \Omega_2$. As explained earlier, for an arbitrary subsequence $\mathcal{R}_3(\omega) \subseteq \mathcal{R}_2$, there exists another subsequence $\mathcal{R}_4(\omega) \subseteq \mathcal{R}_3(\omega)$ such that the convergence (72), namely, $\tilde{y}^r \Rightarrow \tilde{y}$, holds along $\mathcal{R}_4(\omega)$. Then the proof of (95) and (96) will be complete if we can prove the following claims [for the chosen ω , with $r \rightarrow \infty$ along $\mathcal{R}_4(\omega)$].

STEP 1. The limit function \tilde{y} is finite everywhere in $[0, \infty)$.

STEP 2. The function \tilde{y} is continuous and $\tilde{y}(0) = 0$.

STEP 3. If $\tilde{x}(t) > 0$, then t is *not* a point of increase of \tilde{y} .

STEP 4. $\tilde{y} = \tilde{y}^\circ$.

STEP 5. Convergence (95) holds.

STEP 6. Convergence (96) holds.

In this proof, we will use the convention that functions \tilde{x} , \tilde{w} and \tilde{y} are defined additionally at $0-$ as follows: $\tilde{y}(0-) = 0$, $\tilde{w}(0-) = \tilde{x}(0-) = \tilde{w}(0) [= \tilde{x}^\circ(0)]$.

Using this convention, the case $\tilde{y}(0) > 0$ will be viewed as a discontinuity of \tilde{y} (and \tilde{x}) at 0.

We notice that by Lemma 8(ii) and the remark after the lemma, we have

$$(97) \quad (\tilde{q}^{r,(NC)}(t), t \in [0, \infty)) \xrightarrow{\text{u.o.c.}} 0.$$

It follows from (97) and the basic relation $\|\tilde{q}^{r,(C)}(t)\| < \kappa \tilde{x}^r(t)$ that condition (75) with any $\kappa_1 > \kappa$ holds for any sequence $\{\tilde{t}^r, r \in \mathcal{R}'_4 \subseteq \mathcal{R}_4(\omega)\}$, satisfying (73) and (74). For the rest of the proof of Theorem 1(i), we fix an arbitrary $\kappa_1 > \kappa$.

PROOF OF STEP 1. Suppose the statement does not hold. Denote $t^* = \inf\{t \geq 0 \mid \tilde{y}(t) = \infty\}$. The inf is attained because \tilde{y} is RCLL.

Let us fix $\delta > 0$, and denote $\varepsilon^* = \text{Osc}(\tilde{w}; [t^* - 4\delta, t^* + 4\delta] \cap \mathbb{R}_+)$. If $t^* > 0$, let us choose $\Delta \in (0, \delta \wedge t^*)$ and a large C such that $C > \tilde{x}(t^* - \Delta) + 2\varepsilon^*$; otherwise, if $t^* = 0$, we choose $\Delta = 0$ and a large $C > \tilde{x}(0-) + 2\varepsilon^*$. We define

$$\tilde{t}^r = \min\{t \geq t^* - \Delta \mid \tilde{x}^r(t) \geq C\}.$$

Since $\tilde{y}(t) = \infty$ for all $t \geq t^*$, we see (using the facts that $\tilde{y}^r \Rightarrow \tilde{y}$ and \tilde{y} and all \tilde{y}^r are RCLL) that $\limsup \tilde{t}^r \leq t^*$. Our choice of Δ and C also implies $\limsup \tilde{x}^r(t^* - \Delta) < C$ (again, using properties of \tilde{y} and all \tilde{y}^r), which in turn means that $\tilde{t}^r > t^* - \Delta$ for all large r . We see that (for large r) at time \tilde{t}^r , the value of $\tilde{x}(\cdot)$ crosses level C by a positive jump upper bounded as follows:

$$\tilde{x}^r(\tilde{t}^r) - \tilde{x}^r(\tilde{t}^r -) \leq \zeta \cdot [\tilde{f}^r(\tilde{t}^r) - \tilde{f}^r(\tilde{t}^r -)].$$

Then, we must have $\tilde{x}^r(\tilde{t}^r) \rightarrow C$ because, as $r \rightarrow \infty$, by property (71), the jump sizes of all functions $\tilde{f}_n^r(t)$ become arbitrarily small uniformly on compact sets. Let us choose a further subsequence $\mathcal{R}'_4 \subseteq \mathcal{R}_4(\omega)$ along which

$$\tilde{t}^r \rightarrow t' \in [t^* - \Delta, t^*].$$

Note that $[t' - 3\delta, t' + 3\delta] \subset [t^* - 4\delta, t^* + 4\delta]$. The conditions of Lemma 7(i) are satisfied for C , δ and $\{\tilde{t}^r, r \in \mathcal{R}'_4\}$. Therefore, \tilde{y} is finite in $[0, t' + \delta]$ —a contradiction, since $t' + \delta > t^*$. Step 1 has been proved. \square

PROOF OF STEP 2. Suppose the statement does not hold. The contradiction is obtained very similarly to the way it is done in the proof of Step 1. Let t^* be a discontinuity point, that is, $\tilde{y}(t^* -) < \tilde{y}(t^*)$. (The case $t^* = 0$ is not excluded.) Since $\tilde{x} = \tilde{w} + \tilde{y}$ and \tilde{w} is continuous, $\tilde{x}(t^*) - \tilde{x}(t^* -) = \tilde{y}(t^*) - \tilde{y}(t^* -)$. There are two possible cases: (a) $\tilde{x}(t^* -) > 0$ and (b) $\tilde{x}(t^* -) = 0$.

