

On the Flow-level Dynamics of a Packet-switched Network

Ciamac C. Moallemi and Devavrat Shah *

Abstract: The packet is the fundamental unit of transportation in modern communication networks such as the Internet. Physical layer scheduling decisions are made at the level of packets, and packet-level models with exogenous arrival processes have long been employed to study network performance, as well as design scheduling policies that more efficiently utilize network resources. On the other hand, a user of the network is more concerned with end-to-end bandwidth, which is allocated through congestion control policies such as TCP. Utility-based flow-level models have played an important role in understanding congestion control protocols. In summary, these two classes of models have provided separate insights for flow-level and packet-level dynamics of a network.

In this paper, we wish to study these two dynamics together. We propose a joint flow-level and packet-level stochastic model for the dynamics of a network, and an associated policy for congestion control and packet scheduling that is based on α -weighted policies from the literature. We provide a fluid analysis for the model that establishes the throughput optimality of the proposed policy, thus validating prior insights based on separate packet-level and flow-level models. By analyzing a critically scaled fluid model under the proposed policy, we provide constant factor performance bounds on the delay performance and characterize the invariant states of the system.

Keywords and phrases: Flow-level model, Packet-level model, Congestion control, Scheduling, Utility maximization, Back-pressure maximum weight.

1. Introduction

The optimal control of a modern, packet-switched data network can be considered from two distinct vantage points. From the first point of view, the atomic unit of the network is the *packet*. In a packet-level model, the limited resources of a network are allocated via the decisions on the scheduling of packets. Scheduling policies for packet-based networks have been studied across a long line of literature (e.g., [30, 27, 24]). The insights from this literature have enabled the design of scheduling policies that allow for the efficient utilization of the resources of a network, in the sense of maximizing the throughput of packets across the network, while minimizing the delay incurred by packets, or, equivalently, the size of the buffers needed to queue packets in the network.

Packet-level models accurately describe the mechanics of a network at a low level. However, they model the arrival of new packets to the network exogenously. In reality, the arrival of new packets is also under the control of the network designer, via rate allocation or congestion control decisions. Moreover, while efficient utilization of network resources is a reasonable objective, a network designer may also be concerned with the satisfaction of end users of the network. Such objectives cannot directly be addressed in a packet-level model.

*C. C. Moallemi is at the Graduate School of Business, Columbia University. D. Shah is with Laboratory for Information and Decision Systems, Massachusetts Institute of Technology. This work was supported in parts by NSF projects CNS 0546590 and TF 0728554. Authors' email addresses: ciamac@gsb.columbia.edu, devavrat@mit.edu.

Flow-level models (cf. [12, 3]) provide a different point of view by considering the network at a higher level of abstraction or, alternatively, over a longer time horizon. In a flow-level model, the atomic unit of the network is a *flow*, or user, who wishes to transmit data from a source to a destination. Resource allocation decisions are made via the allocation of a transmission rate to each flow. Each flow generates utility as a function of its rate allocation, and rate allocation decisions may be made so as to maximize a global utility function. In this way, a network designer can address end user concerns such as fairness.

Flow-level models typically make two simplifying assumptions. The first assumption is that, as the number of flows evolves stochastically over time, the rates allocated to flows are updated *instantaneously*. The rate allocation decision for a particular flow is made in a manner that requires immediate knowledge of the demands of other flows for the limited transmission resources along the flow's entire path. This assumption, referred to as *time-scale separation*, is based on the idea that flows arrive and depart according to much slower processes than the mechanisms of the rate control algorithm. The second assumption is that, once the rate allocation decision is made, each flow can transmit data *instantaneously* across the network at its given rate. In reality, each flow generates discrete packets, and these packets must travel through queues to traverse the network. Moreover, the packet scheduling decisions within the network must be made in a manner that is consistent with and can sustain the transmission rates allocated to each flow, and the induced packet arrival process must not result in the inefficient allocation of low level network resources.

In this paper, our goal is to develop a stochastic model that jointly captures the packet-level and flow-level dynamics of a network, without any assumption of time-scale separation. The contributions of this paper are as follows:

1. We present a joint model where the dynamic evolution of flows and packets is simultaneous. In our model, it is possible to simultaneously seek efficient allocation of low level network resources (buffers) while maximizing the high-level metric of end-user utility.
2. For our network model, we propose packet scheduling and rate allocation policies where decisions are made via myopic algorithms that combine the distinct insights of prior packet- and flow-level models. Packets are scheduled according to a maximum weight policy. The rate allocation decisions are completely local and distributed. Further, in long term (i.e., under fluid scaling), the rate control policy exhibits the behavior of a primal algorithm for an appropriate utility maximization problem.
3. We provide a fluid analysis of the joint packet- and flow-level model. This analysis allows us to establish stability of the joint model and the throughput optimality of our proposed control policy.
4. We establish, using a fluid model under critical loading, a performance bound on our control policy under the metric of minimizing the outstanding number of packets and flows in the network (or, in other words, minimizing delay). We demonstrate that, for a class of *balanced* networks, our control policy performs to within a constant factor of any other control policy.
5. Under critical loading, we characterize the invariant manifold of the fluid model of our control policy, as well as establishing convergence to this manifold starting from any initial state. These results, along with the method of Bramson [4], lead to the characterization of multiplicative state space collapse under *heavy traffic scaling*. Further, we establish that the invariant states of the fluid model are asymptotically optimal under a limiting control policy.

In summary, our work provides a joint dynamic flow- and packet-level model that captures the microscopic (packet) and macroscopic (fluid, flow) behavior of large packet-based communications network faithfully. The performance analysis of our rate control and scheduling algorithm suggests that the separate insights obtained for dynamic flow-level models [12, 3] and for packet-level models [30, 27, 24] indeed continue to hold in the combined model.

The balance of the paper is organized as follows. In Section 1.1, we survey the related literature on flow- and packet-level models. In Section 2, we introduce our network model. Our network control policy, which combines features of maximum weight scheduling and utility-based rate allocation is described in Section 3. A fluid model is derived in Section 4. Stability (or, throughput optimality) of the network control policy is established in Section 5. The critically scaled fluid model is described in Section 6. In Section 7, we provide performance guarantees for balanced networks. The invariant states of the critically scaled fluid model are described in Section 8. Finally, in Section 9, we conclude.

1.1. Literature Review

Tassiulas and Ephremides [30] considered packet-level network model in the context of wireless radio network and proposed a class of ‘maximum weight’ (MW) or ‘back-pressure’ policies. Such policies assign a weight to every schedule, which is computed by summing the number of packets queued at links that the schedule will serve. At each instant of time, the schedule with the maximum weight will be selected. Tassiulas and Ephremides [30] establish that, in the context of multi-hop wireless networks, MW is throughput optimal. That is, the stability region of MW contains the stability region of any other scheduling algorithm. This work was subsequently extended to a much broader class of queueing networks by others (e.g., [18, 7, 29, 6]).

By allowing for a broader class of weight functions, the MW algorithm can be generalized to the family of so-called MW- α scheduling algorithms. These algorithms are parameterized by a scalar $\alpha \in (0, \infty)$. MW- α can be shown to inherit the throughput optimality of MW [14, 22] for all values of $\alpha \in (0, \infty)$. However, it has been observed experimentally that the average queue length (or, ‘delay’) under MW- α decreases as $\alpha \rightarrow 0^+$ [14]. Certain delay properties of this class of algorithms have been subsequently established under a heavy traffic scaling [27, 6, 24].

Flow-level models have received significant recent attention in the literature. The work of Kelly, Maulloo, and Tan [12] utilizes such model for Internet and suggests link between rate-control algorithm and decentralized solutions of a deterministic utility maximization problem. This optimization problem seeks to maximize the utility generated by a rate allocation, subject to capacity constraints that define a set of feasible rates. This work was subsequently generalized to settings where flows stochastically depart and arrive [17, 8, 3], addressing the question of the stability of the resulting control policies. Fluid and diffusion approximations of the resulting systems have been subsequently developed [13, 11, 32]. Under these stochastic models, flows are assumed to be allocated rate as per the optimal solution of the utility maximization problem instantaneously. Essentially, this *time-scale separation* assumption captures the intuition that the dynamics of the arrivals and departures of flows happens on a much slower time-scale than the dynamics of rate control algorithm.

In reality, flow arrivals/departures and rate control happen on the *same* time-scale. Various authors have considered this issue, in the context of understanding the stability of the stochastic

flow level models without the time-scale separation assumption [15, 10, 28, 20, 26]. Lin, Schroff, and Srikant [15] assume a stochastic model of flow arrivals and departures as well as the operation of a primal-dual algorithm for rate allocation. However, there are no packet dynamics present. Other work [10, 28, 20] has assumed that rate control for each type of flow is a function of a local Lagrange multiplier; and a separate Lagrange multiplier is associated with each link in the network. These multipliers are updated using a maximum weight-type policy. In this line of work, Lagrange multipliers are interpreted as queue lengths, but there are no actual packet-level dynamics present. Further, these models lack flow-level dynamics as well. Thus, while overall this collection of work is closest to the results of this paper, it stops short of offering a complete characterization of a joint flow- and packet-level dynamic model.

We take note of recent work by Walton [31], which presents a simple but insightful model for joint flow- and packet-level dynamics. In this model, each source generates packets by reacting to the acknowledgements from its destination, and at each time instant, each source has at most one packet in flight. Under a many-source scaling for a specific network topology, it is shown that the network operates with rate allocation as per the proportional fair criteria. This work provides important intuition about the relationship between utility maximization and the rate allocation resulting from the packet-level dynamics in a large network. However, it is far from providing a comprehensive joint flow- and packet-level dynamic model as well as efficient control policy.

Some remarks are in order regarding the relation between analytic methods of this paper and those existing in the literature. At the core of the present paper, we utilize fluid model techniques along with an appropriate Lyapunov function in order to establish positive-recurrence of a Markovian description of the system along with the existence and attractiveness of the invariant manifold for the critical fluid model. The latter is an important step towards establishing diffusion approximations. At some level, these approaches are well known for more than a decade now and have been developed for related systems in the literature — positive-recurrence in [15, 28] and critical fluid model analysis [24]. The non-triviality usually lies in the ability to execute this program in the system at hand. In our setting, this requires totally different sets of analytic arguments due to presence of two different dynamics. This, for example, can be seen in one of the key results (cf. Lemma 3) that establishes appropriate ‘drift’ condition for the model of this paper. Similarly, the optimality results we establish for the critical fluid are quite distinct from those known in the literature.

2. Network Model

In this section, we introduce our network model. This model captures both the flow-level and the packet-level aspects of a network, and will allow us to study the interplay between the dynamics at these two levels. In a nutshell, flows of various types arrive according to an exogenous process and seek to transmit some amount of data through the network. As in the standard congestion control algorithm, TCP, the flows generate packets at their ingress to the network. The packets travel to their respective destinations along links in the network, queueing in buffers at intermediate locations. As each packet travels along its route, it is subject to physical layer constraints, such as medium access constraints, switching constraints, or constraints due to limited link capacity. A flow departs once all of its packets have been sent.

In this section, we focus on the mechanics of the network that are independent of the network control policy. In Section 3, we will propose a specific network control policy to be applied in the context of this model.

