SWITCHED NETWORKS WITH MAXIMUM WEIGHT POLICIES: FLUID APPROXIMATION AND MULTIPLICATIVE STATE SPACE COLLAPSE

By Devavrat Shah* and Damon Wischik†

MIT and UCL

We consider a queueing network in which there are constraints on which queues may be served simultaneously; such networks may be used to model input-queued switches and wireless networks. The scheduling policy for such a network specifies which queues to serve at any point in time. We consider a family of scheduling policies, related to the maximum-weight policy of Tassiulas and Ephremides [28], for single-hop and multihop networks. We specify a fluid model, and show that fluid-scaled performance processes can be approximated by fluid model solutions. We study the behaviour of fluid model solutions under critical load, and characterize invariant states as those states which solve a certain network-wide optimization problem. We use fluid model results to prove multiplicative state space collapse. A notable feature of our results is that they do not assume complete resource pooling.

1. Introduction.

A switched network consists of a collection of queues, operating in discrete time. In each timeslot, queues are offered service according to a *service schedule* chosen from a specified finite set. For example, in a three-queue system, the set of allowed schedules might consist of "Serve 3 units of work each from queues A & B" and "Serve 1 unit of work each from queues A & C, and 2 units from queue B". The rule for choosing a schedule is called the *scheduling policy*. New work may arrive in each timeslot; let each queue have a dedicated exogeneous arrival process, with specified mean arrival rates. Once work is served, it may either rejoin one of the queues or leave the network.

Switched networks are special cases of what Harrison [12] calls 'stochastic processing networks'. We believe that switched networks are general enough to model a variety of interesting applications. For example, they have been used to model input-queued switches, the devices at the heart of high-end Internet routers, whose underlying silicon architecture imposes constraints on which traffic streams can be

*Supported by NSF CAREER CNS-0546590.

†Supported by a Royal Society university research fellowship. Collaboration supported by the British Council Researcher Exchange program.

AMS 2000 subject classifications: 60K25, 60K30, 90B36

Keywords and phrases: switched network, maximum weight scheduling, fluid models, state space collapse, heavy traffic, diffusion approximation

transmitted simultaneously [8]. They have also been used to model a multihop wireless network in which interference limits the amount of service that can be given to each host [28].

The main result of this paper is Theorem 7.1, which proves multiplicative state space collapse (as defined in Bramson [3]) for a switched network running a generalized version of the maximum-weight scheduling policy of Tassiulas and Ephremides [28], in the diffusion (or heavy traffic) limit. Whereas previous works on switched networks and stochastic processing networks in the diffusion limit [6, 7, 27] have assumed the 'complete resource pooling' condition, which roughly means that there is a single bottleneck cut constraint, we do not make this assumption. Section 3 discusses further the related work and our reasons for being interested in the case without complete resource pooling.

To prove multiplicative state space collapse we follow the general method laid out by Bramson [3]. In Section 2 we specify a stochastic switched network model and describe the generalized maximum-weight policy. In Section 4 we specify a fluid model, and prove that fluid-scaled performance processes of the switched network are approximated by solutions of this fluid model. Sections 5 and 6 give properties of the solutions of the fluid model for single-hop and multihop networks respectively. Specifically, for non-overloaded fluid model solutions, we characterize the invariant states and prove that fluid model solutions converge towards the set of invariant states. In Section 7 we use these properties to prove multiplicative state space collapse.

We use the cluster-point method of Bramson [3] to prove the fluid model approximation in Section 4, rather than following an approach based on weak convergence. The former yields uniform bounds on the error of fluid model approximations, and these uniform bounds are needed in proving multiplicative state space collapse. However, the assumptions we make on the arrival process are not the same as those of Bramson [3].

In Section 8 we give results concerning the fluid model behaviour of a general single-hop switched network in critical load, and the set of invariant states for the input-queued switch, under a condition that we call 'complete loading'. Motivated by these results, we define a scheduling policy which we conjecture is optimal in the diffusion limit.

Notation. Let $\mathbb{N}$ be the set of natural numbers $\{1, 2, \dots\}$, let $\mathbb{Z}_+ = \{0, 1, 2, \dots\}$, let $\mathbb{R}$ be the set of real numbers, and let $\mathbb{R}_+ = \{x \in \mathbb{R} : x \geq 0\}$. Let $1_{\{\cdot\}}$ be the indicator function, where $1_{\text{true}} = 1$ and $1_{\text{false}} = 0$. Let $x \wedge y = \min(x, y)$, $x \vee y = \max(x, y)$ and $[x]^+ = x \vee 0$. When x is a vector, the maximum is taken componentwise.

We will reserve bold letters for vectors in $\mathbb{R}^N$, where N is the number of queues, for example $\mathbf{x} = [x_n]_{1\le n\le N}$. Superscripts on vectors are used to denote labels, not exponents, except where otherwise noted; thus for example $(\mathbf{x}^0, \mathbf{x}^1, \mathbf{x}^2)$ refers to three arbitrary vectors. Let $\mathbf{0}$ be the vector of all 0s, and $\mathbf{1}$ be the vector of all 1s. Use the norm $|\mathbf{x}| = \max_n |x_n|$. For vectors $\mathbf{u}$ and $\mathbf{v}$ and functions $f : \mathbb{R} \to \mathbb{R}$, let

$$\mathbf{u}\cdot\mathbf{v} = \sum_{n=1}^{N} u_n v_n, \quad \mathbf{u}\mathbf{v} = [u_n v_n]_{1\le n\le N}, \quad \text{and} \quad f(\mathbf{u}) = \big[f(u_n)\big]_{1\le n\le N}$$

and let matrix multiplication take precedence over dot product so that

$$\mathbf{u}\cdot A\mathbf{v} = \sum_{n=1}^{N} u_n \left(\sum_{m=1}^{N} A_{nm} v_m \right).$$

Let A^{T} be the transpose of matrix A. For a set $\mathcal{S} \subset \mathbb{R}^N$, denote its convex hull by $\langle \mathcal{S} \rangle$.

For a fixed $T > 0$, and $I \in \mathbb{N}$, let $C^I(T)$ be the set of continuous functions $[0, T] \to \mathbb{R}^I$, where $\mathbb{R}^I$ is equipped with the norm $|x| = \max_i |x_i|$. Equip $C^I(T)$ with the norm

$$\|f\| = \sup_{t\in[0,T]} |f(t)|.$$

Let $d(f, g) = \|f - g\|$ be the metric induced by this norm. For $E \subset C^I(T)$ and $f \in C^I(T)$, let $d(f, E) = \inf\{d(f, g) : g \in E\}$. Define the modulus of continuity $\mathrm{mc}_\delta(\cdot)$ by

$$\mathrm{mc}_\delta(f) = \sup_{|s-t|<\delta} \big|f(s) - f(t)\big|$$

where $s, t \in [0, T]$. Since $[0, T]$ is compact, each $f \in C^I(T)$ is uniformly continuous, therefore $\mathrm{mc}_\delta(f) \to 0$ as $\delta \to 0$.

2. Switched network model.

We now introduce the switched network model. Section 2.1 describes the general system model, Section 2.2 specifies the class of scheduling policies we are interested in, and Section 2.3 lists the probabilistic assumptions about the arrival process that are needed for the main theorems.

2.1. *Queueing dynamics.* Consider a collection of N queues. Let time be discrete, indexed by $\tau \in \{0, 1, \dots\}$. Let $Q_n(\tau)$ be the amount of work in queue $n \in \{1, \dots, N\}$ at timeslot τ. Following our general notation for vectors, we write $\mathbf{Q}(\tau)$ for $[Q_n(\tau)]_{1\le n\le N}$. The initial queue sizes are $\mathbf{Q}(0)$. Let $A_n(\tau)$ be the total amount of work arriving to queue n, and $B_n(\tau)$ be the cumulative potential service to queue n, up to time τ, with $\mathbf{A}(0) = \mathbf{B}(0) = \mathbf{0}$.

We first define the queueing dynamics for a single-hop switched network. Defining $d\mathbf{A}(\tau) = \mathbf{A}(\tau+1) - \mathbf{A}(\tau)$ and $d\mathbf{B}(\tau) = \mathbf{B}(\tau+1) - \mathbf{B}(\tau)$, the basic Lindley recursion that we will consider is

$$\mathbf{Q}(\tau+1) = \big[\mathbf{Q}(\tau) - d\mathbf{B}(\tau)\big]^+ + d\mathbf{A}(\tau) \tag{1}$$

where the $[\cdot]^+$ is taken componentwise. The fundamental 'switched network' constraint is that there is some finite set $\mathcal{S} \subset \mathbb{R}_+^N$ such that

$$d\mathbf{B}(\tau) \in \mathcal{S} \quad \text{for all } \tau. \tag{2}$$

We will refer to $\boldsymbol{\pi} \in \mathcal{S}$ as a schedule, and $\mathcal{S}$ as the set of allowed schedules. In the applications in this paper, the schedule is chosen based on current queue sizes, which is why it is natural to write the basic Lindley recursion as (1) rather than the more standard $[\mathbf{Q}(\tau) + d\mathbf{A}(\tau) - d\mathbf{B}(\tau)]^+$.

For the analyses in this paper it is useful to keep track of two other quantities. Let $Y_n(\tau)$ be the cumulative amount of idling at queue n, defined by $\mathbf{Y}(0) = \mathbf{0}$ and

$$d\mathbf{Y}(\tau) = \big[d\mathbf{B}(\tau) - \mathbf{Q}(\tau)\big]^+, \tag{3}$$

where $d\mathbf{Y}(\tau) = \mathbf{Y}(\tau+1) - \mathbf{Y}(\tau)$. Then (1) can be rewritten

$$\mathbf{Q}(\tau) = \mathbf{Q}(0) + \mathbf{A}(\tau) - \mathbf{B}(\tau) + \mathbf{Y}(\tau). \tag{4}$$

Also, let $S_{\boldsymbol{\pi}}(\tau)$ be the cumulative time spent on schedule $\boldsymbol{\pi}$ up to time τ, so that

$$\mathbf{B}(\tau) = \sum_{\boldsymbol{\pi} \in \mathcal{S}} S_{\boldsymbol{\pi}}(\tau)\boldsymbol{\pi}. \tag{5}$$

For a multihop switched network, let $R \in \{0,1\}^{N\times N}$ be the routing matrix, $R_{mn} = 1$ if work served from queue m is sent to queue n, and $R_{mn} = 0$ otherwise; if $R_{mn} = 0$ for all n then work served from queue m departs the network. For each m we require $R_{mn} = 1$ for at most one n. (Tassiulas and Ephremides [28] described a network model with routing choice, whereas we have restricted ourselves to fixed routing for the sake of simplicity.) The scheduling constraint (2) is as before, the definition of idling (3) is as before, and the queueing dynamics are now defined by

$$Q_n(\tau+1) = Q_n(\tau) + dA_n(\tau) - \big(dB_n(\tau) - dY_n(\tau)\big) + \sum_m R_{mn}\big(dB_m(\tau) - dY_m(\tau)\big).$$

Equivalently,

$$\mathbf{Q}(\tau) = \mathbf{Q}(0) + \mathbf{A}(\tau) - (I - R^{\mathsf{T}})\big(\mathbf{B}(\tau) - \mathbf{Y}(\tau)\big). \tag{6}$$

Note that $\mathbf{A}$ includes only exogenous arrivals to the network, not internally routed traffic. We will assume that routing is acyclic, i.e. that work served from some queue n never returns to queue n. For example, Border Gateway Protocol (BGP) utilized for routing in the Internet is designed to be acyclic [30]. This implies that the inverse $\vec{R} = (I - R^{\mathsf{T}})^{-1}$ exists; by considering the expansion $\vec{R} = I + R^{\mathsf{T}} + (R^{\mathsf{T}})^2 + \cdots$ it is clear that $\vec{R}_{mn} \in \{0, 1\}$ for all m, n and that $\vec{R}_{mn} = 1$ if work injected at queue n eventually passes through m, and $\vec{R}_{mn} = 0$ otherwise. When $R = 0$ we obtain a single-hop switched network.

A straightforward bound we shall need is

$$Q_n(\tau) \le Q_n(\tau') + A_n(\tau) - A_n(\tau') + \sum_m R_{mn}\big(B_m(\tau) - B_m(\tau')\big) \tag{7}$$

for $\tau' \le \tau$.

2.2. *Scheduling policy.* A policy that decides which schedule to choose at each timeslot $\tau \in \mathbb{Z}_+$ is called a *scheduling policy.* In this paper we will be interested in the Max-Weight scheduling policy, introduced by Tassiulas and Ephremides [28]. We will refer to it as MW.

2.2.1. *Max-weight policy for single-hop network.* We describe the policy first for a single-hop network. Let $\mathbf{Q}(\tau)$ be the vector of queue sizes at time τ. Define the weight of a schedule $\boldsymbol{\pi} \in \mathcal{S}$ to be $\boldsymbol{\pi}\cdot\mathbf{Q}(\tau)$. The MW policy then chooses[1] for timeslot τ a schedule $d\mathbf{B}(\tau)$ with the greatest weight,

$$d\mathbf{B}(\tau) \in \operatorname*{argmax}_{\boldsymbol{\pi}\in\mathcal{S}} \boldsymbol{\pi}\cdot\mathbf{Q}(\tau). \tag{8}$$

This policy can be generalized to choose a schedule which maximizes $\boldsymbol{\pi}\cdot\mathbf{Q}(\tau)^\alpha$, where the exponent is taken componentwise for some $\alpha > 0$; call this the MW-α policy. More generally, one could choose a schedule such that

$$d\mathbf{B}(\tau) \in \operatorname*{argmax}_{\boldsymbol{\pi}\in\mathcal{S}} \boldsymbol{\pi}\cdot f(\mathbf{Q}(\tau)) \tag{9}$$

for some function $f : \mathbb{R}_+ \to \mathbb{R}_+$; call this the MW-$f$ policy. It is assumed that f satisfies the following scale-invariance property:

[1]There may be several schedules which jointly have the greatest weight. To be concrete, we might specify some fixed numbering of schedules, and choose the highest-numbered maximum-weight schedule. Alternatively, we might treat MW not as a policy *per se* but as a constraint on the set of allowed sample paths. For example, in a stochastic setting, we might allow $d\mathbf{B}(\tau)$ to be a random variable, measurable with respect to the underlying probability space, satisfying (8) for every randomness. This permits 'break ties at random'. For the analyses in this paper, it makes no difference which of these two options is used.

Assumption 2.1 *Assume f is differentiable and strictly increasing with $f(0) = 0$. Assume also that for any $\mathbf{q} \in \mathbb{R}_+^N$ and $\boldsymbol{\pi} \in \mathcal{S}$, with $m(\mathbf{q}) = \max_{\boldsymbol{\rho}\in\mathcal{S}} \boldsymbol{\rho}\cdot f(\mathbf{q})$,*

$$\boldsymbol{\pi}\cdot f(\mathbf{q}) = m(\mathbf{q}) \quad\implies\quad \boldsymbol{\pi}\cdot f(\kappa\mathbf{q}) = m(\kappa\mathbf{q}) \text{ for all } \kappa \in \mathbb{R}_+.$$

This is satisfied by $f(x) = x^\alpha$, $\alpha > 0$, but it is not satisfied for example for an input-queued switch with $f(x) = \log(1 + x)$.

2.2.2. *Max-weight policy for multihop network.* Now we define the multihop version of the MW-f scheduling policy. This policy chooses a schedule $d\mathbf{B}(\tau)$ at time τ such that

$$d\mathbf{B}(\tau) \in \operatorname*{argmax}_{\boldsymbol{\pi}\in\mathcal{S}} \boldsymbol{\pi}\cdot(I - R)f\big(\mathbf{Q}(\tau)\big).$$

Recall that matrix multiplication takes precedence over the $\cdot$ operator, so the argmax is of $\boldsymbol{\pi}\cdot\big\{(I - R)f(\mathbf{Q}(\tau))\big\}$; note also that

$$[Rf(\mathbf{Q})]_n = \sum_m R_{nm} f(Q_m) = f\big([R\mathbf{Q}]_n\big)$$

where $[R\mathbf{Q}]_n$ is the queue size at the first queue downstream of n (or 0 if there is no queue downstream). Thus

$$\big[(I - R)f(\mathbf{Q})\big]_n = f(Q_n) - f\big([R\mathbf{Q}]_n\big). \tag{10}$$

The difference $f(Q_n) - f([R\mathbf{Q}]_n)$ is interpreted as the pressure to send work from queue n to the queue downstream of n; if the downstream queue has more work in it than the upstream queue then there is no pressure to send work downstream. For this reason, it is also known as *backpressure* policy.

As before we will assume that f satisfies a scale-invariance property, the multihop equivalent of Assumption 2.1:

Assumption 2.2 *Assume f is differentiable and strictly increasing with $f(0) = 0$. Assume also that for any $\mathbf{q} \in \mathbb{R}_+^N$ and $\boldsymbol{\pi} \in \mathcal{S}$, with $m(\mathbf{q}) = \max_{\boldsymbol{\rho}\in\mathcal{S}} \boldsymbol{\rho}\cdot(I - R)f(\mathbf{q})$,*

$$\boldsymbol{\pi}\cdot(I - R)f(\mathbf{q}) = m(\mathbf{q}) \quad\implies\quad \boldsymbol{\pi}\cdot(I - R)f(\kappa\mathbf{q}) = m(\kappa\mathbf{q}) \text{ for all } \kappa \in \mathbb{R}_+.$$

We further require that the scheduler always have the option of not sending work downstream at any individual queue. Our Lyapunov function, and indeed our whole fluid analysis in Section 6, rely on this assumption.

Assumption 2.3 *For the multihop setting, assume that $\mathcal{S}$ satisfies the following: if $\boldsymbol{\pi} \in \mathcal{S}$ is an allowed schedule, and $\boldsymbol{\rho} \in \mathbb{R}_+^N$ is some other vector with $\rho_n \in \{0, \pi_n\}$ for all n, then $\boldsymbol{\rho} \in \mathcal{S}$.*

In the rest of this paper, whenever we refer to a network running the MW-f backpressure policy, we mean that Assumptions 2.2 and 2.3 are satisfied.

2.3. *Stochastic model.* Some of the results in this paper are about fluid-scaled processes, and others are about multiplicative state space collapse in the diffusion scaling, and the different results make different assumptions about the arrival process.

Assumption 2.4 (Assumptions for the fluid scale) *Let $\mathbf{A}(\cdot)$ be a random process with stationary increments. Assume it has a well-defined mean arrival rate vector $\boldsymbol{\lambda}$, i.e. assume $\lim_{\tau\to\infty} A_n(\tau)/\tau$ exists almost surely and is deterministic for every queue $1 \le n \le N$, and define*

$$\lambda_n = \lim_{\tau\to\infty} \frac{1}{\tau} A_n(\tau). \tag{11}$$

Assume there is a sequence of deviation terms $\delta_r \in \mathbb{R}_+$, $r \in \mathbb{N}$, such that $\delta_r \to 0$ as $r \to \infty$ and

$$\mathbb{P}\Big(\sup_{\tau\le r} \frac{1}{r}\big|\mathbf{A}(\tau) - \boldsymbol{\lambda}\tau\big| \ge \delta_r\Big) \to 0 \quad \text{as } r \to \infty.$$

Assumption 2.5 (Assumptions for multiplicative state space collapse) *Let $\mathbf{A}^r(\cdot)$ be a sequence of random processes indexed by $r \in \mathbb{N}$. For each r, assume that $\mathbf{A}^r$ has stationary increments, and a well-defined mean arrival rate vector $\boldsymbol{\lambda}^r$, and that there is some limiting arrival rate vector $\boldsymbol{\lambda}$ such that*

$$\boldsymbol{\lambda}^r \to \boldsymbol{\lambda} \quad \text{as } r \to \infty.$$

Assume there is a sequence of deviation terms $\delta_z \in \mathbb{R}_+$, $z \in \mathbb{N}$, such that $\delta_z \to 0$ as $z \to \infty$ and

$$z(\log z)^2\, \mathbb{P}\Big(\sup_{\tau\le z} \frac{1}{z}\big|\mathbf{A}^r(\tau) - \boldsymbol{\lambda}^r\tau\big| \ge \delta_z\Big) \to 0 \quad \text{as } z \to \infty, \text{ uniformly in } r.$$

If the arrival process is the same for all r, say $\mathbf{A}^r = \mathbf{A}$ where $\mathbf{A}$ has a well-defined mean arrival rate vector, then Assumption 2.5 reduces to

$$\mathbb{P}\Big(\sup_{\tau\le r} \frac{1}{r}\big|\mathbf{A}(\tau) - \boldsymbol{\lambda}\tau\big| \ge \delta_r\Big) = o\Big(\frac{1}{r(\log r)^2}\Big) \tag{12}$$

and it implies Assumption 2.4. For any arrival process with i.i.d. increments that are uniformly bounded, i.e. such that there is an $A^{\max}$ for which

$$A_n(\tau+1) - A_n(\tau) \in \big[0, A^{\max}\big] \quad \text{for all } n,\ \tau,$$

equation (12) holds with $\delta_r = C\sqrt{\log r}/\sqrt{r}$, with choice of an appropriate constant C that depends on $A^{\max}$, by an application of concentration inequality by Azuma [2] and Hoeffding [14]. More generally, it holds when the increments are not uniformly bounded but instead satisfy a reasonable moment bound. For example, an application of Doob's maximal inequality [10] with bounded fourth moment and $\delta_r = r^{-1/6}$ yields a stronger result than (12); this can be used to show that a Poisson process satisfies that equation. Furthermore (12) holds for a much larger class of stationary arrival processes beyond processes with i.i.d. increments, for example Markov modulated processes (see Dembo and Zeitouni [9]).

2.4. *Motivating example.* An Internet router has several input ports and output ports. A data transmission cable is attached to each of these ports. Packets arrive at the input ports. The function of the router is to work out which output port each packet should go to, and to transfer packets to the correct output ports. This last function is called *switching*. There are a number of possible switch architectures; we will consider the commercially popular input-queued switch architecture.

Figure 1 illustrates an input-queued switch with three input ports and three output ports. Packets arriving at input i destined for output j are stored at input port i, in queue $Q_{i,j}$, thus there are $N = 9$ queues in total. (For this example it is more natural to use double indexing, e.g. $Q_{3,2}$, whereas for general switched networks it is more natural to use single indexing, e.g. Q_n for $1 \leq n \leq N$.)

The switch operates in discrete time. In each timeslot, the switch fabric can transmit a number of packets from input ports to output ports, subject to the two constraints that each input can transmit at most one packet and that each output can receive at most one packet. In other words, at each timeslot the switch can choose a *matching* from inputs to outputs. The schedule $\boldsymbol{\pi} \in \mathbb{R}_+^{3\times 3}$ is given by $\pi_{i,j} = 1$ if input port i is matched to output port j in a given timeslot, and $\pi_{i,j} = 0$ otherwise. Clearly $\boldsymbol{\pi}$ is a permutation matrix, and the set $\mathcal{S}$ of allowed schedules is the set of 3×3 permutation matrices.

Figure 1 shows two possible matchings. In the left hand figure, the matching allows a packet to be transmitted from input port 3 to output port 2, but since $Q_{3,2}$ is empty, no packet is actually transmitted.

3. Related work. Keslassy and McKeown [20] found from extensive simulations of an input-queued switch that the average queueing delay is different under MW-α policies for different values of $\alpha > 0$. They conjecture:

Conjecture 3.1 *For an input-queued switch running the MW-α policy, the average queueing delay decreases as α decreases.*

Though our work is motivated by the desire to establish Conjecture 3.1, we have not been able to prove it. But whereas the two main analytic approaches that have been employed in the literature yield results for the input-queued switch that are insensitive to α, our result about multiplicative state space collapse *is* sensitive, as

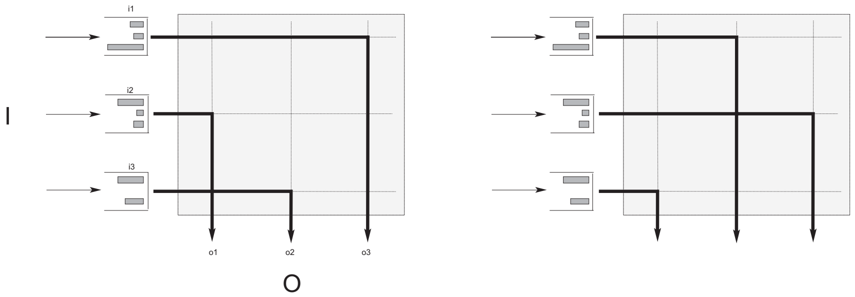

FIG 1. *An input-queued switch, and two example matching of inputs to outputs.*

shown in Section 8. We speculate that our result might eventually form part of a proof of the conjecture.

The two main analytic approaches that have been employed in the literature are stability analysis and heavy traffic analysis. In stability analysis, one calculates the set of arrival rates for which a policy is stable (in the sense of [1, 8, 20, 24, 25, 28]). All the prior work in this context leads to the conclusion that MW-α has the optimal stability region, regardless of α.

In heavy traffic analysis, one looks at queue size behavior under a diffusion (or heavy traffic) scaling. This regime was first described by Kingman [21]; since then a substantial body of theory has developed, and modern treatments can be found in [3, 11, 29, 31]. Stolyar has studied MW-α for a generalized switch model in the diffusion scaling, and obtained a complete characterization of the diffusion approximation for the queue size process, *under a condition known as 'complete resource pooling'*. This condition effectively requires that a clever scheduling policy be able to balance work between all the heavily loaded queues. Stolyar [27] showed in a remarkable paper that the limiting queue size lives in a one-dimensional state space. Operationally, this means that all one needs to keep track of is the one-dimensional total amount of work in the system (called the *workload*), and at any point in time one can assume that the individual queues have all been balanced. Dai and Lin [6, 7] have established that similar result holds in the more general setting of a stochastic processing network.

Under the complete resource pooling condition, the results in [6, 7, 27] imply that the performance of MW-α in an input-queued switch is always optimal (in the diffusion scaling) regardless of the value of $\alpha > 0$. Therefore these results do not help in addressing Conjecture 3.1. This is our motivation for studying switched networks in the absence of complete resource pooling. Technically, the lifting map for a critically-loaded input-queued switch is degenerate and insensitive to α under complete resource pooling, but it is sensitive to α otherwise.

We prove multiplicative state space collapse, following the method of Bramson

[3]. The complement of Bramson's work is by Williams [31], and consists in proving a diffusion approximation, using an appropriate invariance principle along with the multiplicative state space collapse. We do not carry out this complementary aspect. Stolyar [27] and Dai and Lin [6, 7] have proved the diffusion approximation under complete resource pooling condition; and Kang and Williams [17] have made progress towards it in the case without complete resource pooling, for an input-queued switch under the MW-1 policy.

Whereas in heavy traffic models of other systems [3, 11, 27, 31] the lifting map from workloads to queue sizes is linear, we find instead that it is nonlinear—in fact it can be expressed as the solution to an optimization problem. The objective function of the problem is a natural generalization of the Lyapunov function introduced by Tassiulas and Ephremides [28] for proving stability of the MW-1 policy; the constraints of the problem are closely linked to the canonical representation of workload identified by Harrison [12]. The objective function for MW-α depends on α, and this hints that the performance measures might also depend on α.

Finally, we take note of two related results. First, in [26] we have reported some results about a critically loaded input-queued switch without complete resource pooling condition. Second, a sequence of works by Kelly and Williams [19] and Kang et al. [16] has resulted in a diffusion approximation for a bandwidth sharing network model operating under proportionally fair rate allocation, assuming a technical 'local traffic' condition, but without assuming complete resource pooling. They show that the resulting diffusion approximation model has a product form stationary distribution.

4. The fluid approximation. This section introduces the fluid model and establishes it as an approximation to a fluid-scaled descriptor of the switched network. Intuitively, the fluid model describes the dynamics of the system at the 'rate' level rather than at finer granularity. The reader is referred to a recent monography by Bramson [4] and lecture notes by Dai [5] for a detailed account of fluid approximation for multiclass queueing networks. In Section 4.1 we specify the fluid model, in Section 4.2 we state the main result, and in Section 4.3 we prove it.

4.1. *Definition of fluid model.* Let time be measured by $t \in [0, T]$ for some fixed $T > 0$. Let $\mathbf{q}$, $\mathbf{a}$ and $\mathbf{y}$ all be continuous functions mapping $[0, T]$ into $\mathbb{R}_+^N$, and let $s = (s_{\boldsymbol{\pi}})_{\boldsymbol{\pi} \in \mathcal{S}}$ be a collection of continuous functions mapping $[0, T]$ into $\mathbb{R}_+$. Let $x(\cdot) = (\mathbf{q}(\cdot), \mathbf{a}(\cdot), \mathbf{y}(\cdot), s(\cdot))$. This lies in $C^I(T)$ where $I = 3N + |\mathcal{S}|$. The definition below requires these functions to be absolutely continuous; such functions are differentiable almost everywhere, and the time instants where they are differentiable are called regular times. Any equations we write involving derivatives are taken to apply only at regular times.