CASE (a). In this case we must have $t^* > 0$. [Indeed, if $\tilde{w}(0) = \tilde{x}^\circ(0) > 0$, then, by Lemma 7(ii) with $\tilde{t}^r \equiv t' = 0$ and sufficiently small $\delta > 0$, we have $\tilde{x}(0) = \tilde{x}^\circ(0) = \tilde{w}(0)$, which implies that \tilde{y} has no jump at 0. If $\tilde{w}(0) = 0$, then $\tilde{x}(0-) = 0$.] We can always fix a small $\delta \in (0, t^*)$ and small $\Delta \in (0, \delta)$,

such that $t' = t^* - \Delta > 0$ is a point of continuity of \tilde{y} (and \tilde{x}) and $\varepsilon^* = \text{Osc}(\tilde{w}; [t^* - 4\delta, t^* + 4\delta] \cap \mathbb{R}_+) < \tilde{x}(t')/2 = C/2$. We have convergence $\tilde{x}^r(t') \rightarrow \tilde{x}(t') = C$ (since \tilde{x} is continuous at t'). The conditions of Lemma 7(i) are satisfied for C , δ and $\tilde{t}^r \equiv t'$. Therefore, \tilde{y} cannot increase in the interval $(t', t' + \delta)$ which contains t^* . So, \tilde{y} cannot have a jump at t^* .

CASE (b). In this case, let us fix a small $C > 0$ and then a sufficiently small $\delta > 0$ so that

$$C_1^* = \kappa^2 \kappa_1 C + 2\varepsilon^* < \tilde{x}(t^*) \quad \text{and} \quad 2\varepsilon^* < C,$$

where $\varepsilon^* = \text{Osc}(\tilde{w}; [t^* - 4\delta, t^* + 4\delta] \cap \mathbb{R}_+)$. If $t^* > 0$, we fix a small $\Delta \in (0, \delta \wedge t^*)$ such that $\tilde{x}(t^* - \Delta) < C$; otherwise, if $t^* = 0$, we set $\Delta = 0$. Analogously to the way it was done in the proof of Step 1, we define $\tilde{t}^r = \min\{t \geq t^* - \Delta \mid \tilde{x}^r(t) \geq C\}$ and observe that we can choose a further subsequence $\mathcal{R}'_4 \subseteq \mathcal{R}_4(\omega)$ along which $\tilde{t}^r \rightarrow t' \in [t^* - \Delta, t^*]$. The conditions of Lemma 7(i) are satisfied for C , δ and $\{\tilde{t}^r, r \in \mathcal{R}'_4\}$, and so we must have

$$\tilde{x}(t) < C_1 = \kappa^2 \kappa_1 C + 2\text{Osc}(\tilde{w}; [t' - 3\delta, t' + 3\delta] \cap \mathbb{R}_+) \leq C_1^*$$

for all $t \in [t', t' + \delta]$, which is impossible since $C_1^* < \tilde{x}(t^*)$ and $t^* \in [t', t' + \delta]$. Step 2 has been proved. \square

PROOF OF STEP 3. Let $t^* \geq 0$ be such that $\tilde{x}(t^*) > 0$. If $t^* = 0$, then the fact that \tilde{y} does not increase in some interval $(0, \delta]$ follows from Lemma 7(i) with $\tilde{t}^r \equiv t' = 0$, if we choose δ sufficiently small. If $t^* > 0$, then precisely the same construction as in the proof of Step 2(a) shows that \tilde{y} does not increase in some interval $(t', t' + \delta)$ containing t^* . Step 3 has been proved. \square

PROOF OF STEP 4. This follows from the statements of Steps 2 and 3, and Proposition 4 (in the Appendix). \square

PROOF OF STEP 5. This follows from $\tilde{y}^r \Rightarrow \tilde{y}^\circ$, along with the facts that \tilde{y}° is continuous nondecreasing with $\tilde{y}^\circ(0) = 0$, and \tilde{y}^r is nondecreasing with $\tilde{y}^r(0) = 0$. \square

PROOF OF STEP 6. It suffices to show that, for any $t^* \geq 0$ and any $\varepsilon > 0$, there exists $\delta > 0$, such that

$$(98) \quad \limsup_{r \rightarrow \infty} \sup_{\xi \in [t^* - \delta, t^* + \delta] \cap \mathbb{R}_+} \|\tilde{q}^r(\xi) - \tilde{q}^\circ(\xi)\| < \varepsilon.$$

(The u.o.c. convergence will then follow from the Heine–Borel lemma.) Since $\tilde{q}^\circ(t) \equiv \tilde{x}^\circ(t)\nu$ and, by (95), $\tilde{x}^r \rightarrow \tilde{x}^\circ$ (u.o.c.), to prove (98) it suffices to prove

$$(99) \quad \limsup_{r \rightarrow \infty} \sup_{\xi \in [t^* - \delta, t^* + \delta] \cap \mathbb{R}_+} \|\tilde{q}^r(\xi) - \tilde{x}^r(\xi)\nu\| < \varepsilon.$$

If $\tilde{x}^\circ(t^*) = 0$, then property (99) must hold because both functions $\|\tilde{q}^r\|$ and \tilde{x}^r (for large r) are bounded by an arbitrarily small constant in a sufficiently small neighborhood of t^* . [This is implied by the following facts: $\tilde{x}^r \rightarrow \tilde{x}^\circ$ (u.o.c.), $\|\tilde{q}^{r,(C)}(t)\| \leq \kappa \tilde{x}^r(t)$ and (97).] If $\tilde{x}^\circ(t^*) > 0$ and $t^* = 0$, then (99) follows from Lemma 7(ii) with $\tilde{t}^r \equiv t' = 0$ and sufficiently small $\delta > 0$. If $\tilde{x}^\circ(t^*) > 0$ and $t^* > 0$, then to obtain (99) we can repeat the construction of the proof of Step 2(a) and then apply Lemma 7(i). Step 6 has been proved, and this completes the proof of Theorem 1(i). \square

12.3. *Proof of Theorem 1(ii).* The probability space (Ω, \mathcal{F}, P) is the one constructed earlier, same as in the proof of Theorem 1(i). Consider arbitrary scheduling discipline G . For arbitrary fixed $\omega \in \Omega$, consider paths of \tilde{x}_G^r , \tilde{y}_G^r and \tilde{w}_G^r , corresponding to discipline G . Since these paths are constructed on the same probability space as the paths corresponding to the MaxWeight discipline, we have $\tilde{w}_G^r = \tilde{w}^r$ and therefore $\tilde{w}_G^r \rightarrow \tilde{w}_G = \tilde{w}$ u.o.c.