2.1. Network Structure

Consider a network consisting of a set \mathcal{V} of destination nodes, a set \mathcal{L} of links, and a set \mathcal{F} of flow types. Each flow type is identified by a fixed given route starting at the source link $s(f) \in \mathcal{L}$ and ending at the destination node $d(f) \in \mathcal{V}$. At a given time, multiple flows of a given type exist in the network, each flow injects packets into the network.

The network maintains buffers for packets that are in transit across the network. At each link, there is a separate queue for the packets corresponding to each possible destination. Let $\mathcal{E} = \mathcal{L} \times \mathcal{V}$ denote the set of all such queues, with each $e = (\ell, v)$ being the queue at link ℓ for final destination v . Traffic in each queue is transmitted to the next hop along the route to the destination, and leaves the network when it reaches the destination. We define the routing matrix $R \in \{0, 1\}^{\mathcal{E} \times \mathcal{E}}$ by setting $R_{ee'} \triangleq 1$ if the next hop for queue e is queue e' , and $R_{ee'} \triangleq 0$ otherwise. Traffic for a flow of type f enters the network in the ingress queue $\iota(f) \triangleq (s(f), d(f)) \in \mathcal{E}$. Define the ingress matrix $\Gamma \in \{0, 1\}^{\mathcal{E} \times \mathcal{F}}$ by setting $\Gamma_{ef} \triangleq 1$ if $\iota(f) = e$, and $\Gamma_{ef} \triangleq 0$ otherwise. We will assume that the routes are acyclic. In this case, we can define the matrix

$$\Xi \triangleq (I - R^\top)^{-1} = I + R^\top + (R^\top)^2 + \dots \quad (1)$$

Under the acyclic routing assumption, $\Xi_{e'e} = 1$ if and only if a packet arriving at queue e will subsequently eventually pass through queue e' .

2.2. Dynamics: Flow-Level

In this section, we will describe in detail the stochastic model for dynamics of flows in the network. The system evolves in continuous time, with $t \in [0, \infty)$ denoting time, starting at $t = 0$. For each flow type $f \in \mathcal{F}$, let $N_f(t)$ denote the number of flows of type f active at time t . Flows of type f arrive according to an independent Poisson process of rate ν_f . Flows of type f receive an *aggregate* rate of service $X_f(t) \in [0, C]$ at time t . Here, $C > 0$ is the maximal rate of service¹ that can be provided to any flow type. The total rate of service $X_f(t)$ is divided equally among the $N_f(t)$ flows. As flows are serviced, packets are generated. The evolution of packets and flows proceeds according to the following:

- Packets are generated by all the flows of type f , in aggregate, as a time varying Poisson process of rate $X_f(t)$ at time t . If $N_f(t) = 0$, then we require that $X_f(t) = 0$.
- When a packet is generated by a flow of type f , it joins the ingress queue $\iota(f) \in \mathcal{E}$.
- When a packet is generated by a flow of type f , the flow departs from the network with a probability of $0 < \mu_f < 1$, independent of everything else.

¹Our assumption of a maximal service rate is a technical assumption necessary for the mathematical developments later on. This is not meant to model the capacity of links, for example.

Thus, each flow of type f generates a number of packets that is distributed according to an independent geometric random variable with mean² $1/\mu_f$, and the flow departure process for flows of type f is a Poisson process of rate $\mu_f X_f(t)$ at time t . We can summarize the flow count process $N_f(\cdot)$ by the transitions

$$N_f(t) \rightarrow \begin{cases} N_f(t) + 1 & \text{at rate } \nu_f, \\ N_f(t) - 1 & \text{at rate } \mu_f X_f(t). \end{cases}$$

Define the *offered load* vector $\rho \in \mathbb{R}_+^{\mathcal{F}}$ by $\rho_f \triangleq \nu_f/\mu_f$, for each flow type f . Without loss of generality, we will make the following assumptions:³

- $\rho > \mathbf{0}$, i.e., we restrict attention to flows with a non-trivial loading.
- $\rho < C\mathbf{1}$, i.e., we assume that the maximal service rate C is sufficient for the load generated by any single flow type.
- $\Xi\Gamma\rho > \mathbf{0}$, i.e., we restrict attention to queues with a non-trivial loadings.

Denote by $\mathcal{A}_f(t)$ the cumulative number of flows of type f that have arrived in the time interval $[0, t]$. Denote the cumulative number of packets generated by flows of type f in the time interval $[0, t]$ by $A_f(t)$. We suggest that the reader take note of difference between $\mathcal{A}_f(\cdot)$ and $A_f(\cdot)$. Let $D_f(t)$ denote the cumulative number of flows of type f that have departed in the time interval $[0, t]$. The evolution of the flow count for flow type f over time can be written as

$$N_f(t) = N_f(0) + \mathcal{A}_f(t) - D_f(t). \quad (2)$$

2.3. Dynamics: Packet-Level

As we have just described, flows generate packets which are injected into the network. These packets must traverse the links of the network from source to destination. In this section, we describe the dynamics of packets in the network.

We assume that each queue in the network is capable of transmitting at most 1 data packet per unit time. However, the collection of queues that can simultaneously transmit is restricted by a set of scheduling constraints. These scheduling constraints are meant to capture any limitations of the network due to scarce resources (e.g., limited wireless bandwidth, limited link capacity, switching constraints, etc.).

Formally, the scheduling constraints are described by the set $\mathcal{S} \subset \{0, 1\}^{\mathcal{E}}$. Under a permissible schedule $\pi \in \mathcal{S}$, a packet will be transmitted from a queue $e \in \mathcal{E}$ if and only if $\pi_e = 1$. We assume that $\mathbf{0} \in \mathcal{S}$. We require that each queue e be served by *some* schedule, i.e., there exists a $\pi \in \mathcal{S}$ with $\pi_e = 1$. Further, we assume that \mathcal{S} is *monotone*: if $\sigma \in \mathcal{S}$ and $\sigma' \in \{0, 1\}^{\mathcal{E}}$ is such that $\sigma'_e \leq \sigma_e$ for every queue e , then $\sigma' \in \mathcal{S}$. Finally, denote by $\Pi \in \{0, 1\}^{\mathcal{E} \times \mathcal{S}}$ the matrix with columns consisting of the elements of \mathcal{S} . The results of this paper can be extended, using the same arguments, to a the setting where $\mathcal{S} \subset \{0, 1, \dots, L\}^{\mathcal{E}}$, for some finite integer L , and \mathcal{S} is monotone. In this case, π_e would be the total number of packets scheduled per unit time from queue e , and we would allow

²The assumption that $\mu_f < 1$ is equivalent to requiring that $1/\mu_f > 1$, i.e., each flow is expected to generate more than one packet. This is reasonable since we require flows to arrive with at least one packet and for there to be some variability in the number of packets associated with a flow.

³In what follows, inequalities between vectors are to be interpreted component-wise. The vector $\mathbf{0}$ (resp., $\mathbf{1}$) is the vector where every component is 0 (resp., 1), and whose dimension should be inferred from the context.

scheduling of up to L packets. To keep the exposition simple, however, we shall work with the setting of $L = 1$ as described above.

We assume that the scheduling of packets happens at every integer time. At a time $\tau \in \mathbb{Z}_+$, let $\pi(\tau) \in \mathcal{S}$ denote the scheduled queues for the time interval $[\tau, \tau + 1)$. For each queue e , denote by $Q_e(\tau^-)$ the length of the queue e immediately prior to the time τ (i.e., before scheduling happens). The queue length evolves, for times $t \in [\tau, \tau + 1)$ according to⁴

$$Q_e(t) \triangleq Q_e(\tau^-) - \pi_e(\tau) \mathbb{I}_{\{Q_e(\tau^-) > 0\}} + \sum_{f \in \mathcal{F}} \Gamma_{ef} (A_f(t) - A_f(\tau^-)) + \sum_{e' \in \mathcal{E}} R_{e'e} \pi_{e'}(\tau) \mathbb{I}_{\{Q_{e'}(\tau^-) > 0\}}.$$

Here, for each flow type f , $A_f(\tau^-)$ is the cumulative number of packets generated by flows of type f in the time interval $[0, \tau)$. The term $\pi_e(\tau) \mathbb{I}_{\{Q_e(\tau^-) > 0\}}$ enforces an idling constraint, i.e., if queue e is scheduled but empty, no packet departs. Note that, over a time interval $[\tau, \tau + 1)$, we assume the transmission of packets already present in the network occurs instantly at time τ , while the arrival of new packets to the network occurs continuously throughout the entire time interval.

Finally, let $S_\pi(\tau)$ denote the cumulative number of time slots during which the schedule π was employed up to and including time τ . Let $Z_e(\tau)$ denote the cumulative idling time for queue e up to and including time τ . That is,

$$Z_e(\tau) \triangleq \sum_{s=0}^{\tau} \sum_{\pi \in \mathcal{S}} \pi_e(S_\pi(s) - S_\pi(s-1)) \mathbb{I}_{\{Q_e(s-1)=0\}}.$$

Then, the overall queue length evolution can be written in vector form as

$$Q(\tau + 1) = Q(0) - (I - R^\top) \Pi S(\tau) + (I - R^\top) Z(\tau) + \Gamma A(\tau + 1), \quad (3)$$

where we define the vectors

$$\begin{aligned} Q(t) &\triangleq [Q_e(t)]_{e \in \mathcal{E}}, & A(t) &\triangleq [A_f(t)]_{f \in \mathcal{F}}, \\ S(\tau) &\triangleq [S_\pi(\tau)]_{\pi \in \mathcal{S}}, & Z(\tau) &\triangleq [Z_e(\tau)]_{e \in \mathcal{E}}. \end{aligned}$$

3. MWUM Control Policy

A network control policy is a rule that, at each point in time, provides two types of decisions: (a) the rate of service provided to each flow, and (b) the scheduling of packets subject to the physical constraints in the network. In Section 2, we described the stochastic evolution of flows and packets in the network, taking as given the network control policy. In this section, we describe a control policy called the *maximum weight utility maximization- α* (MWUM- α) policy. MWUM- α takes as a parameter a scalar $\alpha \in (0, \infty)$.

The MWUM- α policy is myopic and based only on local information. Specifically, a flow generates packets at rate that is based on the queue length at its ingress, and the scheduling of packets is decided as a function of the effected queue lengths.

At the flow-level, rate allocation decisions are made according to a per flow utility maximization problem. Each flow chooses a rate so as to myopically maximize its utility as a function of rate

⁴ $\mathbb{I}_{\{t\}}$ denotes the indicator function.

consumption, subject to a linear penalty (or, ‘price’) for consuming limited network resources. As in the case of α -fair rate allocation, the utility function is assumed to have a constant relative risk aversion of α . The price charged is a function of the number of packets queued at the ingress queue associated with the flow, raised to the α power.

At the packet-level, packets are scheduled according to a maximum weight- α scheduling algorithm. In particular, each queue is assigned a weight that is a function of number of packets in the queue and in downstream queues, each raised to the α power, and a schedule is picked which maximizes the total weight of all scheduled queues.