Definition 4.1 (Fluid model solution for single-hop network) *Let $f : \mathbb{R}_+ \to \mathbb{R}_+$ satisfy Assumption 2.1. Say that $x(\cdot)$ is a fluid model solution for a single-hop switched network with arrival rate $\boldsymbol{\lambda} \in \mathbb{R}_+^N$ operating under the MW-f policy if it satisfies equations (13)–(20) below. Write* FMS *for the set of all such $x \in C^I(T)$. Additionally, define*

$$\mathrm{FMS}_K = \big\{ x \in \mathrm{FMS} \ : \ |\mathbf{q}(0)| \le K \big\}$$
$$\mathrm{FMS}(\mathbf{q}_0) = \big\{ x \in \mathrm{FMS} \ : \ \mathbf{q}(0) = \mathbf{q}_0 \big\}.$$

The equations are

$$\mathbf{q}(t) = \mathbf{q}(0) + \mathbf{a}(t) - \sum_{\boldsymbol{\pi}} s_{\boldsymbol{\pi}}(t)\boldsymbol{\pi} + \mathbf{y}(t) \tag{13}$$

$$\mathbf{a}(t) = \boldsymbol{\lambda} t \tag{14}$$

$$\sum_{\boldsymbol{\pi}\in\mathcal{S}} s_{\boldsymbol{\pi}}(t) = t \tag{15}$$

$$\mathbf{y}(t) \le \sum_{\boldsymbol{\pi}\in\mathcal{S}} s_{\boldsymbol{\pi}}(t)\boldsymbol{\pi} \tag{16}$$

(17) *each $s_{\boldsymbol{\pi}}(\cdot)$ and $y_n(\cdot)$ is increasing (not necessarily strictly increasing)*

(18) *all the components of $x(\cdot)$ are absolutely continuous*

(19) *for regular times t, all n, $\dot{y}_n(t) = 0$ if $q_n(t) > 0$*

(20) *for regular times t, all $\boldsymbol{\pi} \in \mathcal{S}$, $\dot{s}_{\boldsymbol{\pi}}(t) = 0$ if $\boldsymbol{\pi} \cdot f(\mathbf{q}(t)) < \max_{\boldsymbol{\rho}\in\mathcal{S}} \boldsymbol{\rho} \cdot f(\mathbf{q}(t))$*

Here, $\mathbf{q}(t)$ represents the vector of queue sizes at time t, $\mathbf{a}(t)$ represents the cumulative arrivals up to time t, $\mathbf{y}(t)$ represents the cumulative idleness up to time t, and $s_{\boldsymbol{\pi}}(t)$ represents the total amount of time spent on schedule $\boldsymbol{\pi}$ up to time t. The equation in (13) is the continuous analogue of (4) combined with (5), and the inequality is the analogue of the single-hop version of (7). Equation (14) represents an assumption about the arrival process, related to (11). Equation (15) says that the scheduling policy must choose some schedule at every timestep. Both (16) and (19) derive from the definition of idling, (3). Equation (20) is the continuous analogue of (9).

Definition 4.2 (Fluid model solution for multihop network) *Let $f : \mathbb{R}_+ \to \mathbb{R}_+$ satisfy Assumption 2.2, and let $\mathcal{S}$ satisfy Assumption 2.3. Say that $x(\cdot)$ is a fluid model solution for a multihop switched network operating under the MW-f policy if it satisfies equations (14)–(19), and additionally (21) and (22) below. Let* FMSm *be the set of all such $x \in C^I(T)$. Also, let* FMSm_K *and* $\mathrm{FMSm}(\mathbf{q}_0)$ *be defined* *analogously to the single-hop case. The extra equations are*

$$\mathbf{q}(t) = \mathbf{q}(0) + \mathbf{a}(t) - (I - R^T)\Big(\sum_{\boldsymbol{\pi}} s_{\boldsymbol{\pi}}(t)\boldsymbol{\pi} - \mathbf{y}(t)\Big) \tag{21}$$

(22) *for all regular times t, all $\boldsymbol{\pi} \in \mathcal{S}$,*

$$\dot{s}_{\boldsymbol{\pi}}(t) = 0 \ \textit{ if } \ \boldsymbol{\pi}\cdot(I - R)f(\mathbf{q}(t)) < \max_{\boldsymbol{\rho}\in\mathcal{S}} \boldsymbol{\rho}\cdot(I - R)f(\mathbf{q}(t))$$

When we refer to 'fluid model solutions for any scheduling policy', we mean processes $x(\cdot) \in C^I(T)$ satisfying (13) to (19) in the single-hop case, or satisfying (14) to (19) plus (21) in the multihop case.

4.2. *Main fluid model result.* The development in this section follows the general pattern of Bramson [3]. There is however a difference in presentation that is worth noting. The main result of this section, Theorem 4.3, is a general purpose sample path-wise result: it does not make any probabilistic claim nor does it depend on any probabilistic assumptions. It can be applied to a switched network with stochastic arrivals in two ways: to obtain a result about fluid approximations (Corollary 4.4), and to obtain a result about multiplicative state space collapse (Section 7).

We start by defining the fluid scaling. Consider a switched network of the type described in Section 2.1 running a scheduling policy of the type described in Section 2.2. Write $X(\tau) = (\mathbf{Q}(\tau), \mathbf{A}(\tau), \mathbf{Y}(\tau), S(\tau))$, $\tau \in \mathbb{Z}_+$ to denote its sample path. Given a scaling parameter $z \geq 1$, define the fluid-scaled sample path $\tilde{x}(t) = (\tilde{\mathbf{q}}(t), \tilde{\mathbf{a}}(t), \tilde{\mathbf{y}}(t), \tilde{s}(t))$ for $t \in \mathbb{R}_+$ by

$$\tilde{\mathbf{q}}(t) = \mathbf{Q}(zt)/z \qquad\qquad \tilde{\mathbf{a}}(t) = \mathbf{A}(zt)/z$$
$$\tilde{\mathbf{y}}(t) = \mathbf{Y}(zt)/z \qquad\qquad \tilde{s}_{\boldsymbol{\pi}}(t) = S_{\boldsymbol{\pi}}(zt)/z$$

after extending the domain of $X(\cdot)$ to $\mathbb{R}_+$ by linear interpolation in each interval $(\tau, \tau+1)$. In this section we are interested in the evolution of $\tilde{x}(t)$ over $t \in [0, T]$ for some fixed $T > 0$, therefore we take $\tilde{x}$ to lie in $C^I(T)$ with $I = 3N + |\mathcal{S}|$.

The following theorem concerns uniform convergence of a set of fluid-scaled sample paths. Every fluid-scaled sample path is assumed to relate to some (unscaled) switched network, and all the switched networks are assumed to have the same network data, i.e. the same number of queues N, the same set of allowed schedules $\mathcal{S}$, the same routing matrix R, and the same scheduling policy.

The convergence is indexed by a parameter j lying in some totally ordered countable set. For Corollary 4.4 we will use $j \in \mathbb{N}$, and for Section 7 we will use a subset of $\mathbb{N} \times \mathbb{N}$ as the index set. We are purposefully using the symbol j here as an index, rather than the r used elsewhere, to remind the reader that the index set is interpreted differently in different results.

Theorem 4.3 *Let $\mathcal{X}$ be the set of all possible sample paths for single-hop switched networks with the network data specified above, running the MW-f scheduling policy, where f satisfies Assumption 2.1. Fix $K > 0$ and $\boldsymbol{\lambda} \in \mathbb{R}_+^N$. Let there be sequences $\varepsilon_j \in \mathbb{R}_+$ and $\boldsymbol{\lambda}^j \in \mathbb{R}_+^N$, indexed by j in some totally ordered countable set, such that*

$$\varepsilon_j \to 0 \quad \text{and} \quad \boldsymbol{\lambda}^j \to \boldsymbol{\lambda}, \qquad \text{as } j \to \infty. \tag{23}$$

Consider a sequence of subsets $G_j \subset C^I(T) \times [1, \infty)$ which satisfy the following: for every $(\tilde{x}, z) \in G_j$ there is some unscaled sample path $X \in \mathcal{X}$ such that $\tilde{x}$ is the fluid-scaled version of X with scaling parameter z (here z is permitted to be a function of X); and furthermore

$$\inf\{z : (\tilde{x}, z) \in G_j\} \to \infty \quad \text{as } j \to \infty, \tag{24}$$

$$\sup_{(\tilde{x},z)\in G_j} \sup_{t\in[0,T]} \left|\tilde{\mathbf{a}}(t) - \boldsymbol{\lambda}^j t\right| \le \varepsilon_j \quad \text{for all } j, \text{ and} \tag{25}$$

$$\sup_{(\tilde{x},z)\in G_j} \left|\tilde{\mathbf{q}}(0)\right| \le K \quad \text{for all } j. \tag{26}$$

Then

$$\sup_{(\tilde{x},z)\in G_j} d\big(\tilde{x}, \mathrm{FMS}_K\big) \to 0 \quad \text{as } j \to \infty. \tag{27}$$

Furthermore, fix $\mathbf{q}_0 \in \mathbb{R}_+^N$ and a sequence $\varepsilon'_j \to 0$, and assume that the sets G_j also satisfy

$$\sup_{(\tilde{x},z)\in G_j} \left|\tilde{\mathbf{q}}(0) - \mathbf{q}_0\right| \le \varepsilon'_j \quad \text{for all } j. \tag{28}$$

Then

$$\sup_{(\tilde{x},z)\in G_j} d(\tilde{x}, \mathrm{FMS}(\mathbf{q}_0)) \to 0 \quad \text{as } j \to \infty. \tag{29}$$

Equivalent results to (27) and (29) apply to multihop switched networks, with references to FMS *replaced by* FMSm *and the set $\mathcal{X}$ modified to refer to multihop networks running the MW-f scheduling policy where f satisfies Assumption 2.2 and $\mathcal{S}$ satisfies Assumption 2.3.*

The above theorem as stated applies to the MW-f scheduling policy, but it is clear from the proof that a corresponding limit result holds, relating sample paths of *any* scheduling policy to fluid models defined by equations (13)–(19).

The following corollary is a straightforward application of Theorem 4.3. It specializes the theorem to the case of a single random system X, and the sequence of fluid-scaled versions indexed by $r \in \mathbb{N}$ where the rth version uses scaling parameter r. The arrival process is assumed to satisfy certain stochastic assumptions. This corollary is useful when studying the behaviour of a single switched network with random arrivals, over long timescales.

Corollary 4.4 *Consider a single-hop switched network as described in Section 2.1, running the MW-f policy as described in Section 2.2 where f satisfies Assumption 2.1. Let the arrival process $\mathbf{A}(\cdot)$ satisfy Assumption 2.4, and let the initial queue size $\mathbf{Q}(0) \in \mathbb{R}_+^N$ be random. For $r \in \mathbb{N}$, let*

$$\tilde{\mathbf{q}}^r(t) = \mathbf{Q}(rt)/r \qquad \tilde{\mathbf{a}}^r(t) = \mathbf{A}(rt)/r$$
$$\tilde{\mathbf{y}}^r(t) = \mathbf{Y}(rt)/r \qquad \tilde{s}^r_{\boldsymbol{\pi}}(t) = S_{\boldsymbol{\pi}}(rt)/r,$$

and let $\tilde{x}^r(t) = (\tilde{\mathbf{q}}^r(t), \tilde{\mathbf{a}}^r(t), \tilde{\mathbf{y}}^r(t), \tilde{s}^r(t))$, for $t \in [0, T]$ where $T > 0$ is some fixed time horizon. Then for any $\delta > 0$

$$\mathbb{P}\Big(d\big(\tilde{x}^r(\cdot), \mathrm{FMS}(\mathbf{0})\big) < \delta\Big) \to 1 \quad \text{as } r \to \infty.$$

The same conclusion holds for a multihop switched network running the MW-f back-pressure policy where f satisfies Assumption 2.2 and $\mathcal{S}$ satisfies Assumption 2.3, with FMS *replaced by* FMSm.

Proof. First define the event E_r by

$$E_r = \Big\{\sup_{\tau \le r} \frac{1}{r}\Big|\mathbf{A}(\tau) - \boldsymbol{\lambda}\tau\Big| < \delta_r \text{ and } \big|\mathbf{Q}(0)\big| \le \sqrt{r}\Big\}$$

where $\boldsymbol{\lambda}$ and δ_r are as in Assumption 2.4. By this we mean that E_r is a subset of the probability sample space, and we write $X(\cdot)(\omega)$ etc. for $\omega \in E_r$ to emphasize the dependence on E_r.

We will apply Theorem 4.3 with index set $j \equiv r \in \mathbb{N}$ to the sequence of sets

$$G_j \equiv G_r = \Big\{\big(\tilde{x}^r(\cdot)(\omega), r\big) \;:\; \omega \in E_r\Big\}.$$

In order to apply the theorem we will pick constants as follows. Let $K = 1$, let $\boldsymbol{\lambda}$ be as in Assumption 2.4, $\boldsymbol{\lambda}^j = \boldsymbol{\lambda}$ for all j, $\varepsilon_j \equiv \varepsilon_r = T\delta_r$ where δ_r is as in Assumption 2.4, $\mathbf{q}_0 = \mathbf{0}$ and $\varepsilon'_r = 1/\sqrt{r}$. We now need to verify the conditions of Theorem 4.3. Equation (23) holds by the choice of $\boldsymbol{\lambda}^j$ and by Assumption 2.4. Equation (24) holds automatically by choice of G_j. To see that (25) holds, rewrite event E_r in terms of the fluid scaled arrival process $\tilde{\mathbf{a}}^r$ to see

$$\sup_{t \in [0,T]} \Big|\tilde{\mathbf{a}}^r(t)(\omega) - \boldsymbol{\lambda}t\Big| < T\delta_r \quad \text{for all } r \text{ and } \omega \in E_r$$

which implies (25). Likewise (26) and (28). We conclude that (29) holds. It may be rewritten in terms of E_r as

$$\sup_{\omega\in E_r} d\big(\tilde{x}^r(\cdot)(\omega), \mathrm{FMS}(\mathbf{0})\big) \to 0 \quad \text{as } r\to\infty. \tag{30}$$

We next argue that $\mathbb{P}(E_r) \to 1$ as $r\to\infty$. The event E_r is the intersection of two events, one concerning arrivals and the other concerning initial queue size. The probability of the former $\to 1$ as $r \to \infty$ by Assumption 2.4. For the latter, $\mathbb{P}(|\mathbf{Q}(0)| \le \sqrt{r}) \to 1$ as $r\to\infty$ since $\mathbf{Q}(0)$ is assumed not to be infinite. Therefore $\mathbb{P}(E_r) \to 1$. Combining this with (30) gives the desired result for single-hop networks. The multihop version follows similarly. □

4.3. *Proof of Theorem 4.3.* We shall present the proof of Theorem 4.3 for a single-hop network in detail followed by main ideas required to extend it to multihop networks.

4.3.1. *Cluster points.* Here we are interested in convergence in $C^I(T)$, where $I = 3N + |\mathcal{S}|$ and $T > 0$ is fixed, equipped with the norm $\|\cdot\|$. The appropriate concept for proving convergence is *cluster points*. Consider any metric space E with metric d and a sequence $(E_1, E_2, \dots)$ of subsets of E. Say that $x \in E$ is a *cluster point* of the sequence if $\liminf_{j\to\infty} d(x, E_j) = 0$ where $d(x,E_j) = \inf\{d(x,y) : y\in E_j\}$.

Proposition 4.5 (Cluster points in $C^I(T)$) [2] *Given $K > 0$, $A > 0$ and a sequence $B_j \to 0$, let*

$$K_j = \Big\{x \in C^I(T) \;:\; |x(0)| \le K \;\; \textit{and} \;\; \mathrm{mc}_\delta(x) \le A\delta + B_j \textit{ for all } \delta > 0\Big\}$$

and consider a sequence $(E_1, E_2, \dots)$ of subsets of $C^I(T)$ for which $E_j \subset K_j$. Then $\sup_{y\in E_j} d(y, CP) \to 0$ as $j\to\infty$, where CP is the set of cluster points of $(E_1, E_2, \dots)$.

4.3.2. *Proof of Theorem 4.3.* Let $E_j = \{\tilde{x} : (\tilde{x}, z) \in G_j\}$. Lemma 4.6 below shows that $E_j \subset K_j$, with K_j as defined in Proposition 4.5 for appropriate constants K, A and B_j. By applying that proposition,

$$\sup_{\tilde{x}\in E_j} d\big(\tilde{x}, \mathrm{CP}\big) \to 0 \quad \text{as } j\to\infty$$

where CP is the set of cluster points of the E_j sequence. Lemma 4.7 below shows that all cluster points of the E_j sequence satisfy the fluid model equations. Every cluster point x must also satisfy $|\mathbf{q}(0)| \le K$, by (26). Therefore

$$\sup_{\tilde{x}\in E_j} d\big(\tilde{x}, \mathrm{FMS}_K\big) \to 0 \quad \text{as } j\to\infty.$$

[2] taken from Bramson [3, Proposition 4.1]

If in addition (28) holds then every cluster point x must also satisfy $\mathbf{q}(0) = \mathbf{q}_0$. Therefore

$$\sup_{\tilde{x}\in E_j} d\big(\tilde{x}, \mathrm{FMS}(\mathbf{q}_0)\big) \to 0 \quad \text{as } j \to \infty.$$

□

Lemma 4.6 (Tightness of fluid scaling) *Let K and G_j be as in Theorem 4.3. Then there exist a constant $A > 0$ and a sequence $B_j \to 0$ such that for every $(\tilde{x}, z) \in G_j$, $|\tilde{x}(0)| \le K$ and*

$$\big|\tilde{x}(u) - \tilde{x}(t)\big| \le A\,|u-t| + B_j \quad \textit{for all } 0 \le t, u \le T.$$

Proof. Consider $(\tilde{x}, z) \in G_j$, where $\tilde{x} = (\tilde{\mathbf{q}}, \tilde{\mathbf{a}}, \tilde{\mathbf{y}}, \tilde{s})$. As per the definitions in Section 2.1, the only non-zero component of $\tilde{x}(0)$ is $\tilde{\mathbf{q}}(0)$, and $|\tilde{\mathbf{q}}(0)| \le K$ by choice of G_j, hence $|\tilde{x}(0)| \le K$. For the second inequality, without loss of generality pick any $0 \le t < u \le T$, and let us now look at each component of $|\tilde{x}(u) - \tilde{x}(t)|$ in turn.

For arrivals, let $\lambda^{\max} = \sup_j |\boldsymbol{\lambda}^j|$; this is finite by the assumption that $\boldsymbol{\lambda}^j \to \boldsymbol{\lambda}$ in Theorem 4.3. Then for $(\tilde{x}, z) \in G_j$,

$$\begin{aligned}
\big|\tilde{\mathbf{a}}(u) - \tilde{\mathbf{a}}(t)\big| &\le \big|\tilde{\mathbf{a}}(u) - \boldsymbol{\lambda}^j u\big| + \big|\boldsymbol{\lambda}^j (u-t)\big| + \big|\tilde{\mathbf{a}}(t) - \boldsymbol{\lambda}^j t\big| \\
&\le 2\varepsilon_j + |\boldsymbol{\lambda}^j|(u-t) \quad \text{by (25)} \\
&\le 2\varepsilon_j + \lambda^{\max}(u-t).
\end{aligned}$$

For idling, let $S^{\max} = \max_{\boldsymbol{\pi}\in\mathcal{S}} \max_n \pi_n$. This is the maximum amount of service that can be offered to any queue per unit time, and it must be finite since $|\mathcal{S}|$ is finite. Then, based on (3),

$$\begin{aligned}
\big|\tilde{y}_n(u) - \tilde{y}_n(t)\big| &\le (u-t)\, S^{\max} + 2S^{\max}/z \\
&\le (u-t) S^{\max} + 2S^{\max}/z_j^{\min},
\end{aligned}$$

where $z_j^{\min} = \inf\{z : (\tilde{x}, z) \in G_j\}$. For service, let $S_{\boldsymbol{\pi}}(\cdot)$ be the unscaled process that corresponds to $\tilde{s}_{\boldsymbol{\pi}}(\cdot)$; since $S_{\boldsymbol{\pi}}(\cdot)$ is increasing and since a schedule must be chosen not more than once every timeslot,

$$\big|\tilde{s}_{\boldsymbol{\pi}}(u) - \tilde{s}_{\boldsymbol{\pi}}(t)\big| \le \frac{1}{z}\Big(S_{\boldsymbol{\pi}}(\lceil zu \rceil) - S_{\boldsymbol{\pi}}(\lfloor zt \rfloor)\Big) \le (u-t) + 2/z \le (u-t) + 2/z_j^{\min}.$$

For queue size, note that (4) carries through to the fluid model scaling, i.e.

$$\tilde{\mathbf{q}}(t) = \tilde{\mathbf{q}}(0) + \tilde{\mathbf{a}}(t) - \sum_{\boldsymbol{\pi}} \tilde{s}_{\boldsymbol{\pi}}(t)\boldsymbol{\pi} + \tilde{\mathbf{y}}(t),$$

thus

$$\begin{aligned}\big|\tilde{q}_n(u)-\tilde{q}_n(t)\big| &\le \big|\tilde{a}_n(u)-\tilde{a}_n(t)\big| + \sum_{\boldsymbol{\pi}} \pi_n \big|\tilde{s}_{\boldsymbol{\pi}}(u)-\tilde{s}_{\boldsymbol{\pi}}(t)\big| + \big|\tilde{y}_n(u)-\tilde{y}_n(t)\big| \\ &\le (u-t)\Big(\lambda^{\max} + |\mathcal{S}|S^{\max} + S^{\max}\Big) + \Big(2|\mathcal{S}|S^{\max} + 2S^{\max}\Big)/z_j^{\min} + 2\varepsilon_j.\end{aligned}$$

Putting all these together,

$$|\tilde{x}(u)-\tilde{x}(t)| \le A(u-t) + B_j, \tag{31}$$

where the constants are

$$\begin{aligned} A &= (N+1)\lambda^{\max} + 2NS^{\max} + |\mathcal{S}| + N|\mathcal{S}|S^{\max}, \quad \text{and} \\ B_j &= \Big(4NS^{\max} + 2|\mathcal{S}| + 2N|\mathcal{S}|S^{\max}\Big)\frac{1}{z_j^{\min}} + (2+2N)\varepsilon_j. \end{aligned}$$

By the assumptions of Theorem 4.3, $\varepsilon_j \to 0$ and $z_j^{\min} \to \infty$ as $j \to \infty$, thus $B_j \to 0$ as required. □

Lemma 4.7 (Dynamics at cluster points) *Make the same assumptions as Theorem 4.3 and let $E_j = \{\tilde{x} : (\tilde{x}, z) \in G_j\}$. Then $x \in \mathrm{FMS}_K$ if x is a cluster point of the E_j sequence.*

Proof. From Lemma 4.6 and Proposition 4.5, it follows that $\limsup_{\tilde{x}\in E_j} d(\tilde{x}, \mathrm{CP}) \to 0$ as $j \to \infty$ where CP is the set of cluster points of the sequence E_j. Let x be one such cluster point. That is, there exists a subsequence j_k and a collection $\tilde{x}^{j_k} \in E_{j_k}$ such that $\tilde{x}^{j_k} \to x$. It easily follows that $|x(0)| \le K$ since $|\tilde{x}^{j_k}(0)| \le K$ for all $\tilde{x}^{j_k} \in E_{j_k}$ as argued in Lemma 4.6. Using this, we wish to establish that x satisfies all the fluid model equations to conclude $x \in \mathrm{FMS}_K$. For convenience, we shall omit the subscript k in the rest of the proof, that is we shall use j in place of j_k and $j \to \infty$.

Proof of (13), (15), (17). The discrete (unscaled) system satisfies these properties, therefore the scaled systems $\tilde{x}^j$ do too. Taking the limit yields the fluid equations.

Proof of (16). In (3), $d\mathbf{B}(\tau)$ and $\mathbf{Q}(\tau)$ are both non-negative (component-wise), hence $d\mathbf{Y}(\tau) \le d\mathbf{B}(\tau)$ for all τ. Summing up over τ, we see the discrete (unscaled) system satisfies the equivalent of (16), so as above we obtain the fluid equation.

Proof of (14). Observe that

$$\sup_{t\in[0,T]} \big|\mathbf{a}(t) - \boldsymbol{\lambda} t\big| \le \sup_{t\in[0,T]} \big|\mathbf{a}(t) - \tilde{\mathbf{a}}^j(t)\big| + \sup_{t\in[0,T]} \big|\tilde{\mathbf{a}}^j(t) - \boldsymbol{\lambda}^j t\big| + T\big|\boldsymbol{\lambda}^j - \boldsymbol{\lambda}\big|.$$

Each term converges to 0 as $j \to \infty$: the first because $\tilde{x}^j \to x$, the second because $\tilde{x}^j \in E_j$ so the deviation in $\tilde{\mathbf{a}}^j$ is bounded by ε_j and $\varepsilon_j \to 0$, and the third because $\boldsymbol{\lambda}^j \to \boldsymbol{\lambda}$. Since the left hand side does not depend on j, it must be that $\mathbf{a}(t) = \boldsymbol{\lambda} t$.

Proof of (18). In Lemma 4.6 we found constants A and B_j such that for all $\tilde{x} \in E_j$

$$\left|\tilde{x}(u) - \tilde{x}(t)\right| \leq A|u-t| + B_j,$$

with $B_j \to 0$ as $j \to \infty$. Taking the limit of $|\tilde{x}^j(u) - \tilde{x}^j(t)|$ as $j \to \infty$, we find that $|x(u) - x(t)| \leq A|u-t|$, i.e. x is (globally) Lipschitz continuous (of order 1 with respect to the appropriate metric as defined earlier). This immediately implies that x is absolutely continuous.

Proof of (19). Since $x = (\mathbf{q}, \mathbf{a}, \mathbf{y}, s)$ is absolutely continuous, each component is too, which means that y_n is differentiable for almost all t. Pick some such t, and suppose that $q_n(t) > 0$. Consider some small interval $I = [t, t+\delta]$ about t. Since q_n is continuous, we can choose δ sufficiently small that $\inf_{s\in I} q_n(s) > 0$. Since $\|\tilde{\mathbf{q}}^j - \mathbf{q}\| \to 0$, we can find $c > 0$ such that $\inf_{s \in I} \tilde{q}^j_n(s) > c$ for all j sufficiently large. Since $\tilde{x}^j \in E_j$, there exists a corresponding unscaled version of the system, say X^j, and scaling parameter, say z_j, so that $\tilde{x}^j(\cdot) = X^j(z_j \cdot)/z_j$. Therefore, it must be that the corresponding unscaled queue satisfies $\inf_{s\in I} Q^j_n(z_j s) > z_j c$. That is, the queue size in the entire interval never vanishes to 0 and hence idling in the entire interval is not possible. Therefore after rescaling we find $\tilde{y}^j_n(t+\delta/2) = \tilde{y}^j_n(t)$. (The switch from δ to $\delta/2$ sidesteps any discretization problems.) Therefore the same holds for y_n in the limit. We assumed y_n to be differentiable at t, the derivative must be 0.

Proof of (20). Pick a t at which $s_{\boldsymbol{\pi}}$ is differentiable, and suppose that $\boldsymbol{\pi} \cdot f(\mathbf{q}(t)) < \max_{\boldsymbol{\rho}} \boldsymbol{\rho} \cdot f(\mathbf{q}(t))$. As above, pick some small interval $I = [t, t+\delta]$ and j sufficiently large that

$$\boldsymbol{\pi} \cdot f\big(\tilde{\mathbf{q}}^j(s)\big) < \max_{\boldsymbol{\rho}\in\mathcal{S}} \boldsymbol{\rho} \cdot f\big(\tilde{\mathbf{q}}^j(s)\big) \quad \text{for } s \in I.$$

Writing this in terms of the unscaled system and applying Assumption 2.1,

$$\boldsymbol{\pi} \cdot f\big(\mathbf{Q}^j(z_j s)\big) < \max_{\boldsymbol{\rho}\in\mathcal{S}} \boldsymbol{\rho} \cdot f\big(\mathbf{Q}^j(z_j s)\big) \quad \text{for } s \in I.$$

The MW-f policy ensures by (9) that $\boldsymbol{\pi}$ will not be chosen throughout this entire interval, so after rescaling we find $\tilde{s}^j_{\boldsymbol{\pi}}(t+\delta/2) - \tilde{s}^j_{\boldsymbol{\pi}}(t) = 0$, and taking the limit gives $s_{\boldsymbol{\pi}}(t+\delta/2) = s_{\boldsymbol{\pi}}(t)$. Since $s_{\boldsymbol{\pi}}$ is assumed to be differentiable at t; the derivative must be 0. ∎

4.3.3. *Proof of Theorem 4.3 for multihop networks.* The proof of Theorem 4.3 for single-hop network applies verbatim, except that the two lemmas need to be replaced.