We claim that, along sequence \mathcal{R} ,

$$(100) \quad \liminf_{r \rightarrow \infty} \tilde{x}_G^r(t) \geq \tilde{x}^\circ(t), \quad t \geq 0.$$

Indeed, as explained earlier, for any subsequence $\mathcal{R}_3(\omega) \subseteq \mathcal{R}$, we can choose a further subsequence $\mathcal{R}_4(\omega) \subseteq \mathcal{R}_3(\omega)$ such that $\tilde{y}_G^r \Rightarrow \tilde{y}_G$, where \tilde{y}_G is some nondecreasing nonnegative function in $D([0, \infty), \bar{\mathbb{R}})$. Therefore, for any $t > 0$ where $\tilde{y}_G(\cdot)$ is continuous, as $r \rightarrow \infty$ along $\mathcal{R}_4(\omega)$,

$$\lim \tilde{x}_G^r(t) = \tilde{w}(t) + \tilde{y}_G(t),$$

and $\tilde{w}(t) + \tilde{y}_G(t) \geq 0$ [since $\tilde{x}_G^r(t) \geq 0$]. Then, $\tilde{w}(t) + \tilde{y}_G(t) \geq 0$ for all $t \geq 0$ (by right-continuity). By Proposition 4 (in the Appendix), $\tilde{y}_G(t) \geq \tilde{y}^\circ(t)$ for all $t \geq 0$. This and the continuity of \tilde{y}° implies that, for any $t \geq 0$, (100) holds along the subsequence $\mathcal{R}_4(\omega)$, and therefore along \mathcal{R} [since the subsequence $\mathcal{R}_3(\omega)$ can be arbitrary]. The pathwise lower bound (100) implies (22).

13. Conclusions. The main conclusion of this work is that, even for quite general queueing systems, allowing flexible allocation of service resources and randomness of the service environment, the workload minimization property (and additional optimality properties, like holding cost minimization) in heavy traffic can be achieved by parsimonious dynamic (“on-line”) rules, not requiring information on the mean input rates, and not requiring any “precomputation” of rule parameters or any “preallocation” of the service resources. Our analysis is for the generalized switch model and we prove that the MaxWeight rule possesses such “nice” properties. We believe that main parts of this analysis are quite general, in particular the analysis of convergence properties of fluid sample paths, which relies on a Lyapunov function (Ψ in our case), a key differential inclusion [(57) in our case] showing that the derivative of this Lyapunov function is minimized at

all times [as in (58)], and geometry of the rate region. As a result, the approach and techniques used in this paper may be useful in analysis of MaxWeight-type algorithms for other models, in particular those for which MaxWeight-type algorithms have already been defined in the previous work such as [29, 31, 32].

Finally, we believe that the MaxWeight discipline (or other parsimonious dynamic disciplines having similar optimality properties) is very attractive in applications. The RP condition, which is required for the optimality properties to hold, is common in applications.

APPENDIX

A.1. Linear programming characterization of the RP condition. As we mentioned before, a linear programming characterization can be done along the lines of [14, 16, 20, 21, 36], where such characterization was in fact used as a definition of the heavy traffic regime and resource pooling.

Suppose a vector of mean rates $\lambda \in \mathbb{R}_{++}^N$ is given. The first step is to define the system *load* ρ . In the work cited above, the load ρ is defined as the minimum possible upper bound on the utilization of individual servers, such that the average service rates λ can be achieved. This ρ is determined by a linear program. In our case there are no individual servers (resources). Consequently, a natural definition of load, which is consistent with the previous work (i.e., could be used in the previous work and produce the same answer), is $\rho = 1/c^*$, where c^* is the maximum factor by which vector λ can be scaled until the boundary of the rate region is hit. The linear program to determine c^* is as follows. We will refer to it as the *primal problem*:

$$(101) \quad \max_{c, \phi} c$$

subject to

$$(102) \quad \sum_{m \in \bar{M}} \pi_m \sum_{k \in K(m)} \phi_{mk} \mu_n^m(k) \geq c \lambda_n, \quad n \in \bar{N},$$

$$(103) \quad \sum_{k \in K(m)} \phi_{mk} = 1, \quad m \in \bar{M},$$

$$(104) \quad \phi_{mk} \geq 0, \quad m \in \bar{M}, k \in K(m),$$

$$(105) \quad c \in \mathbb{R}.$$

Note that the set of constraints (102) is equivalent to the condition $v(\phi) \geq c\lambda$.

Let (ϕ^*, c^*) be any fixed optimal solution of this linear program. (We must have $c^* > 0$, since set \bar{V} contains elements with all positive components.) Obviously, $v^* = c^*\lambda \in \bar{V}^*$ is the point of the boundary of \bar{V} which we “hit” by scaling vector λ . Therefore, the condition $\rho = 1/c^* = 1$ is equivalent to the condition that λ is on the rate region boundary, that is, $\lambda = v^*$.