3.1. Control: Rate Allocation

The first control decision we shall consider is that of *rate allocation*, or, the determination of the aggregate rate of service $X_f(t)$, at time t , for flows of type f . We will assume our network is governed by a variant of an α -fair rate allocation policy. This is as follows:

Assume that each flow of type f is allocated a rate $Y_f(t) \geq 0$ at time t by maximizing a (per flow) utility function that depends on the allocated rate, subject to a linear penalty (or, cost) for consuming resources from the limited capacity of the network. In particular, we will assume a utility function given a rate allocation of $y \geq 0$ to an individual flow of type f of the form

$$V_f(y) \triangleq \frac{y^{1-\alpha}}{1-\alpha},$$

for some⁵ $\alpha \in (0, \infty) \setminus \{1\}$. This utility function is popularly known as α -fair in the congestion control literature [19], and has a constant relative risk aversion of α . The individual flow will be assigned capacity according to

$$Y_f(t) \in \operatorname{argmax}_{y \geq 0} V_f(y) - Q_{\iota(f)}^\alpha(t)y.$$

Here, $Q_{\iota(f)}^\alpha(t)$ represents a ‘price’ or ‘congestion signal’. Intuitively, a flow reacts to the congestion in the network (or lack thereof) through the length of the ingress or ‘first-hop’ queue $\iota(f)$. Then, if $N_f(t) > 0$, the aggregate rate $X_f(t)$ allocated to all flows of type f at time t is determined according to

$$X_f(t) = N_f(t)Y_f(t) = \operatorname{argmax}_{x \geq 0} \frac{x^{1-\alpha}N_f^\alpha(t)}{1-\alpha} - Q_{\iota(f)}^\alpha(t)x.$$

If $N_f(t) = 0$, we require that $X_f(t) = 0$. Further, we will constrain the overall rate allocated to flows of type f by the constant C . Thus, rate allocation is determined by the equation

$$X_f(t) = \begin{cases} \operatorname{argmax}_{x \in [0, C]} \frac{x^{1-\alpha}N_f^\alpha(t)}{1-\alpha} - Q_{\iota(f)}^\alpha(t)x & \text{if } N_f(t) > 0, \\ 0 & \text{otherwise.} \end{cases}$$

⁵Our results easily extend to the case where $\alpha = 1$, by adopting the standard definition $V_f(y) \triangleq \log y$ in this case. We will restrict ourselves to the $\alpha \neq 1$ case in what follows, however, to simplify the presentation.

Given the strictly concave nature of the objective in this optimization program, it is clear that the maximizer is unique and is given by

$$X_f(t) = \begin{cases} \max\{N_f(t)/Q_f(t), C\} & \text{if } Q_f(t) > 0, \\ C & \text{otherwise.} \end{cases} \quad (4)$$

Note that while (4) does not explicitly depend on the parameter α , this parameter will impact downstream queue lengths through the scheduling mechanism that we will define momentarily, hence the choice of α will indirectly influence rate allocation decisions.

Denote by $\bar{X}_f(t)$ the cumulative rate allocation to flows of type f in the time interval $[0, t]$, i.e.,

$$\bar{X}_f(t) \triangleq \int_0^t X_f(s) ds.$$

$\bar{X}_f(\cdot)$ is Lipschitz continuous and differentiable, since $X_f(\cdot)$ is always bounded by C .

3.2. Control: Scheduling

The second control decision that must be specified is that of scheduling. We will assume the following variation of the ‘maximum weight’ or ‘back-pressure’ policies introduced by Tassiulas and Ephremides [30].

At the beginning of each discrete-time slot $\tau \in \mathbb{Z}_+$, a schedule $\pi(\tau) \in \mathcal{S}$ is chosen according to the optimization problem

$$\begin{aligned} \pi(\tau) &\in \operatorname{argmax}_{\pi \in \mathcal{S}} \sum_{e \in \mathcal{E}} \pi_e \left[Q_e^\alpha(\tau^-) - \sum_{e' \in \mathcal{E}} R_{ee'} Q_{e'}^\alpha(\tau^-) \right] \\ &= \operatorname{argmax}_{\pi \in \mathcal{S}} \pi^\top (I - R) Q^\alpha(\tau^-), \end{aligned} \quad (5)$$

where $Q^\alpha(\tau^-) \triangleq [Q_e^\alpha(\tau^-)]_{e \in \mathcal{E}}$ is a vector of component-wise powers of queue lengths, immediately prior to time τ . In other words, the schedule $\pi(\tau)$ is chosen so as to maximize the summation of weights of queues served, where weight of a queue $e \in \mathcal{E}$ is given by $[(I - R)Q^\alpha(\tau^-)]_e$. Given that \mathcal{S} is monotone, there exists a $\pi \in \mathcal{S}$ that maximizes this weight and is such that $\pi_e = 0$ if $Q_e(\tau^-) = 0$. We will restrict our attention to such schedules only. From this, it follows that the objective value of the optimization program in (5) is always non-negative.

By the discussion above, it is clear that the following invariants are satisfied:

1. For any schedule π and time τ , $S_\pi(\tau) = S_\pi(\tau - 1)$ if

$$\pi^\top (I - R) Q^\alpha(\tau^-) < \sigma^\top (I - R) Q^\alpha(\tau^-),$$

for some $\sigma \in \mathcal{S}$.

2. For any queue e and time τ , $Z_e(\tau) = 0$. In other words, there is no idling.

4. Fluid Model

In this section, we introduce the fluid model of the our system. As we shall see, the allocation of rate to flows in the fluid model resembles rate allocation of ‘flow-level’ models that has been popular in the literature [17, 3]. In that sense, our model on original time-scale operates at a packet-level granularity and on the fluid-scale operates at a flow-level granularity.

4.1. Fluid Scaling and Fluid Model Equations

In order to introduce the fluid model of our network, we will consider the scaled version of the system. To this end, denote the overall system state at a time $t \geq 0$ by

$$\mathcal{Z}(t) \triangleq \left(Q(t), Z(\lfloor t \rfloor), N(t), S(\lfloor t \rfloor), \bar{X}(t), \mathcal{A}(t), D(t), A(t) \right).$$

Here, the components of the state $\mathcal{Z}(t)$ are the primitives introduced in Sections 2.2 and 2.3. That is, at times $t \in \mathbb{R}_+$, we have,

$$\begin{aligned} Q(t) &\triangleq [Q_e(t)]_{e \in \mathcal{E}}, & \text{where } Q_e(t) \text{ is the length of queue } e, \\ N(t) &\triangleq [N_f(t)]_{f \in \mathcal{F}}, & \text{where } N_f(t) \text{ is the number of flows of type } f, \\ \bar{X}(t) &\triangleq [\bar{X}_f(t)]_{f \in \mathcal{F}}, & \text{where } \bar{X}_f(t) \text{ is the cumulative rate allocated to flow type } f, \\ \mathcal{A}(t) &\triangleq [\mathcal{A}_f(t)]_{f \in \mathcal{F}}, & \text{where } \mathcal{A}_f(t) \text{ is the cumulative arrival count of flow type } f, \\ D(t) &\triangleq [D_f(t)]_{f \in \mathcal{F}}, & \text{where } D_f(t) \text{ is the cumulative departure count of flow type } f, \\ A(t) &\triangleq [A_f(t)]_{f \in \mathcal{F}}, & \text{where } A_f(t) \text{ is the cumulative packet arrival count of flow type } f, \end{aligned}$$

and, at times $\tau \in \mathbb{Z}_+$, we have

$$\begin{aligned} Z(\tau) &\triangleq [Z_e(\tau)]_{e \in \mathcal{E}}, & \text{where } Z_e(\tau) \text{ is the cumulative idleness for queue } e, \\ S(\tau) &\triangleq [S_\pi(\tau)]_{\pi \in \mathcal{S}}, & \text{where } S_\pi(\tau) \text{ is the cumulative time schedule } \pi \text{ is employed.} \end{aligned}$$

Given a scaling parameter $r \in \mathbb{R}$, $r \geq 1$, define the scaled system state as

$$\mathcal{Z}^{(r)}(t) \triangleq \left(Q^{(r)}(t), Z^{(r)}(t), N^{(r)}(t), S^{(r)}(t), \bar{X}^{(r)}(t), \mathcal{A}^{(r)}(t), D^{(r)}(t), A^{(r)}(t) \right). \quad (6)$$

Here, the scaled components are defined as

$$\begin{aligned} Q^{(r)}(t) &\triangleq r^{-1}Q(rt), & N^{(r)}(t) &\triangleq r^{-1}N(rt), \\ \bar{X}^{(r)}(t) &\triangleq r^{-1}\bar{X}(rt), & \mathcal{A}^{(r)}(t) &\triangleq r^{-1}\mathcal{A}(rt), \\ D^{(r)}(t) &\triangleq r^{-1}D(rt), & A^{(r)}(t) &\triangleq r^{-1}A(rt), \\ Z^{(r)}(t) &\triangleq r^{-1}[(rt - \lfloor rt \rfloor)Z(\lceil rt \rceil) + (\lceil rt \rceil - rt)Z(\lfloor rt \rfloor)], \\ S^{(r)}(t) &\triangleq r^{-1}[(rt - \lfloor rt \rfloor)S(\lceil rt \rceil) + (\lceil rt \rceil - rt)S(\lfloor rt \rfloor)]. \end{aligned}$$

In the above, the components $Z(\cdot)$ and $S(\cdot)$ are linearly interpolated for technical convenience only.

Our interest is in understanding the behavior of $\mathcal{Z}^{(r)}(\cdot)$ as $r \rightarrow \infty$. Roughly speaking, in this limiting system the trajectories will satisfy certain deterministic equations, called fluid model equations. Solutions to these equations, which are defined below, will be denoted as fluid model solutions. The formal result is stated in Theorem 1.

Definition 1 (Fluid Model Solution). *Given fixed initial conditions $q(0) \in \mathbb{R}_+^{\mathcal{E}}$ and $n(0) \in \mathbb{R}_+^{\mathcal{F}}$, for every time horizon $T > 0$, let $\text{FMS}(T)$ denote the set of all trajectories*

$$\begin{aligned} \mathfrak{z}(t) &\triangleq (q(t), z(t), n(t), s(t), \bar{x}(t), \mathbf{a}(t), d(t), a(t)) \\ &\in \mathfrak{Z} \triangleq \mathbb{R}_+^{\mathcal{E}} \times \mathbb{R}_+^{\mathcal{E}} \times \mathbb{R}_+^{\mathcal{F}} \times \mathbb{R}_+^{\mathcal{S}} \times \mathbb{R}_+^{\mathcal{F}} \times \mathbb{R}_+^{\mathcal{F}} \times \mathbb{R}_+^{\mathcal{F}} \times \mathbb{R}_+^{\mathcal{F}}, \quad (7) \end{aligned}$$

over the time interval $[0, T]$ such that:

- (F1) All components of $\mathbf{z}(t)$ are uniformly Lipschitz continuous and thus differentiable for almost every $t \in (0, T)$. Such values of t are known as regular points.
- (F2) For all $t \in [0, T]$, $n(t) = n(0) + \mathbf{a}(t) - d(t)$.
- (F3) For all $t \in [0, T]$, $\mathbf{a}(t) = \nu t$.
- (F4) For all $t \in [0, T]$, $d(t) = \text{diag}(\mu)\bar{x}(t)$.
- (F5) For all $t \in [0, T]$,

$$q(t) = q(0) - (I - R^\top)\Pi s(t) + (I - R^\top)z(t) + \Gamma a(t).$$

- (F6) For all $t \in [0, T]$, $a(t) = \bar{x}(t)$.
- (F7) For all $t \in [0, T]$, $\mathbf{1}^\top s(t) = t$.
- (F8) Each component of $z(\cdot)$, $s(\cdot)$, and $\bar{x}(\cdot)$ is non-decreasing.