Lemma 4.6 (Tightness of fluid scaling) $\longrightarrow$ Lemma 4.8
Lemma 4.7 (Dynamics at cluster points) $\longrightarrow$ Lemma 4.9

Lemma 4.8 (Tightness of fluid scaling) *Make the same assumptions as Theorem 4.3, multihop case, and use the same definition of G_j. Then there exist a constant $A > 0$ and a sequence $B_j \to 0$ such that for every $(\tilde{x}, z) \in G_j$, $|\tilde{x}(0)| \le K$ and*

$$\left|\tilde{x}(u) - \tilde{x}(t)\right| \le A\,|u-t| + B_j \quad \text{for all } 0 \le t, u \le T.$$

Proof. Consider $(\tilde{x}, z) \in G_j$, $\tilde{x} = (\tilde{\mathbf{q}}, \tilde{\mathbf{a}}, \tilde{\mathbf{y}}, \tilde{s})$. The bound $|\tilde{x}(0)| \le K$ follows from argument similar to that in the single-hop case. The bounds on the arrival process, the idleness, and service allocation, are as in the single-hop case: for any $0 \le t < u \le T$,

$$\begin{aligned}
\left|\tilde{\mathbf{a}}(u) - \tilde{\mathbf{a}}(t)\right| &\le (u-t)\lambda^{\max} + 2\varepsilon_j,\\
\left|\tilde{y}_n(u) - \tilde{y}_n(t)\right| &\le (u-t)\,S^{\max} + 2S^{\max}/z_j^{\min},\\
\left|\tilde{s}_{\boldsymbol{\pi}}(t) - \tilde{s}_{\boldsymbol{\pi}}(s)\right| &\le (u-t) + 2/z_j^{\min},
\end{aligned}$$

where $z_j^{\min} = \inf\{z : (\tilde{x}, z) \in G_j\}$. The bound on queue size is a little different. Note that (6) carries through to the fluid-scaling, i.e.

$$\tilde{\mathbf{q}}(t) = \tilde{\mathbf{q}}(0) + \tilde{\mathbf{a}}(t) - (I - R^{\mathsf{T}}) \sum_{\boldsymbol{\pi}} \tilde{s}_{\boldsymbol{\pi}}(t)\boldsymbol{\pi} + \tilde{\mathbf{y}}(t),$$

thus

$$\begin{aligned}
\left|\tilde{q}_n(u) - \tilde{q}_n(t)\right| &\le \left|\tilde{a}_n(u) - \tilde{a}_n(t)\right| + \sum_{\boldsymbol{\pi}} \left|[(I - R^{\mathsf{T}})\boldsymbol{\pi}]_n\right| \left|\tilde{s}_{\boldsymbol{\pi}}(u) - \tilde{s}_{\boldsymbol{\pi}}(t)\right| + \left|\tilde{y}_n(u) - \tilde{y}_n(t)\right|\\
&\le (u-t)\Big(\lambda^{\max} + |\mathcal{S}|(NS^{\max})S^{\max} + S^{\max}\Big)\\
&\qquad + \Big(2|\mathcal{S}|(NS^{\max})S^{\max} + 2S^{\max}\Big)/z_j^{\min} + 2\varepsilon_j.
\end{aligned}$$

Putting all these together, for any $(\tilde{x}, z) \in G_j$,

$$|\tilde{x}(u) - \tilde{x}(t)| \le A(t-s) + B_j,$$

where the constants A and B_j are

$$\begin{aligned}
A &= (1+N)\lambda^{\max} + 2NS^{\max} + |\mathcal{S}| + |\mathcal{S}|(NS^{\max})^2\\
B_j &= \Big(4NS^{\max} + 2|\mathcal{S}| + 2|\mathcal{S}|(NS^{\max})^2\Big)\frac{1}{z_j^{\min}} + (2N+2)\varepsilon_j.
\end{aligned} \tag{32}$$

Here $B_j \to 0$ as $j \to \infty$ since $\varepsilon_j \to 0$ by (23) and $z_j^{\min} \to \infty$ as $j \to \infty$ by (24). □

Lemma 4.9 (Dynamics at cluster points) *Under the setup of Theorem 4.3 for a multihop network, let $E_j = \{\tilde{x} : (\tilde{x}, z) \in G_j\}$. Then $x \in \mathrm{FMSm}_K$ if x is a cluster point of the E_j sequence.*

Proof. Given a cluster point $x = (\mathbf{q}, \mathbf{a}, \mathbf{y}, s)$, let there be $(\tilde{x}^j, z_j) \in G_j$ so that $\tilde{x}^j \to x$, as in the proof of Lemma 4.7. Now the bound $|x(0)| \leq K$ and equations (14)–(19) all work exactly as in the single-hop case, as does the queue size equation (21). The only equation that needs further argument is the MW-f backpressure equations (22).

Proof of (22). Pick a t at which $s_{\boldsymbol{\pi}}$ is differentiable, and suppose that $\boldsymbol{\pi} \cdot (I - R) f(\mathbf{q}(t)) < \max_{\boldsymbol{\rho}} \boldsymbol{\rho} \cdot (I - R) f(\mathbf{q}(t))$. As in Lemma 4.7, proof of (20), it must be that there is some small interval $I = [z_j t, z_j t + z_j \delta]$ such that $\boldsymbol{\pi}$ is not chosen for any $\tau \in I$, therefore $\dot{s}_{\boldsymbol{\pi}}(t) = 0$. □

5. Fluid model behaviour (single-hop case).

In this section we prove certain properties of fluid model solutions, which will be needed for the main result of this paper, multiplicative state space collapse. In order to state these properties, we first need some definitions. We then state a portmanteau theorem listing all the properties, and give an example to illustrate the definitions. The rest of the section is given over to proofs and supplementary lemmas.

This section deals with a single-hop switched network; in the next section we give corresponding results for multihop. Our reason for giving separate single-hop and multihop proofs, rather than just treating single-hop as a special case of multihop, is that our multihop results place additional restrictions on the set of allowed schedules (Assumption 2.3) beyond what is required for single-hop networks. This mainly affects the proof; the portmanteau theorem for multihop is nearly identical to that for single-hop.

Definition 5.1 (Admissible region) *Let $\mathcal{S} \subset \mathbb{R}_+^N$ be the set of allowed schedules. Let $\langle \mathcal{S} \rangle$ be the convex hull of $\mathcal{S}$,*

$$\langle \mathcal{S} \rangle = \Big\{ \sum_{\boldsymbol{\pi} \in \mathcal{S}} \alpha_{\boldsymbol{\pi}} \boldsymbol{\pi} : \sum_{\boldsymbol{\pi} \in \mathcal{S}} \alpha_{\boldsymbol{\pi}} = 1, \text{ and } \alpha_{\boldsymbol{\pi}} \geq 0 \text{ for all } \boldsymbol{\pi} \Big\}.$$

Define the admissible region Λ *to be*

$$\Lambda = \big\{ \boldsymbol{\lambda} \in \mathbb{R}_+^N : \boldsymbol{\lambda} \leq \boldsymbol{\sigma} \text{ componentwise, for some } \boldsymbol{\sigma} \in \langle \mathcal{S} \rangle \big\}.$$

Definition 5.2 (Static planning problems and virtual resources) *Define the optimization problem* PRIMAL($\boldsymbol{\lambda}$) *for $\boldsymbol{\lambda} \in \mathbb{R}_+^N$ to be*

$$\begin{aligned} &\text{minimize} && \sum_{\boldsymbol{\pi} \in \mathcal{S}} \alpha_{\boldsymbol{\pi}} \quad \text{over} \quad \alpha_{\boldsymbol{\pi}} \in \mathbb{R}_+ \text{ for all } \boldsymbol{\pi} \in \mathcal{S} \\ &\text{such that} && \boldsymbol{\lambda} \leq \sum_{\boldsymbol{\pi} \in \mathcal{S}} \alpha_{\boldsymbol{\pi}} \boldsymbol{\pi} \text{ componentwise} \end{aligned}$$

Let DUAL($\boldsymbol{\lambda}$) *be the dual to this: it is*

$$\begin{array}{lll} \textit{maximize} & \boldsymbol{\xi}\cdot\boldsymbol{\lambda} \quad \textit{over} \quad \boldsymbol{\xi} \in \mathbb{R}_+^N \\ \textit{such that} & \max_{\boldsymbol{\pi}\in\mathcal{S}} \boldsymbol{\xi}\cdot\boldsymbol{\pi} \leq 1 \end{array}$$

Let E be the set of extreme points of the feasible region of the dual problem; the feasible region is a finite convex polytope so E is finite. Define the set of virtual resources $\mathcal{S}^* \subset \mathbb{R}_+^N$ *to be the set of maximal extreme points,*

$$\mathcal{S}^* = \big\{\boldsymbol{\xi} \in E : \textit{ for all } \boldsymbol{\zeta} \in E,\ \boldsymbol{\xi} \leq \boldsymbol{\zeta} \implies \boldsymbol{\xi} = \boldsymbol{\zeta}\big\}.$$

Define the set of critically loaded virtual resources $\Xi(\boldsymbol{\lambda})$ *to be*

$$\Xi(\boldsymbol{\lambda}) = \big\{\boldsymbol{\xi} \in \mathcal{S}^* : \boldsymbol{\xi}\cdot\boldsymbol{\lambda} = 1\big\}.$$

Both problems are clearly feasible, and the optimum is attained in each. By Slater's condition there is strong duality, i.e. PRIMAL($\boldsymbol{\lambda}$) = DUAL($\boldsymbol{\lambda}$). (When we write PRIMAL($\boldsymbol{\lambda}$) or DUAL($\boldsymbol{\lambda}$) in mathematical expressions, we mean the optimum value, not the optimizer.) Clearly, PRIMAL($\boldsymbol{\lambda}$) ≤ 1 if and only if $\boldsymbol{\lambda}$ is feasible.

Laws [22, 23] and Kelly and Laws [18] used primal and dual problems of this general sort for describing multihop queueing networks with routing choice. Harrison [12] extended the problems for stochastic processing networks.

Definition 5.3 (Lyapunov function and lifting map) *Let the scheduling policy be MW-f, where f satisfies Assumption 2.1. Define the function $L : \mathbb{R}_+^N \to \mathbb{R}_+$ by*

$$L(\mathbf{q}) = F(\mathbf{q})\cdot\mathbf{1}$$

where $F(x) = \int_0^x f(y)\,dy$ for $x \in \mathbb{R}$, and $F(\mathbf{q}) = [F(q_n)]_{1\leq n\leq N}$ as per the notation in Section 1. Define the optimization problem ALGD($\mathbf{q}$) *to be*

$$\begin{array}{lll} \textit{minimize} & L(\mathbf{r}) \quad \textit{over} \quad \mathbf{r} \in \mathbb{R}_+^N \\ \textit{such that} & \boldsymbol{\xi}\cdot\mathbf{r} \geq \boldsymbol{\xi}\cdot\mathbf{q} \textit{ for all } \boldsymbol{\xi} \in \Xi(\boldsymbol{\lambda}) \\ & \textit{and} \quad r_n \leq q_n \textit{ for all } n \textit{ such that } \lambda_n = 0. \end{array}$$

Note that F is strictly convex and increasing, and the feasible region is convex, hence this problem has a unique optimizer. Define the lifting map $\Delta W : \mathbb{R}_+^N \to \mathbb{R}_+^N$ *by setting $\Delta W(\mathbf{q})$ to be the optimizer.*

Note that ALGD and ΔW both depend on $\boldsymbol{\lambda}$ and f, but we will surpress this dependency when the context makes it clear which $\boldsymbol{\lambda}$ and f are meant.

The results in this section apply to any $\boldsymbol{\lambda} \in \Lambda$. However, if $\mathrm{PRIMAL}(\boldsymbol{\lambda}) < 1$ then $\Xi(\boldsymbol{\lambda})$ is empty, so $\Delta W(\mathbf{q}) = \mathbf{0}$ for all $\mathbf{q}$. The results are only interesting when $\mathrm{PRIMAL}(\boldsymbol{\lambda}) = 1$, so we define

$$\partial\Lambda = \big\{\boldsymbol{\lambda} \in \Lambda : \mathrm{PRIMAL}(\boldsymbol{\lambda}) = 1\big\}.$$

We can now state the main result of this section.

Theorem 5.4 (Portmanteau theorem, single-hop version) *Let $\boldsymbol{\lambda} \in \Lambda$. Consider a single-hop switched network running MW-f, where f satisfies Assumption 2.1.*

(i) *For any $K < \infty$, $\{\mathbf{q} \in \mathbb{R}_+^N : L(\mathbf{q}) \leq K\}$ is compact. Also, for any fluid model solution with arrival rate $\boldsymbol{\lambda}$, $L(\mathbf{q}(t)) \leq L(\mathbf{q}(0))$ for all $t \geq 0$.*
(ii) *ΔW is continuous.*
(iii) *If $\mathbf{q} = \Delta W(\mathbf{q})$ then $\Delta W(\kappa\mathbf{q}) = \kappa\,\Delta W(\mathbf{q})$ for all $\kappa > 0$.*
(iv) *Say that $\mathbf{q}^0$ is an* invariant state *if all fluid model solutions $\mathbf{q}(\cdot)$ with arrival rate $\boldsymbol{\lambda}$, starting at $\mathbf{q}(0) = \mathbf{q}^0$, satisfy $\mathbf{q}(t) = \mathbf{q}^0$ for all $t \geq 0$. Then $\mathbf{q}^0$ is an invariant state $\iff \mathbf{q}^0 = \Delta W(\mathbf{q}^0)$.*
(v) *For any $\varepsilon > 0$ there exists some $H_\varepsilon < \infty$ such that, if $\mathbf{q}(\cdot)$ is a fluid model solution with arrival rate $\boldsymbol{\lambda}$, and $|\mathbf{q}(0)| \leq 1$, then $|\mathbf{q}(t) - \Delta W(\mathbf{q}(t))| < \varepsilon$ for all $t \geq H_\varepsilon$.*

A loose interpretation of these results is that the MW-f scheduling policy seeks always to reduce $L(\mathbf{q})$ (part (i)), but it is constrained from reducing it too much, because it is not permitted to reduce the workload at any of the critically loaded virtual resource (the constraints of ALGD). However, it can choose how to allocate work between queues, subject to those constraints. It heads towards a state where it is impossible to reduce $L(\mathbf{q})$ any further (parts (iv) & (v)). In all the examples we have looked at, the fluid model solutions reach an invariant state in finite time, i.e. (v) holds also for $\varepsilon = 0$, but we have not been able to prove this in general.

5.1. *Example to illustrate Λ, $\partial\Lambda$, $\mathcal{S}^*$ and Ξ.* Consider a system with $N = 2$ queues, A and B. Suppose the set $\mathcal{S}$ of possible schedules consists of "serve three packets from queue A" and "serve one packet each from A and B". Write these two schedules as $\boldsymbol{\pi}^1 = (3, 0)$ and $\boldsymbol{\pi}^2 = (1, 1)$ respectively. Let λ_A and λ_B be the arrival rates at the two queues, measured in packets per timeslot.

Determining Λ and $\partial\Lambda$. The arrival rate vector $\boldsymbol{\lambda} = (\lambda_A, \lambda_B)$ is feasible if there is some $\boldsymbol{\sigma} = (1-x)\boldsymbol{\pi}^1 + x\boldsymbol{\pi}^2$ with $0 \leq x \leq 1$ such that $\boldsymbol{\lambda} \leq \boldsymbol{\sigma}$. In words, the arrival rates are feasible if the switch can divide its time between the two possible schedules in such a way that the service rates at the two queues are at least as big as the arrival rates. Schedule $\boldsymbol{\pi}^2$ is the only schedule which serves queue B, so we would need $x \geq \lambda_B$. If $\lambda_B > 1$ then it is impossible to serve all the work that arrives at queue B. Otherwise, we may as well set $x = \lambda_B$. The total amount of service given to queue A is then $3(1-x)+x = 3-2\lambda_B$; if $\lambda_A \leq 3-2\lambda_B$ then it is possible to serve all the work arriving at queue A. We have concluded that

$$\Lambda = \Big\{(\lambda_A, \lambda_B) : \lambda_B \leq 1 \text{ and } \frac{1}{3}\lambda_A + \frac{2}{3}\lambda_B \leq 1\Big\}.$$

Further algebra tells us that

$$\mathrm{PRIMAL}(\boldsymbol{\lambda}) = \max\Big(\lambda_B, \frac{1}{3}\lambda_A + \frac{2}{3}\lambda_B\Big).$$

Hence

$$\partial\Lambda = \Big\{(\lambda_A, \lambda_B) \in \Lambda : \lambda_B = 1 \text{ or } \frac{1}{3}\lambda_A + \frac{2}{3}\lambda_B = 1\Big\}.$$

Determining $\mathcal{S}^*$ and Ξ. The feasible region of $\mathrm{DUAL}(\boldsymbol{\lambda})$ is

$$\Big\{(\xi_1, \xi_2) \in \mathbb{R}^2_+ : 3\xi_1 \leq 1 \text{ and } \xi_1 + \xi_2 \leq 1\Big\}.$$

The extreme points may be found by sketching the feasible region; they are $(0,0)$, $(1/3, 0)$, $(1/3, 2/3)$, and $(0,1)$. Clearly the maximal extreme points, i.e. the virtual resources, are

$$\mathcal{S}^* = \big\{(1/3, 2/3),\ (0,1)\big\}.$$

The set of critically loaded virtual resources depends on λ_A and λ_B: $(0,1) \in \Xi(\boldsymbol{\lambda})$ iff $\lambda_B = 1$, and $(1/3, 2/3) \in \Xi(\boldsymbol{\lambda})$ iff $\lambda_A/3 + 2\lambda_B/3 = 1$.

Interpretation of virtual resources[1]. Each virtual resource $\boldsymbol{\xi} \in \mathcal{S}^*$ may be interpreted as a virtual queue. For example, take $\boldsymbol{\xi} = (1/3, 2/3)$, and define the virtual queue size to be $\boldsymbol{\xi}\cdot\mathbf{Q} = Q_A/3 + 2Q_B/3$. Think of the virtual queue as consisting of tokens: every time a packet arrives to queue A put $1/3$ tokens into the virtual queue, and every time a packet arrives to queue B put in $2/3$ tokens. The schedule $\boldsymbol{\pi}^1$ can remove at most $3 \times 1/3 = 1$ token, and schedule $\boldsymbol{\pi}^2$ can remove at most $1/3 + 2/3 = 1$ token. In order that the total rate at which tokens arrive should be no more than the maximum rate at which we can remove tokens, we need

$$\lambda_A/3 + 2\lambda_B/3 \leq 1$$

i.e. $\boldsymbol{\lambda}\cdot\boldsymbol{\xi} \leq 1$. If $\mathrm{DUAL}(\boldsymbol{\lambda}) = \mathrm{PRIMAL}(\boldsymbol{\lambda}) > 1$, then there is some $\boldsymbol{\xi} \in \mathcal{S}^*$ such that $\boldsymbol{\lambda}\cdot\boldsymbol{\xi} > 1$, which means that the corresponding virtual queue is unstable, hence the original system is unstable.

[3] cf. Laws [22, Example 4.4.3]

5.2. *Proofs for the portmanteau theorem.* Throughout this subsection we consider a single-hop switched network running MW-f with arrival rates $\boldsymbol{\lambda} \in \Lambda$.

The first claim of Theorem 5.4(i), that $\{\mathbf{q} \in \mathbb{R}_+^N : L(\mathbf{q}) \le K\}$ is compact for any $K < \infty$, follows straightforwardly from the facts that $L(\mathbf{q}) \to \infty$ as $|\mathbf{q}| \to \infty$, and $L(\cdot)$ is continuous. The second claim follows from a standard result (first given by Dai and Prabhakar [8], for an input-queued switch), which we include here for the sake of completeness.

Lemma 5.5 *For all* $\mathbf{q} \in \mathbb{R}_+^N$,

$$\boldsymbol{\lambda} \cdot f(\mathbf{q}) - \max_{\boldsymbol{\pi} \in \mathcal{S}} \boldsymbol{\pi} \cdot f(\mathbf{q}) \le 0. \tag{33}$$

Also, every fluid model solution satisfies

$$\frac{d}{dt} L(\mathbf{q}(t)) = \boldsymbol{\lambda} \cdot f(\mathbf{q}(t)) - \max_{\boldsymbol{\pi} \in \mathcal{S}} \boldsymbol{\pi} \cdot f(\mathbf{q}(t)) \le 0.$$

Proof. Since $\boldsymbol{\lambda} \in \Lambda$, we can write $\boldsymbol{\lambda} \le \boldsymbol{\sigma}$ componentwise for some $\boldsymbol{\sigma} = \sum_{\boldsymbol{\pi}} \alpha_{\boldsymbol{\pi}} \boldsymbol{\pi}$ with $\alpha_{\boldsymbol{\pi}} \ge 0$ and $\sum \alpha_{\boldsymbol{\pi}} = 1$. Hence

$$\begin{aligned} \boldsymbol{\lambda} \cdot f(\mathbf{q}) - \max_{\boldsymbol{\rho}} \boldsymbol{\rho} \cdot f(\mathbf{q}) &= \sum_{\boldsymbol{\pi}} \alpha_{\boldsymbol{\pi}} \boldsymbol{\pi} \cdot f(\mathbf{q}) - \max_{\boldsymbol{\rho}} \boldsymbol{\rho} \cdot f(\mathbf{q}) \\ &\le \Big(\sum_{\boldsymbol{\pi}} \alpha_{\boldsymbol{\pi}} - 1 \Big) \max_{\boldsymbol{\rho}} \boldsymbol{\rho} \cdot f(\mathbf{q}) \le 0. \end{aligned}$$

For the claim about fluid model solutions,

$$\begin{aligned} \frac{d}{dt} L(\mathbf{q}(t)) &= \dot{\mathbf{q}}(t) \cdot f(\mathbf{q}(t)) \\ &= \Big(\boldsymbol{\lambda} - \sum_{\boldsymbol{\pi} \in \mathcal{S}} \dot{s}_{\boldsymbol{\pi}}(t) \boldsymbol{\pi} + \dot{\mathbf{y}}(t) \Big) \cdot f(\mathbf{q}(t)) \quad \text{by differentiating (13)} \\ &= \Big(\boldsymbol{\lambda} - \sum_{\boldsymbol{\pi}} \dot{s}_{\boldsymbol{\pi}}(t) \boldsymbol{\pi} \Big) \cdot f(\mathbf{q}(t)) \quad \text{by (19), using } f(0) = 0 \\ &= \boldsymbol{\lambda} \cdot f(\mathbf{q}(t)) - \max_{\boldsymbol{\rho}} \boldsymbol{\rho} \cdot f(\mathbf{q}(t)) \sum_{\boldsymbol{\pi}} \dot{s}_{\boldsymbol{\pi}}(t) \quad \text{by (20)} \\ &= \boldsymbol{\lambda} \cdot f(\mathbf{q}(t)) - \max_{\boldsymbol{\rho}} \boldsymbol{\rho} \cdot f(\mathbf{q}(t)) \quad \text{by (15)} \\ &\le 0 \quad \text{by (33).} \end{aligned}$$

□

To prove Theorem 5.4(ii), it is useful to work with a 'fuller' representation of the lifting map. Let E be the set of extreme feasible solutions of DUAL($\boldsymbol{\lambda}$), and define

$$\Xi^+(\boldsymbol{\lambda}) = \{ \boldsymbol{\xi} \in E : \boldsymbol{\xi} \cdot \boldsymbol{\lambda} = 1 \}. \tag{34}$$

This includes non-maximal extreme points, whereas $\Xi(\boldsymbol{\lambda})$ only includes maximal extreme points.

Lemma 5.6 *The lifting map* $\Delta W(\mathbf{q})$ *is the unique solution to the optimization problem* $\mathrm{ALGD}^+(\mathbf{q})$,

$$\begin{array}{ll} \textit{minimize} & L(\mathbf{r}) \quad \textit{over} \quad \mathbf{r} \in \mathbb{R}_+^N \\ \textit{such that} & \boldsymbol{\xi}\cdot\mathbf{r} \geq \boldsymbol{\xi}\cdot\mathbf{q} \textit{ for all } \boldsymbol{\xi} \in \Xi^+(\boldsymbol{\lambda}) \end{array}$$

Proof. $\mathrm{ALGD}^+(\mathbf{q})$ has a unique minimum for the same reason that $\mathrm{ALGD}(\mathbf{q})$ has a unique minimum.

Next we claim that if $\mathbf{r}$ is feasible for $\mathrm{ALGD}(\mathbf{q})$ then it is feasible for $\mathrm{ALGD}^+(\mathbf{q})$. Pick any $\boldsymbol{\xi} \in \Xi^+(\boldsymbol{\lambda})$. By definition, $\boldsymbol{\xi}$ is an extreme feasible solution of $\mathrm{DUAL}(\boldsymbol{\lambda})$ and $\boldsymbol{\xi}\cdot\boldsymbol{\lambda} = 1$. Since it is an extreme feasible solution, $\boldsymbol{\xi} \leq \boldsymbol{\zeta}$ for some virtual resource $\boldsymbol{\zeta} \in \mathcal{S}^*$. Since $\boldsymbol{\xi}\cdot\boldsymbol{\lambda} = 1$ we know $\boldsymbol{\zeta}\cdot\boldsymbol{\lambda} \geq 1$, but by assumption $\boldsymbol{\lambda} \in \Lambda$; hence $\boldsymbol{\zeta}\cdot\boldsymbol{\lambda} = 1$ and furthermore $\xi_n < \zeta_n$ only for n where $\lambda_n = 0$. Now,

$$\boldsymbol{\xi}\cdot\mathbf{r} - \boldsymbol{\xi}\cdot\mathbf{q} = \big(\boldsymbol{\zeta}\cdot\mathbf{r} - \boldsymbol{\zeta}\cdot\mathbf{q}\big) + (\boldsymbol{\xi} - \boldsymbol{\zeta})\cdot(\mathbf{r} - \mathbf{q}).$$

We assumed that $\mathbf{r}$ is feasible for $\mathrm{ALGD}(\mathbf{q})$; by the first constraint of $\mathrm{ALGD}(\mathbf{q})$ the first term in the preceding equation is positive; by the second constraint the second term is positive. We have shown that $\boldsymbol{\xi}\cdot\mathbf{r} \geq \boldsymbol{\xi}\cdot\mathbf{q}$ for all $\boldsymbol{\xi} \in \Xi^+(\boldsymbol{\lambda})$, hence $\mathbf{r}$ is feasible for $\mathrm{ALGD}^+(\mathbf{q})$.

Next we claim that if $\mathbf{r}$ is optimal for $\mathrm{ALGD}^+(\mathbf{q})$ then it is feasible for $\mathrm{ALGD}(\mathbf{q})$. Clearly it satisfies the first constraint of $\mathrm{ALGD}(\mathbf{q})$. Suppose it does not satisfy the second constraint, i.e. that $r_n > q_n$ for some n where $\lambda_n = 0$, and define $\mathbf{r}'$ by $r'_m = r_m$ if $m \neq n$ and $r'_n = q_n$. Then $\mathbf{r}' < \mathbf{r}$ hence $L(\mathbf{r}') < L(\mathbf{r})$. Also, $\mathbf{r}'$ is feasible for $\mathrm{ALGD}^+(\boldsymbol{\lambda})$. To see this, pick any $\boldsymbol{\zeta} \in \Xi^+(\boldsymbol{\lambda})$, and let $\boldsymbol{\xi} \in \Xi^+(\boldsymbol{\lambda})$ be such that $\zeta_m = \xi_m$ if $m \neq n$ and $\xi_n = 0$. Then

$$\boldsymbol{\zeta}\cdot\mathbf{r}' = \boldsymbol{\xi}\cdot\mathbf{r}' + \zeta_n r'_n = \boldsymbol{\xi}\cdot\mathbf{r} + \zeta_n r'_n \geq \boldsymbol{\xi}\cdot\mathbf{q} + \zeta_n r'_n = \boldsymbol{\zeta}\cdot\mathbf{q}.$$

The inequality is because $\mathbf{r}$ is feasible for $\mathrm{ALGD}^+(\mathbf{q})$. This contradicts optimality of $\mathbf{r}$.

Putting these two claims together completes the proof. □

With this representation, the lifting map ΔW can be split into two parts. Let $\Xi^+(\boldsymbol{\lambda}) = \{\boldsymbol{\xi}^1, \ldots, \boldsymbol{\xi}^V\}$ and define the *workload map* $W : \mathbb{R}_+^N \to \mathbb{R}_+^V$ by $W(\mathbf{q}) = [\boldsymbol{\xi}^v\cdot\mathbf{q}]_{1\leq v\leq V}$. Also define $\Delta : \mathbb{R}_+^V \to \mathbb{R}_+^N$ by

$$\Delta(w) = \operatorname{argmin}\big\{L(\mathbf{r}) : \mathbf{r} \in \mathbb{R}_+^N \text{ and } \boldsymbol{\xi}^v\cdot\mathbf{r} \geq w_v \text{ for } 1 \leq v \leq V\big\}. \tag{35}$$

(This has a unique optimum for the same reason that ALGD^+ and ALGD have.) Then the lifting map is simply the composition of Δ and W. It is clear that W is continuous; to prove Theorem 5.4(ii) we just need to prove that Δ is continuous.

Lemma 5.7 Δ *is continuous.*

Proof. If $\Xi^+(\boldsymbol{\lambda})$ is empty, then Δ is trivial and the result is trivial. In what follows, we shall assume that $\Xi^+(\boldsymbol{\lambda})$ is non-empty, and we will abbreviate it to Ξ^+. Furthermore note that for every $\boldsymbol{\xi} \in \Xi^+$ there is some queue n such that $\xi_n > 0$; this is because $\boldsymbol{\xi} \cdot \boldsymbol{\lambda} = 1$ by definition of Ξ^+.

Pick any sequence $w^k \to w \in \mathbb{R}_+^V$, and let $\mathbf{r}^k = \Delta(w^k)$ and $\mathbf{r} = \Delta(w)$. We want to prove that $\mathbf{r}^k \to \mathbf{r}$. We shall first prove that there is a compact set $[0,h]^N$ such that $\mathbf{r}^k \in [0,h]^N$ for all k. We shall then prove that any convergent subsequence of $\mathbf{r}^k$ converges to $\mathbf{r}$; this establishes continuity of Δ.

First, compactness. A suitable value for h is

$$h = \max_{1 \le v \le V} \max_{n:\xi_n > 0} \sup_k \frac{w_v^k}{\xi_n^v}.$$

Note than the maximums are over a non-empty set, as noted at the beginning of the proof. Note also that h is finite because w is finite. Now, suppose that $\mathbf{r}^k \notin [0,h]^N$ for some k, i.e. that there is some queue n for which $r_n^k > h$, and let $\mathbf{r}' = \mathbf{r}^k$ in each component except for $r'_n = h$. We claim that $\mathbf{r}'$ satisfies the constraints of the optimization problem for $\Delta(w^k)$. To see this, pick any $\boldsymbol{\xi}^v \in \Xi^+$; either $\xi_n^v = 0$ in which case $\boldsymbol{\xi}^v \cdot \mathbf{r}' = \boldsymbol{\xi}^v \cdot \mathbf{r}^k \ge w_v^k$, or $\xi_n^v > 0$ in which case $\boldsymbol{\xi}^v \cdot \mathbf{r}' \ge \xi_n^v h \ge w_v^k$ by construction of h. Applying this repeatedly, if $\mathbf{r}^k \notin [0,h]^N$ then we can reduce it to a queue size vector in $[0,h]^N$, thereby improving on $L(\mathbf{r}^k)$, yet still meeting the constraints of the optimization problem for $\Delta(w^k)$; this contradicts the optimality of $\mathbf{r}^k$. Hence $\mathbf{r}^k \in [0,h]^N$.

Next, convergence on subsequences. With a slight abuse of notation, let $\Delta(w^k) = \mathbf{r}^k \to \mathbf{s}$ be a convergent subsequence, and recall that $\Delta(w) = \mathbf{r}$ and $w^k \to w$. By continuity of the constraints, $\mathbf{s}$ is feasible for the optimization problem for $\Delta(w)$; we shall next show that $L(\mathbf{s}) \le L(\mathbf{r})$. Since $\mathbf{r}$ is the unique optimum, it must be that $\mathbf{s} = \mathbf{r}$.