Theorem 3 shows how the RP condition for v^* and the corresponding outer normal vector ζ are characterized in terms of the following linear program dual to (101)–(105):

$$(106) \quad \min_{\zeta, \alpha} \sum_{m \in \bar{M}} \alpha_m,$$

where $\zeta = (\zeta_1, \dots, \zeta_N)$ and $\alpha = (\alpha_m, m \in \bar{M})$ are subject to constraints

$$(107) \quad \sum_{n \in \bar{N}} \zeta_n \lambda_n = 1,$$

$$(108) \quad \alpha_m \geq \sum_{n \in N} \zeta_n \pi_m \mu_n^m(k), \quad m \in \bar{M}, k \in K(m),$$

$$(109) \quad \zeta_n \geq 0, \quad n \in \bar{N},$$

$$(110) \quad \alpha_m \in \mathbb{R}, \quad m \in \bar{M}.$$

(Basic linear programming facts we use in this section can be found, e.g., in Sections 5.1–5.4 of [19].)

THEOREM 3. *Consider vector $v^* \in \bar{V}^*$ defined as above. Then, vector $\zeta^* \in \mathbb{R}_+^N$ is part of an optimal solution (ζ^*, α^*) of the dual linear program (106)–(110) if and only if ζ^* is an outer normal vector to \bar{V} at point v^* and $\zeta^* \cdot \lambda = 1$.*

Theorem 3 immediately implies that the RP condition for the vector v^* is equivalent to the uniqueness of the vector of dual variables ζ^* across all optimal solutions (ζ^*, α^*) to the dual linear program (106)–(110). And if ζ^* is unique, it is (up to a scaling by a positive constant) the normal vector ζ associated with the RP condition. [Also, as a byproduct of the proof of Theorem 3, it is easy to see that the uniqueness of ζ^* is in fact equivalent to the uniqueness of the entire optimal solution (ζ^*, α^*) .]

PROOF OF THEOREM 3. First, recall that (ϕ^*, c^*) is a fixed optimal solution to the primal problem, $v^* = c^* \lambda$, and therefore the set of constraints (102) for (ϕ^*, c^*) can be compactly written as $v(\phi^*) \geq v^*$. Also, note that the Lagrangian form for the pair of primal problem (101)–(105) and dual problem (106)–(110) is

$$\begin{aligned} \mathcal{L}(\phi, c; \zeta, \alpha) &= c + \sum_{n \in \bar{N}} \zeta_n \left[\sum_{m \in \bar{M}} \pi_m \sum_{k \in K(m)} \phi_{mk} \mu_n^m(k) - c \lambda_n \right] \\ &\quad + \sum_{m \in \bar{M}} \alpha_m \left(1 - \sum_{k \in K(m)} \phi_{mk} \right) \\ &= c + \zeta \cdot [v(\phi) - c \lambda] + \sum_{m \in \bar{M}} \alpha_m \left(1 - \sum_{k \in K(m)} \phi_{mk} \right), \end{aligned}$$

with the domain defined by (104), (105), (109) and (110).

Consider a fixed optimal solution (ζ^*, α^*) of the dual problem (106)–(110). We have $\zeta^* \in \mathbb{R}_+^N$ [by (109)] and $\zeta^* \cdot \lambda = 1$ [by (107)]. Also, by the complementary slackness conditions, we have

$$(111) \quad \zeta^* \cdot [v(\phi^*) - v^*] = 0.$$

Let us prove that ζ^* is normal to \bar{V} in point v^* .

We know that ϕ^* must maximize

$$\mathcal{L}(\phi, c^*; \zeta^*, \alpha^*) = \zeta^* \cdot v(\phi) + c^*[1 - \zeta^* \cdot \lambda] + \sum_{m \in \bar{M}} \alpha_m^* \left(1 - \sum_{k \in K(m)} \phi_{mk} \right)$$

over all ϕ satisfying (104), and in particular over those satisfying in addition (103) (because ϕ^* itself satisfies it). We see that ϕ^* maximizes $\zeta^* \cdot v(\phi)$ over all ϕ satisfying (104) and (103), which implies that $v(\phi^*)$ maximizes $\zeta^* \cdot v$ over all $v \in \bar{V}$. But, according to (111), $\zeta^* \cdot v^* = \zeta^* \cdot v(\phi^*)$, so we have

$$(112) \quad v^* \in \arg \max_{v \in \bar{V}} \zeta^* \cdot v,$$

which means that ζ^* is normal to \bar{V} in point v^* .

To prove the converse, consider an arbitrary outer normal vector ζ^* to \bar{V} in point v^* [which means that (112) holds], such that $\zeta^* \cdot \lambda = 1$. (Necessarily, $\zeta^* \in \mathbb{R}_+^N$, since $v^* \in \bar{V}$.) Since $v(\phi^*) \geq v^*$ and $v(\phi^*) \in \bar{V}$, it follows from (112) that $\zeta^* \cdot v(\phi^*) = \zeta^* \cdot v^*$ and, therefore,

$$(113) \quad v(\phi^*) \in \arg \max_{v \in \bar{V}} \zeta^* \cdot v.$$

In turn, (113) means that ϕ^* maximizes

$$\zeta^* \cdot v(\phi)$$

over ϕ satisfying (104) and (103), and, moreover, the maximum value of the objective (in the last display) is c^* , because $\zeta^* \cdot v(\phi^*) = c^* + \zeta^* \cdot [v(\phi^*) - c^* \lambda]$, $c^* \lambda = v^*$, and $\zeta^* \cdot v(\phi^*) = \zeta^* \cdot v^*$. This is equivalent to the fact that ϕ^* is an optimal solution of the linear program

$$(114) \quad \max_{\phi} \sum_{n \in \bar{N}} \zeta_n^* \left[\sum_{m \in \bar{M}} \pi_m \sum_{k \in K(m)} \phi_{mk} \mu_n^m(k) \right]$$

subject to (103) and (104), and the maximum in (114), attained with $\phi = \phi^*$, is c^* . The linear program dual to (114), (103), (104) is

$$(115) \quad \min_{\alpha} \sum_{m \in \bar{M}} \alpha_m,$$

subject to

$$(116) \quad \alpha_m \geq \sum_{n \in \bar{N}} \zeta_n^* \pi_m \mu_n^m(k), \quad m \in \bar{M}, k \in K(m),$$

$$(117) \quad \alpha_m \in \mathbb{R}, \quad m \in \bar{M}.$$

Let us fix arbitrary optimal solution α^* of the problem (115)–(117). Then (by duality) $\sum_{m \in \bar{M}} \alpha_m^* = c^*$. Also, we see that the pair (ζ^*, α^*) satisfies all the constraints of the problem (106)–(110). Finally, since c^* is the optimal objective function value of the primal problem (101)–(105), then—again by duality— (ζ^*, α^*) is an optimal solution of (106)–(110). \square