In addition, define the set $\text{FMS}^\alpha(T)$ to be the subset of trajectories in $\text{FMS}(T)$ that also satisfy:

- (F9) If $t \in (0, T)$ is a regular point, then for all $f \in \mathcal{F}$,

$$x_f(t) = \begin{cases} \underset{x \in [0, C]}{\text{argmax}} \frac{x^{1-\alpha} n_f^\alpha(t)}{1-\alpha} - q_{i(f)}^\alpha(t) x & \text{if } n_f(t) > 0, \\ \nu_f / \mu_f (= \rho_f) & \text{otherwise,} \end{cases}$$

where⁶ $x_f(t) \triangleq \dot{\bar{x}}_f(t)$.

- (F10) If $t \in (0, T)$ is a regular point, then for all $\pi \in \mathcal{S}$, $\dot{s}_\pi(t) = 0$, if

$$\pi^\top (I - R)q^\alpha(t) < \max_{\sigma \in \mathcal{S}} \sigma^\top (I - R)q^\alpha(t).$$

- (F11) If $t \in (0, T)$ is a regular point and $n_f(t) = 0$ for some $f \in \mathcal{F}$, then $q_{i(f)}(t) = 0$.

- (F12) For all $t \in [0, T]$, $z(t) = \mathbf{0}$.

Note that (F1)–(F8) correspond to fluid model equations that must be satisfied under *any* scheduling policy, and, hence, are *algorithm independent* fluid model equations. On the other hand, (F9)–(F12) are particular to networks controlled under the MWUM- α policy. (F9) captures the long-term effect of the rate allocation mechanism through the α -fair utility maximization based policy. (F10) captures the effect of short-term packet-level behavior induced by the scheduling algorithm. Specifically, the characteristics of the MW- α packet scheduling algorithm are captured by this equation. As explained in [10], in a static utility maximization formulation where the network is trying to achieve allocation of resources to maximize the α -fair utility subject to capacity constraint, the equation (F9) turns out to be the *primal* step and (F10) the *dual* step for a gradient algorithm attempting to solve the Lagrangian dual formulation of the static optimization problem. In that sense, the fluid-model above suggests that the system is attempting to solve such a network-wide utility maximization problem at each time instance.

4.2. Formal Statement

We wish to establish fluid model solutions as limit of the scaled system state process $\mathcal{Z}^{(r)}(\cdot)$ as $r \rightarrow \infty$. To this end, fix a time horizon $T > 0$. Let $\mathbf{D}[0, T]$ denote the space of all functions from

⁶We use the notation $\dot{\theta}(t)$ to denote $\frac{d}{dt}\theta(t)$ for $\theta: [0, T] \rightarrow \mathbb{R}$.

$[0, T]$ to \mathfrak{Z} , as defined in (7), that are right continuous with left limits (RCLL). We will denote the Skorohod metric on this space by $\mathbf{d}(\cdot, \cdot)$ (see Appendix B for details). Given a fixed scaling parameter r , consider the scaled system dynamics over interval $[0, T]$. Each sample path $\mathcal{Z}^{(r)}(\cdot)$ of the system state is RCLL, and hence is contained in the space $\mathbf{D}[0, T]$.

The following theorem formally establishes the convergence of the scaled system process to a fluid model solution of the form specified in Definition 1.

Theorem 1. *Given a fixed time horizon $T > 0$, consider a collection of scaled system state processes $\{\mathcal{Z}^{(r)}(\cdot) : r \geq 1\} \subset \mathbf{D}[0, T]$ under an arbitrary control policy. Suppose the initial conditions*

$$\lim_{r \rightarrow \infty} Q^{(r)}(0) = q(0), \quad \lim_{r \rightarrow \infty} N^{(r)}(0) = n(0), \quad (8)$$

are satisfied with probability 1. Then, for any $\varepsilon > 0$,

$$\liminf_{r \rightarrow \infty} \mathbb{P}\left(\mathcal{Z}^{(r)}(\cdot) \in \text{FMS}_\varepsilon(T)\right) = 1.$$

Here, $\text{FMS}_\varepsilon(T)$ is an ε -flattening of the set $\text{FMS}(T)$ of fluid model solutions, i.e.,

$$\text{FMS}_\varepsilon(T) \triangleq \{\mathbf{x} \in \mathbf{D}[0, T] : \mathbf{d}(\mathbf{x}, \mathbf{y}) < \varepsilon, \mathbf{y} \in \text{FMS}(T)\}.$$

Additionally, under the MWUM- α control policy, we have that

$$\liminf_{r \rightarrow \infty} \mathbb{P}\left(\mathcal{Z}^{(r)}(\cdot) \in \text{FMS}_\varepsilon^\alpha(T)\right) = 1.$$

Here, $\text{FMS}_\varepsilon^\alpha(T)$ is an ε -flattening of the set $\text{FMS}^\alpha(T)$ of MWUM- α fluid model solutions, i.e.,

$$\text{FMS}_\varepsilon^\alpha(T) \triangleq \{\mathbf{x} \in \mathbf{D}[0, T] : \mathbf{d}(\mathbf{x}, \mathbf{y}) < \varepsilon, \mathbf{y} \in \text{FMS}^\alpha(T)\}.$$

Theorem 1 can be established by following a somewhat standard sequence of arguments (cf. [4, 25, 13]). First, the collection of measures corresponding to the collection of random processes $\{\mathcal{Z}^{(r)}(\cdot) : r \geq 1\}$ is shown to be tight. This establishes that limit points must exist. Next, it is established that each limit point must satisfy the conditions of a fluid model solution with probability 1. The tightness argument uses concentration properties of Poisson process along with the Lipschitz property of queue length process. A detailed argument is required to establish that, with probability 1, the conditions of fluid model solution are satisfied. The complete proof of Theorem 1 is presented in Appendix B.

5. System Stability

In this section, we characterize the stability of a network under the MWUM- α policy. In particular, we shall see that the network evolves according to a Markov process is positive recurrent as long as the system is *underloaded*. In other words, the system is maximally stable. In order to construct the stability region under the MWUM- α policy, first define the set of *packet arrival rates* by

$$\Lambda \triangleq \left\{ \lambda \in \mathbb{R}_+^{\mathcal{E}} : \exists s \in \mathbb{R}_+^{\mathcal{S}} \text{ with } \lambda \leq \Pi s, \mathbf{1}^\top s \leq 1 \right\}. \quad (9)$$

Imagine that the network has no packet arrivals from flows, but instead has packets arriving according to exogenous processes. Suppose that $\lambda \in \mathbb{R}_+^{\mathcal{E}}$ is the vector of exogenous arrival rates, so

that packets arrived to each queue e at rate λ_e . Then, it is not difficult to see that the network would not be stable under *any* scheduling policy if $\lambda \notin \Lambda$. This is because there is at least one queue in the network that is loaded beyond its service capacity. Hence, the set Λ represents the raw *scheduling capacity* of the network.

The set Λ can alternatively be described as follows: Given a vector $\lambda \in \mathbb{R}_+^{\mathcal{E}}$, consider the linear program

$$\begin{aligned} \text{PRIMAL}(\lambda) \triangleq \underset{s}{\text{minimize}} \quad & \mathbf{1}^\top s \\ \text{subject to} \quad & \lambda \leq \Pi s, \\ & s \in \mathbb{R}_+^{\mathcal{S}}. \end{aligned}$$

Clearly $\lambda \in \Lambda$ if and only if $\text{PRIMAL}(\lambda) \leq 1$. The quantity $\text{PRIMAL}(\lambda)$ is called the *effective load* of a system with exogenous packet arrivals at rate λ .

Now, in our model, packets arrive to the network not through an exogenous process, but rather, they are generated by flows. As discussed in Section 2.2, each flow type $f \in \mathcal{F}$ generates packets at according to an offered load of ρ_f . The generated packets are injected into the network according to the ingress matrix Γ , and subsequently travel through the network along pre-determined paths specified by the routing matrix R . Let $\lambda \in \mathbb{R}_+^{\mathcal{E}}$ be the vector of *implied loads* on the scheduling network due to the packets generated by flows. It seems reasonable to relate λ and the vector $\rho \in \mathbb{R}_+^{\mathcal{F}}$ of offered loads according to $\lambda = \Gamma\rho + R^\top\lambda$. Equivalently, we define $\lambda \triangleq \Xi\Gamma\rho$, where Ξ is from (1). We define the *effective load* $L(\rho)$ of our network by $L(\rho) \triangleq \text{PRIMAL}(\Xi\Gamma\rho)$.

Given the above discussion, it seems natural to suspect that the network's scheduling capacity allows it to operate effectively as long as $L(\rho) \leq 1$. This motivates the following definition:

Definition 2 (Admissibility). A vector $\rho \in \mathbb{R}_+^{\mathcal{F}}$ of offered loads *admissible* if $L(\rho) \leq 1$. Similarly, ρ is *strictly admissible* if $L(\rho) < 1$. Finally, ρ is *critically admissible* if $L(\rho) = 1$.

We will establish system stability, or, more formally, positive recurrence, when offered load is *strictly admissible*. To this end, recall that the system is completely described by the $\mathcal{Z}(\cdot)$ process. Under the MWUM- α policy, the evolution of all the components of $\mathcal{Z}(\cdot)$ is entirely determined by $(N(\cdot), Q(\cdot))$. The packet-scheduling decisions happen at discrete time, flows are served as per Poisson process with state-dependent rate, and arrival process are fixed rate Poisson process. Therefore, the $\mathcal{Z}(\tau)$ at discrete-time instance $\tau \in \mathbb{Z}_+$ is Markov chain. We shall establish positive-recurrence of this Markov chain when system is under-loaded. The following is the main result of this section:

Theorem 2. Consider a network system with strictly admissible ρ operating under the MWUM- α policy. Then, the Markov chain $(N(\tau), Q(\tau))$, $\tau \in \mathbb{Z}_+$, is positive recurrent.

It is worth noting that if $L(\rho) > 1$, then at least one of the queues in the network must be, on average, loaded beyond its capacity. Hence the network Markov process can not be positive recurrent or stable.

The proof of Theorem 2, provided next in Section 5.1, uses a fluid model approach. Dai [5], for example, developed an early instance of such an approach for a class of queueing networks. However, this result does not apply to the present setting, and a specialized analysis is needed. Conceptually, the fluid model approach involves two steps: (1) derive strong stability of fluid model, and (2) use strong stability to establish the positive recurrence of the original Markov chain. To this end, define Lyapunov function L_α over the vector of flow counts $n = [n_f]_{f \in \mathcal{F}} \in \mathbb{R}_+^{\mathcal{F}}$ and the vector of queue

lengths $q = [q_e]_{e \in \mathcal{E}} \in \mathbb{R}_+^{\mathcal{E}}$ by

$$L_\alpha(n, q) \triangleq \sum_{f \in \mathcal{F}} \frac{n_f^{1+\alpha}}{\mu_f \rho_f^\alpha} + \sum_{e \in \mathcal{E}} q_e^{1+\alpha}. \quad (10)$$

The following lemma, whose proof is found in Section 5.2, provides a central argument towards establishing the strong stability of the fluid model. This lemma will also be of use later in establishing the characterization and attractiveness of invariant manifold under critical loading.