It remains to show that $L(\mathbf{s}) \le L(\mathbf{r})$. Consider the sequence $\mathbf{r} + \varepsilon^k \mathbf{1}$ as candidate solutions to the problem $\Delta(w^k)$ where

$$\varepsilon^k = \max_{1 \le v \le V} \frac{w_v^k - w_v}{\boldsymbol{\xi}^v \cdot \mathbf{1}}.$$

This choice ensures that the candidates are feasible, since

$$\boldsymbol{\xi}^v \cdot (\mathbf{r} + \varepsilon^k \mathbf{1}) = \boldsymbol{\xi}^v \cdot \mathbf{r} + \varepsilon^k \boldsymbol{\xi}^v \cdot \mathbf{1} \ge \boldsymbol{\xi}^v \cdot \mathbf{r} + w_v^k - w_v \ge w_v^k.$$

(If we had used $\boldsymbol{\xi} \in \Xi$ rather than $\boldsymbol{\xi} \in \Xi^+$, it would not necessarily be true that the candidates are feasible; this is why we introduced Lemma 5.6.) Since the candidates are feasible solutions to the problem $\Delta(w^k)$, and $\mathbf{r}^k$ is an optimal solution, it must be that

$$L(\mathbf{r}^k) \le L(\mathbf{r} + \varepsilon^k \mathbf{1}).$$

Taking the limit as $k \to \infty$, and noting that L is continuous and $\varepsilon^k \to 0$, we find

$$L(\mathbf{s}) \le L(\mathbf{r})$$

as required. This completes the proof. $\square$

For the proof of Theorem 5.4(iii), it is useful to work with a different representation of the constraint of ALGD, provided by the following lemma.

Lemma 5.8 *(i)* $\Delta W(\mathbf{q}) = [\mathbf{q} + t(\boldsymbol{\lambda} - \boldsymbol{\sigma})]^+$ *for some* $t \ge 0$ *and* $\boldsymbol{\sigma} \in \langle \mathcal{S} \rangle$.
(ii) $[\mathbf{q} + t(\boldsymbol{\lambda} - \boldsymbol{\sigma})]^+$ *is feasible for* ALGD$(\mathbf{q})$ *for all* $t \ge 0$ *and* $\boldsymbol{\sigma} \in \langle \mathcal{S} \rangle$.

Proof of Lemma 5.8(i). We will shortly prove that the following are equivalent, for all $\mathbf{q}$ and $\mathbf{r} \in \mathbb{R}_+^N$:

$$\mathbf{r} \ge \mathbf{q} + t(\boldsymbol{\lambda} - \boldsymbol{\sigma}) \quad \text{for some} \quad t \ge 0,\ \boldsymbol{\sigma} \in \langle \mathcal{S} \rangle, \tag{36}$$

$$\boldsymbol{\xi} \cdot \mathbf{r} \ge \boldsymbol{\xi} \cdot \mathbf{q} \quad \text{for all} \quad \boldsymbol{\xi} \in \Xi^+(\boldsymbol{\lambda}). \tag{37}$$

We use this equivalence as follows. From Lemma 5.6 we know that $\Delta W(\mathbf{q})$ is the solution of ALGD$^+(\mathbf{q})$. That is, letting $\mathbf{q}' = \Delta W(\mathbf{q})$, equation (37) holds with $\mathbf{q}'$ in the place of $\mathbf{r}$. Hence (36) holds for some $t \ge 0$ and $\boldsymbol{\sigma} \in \langle \mathcal{S} \rangle$; moreover since $\mathbf{q}' \ge \mathbf{0}$ it must be that

$$\mathbf{q}' \ge [\mathbf{q} + t(\boldsymbol{\lambda} - \boldsymbol{\sigma})]^+.$$

We claim that this inequality is in fact an equality. To see this, note that $\mathbf{r} = [\mathbf{q} + t(\boldsymbol{\lambda} - \boldsymbol{\sigma})]^+$ satisfies (36), hence it satisfies (37), hence it is a feasible solution of ALGD$^+(\mathbf{q})$. Note also that $L(\cdot)$ is increasing componentwise, hence $L(\mathbf{q}') \ge L(\mathbf{r})$. But ALGD$^+(\mathbf{q})$ has a unique minimum, hence $\mathbf{q}' = \mathbf{r}$ as required. This completes the proof of Lemma 5.8(i), once we have proved the equivalence between (36) and (37).

Proof that (36) $\Longrightarrow$ (37). Pick any $\boldsymbol{\xi} \in \Xi^+(\boldsymbol{\lambda})$. By definition of $\Xi^+(\boldsymbol{\lambda})$, we know: $\boldsymbol{\xi} \ge \mathbf{0}$; $\boldsymbol{\xi} \cdot \boldsymbol{\pi} \le 1$ for all $\boldsymbol{\pi} \in \mathcal{S}$, hence $\boldsymbol{\xi} \cdot \boldsymbol{\sigma} \le 1$ for all $\boldsymbol{\sigma} \in \langle \mathcal{S} \rangle$; and $\boldsymbol{\xi} \cdot \boldsymbol{\lambda} = 1$. Hence

$$\begin{aligned} \boldsymbol{\xi} \cdot \mathbf{r} &\ge \boldsymbol{\xi} \cdot \mathbf{q} + t\big(\boldsymbol{\xi} \cdot \boldsymbol{\lambda} - \boldsymbol{\xi} \cdot \boldsymbol{\sigma}\big) \quad \text{assuming } \mathbf{q},\ \mathbf{r} \in \mathbb{R}_+^N \text{ satisfying (36)} \\ &= \boldsymbol{\xi} \cdot \mathbf{q} + t\big(1 - \boldsymbol{\xi} \cdot \boldsymbol{\sigma}\big) \ge \boldsymbol{\xi} \cdot \mathbf{q} + t(1 - 1) = \boldsymbol{\xi} \cdot \mathbf{q}. \end{aligned}$$

Proof that (36) $\Longleftarrow$ (37). Let $\mathbf{q}$ and $\mathbf{r}$ satisfy (37), and let $\boldsymbol{\sigma}' = \boldsymbol{\lambda} - (\mathbf{r}-\mathbf{q})/t$ for some sufficiently large $t \in \mathbb{R}_+$. We shortly show that the value of DUAL($\boldsymbol{\sigma}'$) at its optimum is ≤ 1. By strong duality the value of PRIMAL($\boldsymbol{\sigma}'$) at its optimum is likewise ≤ 1, and so by definition of PRIMAL($\boldsymbol{\sigma}'$) we can find some $\boldsymbol{\sigma} \in \langle \mathcal{S} \rangle$ such that $\boldsymbol{\sigma}' \leq \boldsymbol{\sigma}$ componentwise. Then

$$\mathbf{r} = \mathbf{q} + t(\boldsymbol{\lambda} - \boldsymbol{\sigma}') \geq \mathbf{q} + t(\boldsymbol{\lambda} - \boldsymbol{\sigma})$$

i.e. $\mathbf{r}$ satisfies (36).

It remains to show that the value of DUAL($\boldsymbol{\sigma}'$) at its optimum is ≤ 1, i.e. that $\boldsymbol{\zeta} \cdot \boldsymbol{\sigma}' \leq 1$ for all dual-feasible $\boldsymbol{\zeta}$. We have assumed that $\boldsymbol{\lambda} \in \Lambda$, hence $\boldsymbol{\zeta} \cdot \boldsymbol{\lambda} \leq 1$. On one hand, if $\boldsymbol{\zeta} \cdot \boldsymbol{\lambda} = 1$ then it follows from the definition of Ξ^+ that $\boldsymbol{\zeta} \in \langle \Xi^+(\boldsymbol{\lambda}) \rangle$, hence

$$\begin{aligned}
\boldsymbol{\zeta} \cdot \boldsymbol{\sigma}' &= \boldsymbol{\zeta} \cdot \boldsymbol{\lambda} - \boldsymbol{\zeta} \cdot (\mathbf{r}-\mathbf{q})/t \\
&= 1 - \boldsymbol{\zeta} \cdot (\mathbf{r}-\mathbf{q})/t \quad \text{since } \boldsymbol{\zeta} \cdot \boldsymbol{\lambda} = 1 \\
&\leq 1 \quad \text{by (37).}
\end{aligned}$$

On the other hand, if $\boldsymbol{\zeta} \cdot \boldsymbol{\lambda} < 1$ then

$$\boldsymbol{\zeta} \cdot \boldsymbol{\sigma}' < 1 - \boldsymbol{\zeta} \cdot (\mathbf{r}-\mathbf{q})/t$$

and this is < 1 for t sufficiently large. Either way, $\boldsymbol{\zeta} \cdot \boldsymbol{\sigma}' \leq 1$. Therefore the value of DUAL($\boldsymbol{\sigma}'$) at its optimum is ≤ 1.

Proof of Lemma 5.8(ii). For this, we need to check two feasibility conditions of ALGD($\mathbf{q}$). The first feasibility condition is

$$\boldsymbol{\xi} \cdot [\mathbf{q} + t(\boldsymbol{\lambda} - \boldsymbol{\sigma})]^+ \geq \boldsymbol{\xi} \cdot \mathbf{q} \quad \text{for all } \boldsymbol{\xi} \in \Xi(\boldsymbol{\lambda}).$$

Pick any $\boldsymbol{\xi} \in \Xi(\boldsymbol{\lambda})$. By definition of $\Xi(\boldsymbol{\lambda})$, $\boldsymbol{\xi} \geq \mathbf{0}$, $\boldsymbol{\xi} \cdot \boldsymbol{\pi} \leq 1$ for all $\boldsymbol{\pi} \in \mathcal{S}$ hence $\boldsymbol{\xi} \cdot \boldsymbol{\sigma} \leq 1$ for all $\boldsymbol{\sigma} \in \langle \mathcal{S} \rangle$, and $\boldsymbol{\xi} \cdot \boldsymbol{\lambda} = 1$, thus

$$\boldsymbol{\xi} \cdot [\mathbf{q} + t(\boldsymbol{\lambda} - \boldsymbol{\sigma})]^+ \geq \boldsymbol{\xi} \cdot [\mathbf{q} + t(\boldsymbol{\lambda} - \boldsymbol{\sigma})] = \boldsymbol{\xi} \cdot \mathbf{q} + t(1 - \boldsymbol{\xi} \cdot \boldsymbol{\sigma}) \geq \boldsymbol{\xi} \cdot \mathbf{q}$$

as required. The second feasibility condition is that if $\lambda_n = 0$ for some n then

$$[q_n + t(\lambda_n - \sigma_n)]^+ = [q_n - t\sigma_n]^+ \leq q_n.$$

This is true because $\boldsymbol{\sigma} \geq \mathbf{0}$ componentwise for all $\boldsymbol{\sigma} \in \langle \mathcal{S} \rangle$. □

Theorem 5.4(iii) is a corollary of the following lemma.

Lemma 5.9 (Scale-invariance of $\boldsymbol{\Delta W}$) *Let $\mathbf{q} \in \mathbb{R}_+^N$. Then $\Delta W(\kappa \mathbf{q}) = \kappa \Delta W(\mathbf{q})$ for all $\kappa > 0$.*

Proof. We will first establish three preliminary properties of ΔW. Preliminary 1 is used to prove 2, and 2 & 3 are used in the main proof.

Preliminary 1. If $\mathbf{q} = \Delta W(\mathbf{q}')$ for some $\mathbf{q}' \in \mathbb{R}_+^N$ then

$$\boldsymbol{\lambda}\cdot f(\mathbf{q}) = \max_{\boldsymbol{\pi}\in\mathcal{S}} \boldsymbol{\pi}\cdot f(\mathbf{q}). \tag{38}$$

To see this, suppose $\boldsymbol{\pi} \in \mathcal{S}$ has maximal weight and consider $\mathbf{r} = \big[\mathbf{q} + t(\boldsymbol{\lambda} - \boldsymbol{\pi})\big]^+$. This is feasible for ALGD($\mathbf{q}$) by Lemma 5.8. Now, using the fact that $f(0) = 0$,

$$\frac{d}{dt} L\Big(\big[\mathbf{q} + t(\boldsymbol{\lambda} - \boldsymbol{\pi})\big]^+\Big)\Big|_{t=0} = (\boldsymbol{\lambda} - \boldsymbol{\pi})\cdot f(\mathbf{q}).$$

Since $\mathbf{q}$ is optimal for ALGD($\mathbf{q}'$) it is *a forteriori* optimal for ALGD($\mathbf{q}$), hence $\boldsymbol{\lambda}\cdot f(\mathbf{q}) \geq \boldsymbol{\pi}\cdot f(\mathbf{q})$. On the other hand, $\boldsymbol{\lambda} \in \Lambda$ so $\boldsymbol{\lambda} \leq \boldsymbol{\sigma}$ for some $\boldsymbol{\sigma} \in \langle\mathcal{S}\rangle$, hence $\boldsymbol{\lambda}\cdot f(\mathbf{q}) \leq \boldsymbol{\sigma}\cdot f(\mathbf{q}) \leq \boldsymbol{\pi}\cdot f(\mathbf{q})$. Hence the result follows.

Preliminary 2. Suppose that $\mathbf{r} = \Delta W(\mathbf{q})$. From Lemma 5.8, $\mathbf{r} = \big[\mathbf{q} + t(\boldsymbol{\lambda} - \boldsymbol{\sigma})\big]^+$ for some $t \geq 0$ and $\boldsymbol{\sigma} \in \langle\mathcal{S}\rangle$. Then either $t = 0$ or

$$\boldsymbol{\sigma}\cdot f(\mathbf{r}) = \max_{\boldsymbol{\pi}\in\mathcal{S}} \boldsymbol{\pi}\cdot f(\mathbf{r}). \tag{39}$$

This is because t is an optimal choice, so either t is constrained to be 0 or

$$\frac{d}{du} L\Big(\big[\mathbf{q} + u(\boldsymbol{\lambda} - \boldsymbol{\sigma})\big]^+\Big)\Big|_{u=t} = (\boldsymbol{\lambda} - \boldsymbol{\sigma})\cdot f(\mathbf{r}) = 0.$$

In this second case, $\boldsymbol{\lambda}\cdot f(\mathbf{r}) = \max_{\boldsymbol{\pi}} \boldsymbol{\pi}\cdot f(\mathbf{r})$ by (38) so the same is true for $\boldsymbol{\sigma}$.

Preliminary 3. Suppose that $\mathbf{r} = \Delta W(\mathbf{q})$. From Lemma 5.8, we can write it as $\mathbf{r} = \big[\mathbf{q} + t(\boldsymbol{\lambda} - \boldsymbol{\sigma})\big]^+$ for some $\boldsymbol{\sigma} \in \langle\mathcal{S}\rangle$. In fact, for any $T \geq t$ we can write it as

$$\mathbf{r} = \big[\mathbf{q} + T(\boldsymbol{\lambda} - \boldsymbol{\rho})\big]^+ \quad \text{for some } \boldsymbol{\rho} \in \langle\mathcal{S}\rangle. \tag{40}$$

To see this, recall that PRIMAL($\boldsymbol{\lambda}$) ≤ 1, so we can pick some $\bar{\boldsymbol{\lambda}} \in \langle\mathcal{S}\rangle$ such that $\boldsymbol{\lambda} \leq \bar{\boldsymbol{\lambda}}$, whence

$$\begin{aligned}
\mathbf{r} &\geq \big[\mathbf{q} + t(\boldsymbol{\lambda} - \boldsymbol{\sigma}) + (T - t)(\boldsymbol{\lambda} - \bar{\boldsymbol{\lambda}})\big]^+ \\
&= \big[\mathbf{q} + T(\boldsymbol{\lambda} - \boldsymbol{\rho})\big]^+ \quad \text{where } \boldsymbol{\rho} = \frac{t}{T}\boldsymbol{\sigma} + \frac{T - t}{T}\bar{\boldsymbol{\lambda}} \in \langle\mathcal{S}\rangle.
\end{aligned}$$

This last expression is feasible for ALGD($\mathbf{q}$) by Lemma 5.8. Since $\mathbf{r}$ is optimal for ALGD($\mathbf{q}$), and the objective function is increasing pointwise, $\mathbf{r} = [\mathbf{q} + T(\boldsymbol{\lambda} - \boldsymbol{\rho})]^+$ as claimed.

Main proof. Let $\mathbf{r} = \Delta W(\mathbf{q})$ and $\kappa\mathbf{r}' = \Delta W(\kappa\mathbf{q})$. We know that $\kappa\mathbf{r}$ is feasible for $\mathrm{ALGD}(\kappa\mathbf{q})$ because the constraints are linear; we will now show that $L(\kappa\mathbf{r}) \le L(\kappa\mathbf{r}')$; hence $\kappa\mathbf{r}$ is also optimal for $\mathrm{ALGD}(\kappa\mathbf{q})$. By uniqueness of the optimum, $\kappa\mathbf{r} = \kappa\mathbf{r}'$ as required.

It remains to prove that $L(\kappa\mathbf{r}) \le L(\kappa\mathbf{r}')$. Since $\mathbf{r}$ solves $\mathrm{ALGD}(\mathbf{q})$ and $\kappa\mathbf{r}'$ solves $\mathrm{ALGD}(\kappa\mathbf{q})$, we can use Lemma 5.8 to write

$$\mathbf{r} = \left[\mathbf{q} + t(\boldsymbol{\lambda} - \boldsymbol{\sigma})\right]^+, \quad \kappa\mathbf{r}' = \left[\kappa\mathbf{q} + \kappa t'(\boldsymbol{\lambda} - \boldsymbol{\sigma}')\right]^+$$

for $t, t' \in \mathbb{R}_+$ and $\boldsymbol{\sigma}, \boldsymbol{\sigma}' \in \langle\mathcal{S}\rangle$. Indeed, for $T > \max(t, t')$ we can use (40) to write

$$\begin{aligned}
\mathbf{r} &= \mathbf{q} + T(\boldsymbol{\lambda} - \boldsymbol{\rho} + \mathbf{y}) \quad \text{for } \boldsymbol{\rho} \in \langle\mathcal{S}\rangle,\ \mathbf{y} \in \mathbb{R}_+^N, \text{ where } y_n = 0 \text{ if } r_n > 0\\
\mathbf{r}' &= \mathbf{q} + T(\boldsymbol{\lambda} - \boldsymbol{\rho}' + \mathbf{y}') \quad \text{for } \boldsymbol{\rho}' \in \langle\mathcal{S}\rangle,\ \mathbf{y}' \in \mathbb{R}_+^N, \text{ where } y'_n = 0 \text{ if } r'_n > 0.
\end{aligned}$$

Now consider the value of $L(\cdot)$ along the trajectory from $\kappa\mathbf{r}$ to $\kappa\mathbf{r}'$. Along this trajectory,

$$\begin{aligned}
\frac{d}{du} L\Big(\kappa\mathbf{r} + (\mathbf{r}' - \mathbf{r})u/T\Big)\Big|_{u=0} &= (\mathbf{r}' - \mathbf{r}) \cdot f(\kappa\mathbf{r})/T\\
&= \big(\boldsymbol{\rho} - \boldsymbol{\rho}' - \mathbf{y} + \mathbf{y}'\big) \cdot f(\kappa\mathbf{r})\\
&\ge \big(\boldsymbol{\rho} - \boldsymbol{\rho}' - \mathbf{y}\big) \cdot f(\kappa\mathbf{r}) \quad \text{since } \mathbf{y}' \ge \mathbf{0}\\
&= \big(\boldsymbol{\rho} - \boldsymbol{\rho}'\big) \cdot f(\kappa\mathbf{r}) \quad \text{since } y_n = 0 \text{ if } r_n > 0\\
&\ge \boldsymbol{\rho} \cdot f(\kappa\mathbf{r}) - \max_{\boldsymbol{\pi} \in \mathcal{S}} \boldsymbol{\pi} \cdot f(\kappa\mathbf{r}) \quad \text{for any } \boldsymbol{\rho}' \in \langle\mathcal{S}\rangle\\
&= 0.
\end{aligned}$$

The final equality is because $\boldsymbol{\rho} \cdot f(\mathbf{r}) = \max_{\boldsymbol{\pi}} \boldsymbol{\pi} \cdot f(\mathbf{r})$ by (39), so $\boldsymbol{\rho} \cdot f(\kappa\mathbf{r}) = \max_{\boldsymbol{\pi}} \boldsymbol{\pi} \cdot f(\kappa\mathbf{r})$ by Assumption 2.1. Since $L(\cdot)$ is convex, it follows that $L(\kappa\mathbf{r}') \ge L(\kappa\mathbf{r})$. This completes the proof. □

The proof of Theorem 5.4(iv) relies on the following lemma.

Lemma 5.10 (Fluid model trajectories preserve ALGD feasibility) *Consider any fluid model solution, for any scheduling policy, with initial queue size $\mathbf{q}(0)$. Then $\mathbf{q}(t)$ is feasible for $\mathrm{ALGD}\big(\mathbf{q}(0)\big)$ for all $t \ge 0$.*

Proof. Pick any critically loaded virtual resource $\boldsymbol{\xi} \in \Xi(\boldsymbol{\lambda})$. By (13),

$$\begin{aligned}
\boldsymbol{\xi} \cdot \mathbf{q}(t) &= \boldsymbol{\xi} \cdot \mathbf{q}(0) + t\Big(\boldsymbol{\xi} \cdot \boldsymbol{\lambda} - \boldsymbol{\xi} \cdot \boldsymbol{\sigma}(t)\Big) + \boldsymbol{\xi} \cdot \mathbf{y}(t) \quad \text{where } \boldsymbol{\sigma}(t) = \sum \boldsymbol{\pi} s_{\boldsymbol{\pi}}(t)/t\\
&\ge \boldsymbol{\xi} \cdot \mathbf{q}(0) + t\Big(\boldsymbol{\xi} \cdot \boldsymbol{\lambda} - \boldsymbol{\xi} \cdot \boldsymbol{\sigma}(t)\Big) \quad \text{since } \mathbf{y}(t) \ge \mathbf{0}\\
&\ge \boldsymbol{\xi} \cdot \mathbf{q}(0) + t(1 - 1) = \boldsymbol{\xi} \cdot \mathbf{q}(0).
\end{aligned}$$

The last inequality is because $\boldsymbol{\xi} \in \Xi(\boldsymbol{\lambda})$; so $\boldsymbol{\xi}\cdot\boldsymbol{\lambda} = 1$, and $\boldsymbol{\xi}\cdot\boldsymbol{\pi} \leq 1$ for all $\boldsymbol{\pi} \in \mathcal{S}$ hence $\boldsymbol{\xi}\cdot\boldsymbol{\sigma} \leq 1$ for all $\boldsymbol{\sigma} \in \langle\mathcal{S}\rangle$. Finally, $\mathbf{q}(t) \leq \mathbf{q}(0) + t\boldsymbol{\lambda}$ by (13) and (16), and this yields the second constraint of ALGD($\boldsymbol{\lambda}$) for queues n with 0 arrival rate. □

Theorem 5.4(iv) is implied by parts (i) and (ii) of the following lemma.

Lemma 5.11 (Characterization of invariant states of MW-f) *The following are equivalent, for $\mathbf{q}^0 \in \mathbb{R}_+^N$:*

(i) $\mathbf{q}^0 = \Delta W(\mathbf{q}^0)$
(ii) $\mathbf{q}^0$ *is an invariant state*
(iii) there exists a fluid model solution with $\mathbf{q}(t) = \mathbf{q}^0$ *for all* t
(iv) $\boldsymbol{\lambda}\cdot f(\mathbf{q}^0) = \max_{\boldsymbol{\pi}\in\mathcal{S}} \boldsymbol{\pi}\cdot f(\mathbf{q}^0)$

Proof that (i) $\Longrightarrow$ (ii). Suppose that $\mathbf{q}^0 = \Delta W(\mathbf{q}^0)$, i.e. that $\mathbf{q}^0$ is optimal for ALGD($\mathbf{q}^0$), and consider any fluid model solution which starts with $\mathbf{q}(0) = \mathbf{q}^0$. On one hand, Lemma 5.5 says that $L(\mathbf{q}(t)) \leq L(\mathbf{q}^0)$. On the other hand, Lemma 5.10 says that $\mathbf{q}(t)$ is feasible for ALGD($\mathbf{q}^0$). Since ALGD($\mathbf{q}^0$) has a unique solution, it must be that $\mathbf{q}(t) = \mathbf{q}^0$.

Proof that (ii) $\Longrightarrow$ (iii). It is easy to find a fluid model solution which starts at $\mathbf{q}(0) = \mathbf{q}^0$: a limit point of the stochastic model from Theorem 4.3 will do. By (ii), the queue size vector is constant.

Proof that (iii) $\Longrightarrow$ (iv). Suppose there is a fluid model solution with $\mathbf{q}(t) = \mathbf{q}^0$. Since $\mathbf{q}(\cdot)$ is constant, $\dot{L}(\mathbf{q}(t)) = 0$. Lemma 5.5 says that $\dot{L}(\mathbf{q}(t)) \leq 0$, so the inequality in the proof must be tight for all t, i.e.

$$\boldsymbol{\lambda}\cdot f(\mathbf{q}^0) = \max_{\boldsymbol{\pi}\in\mathcal{S}} \boldsymbol{\pi}\cdot f(\mathbf{q}^0). \tag{41}$$

Proof that (iv) $\Longrightarrow$ (i). If $\mathbf{q}^0 = \mathbf{0}$ then the result is trivial. Otherwise, let $\mathbf{r} = \Delta W(\mathbf{q}^0)$. By Lemma 5.8, $\mathbf{r} = [\mathbf{q}^0 + t(\boldsymbol{\lambda}-\boldsymbol{\sigma})]^+$ for some $t \geq 0$ and $\boldsymbol{\sigma} \in \langle\mathcal{S}\rangle$. Consider the value of $L(\cdot)$ along the trajectory from $\mathbf{q}^0$ to $\mathbf{r}$:

$$\begin{aligned}\frac{d}{du} L\Big(\big[\mathbf{q}^0 + (\boldsymbol{\lambda}-\boldsymbol{\sigma})u\big]^+\Big)\Big|_{u=0} &= (\boldsymbol{\lambda}-\boldsymbol{\sigma})\cdot f(\mathbf{q}^0) \quad \text{relying on } f(0)=0\\ &= \Big(\max_{\boldsymbol{\pi}\in\mathcal{S}} \boldsymbol{\pi}\cdot f(\mathbf{q}^0)\Big) - \boldsymbol{\sigma}\cdot f(\mathbf{q}^0) \quad \text{by part (iv)}\\ &\geq 0 \quad \text{because } \boldsymbol{\sigma} \in \langle\mathcal{S}\rangle.\end{aligned}$$

By convexity of L, $L(\mathbf{r}) \geq L(\mathbf{q}^0)$; and $\mathbf{q}^0$ is obviously feasible for ALGD($\mathbf{q}^0$); but we chose $\mathbf{r}$ to be optimal for ALGD($\mathbf{q}^0$) and the optimum is unique. Therefore $\mathbf{q}^0 = \Delta W(\mathbf{q}^0)$. □

Theorem 5.4(v) is given by the following lemma. Recall that we are using the norm $|\mathbf{x}| = \max_n |x_n|$.

Lemma 5.12 *Given* $\boldsymbol{\lambda} \in \Lambda$*, for any* $\varepsilon > 0$ *there exists an* H_ε *such that for every fluid model solution with arrival rate* $\boldsymbol{\lambda}$*, for which* $|\mathbf{q}(0)| \leq 1$*,* $|\mathbf{q}(t) - \Delta W(\mathbf{q}(t))| < \varepsilon$ *for all* $t \geq H_\varepsilon$*.*

Proof. The proof is inspired by Kelly and Williams [19, Theorem 5.2, Lemma 6.3]. We start with some definitions. Let

$$\begin{aligned}
\mathcal{D} &= \big\{\mathbf{q} \in \mathbb{R}_+^N : L(\mathbf{q}) \leq L(\mathbf{1})\big\} \quad \text{for } L(\cdot) \text{ as in Definition 5.3} \\
\mathcal{I} &= \big\{\mathbf{q} \in \mathcal{D} : \Delta W(\mathbf{q}) = \mathbf{q}\big\} \\
\mathcal{I}_\delta &= \big\{\mathbf{q} \in \mathcal{D} : |\mathbf{q} - \mathbf{r}| < \delta \text{ for some } \mathbf{r} \in \mathcal{I}\big\} \\
\mathcal{J}_\varepsilon &= \big\{\mathbf{q} \in \mathbb{R}_+ : \big|\mathbf{q} - \Delta W(\mathbf{q})\big| < \varepsilon\big\} \\
\mathcal{K}_\delta &= \Big\{\mathbf{q} \in \mathcal{D} : L(\mathbf{q}) - L\big(\Delta W(\mathbf{q})\big) < \inf_{\mathbf{r} \in \mathcal{D} \setminus \mathcal{I}_\delta} L(\mathbf{r}) - L\big(\Delta W(\mathbf{r})\big)\Big\}
\end{aligned}$$

We will argue that the function $K(\mathbf{q}) = L(\mathbf{q}) - L\big(\Delta W(\mathbf{q})\big)$ is decreasing along fluid model trajectories, so once you hit $\mathcal{K}_\delta$ you stay there. We will then argue that $\mathcal{I} \subset \mathcal{K}_\delta \subset \mathcal{I}_\delta \subset \mathcal{J}_\varepsilon$ for sufficiently small δ. Finally, we will bound the time it takes to hit $\mathcal{K}_\delta$.

K is decreasing. Lemma 5.5 says that for any fluid model solution, $L(\mathbf{q}(\cdot))$ is decreasing. From Lemma 5.10, the feasible set for ALGD($\mathbf{q}(u)$) is a subset of the feasible set for ALGD($\mathbf{q}(t)$) for any $u \geq t \geq 0$, hence $\Delta W(\mathbf{q}(u)) \geq \Delta W(\mathbf{q}(t))$, i.e. $\Delta W(\mathbf{q}(\cdot))$ is increasing. Therefore K is decreasing (not necessarily strictly).