A.2. Proof of property (71). Since our input processes are i.i.d., according to Bramson’s weak law estimate ([4], Proposition 4.2), for any $T_3 > 0$, any $\varepsilon > 0$ and any $n \in \bar{N}$, for all large r , we have the following estimate, uniformly on integer l , $0 \leq l \leq T_3 r$:

$$P \left\{ \max_{\xi \in \Xi_l} |r(f_n^r(l + \xi) - f_n^r(l)) - \lambda_n \chi(\xi)| \geq \varepsilon \bar{r} \right\} \leq \varepsilon / \bar{r},$$

where Ξ_l is the (finite) subset of $\xi \in (0, 1]$ such that $r(l + \xi)$ is an integer [or, equivalently, where $f_n^r(l + \xi)$ may jump], $\bar{r} = \lceil r \rceil$, and $\chi(\xi) \leq \bar{r}$ is the cardinality of $(0, \xi] \cap \Xi_l$. If we rewrite this as

$$P \left\{ \max_{\xi \in \Xi_l} |f_n^r(l + \xi) - f_n^r(l) - \lambda_n \chi(\xi) / r| \geq \varepsilon \bar{r} / r \right\} \leq \varepsilon / \bar{r},$$

notice that the function $(\chi(\xi) / r, \xi \in [0, 1])$ uniformly converges to the identity function $(\xi, \xi \in [0, 1])$, and recall that $f_n^r(l + \xi) - f_n^r(l)$ can only jump at points Ξ_l , then by rechoosing ε we easily obtain the following extension of the above estimate to “continuous time” ξ :

$$P \left\{ \sup_{0 \leq \xi \leq 1} |f_n^r(l + \xi) - f_n^r(l) - \lambda_n \xi| \geq \varepsilon \right\} < \varepsilon / r.$$

Thus, for any $T_3 > 0$, any $\varepsilon > 0$ and any $n \in \bar{N}$, for all large r , we have

$$(118) \quad P \left\{ \max_{0 \leq l \leq T_3 r} \sup_{0 \leq \xi \leq 1} |f_n^r(l + \xi) - f_n^r(l) - \lambda_n \xi| \geq \varepsilon \right\} < (T_3 + 1)\varepsilon.$$

Let us choose a sequence of pairs $(T_3^{(i)}, \varepsilon^{(i)})$, $i = 0, 1, 2, \dots$, such that $T_3^{(i)} \uparrow \infty$, $\varepsilon^{(i)} > 0$ for all i , and

$$\sum_i (T_3^{(i)} + 1)\varepsilon^{(i)} < \infty.$$

Let us choose a subsequence $\mathcal{R}_2 \subseteq \mathcal{R}_1$ as follows. We pick arbitrary $r(0) \in \mathcal{R}_1$. Then, sequentially, for each index $i \geq 1$, we choose $r(i) > r(i - 1)$, $r(i) \in \mathcal{R}_1$, such that (118) holds for $T_3 = T_3^{(i)}$, $\varepsilon = \varepsilon^{(i)}$ and $r = r(i)$. Property (71) then follows from the Borel–Cantelli lemma.

A.3. Proof of Lemma 8. The proof of Lemma 8 uses constructions and arguments similar to those used in Section 12.1 in the proof of Lemma 7. (At one point Lemma 7 itself is applied.) Consider the family of fluid-scaled paths $\{\bar{z}^{r,l}\}$ defined by (79) in Section 12.1 for each integer $l \in [0, 2\delta r/T - 1]$, where $T > 0$ is a fixed constant.

LEMMA 11. (i) *Suppose the conditions of Lemma 8(i) hold and $C_1 = \kappa^2 \kappa_1 C + 2\varepsilon$ with $\kappa_1 = 2\kappa$. Then, for any sufficiently small $\varepsilon_2 > 0$ and $\varepsilon_4 > 0$, there exists $T > 0$ such that, for all sufficiently large r , the following properties (119)–(122) hold for integer $l \in [0, 2\delta r/T - 1]$ and the property (123) holds for integer $l \in [1, 2\delta r/T - 1]$:*

$$(119) \quad \|\bar{q}^{r,l}(0)\| < 2\kappa C_1 \quad \text{and} \quad \|\bar{q}^{r,l}(T)\| \geq \varepsilon_4 \quad \text{imply}$$

$$G(\bar{q}^{r,l}(T)) < 2\varepsilon_2,$$

$$(120) \quad \varepsilon_4 \leq \|\bar{q}^{r,l}(0)\| < 2\kappa C_1 \quad \text{and} \quad G(\bar{q}^{r,l}(0)) < 2\varepsilon_2 \quad \text{imply}$$

$$G(\bar{q}^{r,l}(u)) < 3\varepsilon_2 \quad \forall u \in [0, T],$$

$$(121) \quad \|\bar{q}^{r,l}(0)\| < 2\kappa C_1 \quad \text{implies}$$

$$\|\bar{q}^{r,l}(u)\| < \kappa \|\bar{q}^{r,l}(0)\| + \varepsilon_4 \quad \forall u \in [0, T],$$

$$(122) \quad \bar{x}^{r,l}(u) < C_1 \quad \text{and} \quad \|\bar{q}^{r,l}(u)\| < 2\kappa C_1 \quad \forall u \in [0, T],$$

$$(123) \quad \|\bar{q}^{r,l}(u)\| > \kappa\varepsilon_4 + \varepsilon_4 \quad \text{implies} \quad G(\bar{q}^{r,l}(u)) < 3\varepsilon_2 \quad \forall u \in [0, T].$$

(ii) *Suppose the conditions of Lemma 8(ii) hold. Then, for any sufficiently small $\varepsilon_2 > 0$ and $\varepsilon_4 > 0$, there exists $T > 0$ such that, for all sufficiently large r , all properties (119)–(123) hold for all integer $l \in [0, 2\delta r/T - 1]$.*

PROOF. (i) We choose $\varepsilon_2 > 0$ and $T > 0$ the same way as in the proof of Lemma 10 in Section 12.1. Let us fix an arbitrary $\varepsilon_4 > 0$ such that $\kappa\varepsilon_4 + \varepsilon_4 < 2\kappa C$.