Lemma 3. *Let $(n(\cdot), q(\cdot))$ be, respectively, the flow count process and the queue length process of a fluid model solution in the set $\text{FMS}^\alpha(T)$. If $\mathsf{L}(\rho) \leq 1$, then for every regular point $t \in (0, T)$,*

$$\frac{d}{dt} L_\alpha(n(t), q(t)) \leq 0.$$

Suppose further that $\mathsf{L}(\rho) < 1$. Then, there exist $\delta_ > 0$ and $T_* > 0$ such that, for all $T > T_*$, if the initial conditions $(n(0), q(0))$ satisfy*

$$L_\alpha(n(0), q(0)) = 1,$$

then

$$L_\alpha(n(t), q(t)) \leq 1 - \delta_*, \quad \text{for all } t \in [T_*, T].$$

5.1. Proof of Theorem 2

The following lemma provides a sufficient condition for positive recurrence:

Lemma 4. [21, Theorem 8.13] *Let $\mathcal{X}(\cdot)$ be an irreducible, aperiodic discrete-time Markov chain on a countable state space \mathfrak{X} . Suppose there exists a function $V: \mathfrak{X} \rightarrow \mathbb{R}_+$, constants A and $\varepsilon > 0$, and an integrable stopping time $\tau > 0$ such that, for all $x \in \mathfrak{X}$ with $V(x) > A$,*

$$\mathbb{E}[V(\mathcal{X}(\tau)) \mid \mathcal{X}(0) = x] \leq V(x) - \varepsilon \mathbb{E}[\tau \mid \mathcal{X}(0) = x].$$

If the set $\{x \in \mathfrak{X} : V(x) \leq A\}$ is finite and $\mathbb{E}[V(\mathcal{X}(1)) \mid \mathcal{X}(0) = x] < \infty$ for all $x \in \mathfrak{X}$, then $\mathcal{X}(\cdot)$ is positive recurrent and ergodic.

Our proof of Theorem 2 relies on establishing the sufficient condition for positive recurrence given by Lemma 4, using the stability of the fluid model (Lemma 3). To this end, note that under the MWUM- α policy, $\mathcal{X}(\tau) \triangleq (N(\tau), Q(\tau)) \in \mathbb{Z}_+^{\mathcal{F}} \times \mathbb{Z}_+^{\mathcal{E}}$ is a discrete-time Markov chain over integer times $\tau \geq 0$, which is irreducible and aperiodic (from every state, the null state is reachable and there is a positive probability of remaining at the null state). We shall use ‘normed’ version of the Lyapunov function L_α , defined as

$$\ell(n, q) \triangleq (L_\alpha(n, q))^{\frac{1}{1+\alpha}},$$

for all $(n, q) \in \mathbb{R}_+^{\mathcal{F}} \times \mathbb{R}_+^{\mathcal{E}}$. The role of V in Lemma 4 will be played by ℓ . Lemma 14 implies that there exists constants $0 < C_1 \leq C_2 < \infty$ so that, for all (n, q) ,

$$C_1 \|(n, q)\|_\infty \leq \ell(n, q) \leq C_2 \|(n, q)\|_\infty. \quad (11)$$

Further, for any $\kappa > 0$ and (n, q) , we have that $\ell(\kappa n, \kappa q) = \kappa \ell(n, q)$.

Now, consider any sequence of initial states

$$\left\{ x^k \triangleq (N^k(0), Q^k(0)) \in \mathbb{Z}_+^{\mathcal{F}} \times \mathbb{Z}_+^{\mathcal{E}} : k \in \mathbb{N} \right\},$$

such that $\|x^k\|_\infty \rightarrow \infty$ as $k \rightarrow \infty$. For each k , consider a system starting at the initial state x^k . Denote the state of the k th system at time $t \geq 0$ by

$$\mathcal{X}^k(t) \triangleq (N^k(t), Q^k(t)).$$

Here, for non-integer values of t , we will utilize the continuous time description of the original Markov process. Indeed, the integer-time sampled version of this original continuous time Markov process is precisely what we are interested in.

For each $k \in \mathbb{N}$, define the scaling factor $r_k \triangleq \ell(x^k)$, and notice that $r_k \rightarrow \infty$. Fix a time horizon $T > 0$, and for the k th system, consider the *scaled* state process, defined for $t \in [0, T]$ as

$$\mathcal{X}^{(r_k)}(t) \triangleq \frac{1}{r_k} \mathcal{X}^k(r_k t) = \frac{1}{r_k} (N^k(r_k t), Q^k(r_k t)),$$

The descriptor of the scaled system is given, for $t \in [0, T]$, by $\mathcal{Z}^{(r_k)}(t)$ from (6). Let $\mu^{(r_k)}$ be its distribution on $\mathbf{D}[0, T]$.

From (11), we have that, for any k ,

$$\|\mathcal{X}^{(r_k)}(0)\|_\infty = \frac{\|\mathcal{X}^k(0)\|_\infty}{r_k} \leq 1/C_1. \quad (12)$$

Since the set of scaled initial conditions is compact, there must exist a limit point and a convergence subsequence. Along this subsequence, by the analysis of Theorem 1 (in particular, Lemma 17) the measures $\{\mu^{(r_k)}\}$ are tight, and therefore, there exists a measure $\mu^{(\infty)}$ that is a limit point. By restricting to a further subsequence, we can assume, without loss of generality, that $\mu^{(r_k)} \Rightarrow \mu^{(\infty)}$ as $k \rightarrow \infty$. That is,

$$\left(\mathcal{Z}^{(r_k)}(t) \right)_{t \in [0, T]} \Rightarrow \left(\mathfrak{z}(t) \right)_{t \in [0, T]}, \quad \text{as } k \rightarrow \infty,$$

with $\mathfrak{z}(\cdot) \in \text{FMS}^\alpha(T)$ satisfying the fluid model equations.

Given a fluid model solution $\mathfrak{z}(\cdot)$ of the form (7), denote by $(n(\cdot), q(\cdot))$ the flow count and queue length components. Note that $\ell(\mathcal{X}^{(r_k)}(0)) = 1$, for all k . By the continuity of ℓ , we have that $\ell(n(0), q(0)) = 1$; that is $L_\alpha(n(0), q(0)) = 1$. Then, from Lemma 3, there exist $\delta_* > 0$ and $T_* > 0$ so that, for sufficiently large T ,

$$\ell(n(t), q(t)) \leq 1 - \delta_*, \quad \text{for all } t \in [T_*, T]. \quad (13)$$

Define the functional $F : \mathbf{D}[0, T] \rightarrow \mathbb{R}_+$ by

$$F \left(\left(\mathfrak{z}(t) \right)_{t \in [0, T]} \right) \triangleq \sup_{t \in [T_*, \frac{1}{2}(T+T_*)]} \ell(n(t), q(t))$$

Since F is a continuous, it follows that

$$F \left(\left(\mathcal{Z}^{(r_k)}(t) \right)_{t \in [0, T]} \right) \Rightarrow F \left(\left(\mathfrak{z}(t) \right)_{t \in [0, T]} \right), \quad \text{as } k \rightarrow \infty. \quad (14)$$

Further, from (12), and using the fact that the arrival processes are Poisson and the boundedness of the rate allocation policies, it follows that for all $t \in [0, T]$ and all $k \in \mathbb{N}$,

$$\mathbb{E} \left[\left\| \left(N^{(r_k)}(t), Q^{(r_k)}(t) \right) \right\|_\infty \right] \leq C_3 T, \quad (15)$$

for some constant $C_3 > 0$. From this, the uniform integrability of $F(\mathcal{Z}^{(r_k)}(\cdot))$ follows. Subsequently, from (13) and (14), it follows that

$$\lim_{k \rightarrow \infty} \mathbb{E} \left[F \left(\left(\mathcal{Z}^{(r_k)}(t) \right)_{t \in [0, T]} \right) \right] \leq 1 - \delta_*.$$

Equivalently, in terms of the *unscaled* state process $\mathcal{X}(\cdot)$, we have that

$$\lim_{k \rightarrow \infty} \frac{1}{\ell(x^k)} \mathbb{E} \left[\sup_{t \in [T_*, \frac{1}{2}(T+T_*)/2]} \ell \left(\mathcal{X} \left(\ell(x^k)t \right) \right) \middle| \mathcal{X}(0) = x^k \right] \leq 1 - \delta_*. \quad (16)$$

To complete the proof of Theorem 2, define stopping time τ that is uniformly distributed over integers in range $[\ell(\mathcal{X}(0))T_*, \ell(\mathcal{X}(0))(T+T_*)/2]$. Then, it follows that τ is integrable and for all initial states $x \in \mathbb{Z}_+^{\mathcal{F}} \times \mathbb{Z}_+^{\mathcal{E}}$ with $\ell(x)$ large enough,

$$\mathbb{E} [\ell(\mathcal{X}(\tau)) \mid \mathcal{X}(0) = x] \leq \ell(x) - \varepsilon \mathbb{E}[\tau \mid \mathcal{X}(0) = x],$$

for appropriately chosen $\varepsilon > 0$. This satisfies the conditions of Lemma 4, and hence completes the proof of Theorem 2.

5.2. Proof of Lemma 3

Suppose $t \in (0, T)$ is a regular point. We will start by establishing the first part of Lemma 3: if $L(\rho) \leq 1$, then $\frac{d}{dt} L_\alpha(n(t), q(t)) \leq 0$. To this end, note that

$$\frac{d}{dt} L_\alpha(n(t), q(t)) = (1 + \alpha) \left(\underbrace{\sum_{f \in \mathcal{F}} \frac{n_f^\alpha(t)}{\mu_f \rho_f^\alpha} \dot{n}_f(t)}_{\Delta_n} + \underbrace{\sum_{e \in \mathcal{E}} q_e^\alpha(t) \dot{q}_e(t)}_{\Delta_q} \right). \quad (17)$$

We consider the terms Δ_n and Δ_q separately.