$\mathcal{I} \subset \mathcal{K}_\delta \subset \mathcal{I}_\delta \subset \mathcal{J}_\varepsilon$. To show $\mathcal{I} \subset \mathcal{K}_\delta$: The map ΔW is continuous by Theorem 5.4(ii), and $L(\cdot)$ is clearly continuous, so $K(\cdot)$ is continuous; also the set $\mathcal{D}$ is compact by Theorem 5.4(i), and $\mathcal{I}_\delta$ is open, so $\mathcal{D} \setminus \mathcal{I}_\delta$ is compact; so the infimum in the definition of $\mathcal{K}_\delta$ is attained at some $\hat{\mathbf{r}} \in \mathcal{D} \setminus \mathcal{I}_\delta$. Now, $K(\mathbf{q}) > 0$ for $\mathbf{q} \in \mathcal{D} \setminus \mathcal{I}$, so $K(\hat{\mathbf{r}}) > 0$. Yet $K(\mathbf{q}) = 0$ for $\mathbf{q} \in \mathcal{I}$. Thus $\mathcal{I} \subset \mathcal{K}_\delta$.

It is clear by construction that $\mathcal{K}_\delta \subset \mathcal{I}_\delta$.

To show $\mathcal{I}_\delta \subset \mathcal{J}_\varepsilon$: The map $\Delta W(\cdot)$ is continuous, hence it is uniformly continuous on the compact set $\mathcal{D}$, so for any $\varepsilon > 0$ there exists a $\delta > 0$ such that

$$|\mathbf{q} - \mathbf{r}| < \delta \implies \big|\Delta W(\mathbf{q}) - \Delta W(\mathbf{r})\big| < \varepsilon/2 \quad \text{for } \mathbf{q}, \mathbf{r} \in \mathcal{D}.$$

If $\mathbf{q} \in \mathcal{I}_\delta$ then it is within δ of some $\mathbf{r} \in \mathcal{I}$, hence

$$\begin{aligned}
\big|\mathbf{q} - \Delta W(\mathbf{q})\big| &\leq |\mathbf{q} - \mathbf{r}| + \big|\mathbf{r} - \Delta W(\mathbf{r})\big| + \big|\Delta W(\mathbf{r}) - \Delta W(\mathbf{q})\big| \\
&< \delta + 0 + \varepsilon/2 \\
&< \varepsilon \quad \text{for } \delta \text{ sufficiently small.}
\end{aligned}$$

Time to hit $\mathcal{K}_\delta$. Consider first the rate of change of $K(\cdot)$ while the process is in $\mathcal{D}\setminus\mathcal{K}_\delta$:

$$\begin{aligned}\dot K(\mathbf{q}(t)) \le \dot L(\mathbf{q}(t)) &= \boldsymbol{\lambda}\cdot f(\mathbf{q}(t)) - \max_{\boldsymbol{\pi}\in\mathcal{S}} \boldsymbol{\pi}\cdot f(\mathbf{q}(t)) \\ &\le \sup_{\mathbf{r}\in\mathcal{D}\setminus\mathcal{K}_\delta}\Big[\boldsymbol{\lambda}\cdot f(\mathbf{r}) - \max_{\boldsymbol{\pi}\in\mathcal{S}}\boldsymbol{\pi}\cdot f(\mathbf{r})\Big] \qquad (42) \\ &\le 0 \quad \text{by Lemma 5.5.}\end{aligned}$$

The supremum in (42) is of a continuous function of $\mathbf{r}$, taken over a compact set, hence the supremum is attained at some $\hat{\mathbf{r}}\in\mathcal{D}\setminus\mathcal{K}_\delta$. If the supremum were equal to 0 then $\boldsymbol{\lambda}\cdot f(\hat{\mathbf{r}}) = \max_{\boldsymbol{\pi}}\boldsymbol{\pi}\cdot f(\hat{\mathbf{r}})$ so $\hat{\mathbf{r}}\in\mathcal{I}$ by Lemma 5.11; but $\hat{\mathbf{r}}\in\mathcal{D}\setminus\mathcal{K}_\delta$ and we just proved that $\mathcal{I}\subset\mathcal{K}_\delta$; hence the supremum is some $-\eta_\delta<0$.

Now consider any fluid model solution starting at $\mathbf{q}(0)$ with $|\mathbf{q}(0)|\le 1$. If $\mathbf{q}(0)\in\mathcal{K}_\delta$ then it remains in $\mathcal{K}_\delta$ so the theorem holds trivially. If not, then $\mathbf{q}(0)\le\mathbf{1}$ componentwise, so $L(\mathbf{q}(0))\le L(\mathbf{1})$, so $\mathbf{q}(0)\in\mathcal{D}$; also $L(\mathbf{q}(t))$ is decreasing so $\mathbf{q}(t)\in\mathcal{D}$ for all $t\ge 0$. Now, $\dot K(\mathbf{q}(t))\le -\eta_\delta$ all the time that $\mathbf{q}(t)\in\mathcal{D}\setminus\mathcal{K}_\delta$, and this can't go on for longer than $H_\varepsilon = K(\mathbf{q}(0))/\eta_\delta \le L(\mathbf{1})/\eta_\delta$. □

6. Fluid model behaviour (multihop case). In this section we describe properties of fluid model solutions for a multihop switched network running MW-f backpressure, as described in Section 2.

Let R be the routing matrix, and $\vec{R} = (I-R^{\mathsf{T}})^{-1}$; recall that $\vec{R}_{mn}=1$ if work injected at queue n eventually passes through m, and 0 otherwise. For a vector $\mathbf{x}\in\mathbb{R}^N$, let $\vec{\mathbf{x}} = \vec{R}\mathbf{x}$: for arrival rate vector $\boldsymbol{\lambda}$, $\vec{\lambda}_n$ is the total arrival rate of work destined to pass through queue n; for a queue size vector $\mathbf{q}$, $\vec{q}_n$ is the total amount of work at queue n and queues upstream of n.

The set Λ, the PRIMAL$(\cdot)$ and DUAL$(\cdot)$ problems, the set $\mathcal{S}^*$ of virtual resources, and $\Xi(\cdot)$ are defined as in the single-hop case. The difference is that we will require $\vec{\boldsymbol{\lambda}}\in\Lambda$, and we will define the set of critically loaded virtual resources to be $\Xi(\vec{\boldsymbol{\lambda}})$. We also need to modify the definition of ALGD and the lifting map:

Definition 6.1 (Lifting map) *With $L:\mathbb{R}^N_+\to\mathbb{R}_+$ as in the single-hop case, define the optimization problem* ALGD$(\mathbf{q})$ *to be*

$$\begin{aligned} &\textit{minimize} && L(\mathbf{r}) \quad \textit{over } \mathbf{r}\in\mathbb{R}^N_+ \\ &\textit{such that} && \boldsymbol{\xi}\cdot\vec{\mathbf{r}} \ge \boldsymbol{\xi}\cdot\vec{\mathbf{q}} \quad \textit{for all } \boldsymbol{\xi}\in\Xi(\vec{\boldsymbol{\lambda}}) \\ & && \textit{and} \quad \vec{r}_n \le \vec{q}_n \quad \textit{for all } n \textit{ such that } \vec{\lambda}_n = 0. \end{aligned}$$

Note that L is strictly convex and increasing componentwise, and the feasible region is convex, hence this problem has a unique optimizer. Define the lifting map ΔW : *$\mathbb{R}^N_+\to\mathbb{R}^N_+$ by setting $\Delta W(\mathbf{q})$ to be the optimizer.*

The main result of this section is the following. Throughout this section we are considering a multihop network with arrival rate vector $\boldsymbol{\lambda} \geq \mathbf{0}$ such that $\vec{\boldsymbol{\lambda}} \in \Lambda$, running MW-$f$ backpressure.

Theorem 6.2 (Portmanteau theorem, multihop version) *The statements of Theorem 5.4 parts (i)–(v) hold, for multihop fluid model solutions and using the multihop definition of ΔW.*

Some of the proofs for the single-hop case carry through to the multihop case. Other proofs rely on the fact that for single-hop networks, $\boldsymbol{\lambda} \in \Lambda \implies \boldsymbol{\lambda} \leq \boldsymbol{\sigma}$ for some $\boldsymbol{\sigma} \in \langle \mathcal{S} \rangle$, and these proofs require modification. We will modify them to use the following result.

Lemma 6.3 *Under Assumption 2.3, if $\boldsymbol{\sigma} \in \langle \mathcal{S} \rangle$ and $\boldsymbol{\sigma}' \in \mathbb{R}_+^N$ is such that $\boldsymbol{\sigma}' \leq \boldsymbol{\sigma}$, then $\boldsymbol{\sigma}' \in \langle \mathcal{S} \rangle$.*

Proof. It is sufficient to establish the result for the case when $\boldsymbol{\sigma}'$ differs from $\boldsymbol{\sigma}$ in only one component, as the repeated application of this will yield the full result. Without loss of generality, assume the queues are numbered such that $0 \leq \sigma'_1 < \sigma_1$ and $\sigma'_n = \sigma_n$ for $n \geq 2$. Since $\boldsymbol{\sigma} \in \langle \mathcal{S} \rangle$ there is a collection of positive constants $(a_{\boldsymbol{\pi}})_{\boldsymbol{\pi} \in \mathcal{S}}$ such that $\sum_{\boldsymbol{\pi}} a_{\boldsymbol{\pi}} = 1$ and $\boldsymbol{\sigma} = \sum_{\boldsymbol{\pi}} a_{\boldsymbol{\pi}} \boldsymbol{\pi}$. By Assumption 2.3, if $\boldsymbol{\pi} \in \mathcal{S}$ then $\boldsymbol{\pi}' \in \mathcal{S}$ where

$$\pi'_n = \begin{cases} 0 & \text{if } n = 1 \\ \pi_n & \text{otherwise,} \end{cases}$$

thus $\boldsymbol{\sigma}'' \in \langle \mathcal{S} \rangle$ where $\boldsymbol{\sigma}'' = \sum_{\boldsymbol{\pi}} a_{\boldsymbol{\pi}} \boldsymbol{\pi}'$. By construction, $\sigma''_1 = 0$ and $\sigma''_n = \sigma_n$ for $n \geq 2$. By choosing the appropriate convex combination

$$\boldsymbol{\sigma}' = (1-x)\boldsymbol{\sigma}'' + x\boldsymbol{\sigma} \quad \text{with} \quad x = \sigma'_1/\sigma_1 \in [0,1]$$

we see $\boldsymbol{\sigma}' \in \langle \mathcal{S} \rangle$. □

Now we proceed towards establishing Theorem 6.2. The proof of the first claim of Theorem 6.2(i) is just as for the single-hop case. The second claim follows from the following lemma.

Lemma 6.4 *For all $\mathbf{q} \in \mathbb{R}_+^N$,*

$$\boldsymbol{\lambda} \cdot f(\mathbf{q}) - \max_{\boldsymbol{\pi} \in \mathcal{S}} \boldsymbol{\pi} \cdot (I-R) f(\mathbf{q}) \leq 0. \tag{43}$$

Also, every fluid model solution satisfies

$$\frac{d}{dt} L(\mathbf{q}(t)) = \boldsymbol{\lambda} \cdot f(\mathbf{q}(t)) - \max_{\boldsymbol{\pi} \in \mathcal{S}} \boldsymbol{\pi} \cdot (I-R) f(\mathbf{q}(t)) \leq 0.$$

Proof. Since $\vec{\boldsymbol{\lambda}} \in \Lambda$, $\vec{R}\boldsymbol{\lambda} \leq \boldsymbol{\sigma}$ componentwise for some $\boldsymbol{\sigma} \in \langle\mathcal{S}\rangle$. Because $\vec{R} \geq 0$ and $\boldsymbol{\lambda} \geq \mathbf{0}$, $\vec{R}\boldsymbol{\lambda} \geq \mathbf{0}$ componentwise. By Lemma 6.3, $\boldsymbol{\lambda} = (I - R^{\mathsf{T}})\boldsymbol{\sigma}'$ for some $\boldsymbol{\sigma}' \in \langle\mathcal{S}\rangle$. Hence

$$\boldsymbol{\lambda}\cdot f(\mathbf{q}) - \max_{\boldsymbol{\pi}\in\mathcal{S}} \boldsymbol{\pi}\cdot(I-R)f(\mathbf{q}) = \boldsymbol{\sigma}'\cdot(I-R)f(\mathbf{q}) - \max_{\boldsymbol{\pi}\in\mathcal{S}} \boldsymbol{\pi}\cdot(I-R)f(\mathbf{q}) \leq 0.$$

For the claim about fluid model solutions,

$$\begin{aligned}
\frac{d}{dt}L(\mathbf{q}(t)) &= \dot{\mathbf{q}}(t)\cdot f(\mathbf{q}(t)) \\
&= \Big(\boldsymbol{\lambda} - (I-R^{\mathsf{T}})\Big[\sum_{\boldsymbol{\pi}} \dot{s}_{\boldsymbol{\pi}}(t)\boldsymbol{\pi} + \dot{\mathbf{y}}(t)\Big]\Big)\cdot f(\mathbf{q}(t)) \quad \text{by differentiating (21)} \\
&= \boldsymbol{\lambda}\cdot f(\mathbf{q}(t)) - \sum_{\boldsymbol{\pi}} \dot{s}_{\boldsymbol{\pi}}(t)\boldsymbol{\pi}\cdot(I-R)f(\mathbf{q}(t)) + \dot{\mathbf{y}}(t)\cdot(I-R)f(\mathbf{q}(t)).
\end{aligned}$$

For the middle term,

$$\begin{aligned}
\sum_{\boldsymbol{\pi}} \dot{s}_{\boldsymbol{\pi}}(t)\boldsymbol{\pi}\cdot(I-R)f(\mathbf{q}(t)) &= \max_{\boldsymbol{\rho}} \boldsymbol{\rho}\cdot(I-R)f(\mathbf{q}(t))\sum_{\boldsymbol{\pi}} \dot{s}_{\boldsymbol{\pi}}(t) \quad \text{by (22)} \\
&= \max_{\boldsymbol{\rho}} \boldsymbol{\rho}\cdot(I-R)f(\mathbf{q}(t)) \quad \text{by (15)}.
\end{aligned}$$

For the last term, we claim that

$$\dot{\mathbf{y}}(t)\cdot(I-R)f(\mathbf{q}(t)) = 0. \tag{44}$$

To see this, consider first a queue n with $[(I-R)f(\mathbf{q}(t))]_n > 0$. As noted in (10), this implies $f(q_n(t)) > f([R\mathbf{q}(t)]_n)$. By Assumption 2.2 it must be that $q_n(t) > 0$, hence $\dot{y}_n(t) = 0$ by (19). Second, consider a queue n with $[(I-R)f(\mathbf{q}(t))]_n < 0$. It must be that all of the active schedules do not serve this queue, i.e. $\dot{s}_{\boldsymbol{\pi}}(t) > 0 \implies \pi_n = 0$, since otherwise by Assumption 2.3 there is another schedule that has bigger weight than $\boldsymbol{\pi}$, contradicting (22). Third, if $[(I-R)f(\mathbf{q}(t))]_n = 0$ then obviously $\dot{y}_n(t)[(I-R)f(\mathbf{q}(t))]_n = 0$. Putting these three together proves (44).

Putting together these findings for the middle and last terms,

$$\frac{d}{dt}L(\mathbf{q}(t)) = \boldsymbol{\lambda}\cdot f(\mathbf{q}(t)) - \max_{\boldsymbol{\rho}} \boldsymbol{\rho}\cdot(I-R)f(\mathbf{q}(t)).$$

Applying (43) this is ≤ 0. □

The proof of Theorem 6.2(ii) is broadly similar to the single-hop case, Lemma 5.7, but the formulae all have to be adjusted to deal with multihop.

Lemma 6.5 *ΔW is continuous.*

Proof. If $\Xi(\vec{\boldsymbol{\lambda}})$ is empty, then the lifting map is trivial and the result is trivial. In what follows, we shall assume that $\Xi(\vec{\boldsymbol{\lambda}})$ is non-empty, and we will abbreviate it to Ξ. Furthermore note that for every $\boldsymbol{\xi} \in \Xi$ we know $\boldsymbol{\xi} \cdot \vec{\boldsymbol{\lambda}} = 1$ by definition of Ξ, and hence there is some queue n such that $\xi_n > 0$ and $\vec{\lambda}_n > 0$.

Pick any sequence $\mathbf{q}^k \to \mathbf{q}$, and let $\mathbf{r}^k = \Delta W(\mathbf{q}^k)$ and $\mathbf{r} = \Delta W(\mathbf{q})$. We want to prove that $\mathbf{r}^k \to \mathbf{r}$. We shall first prove that there is a compact set $[0,h]^N$ such that $\mathbf{r}^k \in [0,h]^N$ for all k. We shall then prove that any convergent subsequence of $\mathbf{r}^k$ converges to $\mathbf{r}$; this establishes continuity of ΔW.

First, compactness. A suitable value for h is

$$h = \max_{\boldsymbol{\xi} \in \Xi} \max_{n:\xi_n > 0} \sup_k \frac{\boldsymbol{\xi} \cdot \vec{\mathbf{q}}^k}{\xi_n}.$$

Note that the maximums are over a non-empty set, as noted at the beginning of the proof. Note also that h is finite because $\mathbf{q}$ is finite. Now, suppose that $\mathbf{r}^k \notin [0,h]^N$ for some k, i.e. that there is some queue n for which $r^k_n > h$, and let $\mathbf{r}' = \mathbf{r}^k$ in every coordinate except for $r'_n = h$. We claim that $\mathbf{r}'$ satisfies the two constraints of ALGD($\mathbf{q}^k$). To see that it satisfies the second constraint, note that $\mathbf{r}' \leq \mathbf{r}^k$ and hence if $\vec{\lambda}_n = 0$ then $\vec{r}'_n \leq \vec{r}^k_n \leq \vec{q}_n$. To see that it satisfies the first constraint, pick any $\boldsymbol{\xi} \in \Xi$. Either $\xi_m = 0$ for all queues m that are downstream of n, i.e. for which $\vec{R}_{mn} = 1$; if this is so then

$$\boldsymbol{\xi} \cdot \vec{\mathbf{r}}' = \boldsymbol{\xi} \cdot \vec{\mathbf{r}}^k + \boldsymbol{\xi} \cdot (\vec{\mathbf{r}}' - \vec{\mathbf{r}}^k) = \boldsymbol{\xi} \cdot \vec{\mathbf{r}}^k + \sum_l (r'_l - r^k_l) \sum_m \xi_m \vec{R}_{ml} = \boldsymbol{\xi} \cdot \vec{\mathbf{r}}^k.$$

Or $\xi_m > 0$ for some queue m that is downstream of n; if this is so then

$$\boldsymbol{\xi} \cdot \vec{\mathbf{r}}' \geq \xi_m \vec{r}'_m \geq \xi_m h \geq \boldsymbol{\xi} \cdot \vec{\mathbf{q}}^k \quad \text{by construction of } h.$$

Applying this repeatedly, if $\mathbf{r}^k \notin [0,h]^N$ then we can reduce it to a queue size vector in $[0,h]^N$, thereby improving on $L(\mathbf{r}^k)$, yet still meeting the constraints of ALGD($\mathbf{q}^k$); this contradicts the optimality of $\mathbf{r}^k$. Hence $\mathbf{r}^k \in [0,h]^N$.

Next, convergence on subsequences. With a slight abuse of notation, let $\Delta W(\mathbf{q}^k) = \mathbf{r}^k \to \mathbf{s}$ be a convergent subsequence, and recall that $\Delta W(\mathbf{q}) = \mathbf{r}$ and $\mathbf{q}^k \to \mathbf{q}$. By continuity of the constraints of ALGD, $\mathbf{s}$ is feasible for ALGD($\mathbf{q}$); we shall next show that $L(\mathbf{s}) \leq L(\mathbf{r})$. Since $\mathbf{r}$ is the unique optimum, it must be that $\mathbf{s} = \mathbf{r}$.

It remains to show that $L(\mathbf{s}) \leq L(\mathbf{r})$. We will construct a sequence $\mathbf{r} - \boldsymbol{\delta}^k + \varepsilon^k \mathbf{P}$ of candidate solutions to ALGD($\mathbf{q}^k$), choosing $\boldsymbol{\delta}^k \geq \mathbf{0}$ and $\varepsilon^k \mathbf{P} \geq \mathbf{0}$ to ensure that the candidate solutions are feasible. Specifically, we define

$$\delta^k_n = \begin{cases} 0 & \text{if } \vec{\lambda}_n > 0 \\ (\vec{q}_n - \vec{q}^k_n)^+ & \text{if } \vec{\lambda}_n = 0 \end{cases}$$

and $P_n = 1_{\vec{\lambda}_n > 0}$, and

$$\varepsilon^k = \max_{\boldsymbol{\xi}\in\Xi} \frac{\left(\boldsymbol{\xi}\cdot\vec{\mathbf{q}}^k - \boldsymbol{\xi}\cdot\vec{\mathbf{q}}\right)^+ + \boldsymbol{\xi}\cdot\vec{\boldsymbol{\delta}}^k}{\boldsymbol{\xi}\cdot\vec{\mathbf{P}}}.$$

We will first deal with the feasibility constraint that pertains when $\vec{\lambda}_n = 0$. Note that this implies $\lambda_m = 0$ for all queues m that are upstream of n, since $\vec{\lambda}_n = \sum_m \vec{R}_{nm}\lambda_m$, and hence that $\vec{\lambda}_m = 0$ for all upstream queues. Using this we find

$$\begin{aligned}
\left[\vec{R}(\mathbf{r} - \boldsymbol{\delta}^k + \varepsilon^k\mathbf{P})\right]_n &= \sum_m \vec{R}_{nm}\left[\mathbf{r} - \boldsymbol{\delta}^k + \varepsilon^k\mathbf{P}\right]_m \\
&= \sum_m \vec{R}_{nm}\left(r_m - (\vec{q}_m - \vec{q}^{\,k}_m)^+\right) \quad \text{since } \vec{\lambda}_m = 0 \text{ when } \vec{R}_{nm} = 1 \\
&= \Big(\sum_m \vec{R}_{nm} r_m\Big) - \Big(\sum_m \vec{R}_{nm}(\vec{q}_m - \vec{q}^{\,k}_m)^+\Big) \\
&\le \Big(\sum_m \vec{R}_{nm} r_m\Big) - (\vec{q}_n - \vec{q}^{\,k}_n)^+ \quad \text{as } \vec{R}_{nn} = 1,\ \vec{R}_{nm} \ge 0 \text{ for all } m \\
&= \vec{r}_n - (\vec{q}_n - \vec{q}^{\,k}_n)^+ \\
&\le \vec{q}_n - (\vec{q}_n - \vec{q}^{\,k}_n)^+ \quad \text{since } \mathbf{r} \text{ is feasible for } \mathrm{ALGD}(\mathbf{q}) \\
&= \min(\vec{q}_n, \vec{q}^{\,k}_n) \le \vec{q}^{\,k}_n.
\end{aligned}$$

Hence $\mathbf{r} - \boldsymbol{\delta}^k + \varepsilon^k\mathbf{P}$ satisfies the second feasibility constraint of $\mathrm{ALGD}(\mathbf{q}^k)$. For the other feasibility constraint of $\mathrm{ALGD}(\mathbf{q}^k)$, pick any $\boldsymbol{\xi} \in \Xi$. Then

$$\begin{aligned}
\boldsymbol{\xi}\cdot\vec{R}(\mathbf{r} - \boldsymbol{\delta}^k + \varepsilon^k\mathbf{P}) &= \boldsymbol{\xi}\cdot(\vec{\mathbf{r}} - \vec{\boldsymbol{\delta}}^k) + \varepsilon^k\boldsymbol{\xi}\cdot\vec{\mathbf{P}} \\
&\ge \boldsymbol{\xi}\cdot(\vec{\mathbf{r}} - \vec{\boldsymbol{\delta}}^k) + \left(\boldsymbol{\xi}\cdot\vec{\mathbf{q}}^k - \boldsymbol{\xi}\cdot\vec{\mathbf{q}}\right)^+ + \boldsymbol{\xi}\cdot\vec{\boldsymbol{\delta}}^k \quad \text{by construction of } \varepsilon^k \\
&\ge \boldsymbol{\xi}\cdot\vec{\mathbf{q}} - \left(\boldsymbol{\xi}\cdot\vec{\mathbf{q}}^k - \boldsymbol{\xi}\cdot\vec{\mathbf{q}}\right)^+ \quad \text{since } \vec{\mathbf{r}} \text{ is feasible for } \mathrm{ALGD}(\vec{\mathbf{q}}) \\
&= \max(\boldsymbol{\xi}\cdot\vec{\mathbf{q}}, \boldsymbol{\xi}\cdot\vec{\mathbf{q}}^k) \ge \boldsymbol{\xi}\cdot\vec{\mathbf{q}}^k.
\end{aligned}$$

Since the candidates are feasible solutions to $\mathrm{ALGD}(\mathbf{q}^k)$, and $\mathbf{r}^k$ is an optimal solution, it must be that

$$L(\mathbf{r}^k) \le L(\mathbf{r} - \boldsymbol{\delta}^k + \varepsilon^k\mathbf{P}).$$

Taking the limit as $k \to \infty$, and noting that L is continuous and $\boldsymbol{\delta}^k \to \mathbf{0}$ and $\varepsilon^k \to 0$, we find

$$L(\mathbf{s}) \le L(\mathbf{r})$$

as required. This completes the proof. □

For the proof of Theorem 6.2(iii), it is useful to work with a different representation of ΔW, provided by the following lemma, which draws on monotonicity of $\mathcal{S}$.

Lemma 6.6 *For any $\mathbf{q} \in \mathbb{R}_+^N$, $\Delta W(\mathbf{q})$ can be written*

$$\Delta W(\mathbf{q}) = \mathbf{q} + t\big(\boldsymbol{\lambda} - (I - R^{\mathsf{T}})\boldsymbol{\sigma}\big) \quad \textit{for some } t \geq 0, \boldsymbol{\sigma} \in \langle \mathcal{S} \rangle.$$

Proof. We will choose $\boldsymbol{\sigma}$ simply by multiplying each side of the desired equation by $\vec{R}$:

$$\vec{\mathbf{r}} = \vec{\mathbf{q}} + t(\vec{\boldsymbol{\lambda}} - \boldsymbol{\sigma}) \quad \text{where} \quad \mathbf{r} = \Delta W(\mathbf{q})$$

or, rearranging,

$$\boldsymbol{\sigma} = \vec{\boldsymbol{\lambda}} - (\vec{\mathbf{r}} - \vec{\mathbf{q}})/t.$$

We will show that $\mathbf{0} \leq \boldsymbol{\sigma} \leq \boldsymbol{\rho}$ for some $\boldsymbol{\rho} \in \langle \mathcal{S} \rangle$, hence by Lemma 6.3 $\boldsymbol{\sigma} \in \langle \mathcal{S} \rangle$.

First, we show $\mathbf{0} \leq \boldsymbol{\sigma}$. If $\vec{\lambda}_n > 0$ this can be achieved by choosing t sufficiently large. If $\vec{\lambda}_n = 0$ then by the second constraint of ALGD($\mathbf{q}$) we know that $\vec{r}_n \leq \vec{q}_n$ so $\sigma_n \geq 0$.

Second, we show $\boldsymbol{\sigma} \cdot \boldsymbol{\xi} \leq 1$ for all $\boldsymbol{\xi}$ that are feasible for DUAL($\vec{\boldsymbol{\lambda}}$). Either $\boldsymbol{\xi} \cdot \vec{\boldsymbol{\lambda}} = 1$, in which case $\boldsymbol{\xi} \in \langle \Xi(\vec{\boldsymbol{\lambda}}) \rangle$ and so by the first constraint of ALGD we know that $\vec{\mathbf{r}} \cdot \boldsymbol{\xi} \geq \vec{\mathbf{q}} \cdot \boldsymbol{\xi}$. Or $\boldsymbol{\xi} \cdot \vec{\boldsymbol{\lambda}} < 1$, in which case we simply need to choose t sufficiently large. Either way, $\boldsymbol{\sigma} \cdot \boldsymbol{\xi} \leq 1$ for all dual-feasible $\boldsymbol{\xi}$, hence DUAL($\boldsymbol{\sigma}$) ≤ 1, hence PRIMAL($\boldsymbol{\sigma}$) ≤ 1, hence $\boldsymbol{\sigma} \leq \boldsymbol{\rho}$ for some $\boldsymbol{\rho} \in \langle \mathcal{S} \rangle$ by the definition of PRIMAL($\boldsymbol{\sigma}$). □

The proof of Theorem 6.2(iii) is given by the following lemma. This proof is similar to the single-hop case, Lemma 5.9, but it is much shorter because the monotonicity assumption gives us a stronger representation of the lifting map, Lemma 6.6. Also, this version makes a weaker claim, namely that the lifting map is scale-invariant at invariant states, whereas the single-hop version shows that the lifting map is invariant everywhere.

Lemma 6.7 (Scale-invariance of the lifting map) *If $\mathbf{q} = \Delta W(\mathbf{q})$ then $\kappa \mathbf{q} = \Delta W(\kappa \mathbf{q})$ for all $\kappa > 0$.*

Proof. Suppose that $\mathbf{q} = \Delta W(\mathbf{q})$, and let $\kappa \mathbf{r} = \Delta W(\kappa \mathbf{q})$. Clearly $\kappa \mathbf{q}$ is feasible for ALGD($\kappa \mathbf{q}$); we shall show that $L(\kappa \mathbf{r}) \geq L(\kappa \mathbf{q})$, whence $\kappa \mathbf{q}$ is also optimal for ALGD($\kappa \mathbf{q}$), whence $\kappa \mathbf{q} = \kappa \mathbf{r}$ by uniqueness of the optimum.