Properties (119)–(121) are proved by contradiction, using Lemma 9 analogously to the argument we employed repeatedly in the proof of Lemma 10. Namely, if (119) would not hold, we would be able to construct an FSP with $\|q(0)\| \leq 2\kappa C_1$ and $G(q(T)) \geq 2\varepsilon_2$ which contradicts our choice of ε_2 . If (120) would not hold, we would be able to construct an FSP violating the property that $G(q(t))$ is nonincreasing. Finally, if (121) would not hold, we would be able to construct an FSP with $\|q(t)\| \geq \kappa\|q(0)\| + \varepsilon_4$ for some $t \geq 0$.

Let us prove (122) by contradiction. Suppose there is an infinite subsequence $\mathcal{R}_5 \subseteq \mathcal{R}_4(\omega)$ along which (122) does not hold for at least one $l \in [0, 2\delta r/T - 1]$

and some $u \in [0, T]$. For each $r \in \mathcal{R}_5$, let l_* be the smallest l such that, for some $u \in [0, T]$, either of the following two inequalities holds:

$$(124) \quad \bar{x}^{r,l}(u) \geq C,$$

$$(125) \quad \|\bar{q}^{r,l}(u)\| \geq 2\kappa C.$$

For large r , we must have $l_* \geq 1$ because, for $l = 0$, we have

$$(126) \quad \limsup_r \sup_{u \in [0, T]} \|\bar{q}^{r,0}(u)\| \leq \kappa C_2 < 2\kappa C < 2\kappa C_1$$

[for otherwise we would be able to construct an FSP with $\|q(0)\| = C_2$ and $\|q(t)\| > \kappa C_2$ for some $t > 0$], which means that, for large r and $u \in [0, T]$, we have

$$(127) \quad \bar{x}^{r,0}(u) < \kappa^2 C_2 < C < C_1.$$

Finally, let u_* denote the smallest value of u for which (124) or (125) holds for the corresponding l_* . (Thus, both l_* and u_* are functions of r .) We must have $u_* > 0$, because $u_* = 0$ would contradict the choice of l_* .

The following property is true for all large $r \in \mathcal{R}_5$:

$$(128) \quad (123) \text{ holds for all integer } l \in [1, l_*].$$

Indeed, from (121) we know that, for $l \leq l_*$ and $u \in [0, T]$, $\|\bar{q}^{r,l}(u)\| > \kappa \varepsilon_4 + \varepsilon_4$ implies $\|\bar{q}^{r,l}(0)\| > \varepsilon_4$ and therefore, by (119), $G(\bar{q}^{r,l}(0)) < 2\varepsilon_2$. Then, by (120), we must have $G(\bar{q}^{r,l}(u)) < 3\varepsilon_2$, which proves (128).

As a consequence of (128) and the choice of ε_2 , for all integer $l \in [1, l_*]$ and all $u \in [0, T]$, $\|\bar{q}^{r,l}(u)\| > \kappa \varepsilon_4 + \varepsilon_4$ implies $\|\bar{q}^{r,l}(u)\| < 2\kappa \bar{x}^{r,l}(u)$. Thus $\|\bar{q}^{r,l_*}(u_*)\| \geq 2\kappa C > \kappa \varepsilon_4 + \varepsilon_4$ always implies $\bar{x}^{r,l_*}(u_*) \geq C$. We see that, in fact, for $l = l_*$ and $u = u_*$, condition (124) must hold. Let us denote $\tilde{t}_*^r = \tilde{t}^r + [l_* T + u_*]/r$. [This \tilde{t}_*^r is a time on the diffusion time scale, i.e., that of the processes $\tilde{x}^r(\cdot)$ and $\tilde{q}^r(\cdot)$.] We must have $\tilde{x}^r(\tilde{t}_*^r) \rightarrow C$ because, by property (71), the positive jump sizes of (diffusion-scaled) functions $\tilde{x}^r(t)$ become arbitrarily small uniformly over compact sets. Let us choose a further subsequence $\mathcal{R}_6 \subseteq \mathcal{R}_5$ along which $\tilde{t}_*^r \rightarrow t'_* \in [t', t' + 2\delta]$. Note that $[t'_* - 3(2\delta), t'_* + 3(2\delta)] \subseteq [t' - 9\delta, t' + 9\delta]$. We see that the conditions of Lemma 7(i) (with $\mathcal{R}_4, \kappa_1, \delta, \tilde{t}^r$ and t' replaced by $\mathcal{R}_6, 2\kappa, 2\delta, \tilde{t}_*^r$ and t'_* , resp.) hold for the δ and ε we consider. Thus, by Lemma 7(i), $\tilde{x}^r(t) < C_1$ and $\|\tilde{q}^r(t)\| < 2\kappa C_1$ for all $t \in [\tilde{t}_*^r, \tilde{t}_*^r + 2\delta]$ and all large $r \in \mathcal{R}_6$. This and the construction of times \tilde{t}_*^r imply that (122) must hold for all large $r \in \mathcal{R}_6$ —a contradiction to the choice of \mathcal{R}_5 . The proof of (122) is complete.