First, consider the term Δ_n . For each flow type f , we wish to show that

$$\frac{n_f^\alpha(t)}{\mu_f \rho_f^\alpha} \dot{n}_f(t) \leq q_{i(f)}^\alpha(t) (\rho_f - x_f(t)). \quad (18)$$

By (F2)–(F4), $\dot{n}_f(t) = \nu_f - \mu_f x_f(t)$. There are two cases. If $n_f(t) = 0$, then by (F9), we have that $x_f(t) = \rho_f$, thus both sides of (18) are 0. If $n_f(t) > 0$, then

$$\frac{n_f^\alpha(t)}{\mu_f \rho_f^\alpha} \dot{n}_f(t) = \frac{n_f^\alpha(t)}{\rho_f^\alpha} (\rho_f - x_f(t)) \leq n_f^\alpha(t) \frac{\rho_f^{1-\alpha}}{1-\alpha} - n_f^\alpha(t) \frac{x_f^{1-\alpha}(t)}{1-\alpha}. \quad (19)$$

Here, the inequality follows from the fact that the function $g(z) \triangleq z^{1-\alpha}/(1-\alpha)$ is concave. For a concave function g , $g'(y)(y-x) \leq g(y) - g(x)$, here we have $x = x_f(t)$ and $y = \rho_f$. Now, since $x_f(t)$ is optimal for the rate allocation problem of (F9) and ρ_f is feasible, we have that

$$n_f^\alpha(t) \frac{x_f^{1-\alpha}(t)}{1-\alpha} - q_{i(f)}^\alpha(t) x_f(t) \geq n_f^\alpha(t) \frac{\rho_f^{1-\alpha}}{1-\alpha} - q_{i(f)}^\alpha(t) \rho_f. \quad (20)$$

Combining (19) and (20), we have established (18). Then, we can sum (18) over all $f \in \mathcal{F}$, to obtain

$$\Delta_n \leq [\Gamma(\rho - x(t))]^\top q^\alpha(t). \quad (21)$$

Now, consider the term Δ_q in (17). (F5), (F6), and (F12) imply that

$$\dot{q}(t) = -(I - R^\top)\Pi\dot{s}(t) + \Gamma x(t). \quad (22)$$

From (F7) and (F10),

$$\left[(I - R^\top)\Pi\dot{s}(t) \right]^\top q^\alpha(t) = \sum_{\pi \in \mathcal{S}} \dot{s}_\pi(t) \pi^\top (I - R) q^\alpha(t) = \max_{\sigma \in \mathcal{S}} \sigma^\top (I - R) q^\alpha(t). \quad (23)$$

Together (22) and (23), imply that

$$\Delta_q = [\Gamma x(t)]^\top q^\alpha(t) - \max_{\sigma \in \mathcal{S}} \sigma^\top (I - R) q^\alpha(t). \quad (24)$$

Now, combining (17), (21), and (24), we have that

$$\frac{d}{dt} L_\alpha(n(t), q(t)) \leq (1 + \alpha) \left([\Gamma \rho]^\top q^\alpha(t) - \max_{\sigma \in \mathcal{S}} \sigma^\top (I - R) q^\alpha(t) \right). \quad (25)$$

By the definition of $L(\rho) \triangleq \text{PRIMAL}(\Xi \Gamma \rho)$, there exists some $s \in \mathbb{R}_+^{\mathcal{S}}$ with $\mathbf{1}^\top s = L(\rho)$ and

$$\Gamma \rho \leq (I - R^\top)\Pi s = \sum_{\pi \in \mathcal{S}} s_\pi (I - R^\top)\pi.$$

This implies that

$$[\Gamma \rho]^\top q^\alpha(t) \leq L(\rho) \max_{\sigma \in \mathcal{S}} \sigma^\top (I - R) q^\alpha(t),$$

Hence, by (25),

$$\frac{d}{dt} L_\alpha(n(t), q(t)) \leq -(1 + \alpha)(1 - L(\rho)) \max_{\sigma \in \mathcal{S}} \sigma^\top (I - R) q^\alpha(t). \quad (26)$$

In order to bound the right hand side of (26), we will argue that

$$\max_{\sigma \in \mathcal{S}} \sigma^\top (I - R) q^\alpha(t) \geq \frac{\mathbf{1}^\top q^\alpha(t)}{|\mathcal{E}|^2}. \quad (27)$$

To see this, consider a queue $e_0 \in \mathcal{E}$ with maximal length at time t , i.e., $q_{e_0}(t) = \max_e q_e(t)$. Define $e_1, \dots, e_J \in \mathcal{E}$ to be a sequence of distinct queues so that, for each $0 \leq i < J$, packets departing from queue e_i go to queue e_{i+1} , and packets departing from queue e_J exist the network. By our assumption of acyclic routing, such a sequence exists, and since each queue is distinct, $J + 1 \leq |\mathcal{E}|$.

Consider the distinct schedules $\pi_0, \dots, \pi_J \in \mathcal{S}$, where, for $0 \leq i \leq J$, the schedule π_i serves exactly the queue e_i — such schedules exist by the monotonicity assumption on the scheduling constraints. Clearly, these schedules have weights given by

$$\pi_i^\top (I - R)q^\alpha(t) = \begin{cases} q_{e_i}^\alpha(t) - q_{e_{i+1}}^\alpha(t) & \text{if } 0 \leq i < J, \\ q_{e_J}^\alpha(t) & \text{if } i = J. \end{cases}$$

Averaging over the $J + 1$ schedules,

$$\frac{1}{J+1} \sum_{i=0}^J \pi_i^\top (I - R)q^\alpha(t) = \frac{q_{e_0}^\alpha(t)}{J+1}.$$

Since at least one schedule must have a weight that exceeds this average,

$$\max_{\sigma \in \mathcal{S}} \sigma^\top (I - R)q^\alpha(t) \geq \frac{q_{e_0}^\alpha(t)}{J+1} \geq \frac{\max_{e \in \mathcal{E}} q_e^\alpha(t)}{|\mathcal{E}|}$$

Since

$$\max_{e \in \mathcal{E}} q_e^\alpha(t) \geq \frac{\mathbf{1}^\top q^\alpha(t)}{|\mathcal{E}|},$$

(27) follows.

Combining (26) and (27), we obtain, when $L(\rho) \leq 1$,

$$\frac{d}{dt} L_\alpha(n(t), q(t)) \leq -(1 + \alpha)(1 - L(\rho)) \frac{\mathbf{1}^\top q^\alpha(t)}{|\mathcal{E}|^2} \leq 0. \quad (28)$$

This establishes the first part of Lemma 3.

To prove the second part of Lemma 3, we will consider two separate cases over initial conditions $(n(0), q(0))$ with $L_\alpha(n(0), q(0)) = 1$:

- (i) $\mathbf{1}^\top q^{1+\alpha}(0) > \varepsilon_1$.
- (ii) $\mathbf{1}^\top q^{1+\alpha}(0) \leq \varepsilon_1$.

Here, $\varepsilon_1 > 0$ is a constant that will be determined shortly.

For case (i), from the norm inequality in part (i) of Lemma 14, we have that

$$\varepsilon_1^{\frac{1}{1+\alpha}} < \|q(0)\|_{1+\alpha} \leq \|q(0)\|_\alpha,$$

Thus, $\mathbf{1}^\top q^\alpha(0) > \varepsilon_2 \triangleq \varepsilon_1^{\frac{\alpha}{1+\alpha}}$. Due to (F1), $q(\cdot)$ is uniformly Lipschitz continuous, there exists $T_1 > 0$ such that $\mathbf{1}^\top q^\alpha(t) \geq \varepsilon_2/2$ for all $t \in [0, T_1]$. From (28), since $L(\rho) < 1$, we have that

$$\frac{d}{dt} L_\alpha(n(t), q(t)) \leq -(1 + \alpha)(1 - L(\rho)) \frac{\varepsilon_2}{2|\mathcal{E}|^2} < 0,$$

for all regular $t \in (0, T_1]$. Therefore, there exists $\delta_1 > 0$ so that

$$L_\alpha(n(t), q(t)) \leq 1 - \delta_1, \quad (29)$$

for all $t \geq T_1$.

For case (ii), the argument is more complicated. The basic insight is as follows: if $\mathbf{1}^\top q^{1+\alpha}(0)$ is *small*, then Δ_q in (17) must be small as well. Further, a good fraction of the flows will be allocated the maximal rate (i.e., C), therefore Δ_n in (17) will be significantly negative. This will lead to strictly negative drift in L . We will now formalize this intuition and make an appropriate choice of $\varepsilon_1 > 0$ along the way.

First, by the Lipschitz continuity of $q(\cdot)$, there exists some $\tau_1 > 0$ such that, for all $t \in [0, \tau_1]$, $\mathbf{1}^\top q^{1+\alpha}(t) \leq 2\varepsilon_1$. Note that from (24), the fact that $\mathbf{0} \in \mathcal{S}$, and $x(t) \leq C\mathbf{1}$, for all regular $t \in (0, \tau_1]$,

$$\Delta_q \leq [\Gamma x(t)]^\top q^\alpha(t) \leq C\mathbf{1}^\top q^\alpha(t).$$

By Jensen's inequality,

$$\left(\frac{1}{|\mathcal{F}|} \mathbf{1}^\top q^\alpha(t) \right)^{\frac{1+\alpha}{\alpha}} \leq \frac{1}{|\mathcal{F}|} \mathbf{1}^\top q^{1+\alpha}(t).$$

Therefore, for regular $t \in (0, \tau_1]$,

$$\Delta_q \leq C|\mathcal{F}|^{\frac{1}{1+\alpha}} (2\varepsilon_1)^{\frac{\alpha}{1+\alpha}}. \quad (30)$$

Again using the fact that $\mathbf{1}^\top q^{1+\alpha}(t) \leq 2\varepsilon_1$ for $t \in [0, \tau_1]$, it follows that $q_e(t) \leq \varepsilon_3 \triangleq (2\varepsilon_1)^{1/1+\alpha}$, for all $e \in \mathcal{E}$. Now, consider the set

$$\mathcal{F}' \triangleq \{f \in \mathcal{F} : n_f(0) \geq 4C\varepsilon_3\}.$$

By Lipschitz continuity, there exists some $0 < T_2 < \tau_1$ such that, for all $f \in \mathcal{F}'$ and all $t \in [0, T_2]$, $n_f(t) \geq 2C\varepsilon_3$. Then, by (F9), $x_f(t) = C$ for regular $t \in (0, T_2]$. Thus, if $f \in \mathcal{F}'$, we have, for regular $t \in (0, T_2]$,

$$\frac{n_f^\alpha(t)}{\mu_f \rho_f^\alpha} \dot{n}_f(t) = \frac{n_f^\alpha(t)}{\mu_f \rho_f^\alpha} (\rho_f - x_f(t)) \leq -\frac{n_f^\alpha(t)}{\mu_f \rho_f^\alpha} (C - \rho_f) \leq -\beta_1 n_f^\alpha(t), \quad (31)$$

where we define

$$\beta_1 \triangleq \min_{f \in \mathcal{F}'} \frac{C - \rho_f}{\mu_f \rho_f^\alpha} > 0.$$

Finally, if $f \notin \mathcal{F}'$, we have, for regular $t \in (0, T_2]$,

$$\frac{n_f^\alpha(t)}{\mu_f \rho_f^\alpha} \dot{n}_f(t) = \frac{n_f^\alpha(t)}{\mu_f \rho_f^\alpha} (\rho_f - x_f(t)) \leq \frac{(2C)^\alpha \rho_f^{1-\alpha} \varepsilon_1^{\frac{\alpha}{1+\alpha}}}{\mu_f} \leq \beta_2 (2C)^\alpha \varepsilon_1^{\frac{\alpha}{1+\alpha}}, \quad (32)$$

where we define

$$\beta_2 \triangleq \max_{f \in \mathcal{F}} \frac{\rho_f^{1-\alpha}}{\mu_f}.$$

Using (17) and (30)–(32), it follows that, for all $t \in [0, T_2]$,

$$\frac{d}{dt} L_\alpha(n(t), q(t)) \leq (1 + \alpha) \left(\left[C|\mathcal{F}|^{\frac{1}{1+\alpha}} + \beta_2 |\mathcal{F}| (2C)^\alpha \right] \varepsilon_1^{\frac{\alpha}{1+\alpha}} - \beta_1 \sum_{f \in \mathcal{F}'} n_f^\alpha(t) \right). \quad (33)$$

From (28), for all $t \geq 0$, $L_\alpha(n(t), q(t)) \leq 1$. For $t \in [0, T_2]$, $\mathbf{1}^\top q(t)^{1+\alpha} \leq 2\varepsilon_1$. Suppose that $\varepsilon_1 < 1/4$. Then,

$$\sum_{f \in \mathcal{F}'} n_f^{1+\alpha}(t) \geq \frac{1 - 2\varepsilon_1}{\max_{f \in \mathcal{F}} \frac{1}{\mu_f \rho_f^\alpha}} \geq \frac{1}{4 \max_{f \in \mathcal{F}} \frac{1}{\mu_f \rho_f^\alpha}}.$$

for $t \in [0, T_2]$. From part (i) of Lemma 14,

$$\sum_{f \in \mathcal{F}'} n_f^\alpha(t) \geq \left(\frac{1}{4 \max_{f \in \mathcal{F}} \frac{1}{\mu_f \rho_f^\alpha}} \right)^{\frac{\alpha}{1+\alpha}} \triangleq \beta_3 > 0.$$

Then, for regular $t \in (0, T_2]$,

$$\frac{d}{dt} L_\alpha(n(t), q(t)) \leq -(1 + \alpha) \beta_1 \beta_3 / 2 < 0.$$

It follows that there exists $\delta_2 > 0$ such that, for all $t \geq T_2$,

$$L_\alpha(n(t), q(t)) \leq 1 - \delta_2. \quad (34)$$

Lemma 3 follows from (29) and (34) with $\delta_* = \min\{\delta_1, \delta_2\}$ and $T_* = \max\{T_1, T_2\}$.