It remains to prove that $L(\kappa \mathbf{r}) \geq L(\kappa \mathbf{q})$. By Lemma 6.6, we can write $\kappa \mathbf{r}$ as

$$\kappa \mathbf{r} = \kappa \mathbf{q} + t\big(\boldsymbol{\lambda} - (I - R^{\mathsf{T}})\boldsymbol{\sigma}\big)$$

for some $t \geq 0$ and some $\boldsymbol{\sigma} \in \Sigma$. Now consider the value of L along a straight-line trajectory from $\kappa\mathbf{q}$ to $\kappa\mathbf{r}$:

$$\begin{aligned}\frac{d}{du} L\Big(\kappa\mathbf{q}+(\boldsymbol{\lambda}-(I-R^{\mathsf{T}})\boldsymbol{\sigma})u\Big)\Big|_{u=0} &= \big(\boldsymbol{\lambda}-(I-R^{\mathsf{T}})\boldsymbol{\sigma}\big)\cdot f(\kappa\mathbf{q}) \\ &= \boldsymbol{\lambda}\cdot f(\kappa\mathbf{q}) - \boldsymbol{\sigma}\cdot(I-R)f(\kappa\mathbf{q}) \\ &\geq \boldsymbol{\lambda}\cdot f(\kappa\mathbf{q}) - \max_{\rho\in\mathcal{S}} \boldsymbol{\rho}\cdot(I-R)f(\kappa\mathbf{q}) \quad \text{for any } \sigma\in\langle\mathcal{S}\rangle \\ &= 0.\end{aligned}$$

The final equality is because

$$\boldsymbol{\lambda}\cdot f(\mathbf{q}) = \vec{\boldsymbol{\lambda}}\cdot(I-R)f(\mathbf{q}) = \max_{\boldsymbol{\pi}\in\mathcal{S}} \boldsymbol{\pi}\cdot(I-R)f(\mathbf{q})$$

by Lemma 6.9(iv) below (the proof of which does not assume the result of this lemma). Hence

$$\boldsymbol{\lambda}\cdot f(\kappa\mathbf{q}) = \vec{\boldsymbol{\lambda}}\cdot(I-R)f(\kappa\mathbf{q}) = \max_{\boldsymbol{\pi}\in\mathcal{S}} \boldsymbol{\pi}\cdot(I-R)f(\kappa\mathbf{q}),$$

using Assumption 2.2 and the fact that $\vec{\boldsymbol{\lambda}} \in \langle\mathcal{S}\rangle$ by Lemma 6.3. □

The proof of Theorem 6.2(iv) relies on the following lemma.

Lemma 6.8 (Fluid model trajectories preserve ALGD feasibility) *Consider any fluid model solution, for any scheduling policy, with initial queue size* $\mathbf{q}(0)$*. Then* $\mathbf{q}(t)$ *is feasible for* ALGD($\mathbf{q}(0)$) *for all* $t \geq 0$*.*

Proof. Feasibility for ALGD($\mathbf{q}(0)$) has two parts. For the first part, pick any critically loaded virtual resource $\boldsymbol{\xi} \in \Xi(\vec{\boldsymbol{\lambda}})$, and multiply each side of (21) by $\vec{R} = (I-R^T)^{-1}$ and then by $\boldsymbol{\xi}$ to get

$$\boldsymbol{\xi}\cdot\vec{\mathbf{q}}(t) = \boldsymbol{\xi}\cdot\vec{\mathbf{q}}(0) + \boldsymbol{\xi}\cdot\vec{R}\mathbf{a}(t) - \boldsymbol{\xi}\cdot\Big(\sum_{\boldsymbol{\pi}} s_{\boldsymbol{\pi}}(t)\boldsymbol{\pi} - \mathbf{y}(t)\Big).$$

Defining $\boldsymbol{\sigma}(t) = \sum \boldsymbol{\pi} s_{\boldsymbol{\pi}}(t)/t$, which is in $\langle\mathcal{S}\rangle$ by (15),

$$\begin{aligned}\boldsymbol{\xi}\cdot\vec{\mathbf{q}}(t) &\geq \boldsymbol{\xi}\cdot\vec{\mathbf{q}}(0) + t\big(\boldsymbol{\xi}\cdot\vec{\boldsymbol{\lambda}} - \boldsymbol{\xi}\cdot\boldsymbol{\sigma}(t)\big) \quad \text{by (14) and because } \mathbf{y}(t) \geq \mathbf{0} \\ &= \boldsymbol{\xi}\cdot\vec{\mathbf{q}}(0) + t\big(1-\boldsymbol{\xi}\cdot\boldsymbol{\sigma}(t)\big) \quad \text{since } \boldsymbol{\xi}\in\Xi(\vec{\boldsymbol{\lambda}}) \\ &\geq \boldsymbol{\xi}\cdot\vec{\mathbf{q}}(0) + t\big(1-1\big) \quad \text{since } \boldsymbol{\xi} \text{ is a virtual resource and } \boldsymbol{\sigma}\in\langle\mathcal{S}\rangle \\ &= \boldsymbol{\xi}\cdot\vec{\mathbf{q}}(0)\end{aligned}$$

as required for the first part of ALGD-feasibility. For the second part, suppose that $\vec{\lambda}_n = 0$ for some queue n. Multiply each side of (21) by $\vec{R}$ to get

$$\vec{\mathbf{q}}(t) = \vec{\mathbf{q}}(0) + \vec{\boldsymbol{\lambda}}t - \sum_{\boldsymbol{\pi}} \boldsymbol{\pi} s_{\boldsymbol{\pi}}(t) + \mathbf{y}(t) \leq \vec{\mathbf{q}}(0) + \vec{\boldsymbol{\lambda}}t$$

where the inequality is by (16). Since we assumed $\vec{\lambda}_n = 0$, $\vec{q}_n(t) \leq \vec{q}_n(0)$. This completes the proof that $\mathbf{q}(t)$ is feasible for ALGD($\mathbf{q}(0)$). □

The proof of Theorem 6.2(iv) is implied by parts (i) and (ii) of the following lemma.

Lemma 6.9 (Characterization of invariant states of MW-f backpressure) *The following are equivalent, for $\mathbf{q}^0 \in \mathbb{R}_+^N$:*

(i) $\mathbf{q}^0 = \Delta W(\mathbf{q}^0)$
(ii) $\mathbf{q}^0$ *is an invariant state*
(iii) there exists a fluid model solution with $\mathbf{q}(t) = \mathbf{q}^0$ *for all* t
(iv) $\boldsymbol{\lambda} \cdot f(\mathbf{q}^0) = \max_{\boldsymbol{\pi} \in \mathcal{S}} \boldsymbol{\pi} \cdot (I - R) f(\mathbf{q}^0)$

Proof. That (i) $\implies$ (ii) $\implies$ (iii) $\implies$ (iv) is proved in the same way as in the single-hop case. We just need to appeal to Lemma 6.4 rather than 5.5 for the fact that $L(\mathbf{q}(t))$ is decreasing, and to Lemma 6.8 rather than 5.10 for the fact that $\mathbf{q}(t)$ remains feasible.

Proof that (iv) $\implies$ (i). Let $\mathbf{r} = \Delta W(\mathbf{q}^0)$. By Lemma 6.6, $\mathbf{r} = \mathbf{q}^0 + t(\boldsymbol{\lambda} - (I - R^\mathsf{T})\boldsymbol{\sigma})$ for some $t \geq 0$ and $\boldsymbol{\sigma} \in \langle \mathcal{S} \rangle$. By considering the value of $L(\cdot)$ along the trajectory from $\mathbf{q}^0$ to $\mathbf{r}$, and using (iv), we conclude that $L(\mathbf{r}) \geq L(\mathbf{q}^0)$. By the same argument as in the single-hop case, $\mathbf{q}^0 = \Delta W(\mathbf{q}^0)$. □

The proof of Theorem 6.2(v) is given by the following lemma.

Lemma 6.10 *Given $\vec{\boldsymbol{\lambda}} \in \Lambda$, for any $\varepsilon > 0$ there exists an $H_\varepsilon > 0$ such that for every fluid model solution with arrival rate $\boldsymbol{\lambda}$, for which $|\mathbf{q}(0)| \leq 1$, $|\mathbf{q}(t) - \Delta W(\mathbf{q}(t))| < \varepsilon$ for all $t \geq H_\varepsilon$.*

Proof. The proof of Lemma 5.12 goes through almost verbatim. The only changes are in the penultimate paragraph, which should be replaced by the following:

Time to hit $\mathcal{K}_\delta$. Consider first the rate of change of $K(\cdot)$ while the process is in $\mathcal{D} \setminus \mathcal{K}_\delta$:

$$\begin{aligned}
\dot{K}(\mathbf{q}(t)) \leq \dot{L}(\mathbf{q}(t)) &= \boldsymbol{\lambda} \cdot f(\mathbf{q}(t)) - \max_{\boldsymbol{\pi} \in \mathcal{S}} \boldsymbol{\pi} \cdot (I - R) f(\mathbf{q}(t)) \\
&\leq \sup_{\mathbf{r} \in \mathcal{D} \setminus \mathcal{K}_\delta} \Big[\boldsymbol{\lambda} \cdot f(\mathbf{r}) - \max_{\boldsymbol{\pi} \in \mathcal{S}} \boldsymbol{\pi} \cdot (I - R) f(\mathbf{r}) \Big] \\
&\leq 0 \quad \text{by Lemma 6.4.}
\end{aligned}$$

This supremum is of a continuous function of $\mathbf{r}$, taken over a closed and bounded set, hence the supremum is attained at some $\hat{\mathbf{r}} \in \mathcal{D} \setminus \mathcal{K}_\delta$. If the supremum were equal to 0 then $\boldsymbol{\lambda} \cdot f(\hat{\mathbf{r}}) = \max_{\boldsymbol{\pi}} \boldsymbol{\pi} \cdot (I - R) f(\hat{\mathbf{r}})$ so $\hat{\mathbf{r}} \in \mathcal{I}$ by Lemma 6.9; but $\hat{\mathbf{r}} \in \mathcal{D} \setminus \mathcal{K}_\delta$ and we just proved that $\mathcal{I} \subset \mathcal{K}_\delta$; hence the supremum is some $-\eta_\delta < 0$. □

7. Multiplicative state-space collapse. This section establishes multiplicative state space collapse of queue size. It shows that under the MW-f policy and with suitable initial conditions when the network is not overloaded (i.e. when $\boldsymbol{\lambda} \in \Lambda$), the appropriately normalized queue size vector is constrained to lie in or close to the *set of invariant states*

$$\mathcal{I} = \{\mathbf{q} \in \mathbb{R}_+^N : \mathbf{q} = \Delta W(\mathbf{q})\}.$$

We assume that arrivals satisfy Assumption 2.5, and let the arrival rate vector $\boldsymbol{\lambda}$ be as specified in that assumption. The function ΔW depends on $\boldsymbol{\lambda}$ and f, as specified in Sections 5 and 6 for single-hop and multihop networks respectively, and the interesting case is where $\boldsymbol{\lambda} \in \partial\Lambda$ (since otherwise ΔW is trivial).

This section mostly follows the method developed by Bramson [3], except that our proof avoids the need for regenerative assumptions on the arrival process by imposing slightly tighter bounds on the uniformity of their convergence, as expressed by Assumption 2.5.

Consider a sequence of systems of the type described in Section 2.1 running a scheduling policy of the type described in Section 2.2. Let the systems all have the same number of queues N, the same set of allowed schedules $\mathcal{S}$, the same routing matrix R, and the same scheduling policy. Let the sequence of systems be indexed by $r \in \mathbb{N}$. Write

$$X^r(\tau) = (\mathbf{Q}^r(\tau), \mathbf{A}^r(\tau), \mathbf{Z}^r(\tau), S^r(\tau)), \quad \tau \in \mathbb{Z}_+,$$

for the rth system. Define the scaled system $\hat{x}^r(t) = (\hat{\mathbf{q}}^r(t), \hat{\mathbf{a}}^r(t), \hat{\mathbf{z}}^r(t), \hat{s}^r(t))$ for $t \in \mathbb{R}_+$ by

$$\hat{\mathbf{q}}^r(t) = \mathbf{Q}^r(r^2 t)/r \qquad \hat{\mathbf{a}}^r(t) = \mathbf{A}^r(r^2 t)/r$$
$$\hat{\mathbf{z}}^r(t) = \mathbf{Z}^r(r^2 t)/r \qquad \hat{s}^r_{\boldsymbol{\pi}}(t) = S^r_{\boldsymbol{\pi}}(r^2 t)/r$$

after extending the domain of $X^r(\cdot)$ to $\mathbb{R}_+$ by linear interpolation in each interval $(\tau, \tau+1)$. Note that each sample path of a scaled system $\hat{x}^r(t)$ over the interval $t \in [0, T]$ lies in $C^I(T)$ with $I = 3N + |\mathcal{S}|$. $T > 0$ will be fixed for the remainder of this section. Recall the norm $\|x\| = \sup_{0 \le t \le T} |x(t)|$. The main result of this paper is the following.

Theorem 7.1 (Multiplicative state-space collapse) *Consider a sequence of (single-hop or multihop) switched networks indexed by $r \in \mathbb{N}$, operating under the MW-f policy (with f satisfying Assumption 2.1 or 2.2 and $\mathcal{S}$ with 2.3), as described above. Assume that the arrival processes satisfy Assumption 2.5 with $\boldsymbol{\lambda} \in \Lambda$. Also assume that the initial queue sizes are non-random, and satisfy $\lim_{r\to\infty} \hat{\mathbf{q}}^r(0) = \hat{\mathbf{q}}_0$ for some $\hat{\mathbf{q}}_0 \in \mathcal{I}$. Then for any $\delta > 0$,*

$$\mathbb{P}\left(\frac{\|\hat{\mathbf{q}}^r(\cdot) - \Delta W(\hat{\mathbf{q}}^r(\cdot))\|}{\|\hat{\mathbf{q}}^r(\cdot)\| \vee 1} < \delta\right) \to 1 \quad \textit{as } r \to \infty. \tag{45}$$

Simulations suggest that a stronger result holds in the widely-studied *diffusion* or *heavy traffic* scaling: $\boldsymbol{\lambda}^r = \boldsymbol{\lambda} - \Gamma/r$ for some non-trivial $\Gamma \in \mathbb{R}_+^N$ and $\boldsymbol{\lambda} \in \partial\Lambda$. We conjecture the following.

Conjecture 7.2 *Under the assumptions of Theorem 7.1 and the additional assumption that increments in the arrival process are i.i.d. and uniformly bounded, under the diffusion scaling for any $\delta > 0$*

$$\mathbb{P}\Big(\big\|\hat{\mathbf{q}}^r(\cdot) - \Delta W(\hat{\mathbf{q}}^r(\cdot))\big\| < \delta\Big) \to 1 \quad \text{as } r \to \infty. \tag{46}$$

7.1. *Outline of the proof of Theorem 7.1.* The outline of the proof of Theorem 7.1 is as follows. We are interested in the dynamics of $\hat{\mathbf{q}}^r(t)$ over $t \in [0, T]$, i.e. of $\mathbf{Q}^r(\tau)$ over $\tau \in [0, r^2T]$. We will split this time interval into $\lfloor rT \rfloor + 1$ pieces starting at $0, r, 2r, \dots,$ and look at each piece under a fluid scaling. We will define a 'good event' $\hat{E}_r$ under which the arrivals in all of the pieces are well-behaved (Section 7.1.1). We then apply Theorem 4.3 to deduce that, under this event, the queue size process in each of the pieces can be (uniformly) approximated by a fluid model solution (Lemma 7.3). We then use the properties of the fluid model solution stated in Theorem 5.4 to show that in each of the pieces, the queue size is (uniformly) close to the set of invariant states (Lemmas 7.4 and 7.5). Figure 2 depicts the idea. Finally we show that $\mathbb{P}(\hat{E}_r) \to 1$ (Lemma 7.6). The formal proof is given in Section 7.1.2.

Note that Lemmas 7.3–7.5 are all sample path-wise results that hold for every $\omega \in \hat{E}_r$, and so questions of independence etc. do not arise. The only part of the proof where probability comes in is Lemma 7.6.

The proof is written out for a single-hop switched network. For the multihop case, the argument holds verbatim; simply replace all references to the single-hop fluid limit Theorem 4.3 by references to the equivalent multihop result, and replace all references to the description of single-hop fluid model solutions in Theorem 5.4 by references to the multihop version Theorem 6.2.

7.1.1. *The good event, and the fluid-scaled pieces.* Define the fluid-scaled pieces $\tilde{x}^{r,m,z}(u) = (\tilde{\mathbf{q}}^{r,m,z}(u), \tilde{\mathbf{a}}^{r,m,z}(u), \tilde{\mathbf{y}}^{r,m,z}(u), \tilde{s}^{r,m,z}(u))$ of the original process by

$$\begin{aligned}
\tilde{\mathbf{q}}^{r,m,z}(u) &= \mathbf{Q}^r(rm + zu)/z \\
\tilde{\mathbf{a}}^{r,m,z}(u) &= \big(\mathbf{A}^r(rm + zu) - \mathbf{A}^r(rm)\big)/z \\
\tilde{\mathbf{y}}^{r,m,z}(u) &= \big(\mathbf{Y}^r(rm + zu) - \mathbf{Y}^r(rm)\big)/z \\
\tilde{s}^{r,m,z}(u) &= \big(S^r(rm + zu) - S^r(rm)\big)/z
\end{aligned}$$

for $0 \le m \le \lfloor rT \rfloor$, $z \ge r$ and $u \ge 0$. Here r indicates which process we are considering, m indicates the piece, and z indicates the fluid-scaling parameter. The

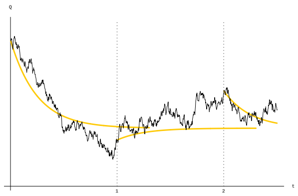

FIG 2. *Splitting the process into fluid-scaled parts, starting at* $0, r, 2r, \ldots$.

scaling parameter $z_{r,m} = |\mathbf{Q}^r(rm)| \vee r$ is particularly important, and for convenience we will define $\tilde{x}^{r,m}(u) = (\tilde{\mathbf{q}}^{r,m}(u), \tilde{\mathbf{a}}^{r,m}(u), \tilde{\mathbf{y}}^{r,m}(u), \tilde{s}^{r,m}(u))$ by

$$\tilde{\mathbf{q}}^{r,m}(u) = \tilde{\mathbf{q}}^{r,m,z_{r,m}}(u) \qquad \tilde{\mathbf{a}}^{r,m}(u) = \tilde{\mathbf{a}}^{r,m,z_{r,m}}(u)$$
$$\tilde{\mathbf{y}}^{r,m}(u) = \tilde{\mathbf{y}}^{r,m,z_{r,m}}(u) \qquad \tilde{s}^{r,m}(u) = \tilde{s}^{r,m,z_{r,m}}(u).$$

The good event is defined to be

$$\hat{E}_r = \left\{ \sup_{u \in [0,T^{\text{fluid}}]} \left| \tilde{\mathbf{a}}^{r,m,w_{r,k}}(u) - \boldsymbol{\lambda}^r u \right| < \eta_r \quad \text{for all } 0 \le m \le \lfloor rT \rfloor \right. \\ \left. \text{and } 0 \le k \le \lfloor Lr \log r \rfloor, \text{ where } w_{r,k} = r(1 + k/\log r) \right\}. \tag{47}$$

By this, we mean that $\hat{E}_r$ is a subset of the sample space for the rth system, and we write $\tilde{x}^{r,m,w_{r,k}}(\cdot)(\omega)$ etc. for $\omega \in \hat{E}_r$ when we wish to emphasize the dependence on $\hat{E}_r$. The constants here are $T^{\text{fluid}} = (2 + \lambda^{\max} + NS^{\max})(H_\zeta + 1)$, $\lambda^{\max} = \sup_r |\boldsymbol{\lambda}^r|$, $\zeta > 0$ is chosen as specified in Section 7.1.2 below, H_ζ is chosen as in Theorem 5.4(v), $L = 1 + T(1 + \lambda^{\max} + NS^{\max})$, $S^{\max} = \max_{\boldsymbol{\pi} \in \mathcal{S}} |\boldsymbol{\pi}|$, and the sequence of deviation terms $\eta_r \in [0,1]$ is chosen as specified in Lemma 7.6 such that $\eta_r \to 0$ as $r \to \infty$.

Lemma 7.3 *Let* FMS *be the set of fluid model solutions over time horizon* $[0, T^{\text{fluid}}]$ *for arrival rate vector* $\boldsymbol{\lambda}$, *and let* $\text{FMS}(\mathbf{q}_0)$ *and* FMS_1 *be as specified in Definition 4.1. Then*

$$\sup_{\omega \in \hat{E}_r} \max_{0 \le m \le \lfloor rT \rfloor} d\big(\tilde{x}^{r,m}(\cdot)(\omega), \text{FMS}_1\big) \to 0 \quad \textit{as } r \to \infty \tag{48}$$

and

$$\sup_{\omega \in \hat{E}_r} d\big(\tilde{x}^{r,0}(\cdot)(\omega), \text{FMS}(\mathbf{q}_0)\big) \to 0 \quad \textit{as } r \to \infty \tag{49}$$

where $\mathbf{q}_0 = \hat{\mathbf{q}}_0 / (|\hat{\mathbf{q}}_0| \vee 1)$.

Proof. The proof of each equation will use Theorem 4.3. We start with (48). The theorem requires the use of an index j in some totally ordered countable set; here we shall use the pair $j \equiv (r, m)$ ordered lexicographically, where $r \in \mathbb{N}$ and $0 \leq m \leq \lfloor rT \rfloor$. Lexicographic ordering means $(r, m) \geq (r', m')$ iff either $r > r'$ or both $r = r'$ and $m \geq m'$. Note that $j \to \infty$ implies $r \to \infty$ (and vice versa).

To apply the theorem, we first need to pick constants. Let $K = 1$, let $\boldsymbol{\lambda}$ and $\boldsymbol{\lambda}^j = \boldsymbol{\lambda}^r \to \boldsymbol{\lambda}$ as per Assumption 2.5, and let $\varepsilon_j = \eta_r(1 + 1/\log r)$ so that $\varepsilon_j \to 0$. Thus condition (23) of Theorem 4.3 is satisfied. Now let

$$G_j \equiv G_{r,m} = \Big\{ \big(\tilde{x}^{r,m}(\cdot)(\omega), z_{r,m}(\omega)\big) \;:\; \omega \in \hat{E}_r \Big\}.$$

It is worth stressing that $G_j \equiv G_{r,m}$ is a set of sample paths and associated scaling parameters, not a probabilistic event, and so any questions about the lack of independence between $\tilde{x}^{r,m}(\cdot)(\omega)$ and $z_{r,m}(\omega)$ are void. Note also that although the events $\hat{E}_r$ lie in different probability spaces for each r, this has no bearing on the definition of G_j nor on the application of Theorem 4.3.

We next show that G_j satisfies conditions (24)–(26) of Theorem 4.3, for j sufficiently large. Equation (24) follows straightforwardly from the fact that $z_{r,m} \geq r$, hence $\inf\{z : (\tilde{x}, z) \in G_{r,m}\} \geq r$, hence $\inf\{z : (\tilde{x}, z) \in G_j\} \to \infty$. For equation (25): later in the proof we will establish that, under $\hat{E}_r$ for r large enough,

$$\sup_{t \in [0, T^{\text{fluid}}]} \left| \tilde{\mathbf{a}}^{r,m}(t) - \boldsymbol{\lambda}^r t \right| < \eta_r \Big(1 + \frac{1}{\log r}\Big) \quad \text{for all } 0 \leq m \leq \lfloor rT \rfloor \tag{50}$$

which implies that for all $(\tilde{x}, z) \in G_j \equiv G_{r,m}$, $\sup_{t \in [0, T^{\text{fluid}}]} |\tilde{\mathbf{a}}^j(t) - \boldsymbol{\lambda}^j t| < \varepsilon_j$ as required. Equation (26) follows straightforwardly from the scaling used to define $\tilde{\mathbf{q}}^j(0) \equiv \tilde{\mathbf{q}}^{r,m}(0)$: for every ω, not merely $\omega \in \hat{E}_r$,

$$\left|\tilde{\mathbf{q}}^{r,m}(0)\right| = \left| \frac{\mathbf{Q}^r(rm)}{z_{r,m}} \right| = \left| \frac{\mathbf{Q}^r(rm)}{|\mathbf{Q}^r(rm)| \vee r} \right| \leq 1.$$

Since G_j satisfies the conditions of Theorem 4.3 for sufficiently large j, we can apply that theorem to deduce

$$\sup_{(\tilde{x}, z) \in G_j} d(\tilde{x}, \text{FMS}_1) \to 0 \quad \text{as } j \to \infty.$$

Rewriting j as (r, m), and turning the limit statement into a lim sup statement,

$$\sup_{(r', m) \geq (r, 0)} \; \sup_{(\tilde{x}, z) \in G_{r', m}} d(\tilde{x}, \text{FMS}_1) \to 0 \quad \text{as } r \to \infty$$

and in particular

$$\max_{0\le m\le\lfloor rT\rfloor}\ \sup_{(\tilde{x},z)\in G_{r,m}} d(\tilde{x},\mathrm{FMS}_1)\to 0 \quad \text{as } r\to\infty.$$

Rewriting $(\tilde{x},z)\in G_{r,m}$ in terms of $\omega\in\hat{E}_r$, as per the definition of $G_{r,m}$,

$$\max_{0\le m\le\lfloor rT\rfloor}\ \sup_{\omega\in\hat{E}_r} d(\tilde{x}^{r,m}(\cdot)(\omega),\mathrm{FMS}_1)\to 0 \quad \text{as } r\to\infty.$$

Interchanging the max and the sup gives (48).

To establish (49), we will again apply Theorem 4.3 but this time using the index $j\equiv r$, $\boldsymbol{\lambda}^j=\boldsymbol{\lambda}^r\to\boldsymbol{\lambda}$ and $\varepsilon_j=\eta_r(1+1/\log r)$ as above, and define

$$G_j\equiv G_r=\Big\{\big(\tilde{x}^{r,0}(\cdot)(\omega),z_{r,0}(\omega)\big)\ :\ \omega\in\hat{E}_r\Big\}.$$

Equations (23)–(26) hold just as before. For equation (28), we will use $\mathbf{q}_0$ as in the statement of this lemma, and $\varepsilon_j'=|\tilde{\mathbf{q}}^{r,0}(0)-\mathbf{q}_0|$. This is a well-defined constant (i.e. it does not depend on the randomness ω), because we assumed in Theorem 7.1 that the initial queue sizes $\mathbf{Q}^r(0)$ are non-random, and by definition $\tilde{\mathbf{q}}^{r,0}(0)=\mathbf{Q}^r(0)/(|\mathbf{Q}^r(0)|\vee r)$. Furthermore, Theorem 7.1 assumes $\hat{\mathbf{q}}^r(0)\to\hat{\mathbf{q}}_0$, which implies $\tilde{\mathbf{q}}^{r,0}(0)\to\mathbf{q}_0$ hence $\varepsilon_j'\to 0$. Equation (28) then follows straightforwardly, for every ω not merely $\omega\in\hat{E}_r$. Applying Theorem 4.3, we deduce that

$$\sup_{(\tilde{x},z)\in G_j} d\big(\tilde{x},\mathrm{FMS}(\mathbf{q}_0)\big)\to 0 \quad \text{as } j\to\infty.$$

Equivalently,

$$\sup_{\omega\in\hat{E}_r} d\big(\tilde{x}^{r,0}(\cdot)(\omega),\mathrm{FMS}(\mathbf{q}_0)\big)\to 0 \quad \text{as } r\to\infty$$

as required.

To complete the proof of Lemma 7.3, the only remaining claim that needs to be established is (50). We will proceed in two steps. First we prove that $|\mathbf{Q}^r(rm)|\le Lr^2$ under $\hat{E}_r$, for r sufficiently large and for all $0\le m\le\lfloor rT\rfloor$. To see this, note from (7) that

$$\mathbf{Q}^r(rm)\le\mathbf{Q}^r(0)+\mathbf{A}^r(rm)+NrmS^{\max}.$$

Now $\hat{E}_r$ gives a suitable bound on arrivals: for all $0\le m'\le\lfloor rT\rfloor$, and using the fact that $1\le T^{\text{fluid}}$,

$$\left|\frac{\mathbf{A}^r(rm'+r)-\mathbf{A}^r(rm')}{r}-\boldsymbol{\lambda}^r\right|=\left|\tilde{\mathbf{a}}^{r,m',w_{r,0}}(1)-\boldsymbol{\lambda}^r\right|<\eta_r$$

and by applying this from $m' = 0$ to $m' = m-1$ we find $|\mathbf{A}^r(rm)| \le rm(\lambda^{\max} + \eta_r)$. The assumptions of Theorem 7.1 tell us that $\mathbf{Q}^r(0)/r \to \hat{\mathbf{q}}_0$ for some $\hat{\mathbf{q}}_0 \in \mathbb{R}_+^N$. Putting all this together, we find that for sufficiently large r

$$|\mathbf{Q}^r(rm)| \le r^2\big(1 + T(1 + \lambda^{\max} + NS^{\max})\big) = Lr^2, \quad \text{for any } 0 \le m \le \lfloor rT \rfloor.$$

Now we proceed to prove (50), under $\hat{E}_r$ for r sufficiently large. Observe that (for $r > 2$) there exists $k \in \{1, \dots, \lfloor Lr \log r \rfloor\}$ such that $w_{r,k-1} \le z_{r,m} \le w_{r,k}$; this follows from $r \le z_{r,m} = |\mathbf{Q}^r(rm)| \vee r \le Lr^2$ and the definition of $w_{r,k}$ in (47). Hence for any $t \in [0, T^{\text{fluid}}]$,

$$\begin{aligned}
\big|\tilde{\mathbf{a}}^{r,m}(t) - \boldsymbol{\lambda}^r t\big| &= \left|\frac{\mathbf{A}^r(rm + z_{r,m}t) - \mathbf{A}^r(rm)}{z_{r,m}} - \boldsymbol{\lambda}^r t\right| \\
&= \left|\frac{\mathbf{A}^r(rm + w_{r,k}u) - \mathbf{A}^r(rm)}{w_{r,k}} - \boldsymbol{\lambda}^r u\right| \Big(\frac{w_{r,k}}{z_{r,m}}\Big) \quad \text{where } u = tz_{r,m}/w_{r,k} \\
&= \big|\tilde{\mathbf{a}}^{r,m,w_{r,k}}(u) - \boldsymbol{\lambda}^r u\big| \Big(\frac{w_{r,k}}{z_{r,m}}\Big) \\
&< \eta_r \frac{w_{r,k}}{z_{r,m}} \quad \text{since } \hat{E}_r \text{ holds and } u \le t \le T^{\text{fluid}} \\
&\le \eta_r \frac{w_{r,k}}{w_{r,k-1}} \quad \text{since } z_{r,m} \ge w_{r,k-1} \\
&= \eta_r \Big(1 + \frac{1}{k - 1 + \log r}\Big) \le \eta_r \Big(1 + \frac{1}{\log r}\Big).
\end{aligned}$$

This establishes (50) and completes the proof. □

Lemma 7.4 (Choice of approximating piece) *Given $t \in [0, T]$ and $r \in \mathbb{N}$, define $m^* = m^*(r,t)$ and $u^* = u^*(r,t)$ by*

$$m^* = \min\Big\{m \in \mathbb{Z}_+ : rm \le r^2 t \le rm + T^{\text{fluid}} z_{r,m}\Big\}, \quad u^* = \frac{r^2 t - rm^*}{z_{r,m^*}}.$$

This is a sound definition (i.e. the set for m^* is non-empty). Further, under event $\hat{E}_r$, either $m^* = 0$ and $0 \le u^* \le T^{\text{fluid}}$, or $0 < m^* \le \lfloor rT \rfloor$ and $H_\zeta < u^* \le T^{\text{fluid}}$.