Now, given that (122) provides uniform bound $\|\bar{q}^{r,l}(u)\| < 2\kappa C_1$ for all $u \in [0, T]$ and all large $r \in \mathcal{R}_4(\omega)$, we obtain property (123) by repeating (for arbitrary integer $l \in [1, 2\delta r/T - 1]$) the argument we used above in the proof of (128). The proof of Lemma 11(i) is complete.

(ii) In addition to the statement of (i), we only need to show that (123) holds also for $l = 0$ for all large r . Given the additional condition that either $\bar{q}^{r,0}(0) \rightarrow 0$ (if $C_2 = 0$) or $G(\bar{q}^{r,0}(0)) \rightarrow 0$ (if $C_2 > 0$), we see that the argument used in the proof of (123) for $l \geq 1$ applies to the case $l = 0$ as well. \square

PROOF OF LEMMA 8. (i) We apply Lemma 11(i). Statement (a) follows from (122), $\tilde{x}^r \Rightarrow \tilde{x}$ and right-continuity of \tilde{x} . Property (123), which holds for $l \in [1, 2\delta r/T - 1]$, means that, for all large r and any $t \in [t' + \delta', t' + \delta]$, either $\|\tilde{q}^r(t)\| \leq \kappa\varepsilon_4 + \varepsilon_4$ or $G(\tilde{q}^r(t)) < 3\varepsilon_2$. Since we have $\|\tilde{q}^{r,(NC)}(t)\| \leq \|\tilde{q}^r(t)\|$, $\|\tilde{q}^{r,(NC)}(t)\| \leq \|\tilde{q}^r(t) - \tilde{x}^r(t)v\|$, property (52) and the uniform bound $\tilde{x}^r(t) < C_1$, we can make $\|\tilde{q}^{r,(NC)}(t)\|$ arbitrarily small uniformly in $[t' + \delta', t' + \delta]$ for large r , if we choose sufficiently small $\varepsilon_2 > 0$ and $\varepsilon_4 > 0$. This proves (b).

(ii) To prove (b') we use Lemma 11(ii) and the same argument as in the proof of (i)(b), except this argument now applies to all $t \in [t', t' + \delta]$. \square

A.4. One-dimensional Skorohod problem. The following proposition describes standard properties of solutions of the one-dimensional Skorohod problem.

PROPOSITION 4. *Let $w = (w(t), t \geq 0)$ be a continuous function in $D([0, \infty), \mathbb{R})$ such that $w(0) \geq 0$. Then the following hold:*

- (i) *There exists a unique pair (x, y) of functions in $D([0, \infty), \bar{\mathbb{R}})$, such that:*
 - (a) $x(t) = w(t) + y(t) \geq 0, t \geq 0$,
 - (b) y is nondecreasing and nonnegative,
 - (c) $y(0) = 0$,
 - (d) for any $t \geq 0$, if $x(t) > 0$, then t is not a point of increase of y .

This unique pair is (x°, y°) , where

$$y^\circ(t) \doteq -\left[0 \wedge \inf_{0 \leq u \leq t} w(u)\right], \quad x^\circ(t) = w(t) + y^\circ(t), \quad t \geq 0.$$

(ii) *For any pair (x, y) of functions in $D([0, \infty), \bar{\mathbb{R}})$ satisfying (a) and (b), we have*

$$y(t) \geq y^\circ(t), \quad x(t) \geq x^\circ(t), \quad t \geq 0.$$

The proof of Proposition 4 can be found, for example, in [8]. (It is also contained in the proof of Theorem 5.1 of [34]. More precisely, it can be obtained using the argument proving inequality (8) in [34].) Note that the formulation of Proposition 4(ii) is formally more general than a more conventional formulation, in which $x, y \in D([0, \infty), \mathbb{R})$ (i.e., x and y required to be finite) and condition $y(0) = 0$ is included into (b). However, this more general statement is proved the same way as the conventional one, with straightforward adjustments.

Acknowledgments. I would like to thank Mike Harrison, Avi Mandelbaum and Marty Reiman for useful discussions of the earlier versions of this paper. I am grateful to the Associate Editor and all referees for extremely helpful comments. Most notably, the importance of relaxing the CRP condition (which was employed in the original version) to the RP was pointed out by one of the referees. Also, this version of the paper incorporates numerous presentation and terminology improvements suggested by the referees.