6. Critical Loading

We have established the throughput optimality of the system under the MWUM- α control policy, for any $\alpha \in (0, \infty) \setminus \{1\}$. Thus, this entire family of policies possesses good *first order* characteristics. Further, there may be many other throughput optimal policies outside the class of MWUM- α policies. This naturally raises the question of whether there is a ‘best’ choice of α , and how the resulting MWUM- α policy might compare against the universe of all other policies.

In order to answer these questions, we desire a more refined analysis of policy performance than throughput optimality. One way to obtain such an analysis is via the study of a *critically loaded system*, i.e., a system with critically admissible arrival rates. Under a critical loading, fluid model solutions take non-trivial values over entire horizon. In contrast, for strictly admissible systems under throughput optimal policies, all fluid trajectories go to $\mathbf{0}$ (cf. Lemma 3). We will employ the study of the fluid model solutions of critically loaded systems as a tool for the comparative analysis of network control policies.

In particular, given a vector of flow counts, $n \in \mathbb{R}_+^{\mathcal{F}}$, and the vector of queue lengths, $q \in \mathbb{R}_+^{\mathcal{E}}$, define the linear cost function

$$c(n, q) \triangleq \sum_{f \in \mathcal{F}} \frac{n_f}{\mu_f} + \sum_{e \in \mathcal{E}} q_e = \mathbf{1}^\top [\Gamma \text{diag}(\mu^{-1})n + q]. \quad (35)$$

This cost function is analogous to a ‘minimum delay’ objective in a packet-level queueing network: a cost is incurred for each queued packet, and a cost is also incurred for each outstanding flow, proportional to the number of packets that it is expected to generate.

In this section, we establish fundamental lower bounds that apply to the cost incurred in a critically loaded fluid model under *any* scheduling policy. In Sections 7 and 8, we will compare these with the costs incurred by MWUM- α control policies. We shall find that as $\alpha \rightarrow 0^+$, the cost induced by the MWUM- α algorithms improves and becomes close to the algorithm independent lower bound

6.1. Virtual Resources and Workload

We begin with some definitions. First, consider the dual of the LP PRIMAL(λ),

$$\begin{aligned} \text{DUAL}(\lambda) \triangleq \underset{\zeta}{\text{maximize}} \quad & \lambda^\top \zeta \\ \text{subject to} \quad & \Pi^\top \zeta \leq \mathbf{1}, \\ & \zeta \in \mathcal{R}_+^\mathcal{E}. \end{aligned}$$

Note that there is no duality gap, thus the value of PRIMAL(λ) is equal to the value of DUAL(λ).

Definition 3 (Virtual Resource). *We will call any feasible solution $\zeta \in \mathbb{R}_+^\mathcal{E}$ of dual optimization problem DUAL($\Xi\Gamma\rho$) a **virtual resource**. Suppose the system is critically loaded, i.e., the offered load vector ρ satisfies*

$$L(\rho) = \text{PRIMAL}(\Xi\Gamma\rho) = \text{DUAL}(\Xi\Gamma\rho) = 1.$$

*Then, we call a virtual resource that is an optimal solution of DUAL($\Xi\Gamma\rho$) a **critical virtual resource**.*

For a critically loaded system with offered load vector ρ , let $\text{CR}(\rho)$ denote the set of all critical virtual resources. Note that $\text{CR}(\rho)$ is a bounded polytope and hence possesses finitely many extreme points. Let $\text{CR}^*(\rho)$ denote the set of extreme points of $\text{CR}(\rho)$. Informally, a critical virtual resource identifies collection of queues with associated weights which form a bottleneck in a critically loaded system. In particular, the amount of ‘work’ associated with a critical virtual resource cannot be reduced under *any* scheduling policy. To see this, we first make the following definition of the work associated with a critical virtual resource, as a function of the current state of the system.

Definition 4 (Workload). *Consider a critically loaded system with an offered load vector ρ and a critical virtual resource $\zeta \in \text{CR}(\rho)$. If the flow count and queue length vectors are given by (n, q) , the **workload** associated with the resource ζ is defined to be*

$$w_\zeta(n, q) \triangleq \zeta^\top \Xi [q + \Gamma \text{diag}(\mu)^{-1}n].$$

6.2. A Lower Bound on Fluid Trajectories

Consider a critically loaded system with offered load vector ρ . We claim that the following fundamental lower bound holds on the fluid trajectory under *any* algorithm. This bound can be thought of as a minimal work-conservation requirement.

Lemma 5. *Consider the fluid model trajectory of system under any scheduling and rate allocation policy, with flow count and queue length processes given by $(n(\cdot), q(\cdot))$. Then, for any time $t \geq 0$ and any critical virtual resource $\zeta \in \text{CR}(\rho)$,*

$$w_\zeta(n(0), q(0)) \leq w_\zeta(n(t), q(t)). \quad (36)$$

Proof. Given a time interval $[0, T]$, for any $T > 0$, consider the fluid model trajectory $\mathfrak{z}(\cdot)$ of the form (7). By Theorem 1, this fluid trajectory must satisfy the algorithm independent fluid model equations (F1)–(F8). By (F1), the trajectory is Lipschitz continuous and differentiable for almost all $t \in (0, T)$. For any such regular point t , by (F2)–(F4), we have

$$\dot{n}(t) = \nu - \text{diag}(\mu)\dot{x}(t).$$

Thus,

$$\Gamma \operatorname{diag}(\mu^{-1})\dot{n}(t) = \Gamma\rho - \Gamma\dot{x}(t). \quad (37)$$

From (F5)–(F6), we obtain

$$\dot{q}(t) = (I - R^\top) \dot{z}(t) - (I - R^\top) \Pi \dot{s}(t) + \Gamma \dot{x}(t). \quad (38)$$

Adding (37) and (38), we obtain

$$\dot{q}(t) + \Gamma \operatorname{diag}(\mu^{-1})\dot{n}(t) = \Gamma\rho + (I - R^\top) (\dot{z}(t) - \Pi \dot{s}(t)).$$

Now, multiplying both sides by $\Xi \triangleq (I - R^\top)^{-1}$, we obtain

$$\Xi [\dot{q}(t) + \Gamma \operatorname{diag}(\mu^{-1})\dot{n}(t)] = \Xi\Gamma\rho + \dot{z}(t) - \Pi \dot{s}(t). \quad (39)$$

Consider a critical virtual resource $\zeta \in \mathbf{CR}(\rho)$. Since ζ is $\mathbf{DUAL}(\Xi\Gamma\rho)$ optimal, $\zeta^\top \Xi\Gamma\rho = 1$. Taking an inner product of (39) with ζ , we obtain

$$\zeta^\top \Xi [\dot{q}(t) + \Gamma \operatorname{diag}(\mu^{-1})\dot{n}(t)] = 1 + \zeta^\top \dot{z}(t) - \zeta^\top \Pi \dot{s}(t). \quad (40)$$

Now, by (F8), $z(\cdot)$ is non-decreasing. Since ζ is non-negative, then $\zeta^\top \dot{z}(t) \geq 0$. By (F8), $\dot{s}(t)$ is non-negative. Since ζ is $\mathbf{DUAL}(\Xi\Gamma\rho)$ feasible and from (F7), it follows that $\zeta^\top \Pi \dot{s}(t) \leq \mathbf{1}^\top \dot{s}(t) = 1$. Applying these observations to (40), it follows that

$$\frac{d}{dt} w_\zeta(n(t), q(t)) = \zeta^\top \Xi [\dot{q}(t) + \Gamma \operatorname{diag}(\mu^{-1})\dot{n}(t)] \geq 0.$$

Given that $(n(\cdot), q(\cdot))$ is Lipschitz continuous, the desired result follows immediately. \blacksquare

Lemma 5 guarantees the conservation of workload under any policy. This motivates the *effective cost* of a state $(n, q) \in \mathbb{R}_+^{\mathcal{F}} \times \mathbb{R}_+^{\mathcal{E}}$, defined by the linear program

$$\begin{aligned} c^*(n, q) &\triangleq \underset{n', q'}{\text{minimize}} && c(n', q') \\ &\text{subject to} && w_\zeta(n', q') \geq w_\zeta(n, q), \quad \forall \zeta \in \mathbf{CR}^*(\rho), \\ &&& n \in \mathbb{R}_+^{\mathcal{F}}, \quad q \in \mathbb{R}_+^{\mathcal{E}}. \end{aligned} \quad (41)$$

The effective cost is the lowest cost of any state with at least as much workload as (n, q) . We have the following lower bound on the cost achieved under any fluid trajectory:

Theorem 6. *Consider fluid model trajectory of system under any scheduling and rate allocation policy, with flow count and queue length processes given by $(n(\cdot), q(\cdot))$. Then, for any time $t \geq 0$, the instantaneous cost $c(n(t), q(t))$ is bounded below according to*

$$c^*(n(0), q(0)) \leq c(n(t), q(t)). \quad (42)$$

Proof. By Lemma 5, if the initial condition of a fluid trajectory satisfies $(n(0), q(0)) = (n, q)$, then $(n(t), q(t))$ is feasible for (41) for every $t \geq 0$. The result immediately follows. \blacksquare

7. Balanced Systems

In this section, we will develop a bound on the cost achieved in a fluid model solution under the MWUM- α policy. In particular, we will establish that this cost, at any instant of time, is within a constant factor of the cost achievable under *any* policy. The constant factor is uniform across the entire fluid trajectory, and relates to a notion of ‘balance’ of the critical resources of the network that we will describe shortly.

We begin with a preliminary lemma, that provides an upper bound on the cost under the MWUM- α policy. This upper bound is closely related to the Lyapunov function introduced earlier for studying the system stability.

Lemma 7. *Consider a fluid model trajectory of system under the MWUM- α policy, and denote the flow count and queue length processes by $(n(\cdot), q(\cdot))$. Suppose that the offered load vector ρ satisfies $L(\rho) \leq 1$. Then, at any time $t \geq 0$, it must be that*

$$c(n(t), q(t)) \leq (1 + \beta(\alpha))c(n(0), q(0)), \quad (43)$$

where $\beta(\alpha) \rightarrow 0$ as $\alpha \rightarrow 0^+$.