Proof. The set for m^* is non-empty because $z_{r,m} \ge r$ and $T^{\text{fluid}} \ge 1$. The upper bound for m^* is trivial. The upper bound for u^* in either case is trivial. To prove the lower bound for u^* when $m^* > 0$, $r^2 t > r(m^* - 1) + T^{\text{fluid}} z_{r,m^*-1}$ due to the minimality of m^*. Hence

$$u^* = \frac{r^2 t - rm^*}{z_{r,m^*}} > \frac{T^{\text{fluid}} z_{r,m^*-1} - r}{z_{r,m^*}} \ge T^{\text{fluid}} \frac{z_{r,m^*-1}}{z_{r,m^*}} - 1.$$

To bound $z_{r,m^*-1}/z_{r,m^*}$, we can use (7) and the bound on $\mathbf{a}^{r,m,w_{r,0}}(1)$ provided by $\hat{E}_r$ to show that for any m

$$\begin{aligned} z_{r,m} &= \big|\mathbf{Q}^r(rm)\big| \vee r \\ &\le \Big(\big|\mathbf{Q}^r(rm-r)\big| + r\big(\lambda^{\max} + \eta_r + NS^{\max}\big)\Big) \vee r \\ &\le \big|\mathbf{Q}^r(rm-r)\big| \vee r + r\big(\lambda^{\max} + 1 + NS^{\max}\big) \quad \text{since } \eta_r \le 1 \\ &\le z_{r,m-1}\big(2 + \lambda^{\max} + NS^{\max}\big) \quad \text{since } z_{r,m-1} \ge r. \end{aligned}$$

Substituting this back into the earlier bound for u^*,

$$u^* > \frac{T^{\text{fluid}}}{2 + \lambda^{\max} + NS^{\max}} - 1$$

and this is equal to H_ζ by choice of T^{fluid}. □

Lemma 7.5 (Pathwise multiplicative state space collapse) *Let* $0 < \zeta < 1$, $t \in [0,T]$ *and* $r \in \mathbb{N}$ *be given. Suppose there exist* $m \in \{0,\dots,\lfloor rT \rfloor\}$, $u \in [0, T^{\text{fluid}}]$ *and* $x \in \mathrm{FMS}$ *such that* $r^2 t = rm + z_{r,m}u$ *and* $\|\tilde{x}^{r,m} - x\| < \zeta$, *and furthermore either (i)* $m > 0$ *and* $u > H_\zeta$ *and* $x \in \mathrm{FMS}_1$, *or (ii)* $m = 0$ *and* $x \in \mathrm{FMS}(\mathbf{q}_0)$ *where* $\mathbf{q}_0$ *is as defined in Lemma 7.3. Then*

$$\frac{\big|\hat{\mathbf{q}}^r(t) - \Delta W(\hat{\mathbf{q}}^r(t))\big|}{\big\|\hat{\mathbf{q}}^r(\cdot)\big\| \vee 1} \le \frac{\big|\hat{\mathbf{q}}^r(t) - \Delta W(\hat{\mathbf{q}}^r(t))\big|}{z_{r,m}/r} < 2\zeta + \mathrm{mc}_\zeta(\Delta W) \tag{51}$$

where $\mathrm{mc}_\zeta(\Delta W)$ *is the modulus of continuity of the map* $\mathbf{q} \mapsto \Delta W(\mathbf{q})$ *over*

$$\mathcal{D} = \big\{\mathbf{q}' \in \mathbb{R}_+^N : |\mathbf{q}' - \mathbf{q}| \le 1 \text{ for some } \mathbf{q} \text{ such that } L(\mathbf{q}) \le L(\mathbf{1})\big\}$$

for $L(\cdot)$ *as in Definition 5.3.*

Proof. The first inequality is trivially true because

$$\frac{z_{r,m}}{r} = \frac{|\mathbf{Q}^r(rm)| \vee r}{r} \le \Big(\sup_{u\in[0,T]} \hat{\mathbf{q}}^r(u)\Big) \vee 1.$$

For the second inequality, note that after unwrapping the $\hat{\mathbf{q}}^r(\cdot)$ scaling and wrapping it up again in the $\tilde{\mathbf{q}}^{r,m}$ scaling, the middle term in (51) is

$$\mathrm{MT} = \big|\tilde{\mathbf{q}}^{r,m}(u) - \Delta W(\tilde{\mathbf{q}}^{r,m}(u))\big|.$$

Writing $\mathbf{q}$ for the queue component of x,

$$\begin{aligned} \mathrm{MT} &\le \Big|\tilde{\mathbf{q}}^{r,m}(u) - \mathbf{q}(u)\Big| + \Big|\mathbf{q}(u) - \Delta W(\mathbf{q}(u))\Big| + \Big|\Delta W(\mathbf{q}(u)) - \Delta W(\tilde{\mathbf{q}}^{r,m}(u))\Big| \\ &= (52a) + (52b) + (52c) \quad \text{respectively.} \end{aligned}$$

We can bound each term as follows:

(52a) is $<\zeta$ since $\|\tilde{x}^{r,m}-x\|<\zeta$ by an assumption of the lemma.

(52b) is $<\zeta$ in the case $m>0$: By the assumptions of the lemma, $x\in\mathrm{FMS}_1$ so $|\mathbf{q}(0)|\leq 1$, and also $u>H_\zeta$. The requirements of Theorem 5.4(v) are met, hence we obtain the inequality.

(52b) is $=0$ in the case $m^*=0$: In this case, by assumption of the lemma $x\in\mathrm{FMS}(\mathbf{q}_0)$ i.e. $\mathbf{q}(0)=\mathbf{q}_0$. By assumption of Theorem 7.1 $\hat{\mathbf{q}}_0\in\mathcal{I}$, that is $\hat{\mathbf{q}}_0=\Delta W(\hat{\mathbf{q}}_0)$, therefore by Theorem 5.4(iii) $\mathbf{q}_0\in\mathcal{I}$, therefore by Theorem 5.4(iv) the fluid model solution $\mathbf{q}(\cdot)$ stays constant at $\mathbf{q}_0$ and so (52b)$=0$.

(52c) is $\leq \mathrm{mc}_\zeta(\Delta W)$: By assumption of the lemma, either $m>0$ and $x\in\mathrm{FMS}_1$, or $m=0$ and $x\in\mathrm{FMS}(\mathbf{q}_0)$ where $\mathbf{q}_0\leq\mathbf{1}$ componentwise; either way $\mathbf{q}(0)\leq\mathbf{1}$ componentwise, so $L(\mathbf{q}(0))\leq L(\mathbf{1})$. By Theorem 5.4(i) $L(\mathbf{q}(u))\leq L(\mathbf{1})$ so $\mathbf{q}(u)\in\mathcal{D}$. Furthermore, since $\|\tilde{x}^{r,m}-x\|<\zeta$ by assumption of the lemma, $|\tilde{\mathbf{q}}^{r,m}(u)-\mathbf{q}(u)|<\zeta<1$ and so $\tilde{\mathbf{q}}^{r,m}(u)\in\mathcal{D}$. The inequality then follows from the definition of the modulus of continuity. □

Lemma 7.6 (The good event has high probability) *Under the assumptions of Theorem 7.1, $\mathbb{P}(\hat{E}_r)\to 1$ as $r\to\infty$. The deviation terms are given by $\eta_r=\min\big(1,\sup_{z\geq r}T^{\text{fluid}}\delta_{\lfloor zT^{\text{fluid}}\rfloor}\big)$, and $\eta_r\to 0$ as $r\to\infty$.*

Proof. By a simple union bound, and then using the fact that the arrival process has stationary increments,

$$\begin{aligned}
\mathbb{P}(\hat{E}_r) &\geq 1-\sum_{m=0}^{\lfloor rT\rfloor}\sum_{k=0}^{\lfloor Lr\log r\rfloor}\mathbb{P}\Big(\sup_{u\in[0,T^{\text{fluid}}]}\big|\tilde{\mathbf{a}}^{r,m,w_{r,k}}(u)-\boldsymbol{\lambda}^r u\big|\geq\eta_r\Big)\\
&=1-(1+\lfloor rT\rfloor)\sum_{k=0}^{\lfloor Lr\log r\rfloor}\mathbb{P}\Big(\sup_{u\in[0,T^{\text{fluid}}]}\big|\tilde{\mathbf{a}}^{r,0,w_{r,k}}(u)-\boldsymbol{\lambda}^r u\big|\geq\eta_r\Big)\\
&=1-(1+rT)\sum_{k=0}^{\lfloor Lr\log r\rfloor}\mathbb{P}\Big(\sup_{u\in[0,T^{\text{fluid}}]}\Big|\frac{A^r(w_{r,k}u)}{w_{r,k}}-\boldsymbol{\lambda}^r u\Big|\geq\eta_r\Big).
\end{aligned}$$

To bound this we will use Assumption 2.5, which says that

$$z(\log z)^2\mathbb{P}\Big(\sup_{\tau\leq z}\frac{1}{z}\big|\mathbf{A}^r(\tau)-\boldsymbol{\lambda}^r\tau\big|\geq\delta_z\Big)\to 0\quad\text{as } z\to\infty$$

uniformly in r. After extending the domain of $\mathbf{A}^r$ to $\mathbb{R}_+$ by linear interpolation in each interval $(\tau,\tau+1)$, and extending the domain of δ_z to $z\in\mathbb{R}_+$ by $\delta(z)=\delta_{\lfloor z\rfloor}\vee\delta_{\lceil z\rceil}$, and rescaling z by T^{fluid},

$$z(\log z)^2\mathbb{P}\Big(\sup_{u\in[0,T^{\text{fluid}}]}\Big|\frac{\mathbf{A}^r(zu)}{z}-\boldsymbol{\lambda}^r u\Big|\geq T^{\text{fluid}}\delta(zT^{\text{fluid}})\Big)\to 0\quad\text{as } z\to\infty$$

uniformly in r. In other words, for any $\phi > 0$ there exists z_0 such that for all $z \ge z_0$ and all r,

$$\mathbb{P}\left(\sup_{u\in[0,T^{\text{fluid}}]} \left| \frac{\mathbf{A}^r(zu)}{z} - \boldsymbol{\lambda}^r u \right| \ge T^{\text{fluid}} \delta(z T^{\text{fluid}}) \right) < \frac{\phi}{z(\log z)^2}.$$

Now pick r_0 sufficiently large that $r_0 \ge z_0$ and $\sup_{z\ge r_0} T^{\text{fluid}}\delta(zT^{\text{fluid}}) < 1$, which we can do since $\delta(z) \to 0$ as $z \to \infty$ by Assumption 2.5. This choice implies that for any $r \ge r_0$ and $z \ge r$, $T^{\text{fluid}}\delta(zT^{\text{fluid}}) \le \eta_r$ (recall that $\eta_r = \min\big(1, \sup_{z\ge r} T^{\text{fluid}}\delta_{\lfloor z T^{\text{fluid}}\rfloor}\big)$). Hence, for any $r \ge r_0$ and $z \ge r$,

$$\mathbb{P}\left(\sup_{u\in[0,T^{\text{fluid}}]} \left| \frac{\mathbf{A}^r(zu)}{z} - \boldsymbol{\lambda}^r u \right| \ge \eta_r \right) < \frac{\phi}{z(\log z)^2}.$$

Applying this bound to $\mathbb{P}(\hat{E}_r)$, and using the facts that $w_{r,k} \ge r$ and $w_{r,k}(\log w_{r,k})^2 \ge r(1 + k/\log r)(\log r)^2$,

$$\begin{aligned}
1 - \mathbb{P}(\hat{E}_r) &< \phi\Big(\frac{1+\lfloor rT\rfloor}{r(\log r)^2}\Big) \sum_{k=0}^{\lfloor Lr\log r\rfloor} \frac{1}{1+k/\log r} = \phi\Big(\frac{1+\lfloor rT\rfloor}{r\log r}\Big) \sum_{k=0}^{\lfloor Lr\log r\rfloor} \frac{1}{k+\log r} \\
&\le \phi\Big(\frac{1+\lfloor rT\rfloor}{r\log r}\Big) \int_{\ell=0}^{\lfloor Lr\log r\rfloor} \frac{1}{\ell - 1 + \log r}\, d\ell \\
&= \phi\Big(\frac{1+\lfloor rT\rfloor}{r\log r}\Big) \log\Big(1 + \frac{\lfloor Lr\log r\rfloor}{\log r - 1}\Big).
\end{aligned}$$

The final expression converges to ϕT as $r \to \infty$. Since ϕ can be chosen arbitrarily small, $\mathbb{P}(\hat{E}_r) \to 1$ as $r \to \infty$. □

7.1.2. *Proof of Theorem 7.1.* Given $\delta > 0$, pick $\zeta > 0$ such that $2\zeta + \mathrm{mc}_\zeta(\Delta W) < \delta \wedge 1$ where $\mathrm{mc}_\zeta(\Delta W)$ is the modulus of continuity of ΔW over the set $\mathcal{D}$ specified in Lemma 7.5. We can achieve the desired bound by making ζ sufficiently small; this is because ΔW is continuous, hence uniformly continuous on compact sets, and $\mathcal{D}$ is compact as a consequence of Theorem 5.4(i), hence $\mathrm{mc}_\zeta(\Delta W) \to 0$ as $\zeta \to 0$. With this choice of ζ, define the good sets $\hat{E}_r$ and the constants T^{fluid} and H_ζ as specified by (47).

By Lemma 7.3, there exists r_0 such that for $r \ge r_0$ and for all $\omega \in \hat{E}_r$ and all m, $d(\tilde{x}^{r,m}, \mathrm{FMS}_1) < \zeta$ and $d(\tilde{x}^{r,0}, \mathrm{FMS}(\mathbf{q}_0)) < \zeta$, where $\mathbf{q}_0$ is defined in the statement of the lemma.

Now, pick any $t \in [0, T^{\text{fluid}}]$ and $r \ge r_0$, and assume $\hat{E}_r$ holds. Lemma 7.4 says that we can choose $m \in \{0, \dots, \lfloor rT\rfloor\}$ and $u \in [0, T^{\text{fluid}}]$ such that $r^2 t = rm + z_{r,m}u$, and furthermore either (i) $m > 0$ and $u > H_\zeta$ or (ii) $m = 0$. By Lemma 7.3, we can pick $x \in \mathrm{FMS}$ (depending on r, t and the ω) such that $\|\tilde{x}^{r,m} - x\| < \zeta$ and furthermore either (i) $m > 0$ and $x \in \mathrm{FMS}_1$ or (ii) $m = 0$ and $x \in \mathrm{FMS}(\mathbf{q}_0)$. Then, by Lemma 7.5,

$$\frac{\left|\hat{\mathbf{q}}^r(t) - \Delta W(\hat{\mathbf{q}}^r(t))\right|}{\left\|\hat{\mathbf{q}}^r(\cdot)\right\| \vee 1} < \delta.$$

This bound holds for every $t \in [0, T]$ and $r \geq r_0$, in a sample path-wise sense, whenever $\omega \in \hat{E}_r$.

Finally, Lemma 7.6 says that $\mathbb{P}(\hat{E}_r) \to 1$. This completes the proof. □

8. An optimal policy? Our motivation for this work was Conjecture 3.1, which says that for an input-queued switch the performance of MW-α improves as $\alpha \searrow 0$. We have not been able to prove this. However, under a condition on the arrival rate, we can show (i) that the critically-loaded fluid model solutions for a single-hop switched network approach optimal (in the sense of minimizing total amount of work in the network) as $\alpha \searrow 0$; and (ii) that for an input-queued switch the set of invariant states $\mathcal{I}$ defined in Section 7 is sensitive to α. We speculate that these findings might eventually form part of a proof of a heavy traffic limit theorem supporting Conjecture 3.1, given that critically loaded fluid models and invariant states play an important role in heavy traffic theorems.

In this section we state the condition on the arrival rates, and give the results (i) and (ii). Motivated by these results, we make a conjecture about an optimal scheduling policy.

Definition 8.1 (Complete loading) *Consider a switched network with arrival rate vector $\boldsymbol{\lambda}$. Say that $\boldsymbol{\lambda}$ satisfies the* complete loading condition *if $\boldsymbol{\lambda} \in \Lambda$ and there is a convex combination of critically loaded virtual resources that gives equal weight to each queue; in other words if*

$$\frac{\mathbf{1}}{\max_{\boldsymbol{\pi} \in \mathcal{S}} \mathbf{1} \cdot \boldsymbol{\pi}} \in \left\langle \Xi(\boldsymbol{\lambda}) \right\rangle.$$

Theorem 8.2 (Near-optimality of fluid models under complete loading) *Consider a single-hop switched network with arrival rate vector $\boldsymbol{\lambda} \in \Lambda$.*

- *(i) For any fluid model solution for the MW-α policy, $\mathbf{1} \cdot \mathbf{q}(t) \leq N^{\alpha/(1+\alpha)} \mathbf{1} \cdot \mathbf{q}(0)$*
- *(ii) For any fluid model solution for any scheduling policy, if $\boldsymbol{\lambda}$ satisfies the complete loading condition then $\mathbf{1} \cdot \mathbf{q}(t) \geq \mathbf{1} \cdot \mathbf{q}(0)$*

Proof. The first claim relies on the standard result that for any $\mathbf{x} \in \mathbb{R}_+^N$ and $\beta > 1$,

$$\frac{1}{N^{1-1/\beta}} \sum_n x_n \leq \Big(\sum_n x_n^\beta\Big)^{1/\beta} \leq \sum_n x_n. \tag{53}$$

Using the Lyapunov function from Definition 5.3,

$$\begin{aligned}\mathbf{1}\cdot\mathbf{q}(t) &\leq N^{1-1/(1+\alpha)}\Big(\sum_n q_n(t)^{1+\alpha}\Big)^{1/(1+\alpha)} \quad \text{by the first inequality in (53)}\\ &= N^{\alpha/(1+\alpha)}L\big(\mathbf{q}(t)\big)^{1/(1+\alpha)} \quad \text{by definition of } L(\cdot)\\ &\leq N^{\alpha/(1+\alpha)}L\big(\mathbf{q}(0)\big)^{1/(1+\alpha)} \quad \text{since } \dot{L}(\mathbf{q}(t))\leq 0 \text{ by Theorem 5.4(i)}\\ &\leq N^{\alpha/(1+\alpha)}\mathbf{1}\cdot\mathbf{q}(0) \quad \text{by the second inequality in (53).}\end{aligned}$$

The second claim is a simple consequence of Lemma 5.10. (This lemma is for a single-hop network. The multihop version, Lemma 6.8, does not have such a simple interpretation.) □

Theorem 8.3 ($\mathcal{I}$ is sensitive to α for an input-queued switch) *Consider an $M \times M$ input-queued switch running MW-α, as introduced in Section 2.4. Let λ_{ij} be the arrival rate at the queue at input port i of packets destined for output port j, $\boldsymbol{\lambda} \in \mathbb{R}_+^{M\times M}$. Suppose that $\boldsymbol{\lambda} > \mathbf{0}$ componentwise, and furthermore that every input port and every output port is critically loaded, i.e.*

$$\sum_{j=1}^{M}\lambda_{\hat{\imath}j} = 1 \text{ and } \sum_{i=1}^{M}\lambda_{i\hat{\jmath}} = 1 \quad \text{for every } 1 \leq \hat{\imath},\hat{\jmath} \leq M. \tag{54}$$

Then $\boldsymbol{\lambda}$ satisfies the complete loading condition, and the critically loaded virtual resources are

$$\Xi(\boldsymbol{\lambda}) = \big\{\mathbf{r}_{\hat{\imath}} \text{ for all } 1 \leq \hat{\imath} \leq M\big\} \cup \big\{\mathbf{c}_{\hat{\jmath}} \text{ for all } 1 \leq \hat{\jmath} \leq M\big\}$$

where $\mathbf{r}_{\hat{\imath}}$ and $\mathbf{c}_{\hat{\jmath}}$ are the row and column indicator matrices, $(\mathbf{r}_{\hat{\imath}})_{i,j} = 1_{i=\hat{\imath}}$ and $(\mathbf{c}_{\hat{\jmath}})_{i,j} = 1_{j=\hat{\jmath}}$. Define the workload map $W : \mathbb{R}_+^{M\times M} \to \mathbb{R}_+^{2M}$ by $W(\mathbf{q}) = [\boldsymbol{\xi}\cdot\mathbf{q}]_{\boldsymbol{\xi}\in\Xi(\boldsymbol{\lambda})}$. Denoting the invariant set by $\mathcal{I}(\alpha)$,

(i) If w is in the relative interior of $\{W(\mathbf{q}) : \mathbf{q} \in \mathbb{R}_+^{M\times M}\}$ then w is in $\{W(\mathbf{q}) : \mathbf{q} \in \mathcal{I}(\alpha)\}$ for sufficiently small $\alpha > 0$

(ii) For a 2×2 input-queued switch, $\{W(\mathbf{q}) : \mathbf{q} \in \mathcal{I}(\alpha)\}$ is strictly increasing as $\alpha \searrow 0$

Item (i) essentially says that $W(\mathcal{I}(\alpha))$ becomes as large as possible as $\alpha \searrow 0$, except for some possible issues at the boundary. We have only been able to prove (ii) for a 2×2 switch, but we believe it holds for any $M \times M$ switch. The proofs are rather long, and depend on the specific structure of the input-queued switch, so they are left to the appendix.

Conjecture 3.1 claims that, for an input-queued switch, performance improves as $\alpha \searrow 0$. Examples due to Ji, Athanasopoulou, and Srikant [15] and Stolyar (personal communication) show that this is not true for general switched networks. However, Theorem 8.2 suggests that the conjecture might apply not just to input-queued switches but also to generalized switches under the complete loading condition; the examples of Ji et al. and Stolyar do not satisfy this condition. We therefore extend Conjecture 3.1 as follows.

Conjecture 8.4 *Consider a general single-hop switched network as described in Section 2, running MW-α. Consider the diffusion scaling limit (described in Conjecture 7.2) and let $\boldsymbol{\lambda}$ be the limiting arrival rates; assume $\boldsymbol{\lambda}$ satisfies the complete loading condition. For every $\alpha > 0$ there is a limiting stationary queue size distribution. The expected value of the sum of queue sizes under this distribution is non-increasing as $\alpha \searrow 0$.*

Theorem 8.2 says that MW-α approaches optimal as $\alpha \searrow 0$, under the complete loading condition. It is natural to ask if there is a policy that is optimal, rather than just a sequence of policies that approach optimal. Given that MW-α chooses a schedule $\boldsymbol{\pi}$ to maximize $\boldsymbol{\pi} \cdot \mathbf{q}^\alpha$ (where the exponent is taken componentwise), and since

$$x^\alpha = \begin{cases} 1 + \alpha \log x + O(\alpha^2) & \text{if } x > 0 \\ 0 & \text{if } x = 0 \end{cases}$$

we propose the following formal limit policy, which we call MSMW-log: *At each timestep, look at all maximum-size schedules, i.e. those $\boldsymbol{\pi} \in \mathcal{S}$ for which $\sum_n \pi_n 1_{Q_n > 0}$ is maximal. Among these, pick one which has maximal log-weight, i.e. for which $\sum_{n:Q_n>0} \pi_n \log Q_n$ is maximal, breaking ties randomly.*

Conjecture 8.5 *Consider a general single-hop switched network running MSMW-log. Consider the diffusion scaling limit and let $\boldsymbol{\lambda}$ be the limiting arrival rates; assume $\boldsymbol{\lambda}$ satisfies the complete loading condition. There is a limiting stationary queue size distribution. This distribution minimizes the expected value of the sum of the queue sizes, over all scheduling policies for which this expected value is defined.*

Scheduling policies based on MW are computationally difficult to implement, because there are so many comparisons to be made. In future work we plan to investigate whether the techniques described in this paper can be applied to policies that may have worse performance but simpler implementation.

Acknowledgements. We would like to thank Mike Harrison, Frank Kelly, Balaji Prabhakar and Ruth Williams for helpful discussions. We thank Ciamac Moallemi for suggesting Lemma 5.8. We also thank the anonymous reviewers for their help in improving readability of the manuscript.

References

[1] Andrews, M., Kumaran, K., Ramanan, K., Stolyar, S., Vijayakumar, R., and Whiting, P. (2001). Scheduling in a queueing system with asynchronously varying service rates. *Probability in the Engineering and Informational Sciences*.

[2] Azuma, K. (1967). Weighted sums of certain dependent random variables. *Tohoku Mathematics Journal 19*, 357–367.

[3] Bramson, M. (1998). State space collapse with application to heavy traffic limits for multiclass queueing networks. *Queueing Systems 30*, 89–148.

[4] Bramson, M. (2008). *Stability of Queueing Networks*. Springer Lecture Notes in Mathematics 1950, Ecole d'Ete de Probabilites de Saint-Flour XXXVI-2006.

[5] Dai, J. G. (1999). Stability of fluid and stochastic processing networks. MaPhySto Lecture Notes. http://www.maphysto.dk/cgi-bin/gp.cgi?publ=70.

[6] Dai, J. G. and Lin, W. (2005). Maximum pressure policies in stochastic processing networks. *Operations Research* **53**, 2.

[7] Dai, J. G. and Lin, W. (2008). Asymptotic optimality of maximum pressure policies in stochastic processing networks. *The Annals of Applied Probability* **18**, 6.

[8] Dai, J. G. and Prabhakar, B. (2000). The throughput of switches with and without speed-up. In *Proceedings of IEEE Infocom*. 556–564.

[9] Dembo, A. and Zeitouni, O. (2009). *Large deviations techniques and applications*. Springer Verlag.

[10] Doob, J. (1951). *Stochastic processes*. Wiley New York.

[11] Harrison, J. M. (1995). Balanced fluid models of multiclass queueing networks: A heavy traffic conjecture. *Sotchastic Networks 71*, 1–20. IMA Volumes in Mathematics and Its Applications.

[12] Harrison, J. M. (2000). Brownian models of open processing networks: canonical representation of workload. *The Annals of Applied Probability 10*, 75–103. Also see [13], http://projecteuclid.org/euclid.aoap/1019737665.

[13] Harrison, J. M. (2003). Correction to [12]. *The Annals of Applied Probability 13*, 390–393.

[14] Hoeffding, W. (1963). Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association* **58**, 301, 13–30.

[15] Ji, T., Athanasopoulou, E., and Srikant, R. (2009). Optimal scheduling policies in small generalized switches. In *Proceedings of IEEE Infocom*.

[16] Kang, W. N., Kelly, F. P., Lee, N. H., and Williams, R. J. (2009). State space collapse and diffusion approximation for a network operating under a fair bandwidth sharing policy. *The Annals of Applied Probability*. http://www.math.ucsd.edu/~williams/bandwidth/kklw.html.

[17] Kang, W. N. and Williams, R. J. (2006). An input-queued switch in heavy traffic. Presented at the Stochastic Networks Conference, Urbana-Champaign.

[18] 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.

[19] Kelly, F. P. and Williams, R. J. (2004). Fluid model for a network operating under a fair bandwidth-sharing policy. *The Annals of Applied Probability 14*, 1055–1083.

[20] Keslassy, I. and McKeown, N. (2001). Analysis of scheduling algorithms that provide 100% throughput in input-queued switches. In *Proceedings of Allerton Conference on Communication, Control and Computing*.

[21] Kingman, J. F. C. (1962). On queues in heavy traffic. *Journal of the Royal Statistical Society, series B* **24**, 2, 383–392.

[22] Laws, C. N. (1990). Dynamic routing in queueing networks. Ph.D. thesis, University of Cambridge.

[23] Laws, C. N. (1992). Resource pooling in queueing networks with dynamic routing. *Advances*

in Applied Probability 24, 699–726.

[24] McKeown, N., Anantharam, V., and Walrand, J. (1996). Achieving 100% throughput in an input-queued switch. In *Proceedings of IEEE Infocom*. 296–302.

[25] Shah, D. (2001). Stable algorithms for input queued switches. In *Proceedings of Allerton Conference on Communication, Control and Computing*.

[26] Shah, D. and Wischik, D. J. (2006). Optimal scheduling algorithms for input-queued switches. In *Proceedings of IEEE Infocom*.

[27] Stolyar, A. L. (2004). MaxWeight scheduling in a generalized switch: State space collapse and workload minimization in heavy traffic. *The Annals of Applied Probability* **14**, 1, 1–53.