REFERENCES

- [1] ANDREWS, M., KUMARAN, K., RAMANAN, K., STOLYAR, A. L., VIJAYAKUMAR, R. and WHITING, P. (2004). Scheduling in a queueing system with asynchronously varying service rates. *Probab. Engrg. Inform. Sci.* To appear.
- [2] ARMONY, M. and BAMBOS, N. (1999). Queueing networks with interacting service resources. In *Proc. 37th Annual Allerton Conference on Communication, Control, and Computing* 42–51. Univ. of Illinois, Urbana-Champaign.
- [3] BELL, S. L. and WILLIAMS, R. J. (2001). Dynamic scheduling of a system with two parallel servers in heavy traffic with complete resource pooling: Asymptotic optimality of a continuous review threshold policy. *Ann. Appl. Probab.* **11** 608–649.
- [4] BRAMSON, M. (1998). State space collapse with applications to heavy traffic limits for multiclass queueing networks. *Queueing Systems Theory Appl.* **30** 89–148.
- [5] BRAMSON, M. and WILLIAMS, R. (2000). On dynamic scheduling of stochastic networks in heavy traffic and some new results for the workload process. In *Proc. 39th IEEE Conference on Decision and Control* 516–521. IEEE, New York.
- [6] BRAMSON, M. and DAI, J. G. (2001). Heavy traffic limits for some queueing networks. *Ann. Appl. Probab.* **11** 49–88.
- [7] BUCHE, R. and KUSHNER, H. J. (2002). Control of mobile communications with time-varying channels in heavy traffic. *IEEE Trans. Automat. Control* **47** 992–1003.
- [8] CHEN, H. and MANDELBAUM, A. (1991). Leontief systems, RBV's and RBM's. In *Applied Stochastic Analysis* (M. H. A. Davis and R. J. Elliott, eds.) 1–43. Gordon and Breach, New York.
- [9] DAI, J. G. and PRABHAKAR, B. (2000). The throughput of data switches with and without speedup. In *Proc. IEEE INFOCOM 2000* 556–564.
- [10] ETHIER, S. N. and KURTZ, T. G. (1986). *Markov Process: Characterization and Convergence*. Wiley, New York.
- [11] GANS, N. and VAN RYZIN, G. (1997). Optimal control of a multiclass, flexible queueing system. *Oper. Res.* **45** 677–693.
- [12] GANS, N. and VAN RYZIN, G. (1998). Optimal dynamic scheduling of a general class of parallel-processing queueing systems. *Adv. in Appl. Probab.* **30** 1130–1156.
- [13] HARRISON, J. M. (1998). Heavy traffic analysis of a system with parallel servers: Asymptotic optimality of discrete review policies. *Ann. Appl. Probab.* **8** 822–848.
- [14] HARRISON, J. M. (2000). Brownian models of open processing networks: Canonical representation of workload. *Ann. Appl. Probab.* **10** 75–103.
- [15] HARRISON, J. M. and VAN MIEGHEM, J. A. (1997). Dynamic control of Brownian networks: State space collapse and equivalent workload formulations. *Ann. Appl. Probab.* **7** 747–771.
- [16] HARRISON, J. M. and LOPEZ, M. J. (1999). Heavy traffic resource pooling in parallel-server systems. *Queueing Systems* **33** 339–368.

- [17] JALALI, A., PADOVANI, R. and PANKAJ, R. (2000). Data throughput of CDMA-HDR, a high efficiency-high data rate personal communication wireless system. In *Proc. IEEE VTC 2000 Spring* 1854–1858. IEEE, New York.
- [18] KAHALE, N. and WRIGHT, P. E. (1997). Dynamic global packet routing in wireless networks. In *Proc. IEEE INFOCOM'97* 1414–1421.
- [19] KARLIN, S. (1992). *Mathematical Methods in Games, Programming, and Economics*. Dover, New York.
- [20] KELLY, F. P. and LAWS, C. N. (1993). Dynamic routing in open queueing networks: Brownian models, cut constraints and resource pooling. *Queueing Systems* **13** 47–86.
- [21] LAWS, C. N. (1992). Resource pooling in queueing networks with dynamic routing. *Adv. in Appl. Probab.* **24** 699–726.
- [22] MANDELBAUM, A. and STOLYAR, A. L. (2004). Scheduling flexible servers with convex delay costs: Heavy-traffic optimality of the generalized $c\mu$ -rule. *Oper. Res.* To appear.
- [23] MCKEOWN, N., ANANTHARAM, V. and WALRAND, J. (1996). Achieving 100% throughput in an input-queued switch. In *Proc. IEEE INFOCOM'96* 296–302.
- [24] MEKKITTIKUL, A. and MCKEOWN, N. (1996). A starvation free algorithm for achieving 100% throughput in an input-queued switch. In *Proc. IEEE ICCCN'96* 226–231.
- [25] REIMAN, M. I. (1988). A multiclass feedback queue in heavy traffic. *Adv. in Appl. Probab.* **20** 179–207.
- [26] REIMAN, M. I. (1984). Some diffusion approximations with state space collapse. *Proc. International Seminar on Modeling and Performance Evaluation Methodology. Lecture Notes in Control and Inform. Sci.* 209–240. Springer, New York.
- [27] SHAKKOTTAI, S. and STOLYAR, A. (2002). Scheduling for multiple flows sharing a time-varying channel: The exponential rule. In *Analytic Methods in Applied Probability. In Memory of Fridrih Karpelevich* (Yu. M. Suhov, ed.) 185–202. AMS, Providence, RI.
- [28] STOLYAR, A. L. (1995). On the stability of multiclass queueing networks: A relaxed sufficient condition via limiting fluid processes. *Markov Process. Related Fields* **1** 491–512.
- [29] TASSIULAS, L. (1995). Adaptive back-pressure congestion control based on local information. *IEEE Trans. Automat. Control* **40** 236–250.
- [30] TASSIULAS, L. and EPHREMIDES, A. (1992). Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks. *IEEE Trans. Automat. Control* **37** 1936–1948.
- [31] TASSIULAS, L. and EPHREMIDES, A. (1993). Dynamic server allocation to parallel queues with randomly varying connectivity. *IEEE Trans. Inform. Theory* **39** 466–478.
- [32] TASSIULAS, L. and BHATTACHARYA, P. P. (2000). Allocation of interdependent resources for maximal throughput. *Comm. Statist. Stochastic Models* **16** 27–48.
- [33] WHITT, W. (1971). Weak convergence theorems for priority queues: Preemptive resume discipline. *J. Appl. Probab.* **8** 74–94.
- [34] WILLIAMS, R. J. (1998). An invariance principle for semimartingale reflecting Brownian motions in an orthant. *Queueing Systems* **30** 5–25.
- [35] WILLIAMS, R. J. (1998). Diffusion approximations for open multiclass queueing networks: Sufficient conditions involving state space collapse. *Queueing Systems* **30** 27–88.
- [36] WILLIAMS, R. J. (2000). On dynamic scheduling of a parallel server system with complete resource pooling. *Fields Institute Communications* **28** 49–71.

BELL LABS, LUCENT TECHNOLOGIES
600 MOUNTAIN AVENUE, 2C-322
MURRAY HILL, NEW JERSEY 07974
USA
E-MAIL: stolyar@research.bell-labs.com