[28] Tassiulas, L. and Ephremides, A. (1992). Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks. *IEEE Transactions on Automatic Control 37*, 1936–1948.

[29] Whitt, W. (2001). *Stochastic-Process Limits*. Springer.

[30] Wikipedia. Border gateway protocol. http://en.wikipedia.org/wiki/Border_Gateway_Protocol.

[31] Williams, R. J. (1998). Diffusion approximations for open multiclass queueing networks: sufficient conditions involving state space collapse. *Queueing Systems 30*, 27–88.

APPENDIX A: RESULTS FOR INPUT-QUEUED SWITCHES

In this section we prove Theorem 8.3. Throughout this appendix we are considering a $M \times M$ input-queued switch. The set of schedules $\mathcal{S}$ consists of all $M \times M$ permutation matrices. We assume the arrival rate matrix $\boldsymbol{\lambda}$ satisfies the complete loading condition (54) and that $\boldsymbol{\lambda} > \mathbf{0}$ componentwise. We let the scheduling algorithm be MW-α, and define $\mathcal{I}(\alpha)$ to be the set of invariant states.

A.1. Identifying $\boldsymbol{\Lambda}$, $\boldsymbol{\mathcal{S}^*}$, $\boldsymbol{\Xi}$, and $\boldsymbol{\Xi^+}$. The Birkhoff–von-Neumann decomposition result says that a matrix is doubly substochastic if and only if it is less than or equal to a convex combination of permutation matrices, which yields

$$\Lambda = \Big\{ \boldsymbol{\lambda} \in [0,1]^{M\times M} \ : \ \sum_{j=1}^{M} \lambda_{\hat{\imath},j} \leq 1 \text{ and } \sum_{i=1}^{M} \lambda_{i,\hat{\jmath}} \leq 1 \quad \text{for all } \hat{\imath}, \hat{\jmath} \Big\}.$$

Since $\boldsymbol{\lambda}$ satisfies the complete loading condition (54), $\boldsymbol{\lambda} \in \Lambda$.

Lemma A.1 below gives $\mathcal{S}^*$, the set of virtual resources i.e. maximal extreme points of the set of feasible solutions to DUAL($\boldsymbol{\lambda}$). From the complete loading condition, it is clear that $\Xi(\boldsymbol{\lambda}) = \mathcal{S}^*$ as claimed in the theorem. It will also be useful, for the proof of Theorem 8.3(i), to identify $\Xi^+(\boldsymbol{\lambda})$ as defined by (34). We claim that

$$\Xi^+(\boldsymbol{\lambda}) = \Xi(\boldsymbol{\lambda}). \tag{55}$$

To see this, suppose $\boldsymbol{\xi}$ is a non-maximal extreme point of the set of feasible solutions to DUAL($\boldsymbol{\lambda}$); then there exists some other extreme point $\boldsymbol{\zeta}$ such that $\boldsymbol{\xi} \leq \boldsymbol{\zeta}$ and $\boldsymbol{\xi} \neq \boldsymbol{\zeta}$; but because $\boldsymbol{\lambda} > \mathbf{0}$ componentwise it must be that $\boldsymbol{\xi}\cdot\boldsymbol{\lambda} < \boldsymbol{\zeta}\cdot\boldsymbol{\lambda}$. We have found that $\boldsymbol{\lambda} \in \Lambda$, so the solution to DUAL($\boldsymbol{\lambda}$) is ≤ 1, hence $\boldsymbol{\xi}\cdot\boldsymbol{\lambda} < 1$. Therefore $\boldsymbol{\xi} \notin \Xi^+(\boldsymbol{\lambda})$, i.e. $\Xi^+(\boldsymbol{\lambda})$ consists only of maximal extreme points.

Lemma A.1 *The set of maximal extreme points of the set*

$$F = \big\{ \boldsymbol{\xi} \in \mathbb{R}_+^{M\times M} \ : \ \boldsymbol{\xi}\cdot\boldsymbol{\pi} \leq 1 \text{ for all } \boldsymbol{\pi} \in \mathcal{S} \big\}.$$

is given by

$$\mathcal{S}^* = \big\{ \mathbf{r}_{\hat{\imath}} \text{ for all } 1 \leq \hat{\imath} \leq M \big\} \cup \big\{ \mathbf{c}_{\hat{\jmath}} \text{ for all } 1 \leq \hat{\jmath} \leq M \big\}$$

where the row and column indicator matrices $\mathbf{r}_{\hat{\imath}}$ and $\mathbf{c}_{\hat{\jmath}}$ are defined by $(\mathbf{r}_{\hat{\imath}})_{i,j} = 1_{i=\hat{\imath}}$ and $(\mathbf{c}_{\hat{\jmath}})_{i,j} = 1_{j=\hat{\jmath}}$.

Proof. First we argue that every $\boldsymbol{\xi} \in \mathcal{S}^*$ is a maximal extreme point of F. It is simple to check that $\boldsymbol{\xi} \in F$. Also, $\boldsymbol{\xi}$ is extreme because $F \subset [0,1]^{M\times M}$. Finally, $\boldsymbol{\xi}$ is maximal, because if it were not then there would be some $\boldsymbol{\varepsilon} \geq \mathbf{0}$, $\boldsymbol{\varepsilon} \neq \mathbf{0}$, such that $\boldsymbol{\xi}+\boldsymbol{\varepsilon}\in F$; but for any such $\boldsymbol{\varepsilon}$ there is a matching $\boldsymbol{\pi}$ such that $\boldsymbol{\varepsilon}\cdot\boldsymbol{\pi}>0$ hence $\boldsymbol{\xi}+\boldsymbol{\varepsilon}\notin F$.

Next we argue the converse, that all maximal extreme points of F are in $\mathcal{S}^*$. The first step is to characterize the extreme points of F. We claim that if $\boldsymbol{\zeta}\in F$ then it can be written $\boldsymbol{\zeta}\leq\boldsymbol{\xi}$ for some $\boldsymbol{\xi}\in\langle\mathcal{S}^*\rangle$. Consider the optimization problem

$$\begin{aligned} &\text{(56)} \qquad \text{minimize} \quad \sum_{\hat{\imath}=1}^{M} x_{\hat{\imath}}+\sum_{\hat{\jmath}=1}^{M} y_{\hat{\jmath}} \quad \text{over} \quad x_{\hat{\imath}}\geq 0,\ y_{\hat{\jmath}}\geq 0 \ \text{ for all } \hat{\imath},\hat{\jmath} \\ &\qquad\qquad \text{such that} \quad \boldsymbol{\zeta}\leq\sum_{\hat{\imath}} x_{\hat{\imath}}\mathbf{r}_{\hat{\imath}}+\sum_{\hat{\jmath}} y_{\hat{\jmath}}\mathbf{c}_{\hat{\jmath}}. \end{aligned}$$

The dual of this problem is

$$\begin{aligned} &\text{(57)} \qquad \text{maximize} \quad \mathbf{a}\cdot\boldsymbol{\zeta} \quad \text{over} \quad \mathbf{a}\in\mathbb{R}_+^{M\times M} \\ &\qquad\qquad \text{such that} \quad \mathbf{a}\cdot\mathbf{r}_{\hat{\imath}}\leq 1,\ \mathbf{a}\cdot\mathbf{c}_{\hat{\jmath}}\leq 1 \ \text{ for all } \hat{\imath},\hat{\jmath}. \end{aligned}$$

(These problems are just PRIMAL($\boldsymbol{\zeta}$) and DUAL($\boldsymbol{\zeta}$) respectively, but with respect to the 'virtual' schedule set $\mathcal{S}^*$ rather than the actual schedule set $\mathcal{S}$.) By Slater's condition, strong duality holds. Now, any matrix $\mathbf{a}$ that is feasible for (57) is non-negative and doubly substochastic, hence by the Birkhoff–von-Neumann decomposition result it can be written as a $\mathbf{a}\leq\mathbf{b}$ where $\mathbf{b}$ is a convex combination of permutation matrices, i.e. $\mathbf{b}\in\langle\mathcal{S}\rangle$. But by the assumption that $\boldsymbol{\zeta}\in F$, $\boldsymbol{\zeta}\cdot\boldsymbol{\pi}\leq 1$ for all $\boldsymbol{\pi}\in\mathcal{S}$, hence $\mathbf{b}\cdot\boldsymbol{\zeta}\leq 1$, hence $\mathbf{a}\cdot\boldsymbol{\zeta}\leq 1$, so the value of the optimization problem (57) is ≤ 1. By strong duality, the value of the optimization problem (56) is also ≤ 1. Therefore $\boldsymbol{\zeta}\leq\boldsymbol{\xi}$ for some $\boldsymbol{\xi}\in\langle\mathcal{S}\rangle$.

We now claim that if $\boldsymbol{\zeta}\in F$ then it can be written $\boldsymbol{\zeta}=\sum_{\boldsymbol{\xi}} a_{\boldsymbol{\xi}}\boldsymbol{\xi}$ where the sum is over some finite collection of values drawn from the set

$$E=\big\{\boldsymbol{\xi}\in\mathbb{R}_+^{M\times M} \ :\ \boldsymbol{\xi}\leq\mathbf{r}_{\hat{\imath}} \text{ or } \boldsymbol{\xi}\leq\mathbf{c}_{\hat{\jmath}} \text{ for some } \hat{\imath},\,\hat{\jmath}\big\}$$

and all $a_{\boldsymbol{\xi}}$ are ≥ 0. We have just shown that

$$\boldsymbol{\zeta}=\sum_{\boldsymbol{\xi}} a_{\boldsymbol{\xi}}\boldsymbol{\xi}-\mathbf{z} \quad \text{for some } \mathbf{z}\geq\mathbf{0} \text{ and } a_{\boldsymbol{\xi}}\geq 0,\ \boldsymbol{\xi}\in E. \tag{58}$$

If $\mathbf{z}=\mathbf{0}$ we are done. Otherwise, pick some (k,l) such that $z_{k,l}>0$ and define $\mathbf{n}^{k,l}$ by $(\mathbf{n}^{k,l})_{i,j}=1-1_{i=k \text{ and } j=l}$. Noting that $\sum a_{\boldsymbol{\xi}}\xi_{k,l}\geq z_{k,l}>0$, we can rewrite $\boldsymbol{\zeta}$ as

$$\boldsymbol{\zeta}=\frac{z_{k,l}}{\sum a_{\boldsymbol{\xi}}\xi_{k,l}}\sum a_{\boldsymbol{\xi}}\boldsymbol{\xi}+\big(1-\frac{z_{k,l}}{\sum a_{\boldsymbol{\xi}}\xi_{k,l}}\big)\sum a_{\boldsymbol{\xi}}\boldsymbol{\xi}\mathbf{n}^{k,l}-\mathbf{z}\mathbf{n}^{k,l}$$

where matrix multiplication is componentwise as per the notation specified in Section 1. We have thus rewritten $\boldsymbol{\zeta}$ in the form (58), but now $\mathbf{z}$ has one fewer non-zero element. Continuing in this way we can remove all non-zero elements of $\mathbf{z}$, until we are left with $\boldsymbol{\zeta} \in \langle E \rangle$.

We have therefore shown that all extreme points of F are in E. Clearly, all maximal points of E are in $\mathcal{S}^*$. Therefore, all the maximal extreme points of F are in $\mathcal{S}^*$ as claimed. $\square$

A.2. Proof of Theorem 8.3(i). We first state two lemmas which will be needed in the proof. The first is a general closure property of permutation matrices, and the second is a property of invariant states of MW-α. We then prove the theorem and the two lemmas.

Lemma A.2 *Let $\mathbf{x} \in \mathbb{R}_+^{M\times M}$, and define $\mathbf{A} \in \mathbb{R}_+^{M\times M}$ by $A_{i,j} = 1$ if there is some matching $\boldsymbol{\rho} \in \mathcal{S}$ whose weight $\boldsymbol{\rho}\cdot\mathbf{x}$ is maximal (i.e. $\boldsymbol{\rho}\cdot\mathbf{x} = \max_{\boldsymbol{\sigma}\in\mathcal{S}} \boldsymbol{\sigma}\cdot\mathbf{x}$) and for which $\rho_{i,j} = 1$; and $A_{i,j} = 0$ otherwise. Then, for any matching $\boldsymbol{\pi}$ such that*

$$\pi_{i,j} = 1 \implies A_{i,j} = 1,$$

$\boldsymbol{\pi}$ is itself a maximum weight matching.

Lemma A.3 *Fix any $\boldsymbol{\lambda} \in \Lambda$, and any $\mathbf{q} \in \mathcal{I}(\alpha)$. For every $1 \le i,j \le M$ such that $\lambda_{i,j} > 0$, there exists a matching $\boldsymbol{\pi} \in \mathcal{S}$ whose weight $\boldsymbol{\pi}\cdot\mathbf{q}^\alpha$ is maximal (i.e. $\boldsymbol{\pi}\cdot\mathbf{q}^\alpha = \max_{\boldsymbol{\sigma}\in\mathcal{S}} \boldsymbol{\sigma}\cdot\mathbf{q}^\alpha$, where the exponent is taken componentwise) and for which $\pi_{i,j} = 1$.*

Proof of Theorem 8.3(i). Suppose the claim of the theorem is not true, i.e. that there exists a sequence $\alpha \searrow 0$ such that $w \notin \mathcal{W}(\alpha) = \{W(\mathbf{q}) : \mathbf{q} \in \mathcal{I}(\alpha)\}$ for each α in the sequence.

Write $w = (w_{1\cdot}, \dots, w_{M\cdot},\ w_{\cdot 1}, \dots, w_{\cdot M})$, and define the function $\Delta^\alpha : \mathbb{R}_+^{2M} \to \mathbb{R}_+^{M\times M}$ to give the (unique) solution to the optimization problem

$$\begin{aligned} &\text{minimize} && \frac{1}{1+\alpha}\sum_{i,j} q_{i,j}^{1+\alpha} \quad \text{over} \quad \mathbf{q} \in \mathbb{R}_+^{M\times M} \\ &\text{such that} && \mathbf{r}_{\hat{\imath}}\cdot\mathbf{q} \ge w_{\hat{\imath}\cdot} \ \text{ and } \ \mathbf{c}_{\hat{\jmath}}\cdot\mathbf{q} \ge w_{\cdot\hat{\jmath}} \quad \text{for all } 1 \le \hat{\imath},\hat{\jmath} \le M. \end{aligned}$$

This is the optimization problem defined in (35); we have simply written out $\Xi^+(\boldsymbol{\lambda})$ explicitly using (55). By Lemma 5.6, the map $\Delta W : \mathbb{R}_+^{M\times M} \to \mathbb{R}_+^{M\times M}$ that defines $\mathcal{I}(\alpha)$ is simply the composition of $W : \mathbb{R}_+^{M\times M} \to \mathbb{R}_+^{2M}$ and $\Delta^\alpha : \mathbb{R}_+^{2M} \to \mathbb{R}_+^{M\times M}$. Let $\mathbf{q}(\alpha) = \Delta^\alpha(w)$. Note that $\mathbf{q}(\alpha) \in \mathcal{I}(\alpha)$; this is because $\mathbf{q}(\alpha)$ is optimal for $\Delta^\alpha(w)$, therefore it is optimal for $\Delta^\alpha(W(\mathbf{q}(\alpha)))$ which has a smaller feasible region, therefore $\mathbf{q}(\alpha) = \Delta^\alpha(W(\mathbf{q}(\alpha))) = \Delta W(\mathbf{q}(\alpha))$ i.e. $\mathbf{q}(\alpha) \in \mathcal{I}(\alpha)$.

We next establish this claim: that *for each α in the sequence, there exists i, j, i' and j' such that $q_{i,j}(\alpha) = 0$, $q_{i',j}(\alpha) \ge w_{\cdot j}/M$ and $q_{i,j'}(\alpha) \ge w_{i\cdot}/M$.* To prove this claim, observe that $W(\mathbf{q}(\alpha)) \ge w$ by the constraints of the optimization problem Δ^α;

and that $\mathbf{q}(\alpha) \in \mathcal{I}(\alpha)$ hence $W(\mathbf{q}(\alpha)) \in \mathcal{W}(\alpha)$; hence $W(\mathbf{q}(\alpha)) \neq w$ by assumption that $w \notin \mathcal{W}(\alpha)$. Therefore $W(\mathbf{q}(\alpha)) > w$ in some component, i.e. there is some i or j such that $\mathbf{r}_i \cdot \mathbf{q}(\alpha) > w_{i\cdot}$ or $\mathbf{c}_j \cdot \mathbf{q}(\alpha) > w_{\cdot j}$. Indeed, there must be both such an i and j, since otherwise the sum of row workloads and column workloads would not be equal. Therefore $q_{i,j}(\alpha) = 0$, since if $q_{i,j}(\alpha) > 0$ then we could reduce $q_{i,j}(\alpha)$ and still have a feasible solution to the problem that defines $\Delta^\alpha(w)$ but with a smaller value of the objective function, which contradicts optimality of $\mathbf{q}(\alpha)$. There must also be a j' such that $q_{i,j'}(\alpha) \geq w_{i\cdot}/M$ since otherwise the workload constraint $\mathbf{r}_i \cdot \mathbf{q}(\alpha) \geq w_{i\cdot}$ would not be met. Likewise for i'.

We assumed that $w \in \mathbb{R}_+^{2M}$ is in the relative interior of $\mathcal{W}^{\max} = \{W(\mathbf{q}) : \mathbf{q} \in \mathbb{R}_+^{M \times M}\}$. This set is clearly convex, and from the characterization of relative interior for convex sets, for all $x \in \mathcal{W}^{\max}$ there exists $y \in \mathcal{W}^{\max}$ and $0 < a < 1$ such that $w = ax + (1-a)y$. In particular, by choosing $x = W(\mathbf{1})$, we find that $w > 0$ componentwise.

For each α in the sequence we can find indices $(i(\alpha), j(\alpha), i'(\alpha), j'(\alpha))$ as above. Some set of indices (i, j, i', j') must be repeated infinitely often, since there are only finitely many choices. Restrict attention to the subsequence of α for which $i(\alpha) = i$, $j(\alpha) = j$, $i'(\alpha) = i'$, $j'(\alpha) = j'$.

Now, consider the submatrix

$$\begin{pmatrix} q_{i,j}(\alpha) & q_{i,j'}(\alpha) \\ q_{i',j}(\alpha) & q_{i',j'}(\alpha) \end{pmatrix}$$

Recall that $\mathbf{q}(\alpha) \in \mathcal{I}(\alpha)$. By Lemma A.3, and the assumption that $\boldsymbol{\lambda} > \mathbf{0}$ componentwise, every queue is involved in some maximum weight matching. By Lemma A.2, every matching is a maximum weight matching. Let $\boldsymbol{\pi}$ be any matching with $\pi_{i,j} = \pi_{i',j'} = 1$, and let $\boldsymbol{\rho}$ be like $\boldsymbol{\pi}$ except with (i,j) and (i',j') flipped, i.e. $\rho_{i,j} = \rho_{i',j'} = 0$ and $\rho_{i,j'} = \rho_{i',j} = 1$. We can write out explicitly the difference in weight between these two matchings:

$$\boldsymbol{\rho} \cdot \mathbf{q}(\alpha)^\alpha - \boldsymbol{\pi} \cdot \mathbf{q}(\alpha)^\alpha = q_{i',j}(\alpha)^\alpha + q_{i,j'}(\alpha)^\alpha - q_{i',j'}(\alpha)^\alpha - q_{i,j}(\alpha)^\alpha.$$

Here, $\mathbf{q}(\alpha)^\alpha$ denotes component-wise exponentiation. Recall that along the subsequence we have chosen, $q_{i,j}(\alpha) = 0$, $q_{i,j'}(\alpha) \geq w_{i\cdot}/M$ and $q_{i',j}(\alpha) \geq w_{\cdot j}/M$. Therefore $\liminf_{\alpha \to 0} \boldsymbol{\rho} \cdot \mathbf{q}(\alpha)^\alpha - \boldsymbol{\pi} \cdot \mathbf{q}(\alpha)^\alpha > 0$. This contradicts the finding that every matching is a maximum weight matching for $\mathbf{q}(\alpha)$.

Thus, we have contradicted our original assumption that there exists a sequence $\alpha \searrow 0$ with $w \notin \mathcal{W}(\alpha)$. This completes the proof. □

Proof of Lemma A.2. Let $a = \max_{\boldsymbol{\rho}} \boldsymbol{\rho} \cdot \mathbf{x}$ be the weight of the maximum weight matching, let $\mathcal{A} = \{\boldsymbol{\rho} : \boldsymbol{\rho} \cdot \mathbf{x} = a\}$, and let $\boldsymbol{\sigma} = \sum_{\boldsymbol{\rho} \in \mathcal{A}} \boldsymbol{\rho}$. Observe that $\boldsymbol{\sigma} \in \mathbb{Z}_+^{M \times M}$ and $A_{i,j} = 1_{\{\sigma_{i,j} > 0\}}$. Therefore, $\boldsymbol{\sigma} \geq \mathbf{A}$ componentwise. Further, by definition $\mathbf{A} \geq \boldsymbol{\pi}$.

Therefore, the matrix $\boldsymbol{\sigma} - \boldsymbol{\pi}$ is non-negative. Since $\boldsymbol{\sigma}$ is sum of $|\mathcal{A}|$ permutation matrices, all its row sums and column sums are equal to $|\mathcal{A}|$. And since $\boldsymbol{\pi}$ is a permutation matrix as well, the matrix $\boldsymbol{\sigma} - \boldsymbol{\pi}$ has all its row sums and column sums equal to $|\mathcal{A}| - 1$; therefore by the Birkhoff–von-Neumann decomposition

$$\boldsymbol{\sigma} = \boldsymbol{\pi} + \sum_{\boldsymbol{\rho} \in \mathcal{S}} \alpha_{\boldsymbol{\rho}} \boldsymbol{\rho}$$

where each $\alpha_{\boldsymbol{\rho}} \geq 0$ and $\sum \alpha_{\boldsymbol{\rho}} = |\mathcal{A}| - 1$. Now, $\boldsymbol{\sigma} \cdot \mathbf{x} = |\mathcal{A}| a$ by construction of $\boldsymbol{\sigma}$. Therefore

$$\begin{aligned} |\mathcal{A}| a = \boldsymbol{\sigma} \cdot \mathbf{x} &= \boldsymbol{\pi} \cdot \mathbf{x} + \sum_{\boldsymbol{\rho} \in \mathcal{S}} \alpha_{\boldsymbol{\rho}} \boldsymbol{\rho} \cdot \mathbf{x} \\ &\leq \boldsymbol{\pi} \cdot \mathbf{x} + \big(|\mathcal{A}| - 1\big) a \quad \text{because } a = \max_{\boldsymbol{\rho} \in \mathcal{S}} \boldsymbol{\rho} \cdot \mathbf{x}. \end{aligned}$$

Rearranging, $\boldsymbol{\pi} \cdot \mathbf{x} \geq a$. But a is the weight of the maximum weight matching, thus $\boldsymbol{\pi} \cdot \mathbf{x} = a$. □

Proof of Lemma A.3. Since $\boldsymbol{\lambda} \in \Lambda$, $\boldsymbol{\lambda} \leq \boldsymbol{\sigma}$ for some $\boldsymbol{\sigma} \in \langle \mathcal{S} \rangle$ i.e. $\boldsymbol{\sigma} = \sum_{\boldsymbol{\pi} \in \mathcal{S}} a_{\boldsymbol{\pi}} \boldsymbol{\pi}$ where $\sum a_{\boldsymbol{\pi}} = 1$ and each $a_{\boldsymbol{\pi}} \geq 0$. Therefore

$$\boldsymbol{\lambda} \cdot \mathbf{q}^{\alpha} \leq \boldsymbol{\sigma} \cdot \mathbf{q}^{\alpha} = \sum a_{\boldsymbol{\pi}} \boldsymbol{\pi} \cdot \mathbf{q}^{\alpha} \leq \Big(\sum a_{\boldsymbol{\pi}} \Big) \max_{\boldsymbol{\pi} \in \mathcal{S}} \boldsymbol{\pi} \cdot \mathbf{q}^{\alpha} = \max_{\boldsymbol{\pi} \in \mathcal{S}} \boldsymbol{\pi} \cdot \mathbf{q}^{\alpha}.$$

But by Lemma 5.11(iv), $\boldsymbol{\lambda} \cdot \mathbf{q}^{\alpha} = \max_{\boldsymbol{\pi}} \boldsymbol{\pi} \cdot \mathbf{q}^{\alpha}$, therefore both the inequalities in the above must be equalities. In particular, all matchings $\boldsymbol{\pi}$ for which $a_{\boldsymbol{\pi}} > 0$ are maximum weight matchings. If $\lambda_{i,j} > 0$ then at least one of these matchings has $\boldsymbol{\pi}_{i,j} = 1$. □

A.3. Proof of Theorem 8.3(ii). Consider a 2×2 switch with arrival rate matrix $\boldsymbol{\lambda}$. Since $\boldsymbol{\lambda}$ satisfies (54), we may write it

$$\boldsymbol{\lambda} = \begin{pmatrix} \lambda_{1,1} & 1 - \lambda_{1,1} \\ 1 - \lambda_{1,1} & \lambda_{1,1} \end{pmatrix}$$

for some $\lambda_{1,1} \in (0,1)$. To find $\mathcal{I}(\alpha)$ we use the characterization from Lemma 5.11(iv), which says that $\mathbf{q} \in \mathcal{I}(\alpha)$ if and only if $\boldsymbol{\lambda} \cdot \mathbf{q}^{\alpha} = \max_{\boldsymbol{\pi}} \boldsymbol{\pi} \cdot \mathbf{q}^{\alpha}$, i.e. if and only if

$$\lambda_{1,1} \big(q_{1,1}^{\alpha} + q_{2,2}^{\alpha} \big) + (1 - \lambda_{1,1}) \big(q_{1,2}^{\alpha} + q_{2,1}^{\alpha} \big) = \big(q_{1,1}^{\alpha} + q_{2,2}^{\alpha} \big) \vee \big(q_{1,2}^{\alpha} + q_{2,1}^{\alpha} \big).$$

Now, the equation $\lambda_{1,1} x + (1 - \lambda_{1,1}) y = x \vee y$ is satisfied if and only if $x = y$, given $0 < \lambda_{1,1} < 1$. Therefore

$$\mathcal{I}(\alpha) = \Big\{ \mathbf{q} \in \mathbb{R}_{+}^{2 \times 2} : q_{1,1}^{\alpha} + q_{2,2}^{\alpha} = q_{1,2}^{\alpha} + q_{2,1}^{\alpha} \Big\}.$$

We wish show that $\{W(\mathbf{q}) : \mathbf{q} \in \mathcal{I}(\alpha)\}$ is strictly increasing as $\alpha \searrow 0$, where $W(\mathbf{q}) = \big(\mathbf{r}_1 \cdot \mathbf{q},\ \mathbf{r}_2 \cdot \mathbf{q},\ \mathbf{c}_1 \cdot \mathbf{q},\ \mathbf{c}_2 \cdot \mathbf{q}\big)$. It suffices to show that $\hat{\mathcal{W}}(\alpha) = \{\hat{W}(\mathbf{q}) : \mathbf{q} \in \mathcal{I}(\alpha)\}$ is strictly increasing, where $\hat{W}(\mathbf{q}) = \big(\mathbf{r}_1 \cdot \mathbf{q},\ \mathbf{c}_1 \cdot \mathbf{q},\ \mathbf{1} \cdot \mathbf{q}\big)$, since there is a straightforward bijection between $W(\mathbf{q})$ and $\hat{W}(\mathbf{q})$. Now, $(w_{1\cdot}, w_{\cdot 1}, w_{\cdot\cdot}) \in \mathbb{R}^3_+$ is in $\hat{\mathcal{W}}(\alpha)$ iff there exists $\mathbf{q} \in \mathbb{R}^{2\times 2}_+$ such that

$$q_{1,1}^\alpha + q_{2,2}^\alpha = q_{1,2}^\alpha + q_{2,1}^\alpha, \quad \mathbf{r}_1 \cdot \mathbf{q} = w_{1\cdot}, \quad \mathbf{c}_1 \cdot \mathbf{q} = w_{\cdot 1}, \quad \mathbf{1} \cdot \mathbf{q} = w_{\cdot\cdot}$$

i.e. iff there exists $x \in \mathbb{R}$ such that

$$x^\alpha + (w_{\cdot\cdot} - w_{1\cdot} - w_{\cdot 1} + x)^\alpha - (w_{1\cdot} - x)^\alpha - (w_{\cdot 1} - x)^\alpha = 0, \tag{59}$$

$$\max\big(0, w_{1\cdot} + w_{\cdot 1} - w_{\cdot\cdot}\big) \le x \le \min\big(w_{1\cdot}, w_{\cdot 1}\big). \tag{60}$$

Write $\theta(x)$ for the left hand side of (59). Since $\theta(x)$ is increasing in x, there exists a solution to (59) iff $\theta(x) \le 0$ at the lower bound in (60) and $\theta(x) \ge 0$ at the upper bound. By considering four separate cases of which of the bounds in (60) are tight, and after some algebra, we find that there exists a solution iff

$$w_{i\cdot} + w_{\cdot j} + \big(w_{i\cdot}^\alpha + w_{\cdot j}^\alpha\big)^{1/\alpha} \ge w_{\cdot\cdot} \quad \text{for each } i, j \in \{1, 2\} \tag{61}$$

where $w_{2\cdot} = w_{\cdot\cdot} - w_{1\cdot}$ and $w_{\cdot 2} = w_{\cdot\cdot} - w_{\cdot 1}$. Now, it is a standard inequality that for any $x > 0$ and $y > 0$, and any $0 < \alpha < \beta$,

$$(x^\alpha + y^\alpha)^{1/\alpha} > (x^\beta + y^\beta)^{1/\beta}.$$

Applying this inequality to (61), it follows that $\hat{\mathcal{W}}(\alpha)$ is strictly increasing as $\alpha \searrow 0$.