Proof. Recall the Lyapunov function L_α from (10). It follows from Lemma 3 that, so long as $L(\rho) \leq 1$,

$$L_\alpha(n(t), q(t)) \leq L_\alpha(n(0), q(0)). \quad (44)$$

Applying Lemma 14, with $p \triangleq 1 + \alpha$ and $d \triangleq |\mathcal{E}| + |\mathcal{F}|$,

$$\sum_{f \in \mathcal{F}} \frac{n_f(t)}{\mu_f} \left(\frac{\mu_f}{\rho_f} \right)^{\frac{\alpha}{1+\alpha}} + \sum_{e \in \mathcal{E}} q_e(t) \leq d^{\frac{\alpha}{1+\alpha}} \left[\sum_{f \in \mathcal{F}} \frac{n_f(0)}{\mu_f} \left(\frac{\mu_f}{\rho_f} \right)^{\frac{\alpha}{1+\alpha}} + \sum_{e \in \mathcal{E}} q_e(0) \right].$$

Now, as $\alpha \rightarrow 0^+$, $d^{\frac{\alpha}{1+\alpha}} \rightarrow 1$. Also, for each f ,

$$\left(\min_{f'} \frac{\mu_{f'}}{\rho_{f'}} \right)^{\frac{\alpha}{1+\alpha}} \leq \left(\frac{\mu_f}{\rho_f} \right)^{\frac{\alpha}{1+\alpha}} \leq \left(\max_{f'} \frac{\mu_{f'}}{\rho_{f'}} \right)^{\frac{\alpha}{1+\alpha}}. \quad (45)$$

Thus, as $\alpha \rightarrow 0^+$, $(\mu_f/\rho_f)^{\frac{\alpha}{1+\alpha}} \rightarrow 1$ uniformly over f . The result then follows. \blacksquare

The following definition is central to our performance guarantee:

Definition 5 (Balance Factor). *Given a system that is critically loaded with offered load vector ρ , define the **balance factor** as the value of the optimization problem*

$$\begin{aligned} \gamma(\rho) \triangleq & \underset{n, q, n', q'}{\text{minimize}} && c(n', q') \\ & \text{subject to} && w_\zeta(n', q') \geq w_\zeta(n, q), \quad \forall \zeta \in \text{CR}^*(\rho), \\ & && c(n, q) = 1, \\ & && n, n' \in \mathbb{R}_+^{\mathcal{F}}, \quad q, q' \in \mathbb{R}_+^{\mathcal{E}}. \end{aligned}$$

We can interpret the balance factor as follows: It is clear that $\gamma(\rho) \geq 0$, since $n', q' \geq \mathbf{0}$. Since there are feasible solutions with $(n, q) = (n', q')$, it is also true that $\gamma(\rho) \leq 1$. In order to interpret

$\gamma(\rho)$, assume for the moment that there is only a single critical extreme resource $\zeta \in \text{CR}^*(\rho)$. If we define $v \triangleq \Xi^\top \zeta$, then the constraint that $w_\zeta(n', q') \geq w_\zeta(n, q)$ is equivalent to

$$v^\top [\Gamma \text{diag}(\mu)^{-1} n' + q'] \geq v^\top [\Gamma \text{diag}(\mu)^{-1} n + q].$$

In this case, it is clear that the solution to the LP defining $\gamma(\rho)$ is given by

$$\gamma(\rho) = (\min_e v_e) / (\max_e v_e).$$

Hence, $\gamma(\rho)$ is the measure of the degree of ‘balance’ of the influence of the critical resource ζ across buffers in the network. In other words, in networks that are more balanced, queues will participate more equally in the critical resource.

In the more general case, there may be multiple critical extreme resources (i.e., $|\text{CR}^*(\rho)| \geq 1$). Define the set $\mathcal{V} \triangleq \text{span} \{\Xi^\top \zeta : \zeta \in \text{CR}^*(\rho)\}$. It is not difficult to see that $\gamma(\rho) > 0$ if and only if, for each queue $e \in \mathcal{V}$, there exists some $v \in \mathcal{V}$ with $v_e > 0$, i.e., if every queue is influenced by *some* critical resource. We call networks where $\gamma(\rho) > 0$ *balanced*. In the extreme, if $\mathbf{1} \in \mathcal{V}$, then $\gamma(\rho) = 1$.

The following is the main theorem of this section. It offers a bound on the cost incurred at any instant in time under the MWUM- α policy, relative that incurred under *any* other policy. This bound is a function of the balance factor.

Theorem 8. *Consider fluid model trajectory of a critically loaded system under the MWUM- α policy and denote the flow count and queue length processes by $(n(\cdot), q(\cdot))$. Suppose that $\gamma(\rho) > 0$. Let $(n'(\cdot), q'(\cdot))$ be the flow count and queue length policies under an arbitrary policy given the same initial conditions, i.e., $n(0) = n'(0)$ and $q(0) = q'(0)$. Then, at any time $t \geq 0$, it must be that*

$$c(n(t), q(t)) \leq \frac{1 + \beta(\alpha)}{\gamma(\rho)} c(n'(t), q'(t)), \quad (46)$$

where $\beta(\alpha) \rightarrow 0$ as $\alpha \rightarrow 0^+$.

Proof. First, note that if $(n(0), q(0)) = \mathbf{0}$, i.e., the system is empty, then this holds for all $t \geq 0$ (cf. Theorem 10). In this case, (46) is immediate. Otherwise, fix $t \geq 0$, and set $\bar{c} \triangleq c(n(0), q(0)) > 0$. Define

$$(n', q') \triangleq (n'(t), q'(t)) / \bar{c}, \quad (n, q) \triangleq (n(0), q(0)) / \bar{c}.$$

Using Lemma 5, it is clear that (n, q, n', q') is feasible for the LP defining $\gamma(\rho)$. Thus,

$$c(n(0), q(0)) \leq \frac{1}{\gamma(\rho)} c(n'(t), q'(t)).$$

The result then follows by applying Lemma 7. ■

8. Invariant Manifold

In Section 7, we proved a constant factor guarantee on the cost of the MWUM- α policy, relative to the cost achieved under any other policy. Our bound held point-wise, at *every* instant of time. However, the constant factor of the bound depends on the balance factor, and this could be very large.

In this section, we consider a different type of analysis. Instead of considering the evolution of the fluid model for *every* time t , we instead examine the asymptotically limiting states of the fluid model as $t \rightarrow \infty$. In particular, we characterize these *invariant* states as fixed points in the solution space of an optimization problem. We shall also show that these fixed points are attractive, i.e., starting from any initial condition, the fluid trajectory reaches an invariant state. We will quantify time to converge to the invariant manifold as a function of the initial conditions of the fluid trajectory.

This characterization of invariant states is key towards establishing the *state space collapse* property of the system under a heavy traffic limit [4]. Moreover, we shall demonstrate that these invariant states are cost optimal as $\alpha \rightarrow 0^+$. In other words, the cost of an invariant state cannot be improved by *any* policy.

8.1. Optimization Problems

We start with two useful optimization problems that will be useful in characterizing invariant states of the fluid trajectory. We assume that the system is critically loaded.

Suppose we are given a state $(n, q) \in \mathbb{R}_+^{\mathcal{F}} \times \mathbb{R}_+^{\mathcal{E}}$ of, respectively, flow counts and queue lengths. Define the optimization problem

$$\begin{aligned} \text{ALGP}(n, q) \triangleq & \underset{n', q', t, x, \sigma}{\text{minimize}} && L_\alpha(n', q') \\ & \text{subject to} && n' = n + t[\nu - \text{diag}(\mu)x], \\ & && q' = q + t[\Gamma x - (I - R^\top)\sigma], \\ & && n' \in \mathbb{R}_+^{\mathcal{F}}, \quad q' \in \mathbb{R}_+^{\mathcal{E}}, \quad t \in \mathbb{R}_+, \\ & && x \in [0, C]^{\mathcal{F}}, \quad \sigma \in \Lambda. \end{aligned}$$

Here, recall that Λ is the scheduling capacity region of the network, defined by (9). Similarly, define the optimization problem

$$\begin{aligned} \text{bALGD}(n, q) \triangleq & \underset{n', q'}{\text{minimize}} && L_\alpha(n', q') \\ & \text{subject to} && w_\zeta(n', q') \geq w_\zeta(n, q), \\ & && \forall \zeta \in \text{CR}^*(\rho), \\ & && n' \in \mathbb{R}_+^{\mathcal{F}}, \quad q' \in \mathbb{R}_+^{\mathcal{E}}. \end{aligned}$$

Intuitively, given a state (n, q) , $\text{ALGP}(n, q)$ finds a state (n', q') which minimizes the Lyapunov function L_α and can be reached starting from (n, q) , using feasible scheduling and rate allocation decisions. $\text{ALGP}(n, q)$, on the other hand, finds a state (n', q') which minimizes the Lyapunov function and has at least as much workload as (n, q) . The following result states that $\text{ALGP}(n, q)$ and $\text{bALGD}(n, q)$ are equivalent optimization problems:

Lemma 9. *A state $(n', q') \in \mathbb{R}_+^{\mathcal{F}} \times \mathbb{R}_+^{\mathcal{E}}$ is feasible for the optimization problem $\text{ALGP}(n, q)$ if and only if it is feasible for the optimization problem $\text{bALGD}(n, q)$.*

Proof. First, consider any (n', q', t, x, σ) that is feasible for $\text{ALGP}(n, q)$. Note that feasibility for $\text{ALGP}(n, q)$ implies that

$$\Gamma \text{diag}(\mu)^{-1}n' + q' \geq \Gamma \text{diag}(\mu)^{-1}n + q + t[\Gamma \text{diag}(\mu)^{-1}\nu - \Gamma x] + t[\Gamma x - (I - R^\top)\sigma].$$

Therefore, if $\zeta \in \text{CR}^*(\rho)$, we have that

$$w_\zeta(n', q') = w_\zeta(n, q) + t \left[\zeta^\top \Xi \Gamma \rho - \zeta^\top \sigma \right].$$

Since $\sigma \in \Lambda$ and ζ is feasible for $\text{DUAL}(\Xi \Gamma \rho)$, we have $\zeta^\top \sigma \leq 1$. Since $\zeta \in \text{CR}^*(\rho)$, we have $\zeta^\top \Xi \Gamma \rho = 1$. Therefore, as $t \geq 0$, it follows that

$$w_\zeta(n', q') \geq w_\zeta(n, q).$$

That is, (n', q') is $\text{bALGD}(n, q)$ feasible.

Next, assume that (n', q') is feasible for $\text{bALGD}(n, q)$. Given some $t \geq 0$, define

$$x \triangleq \text{diag}(\mu)^{-1} [\nu - t^{-1}(n' - n)], \quad \sigma \triangleq \Xi [\Gamma x - t^{-1}(q' - q)].$$

With these definitions, if we establish existence of $t \geq 0$ so that $\mathbf{0} \leq x \leq C\mathbf{1}$ and $\sigma \in \Lambda$, then (n', q', t, x, σ) is feasible for $\text{ALGP}(n, q)$ feasible.

Note that as $t \rightarrow \infty$, $x \rightarrow \rho$. By assumption, $0 < \rho_f < C$, for all $f \in \mathcal{F}$. Therefore, for t sufficiently large, $\mathbf{0} \leq x \leq C\mathbf{1}$.

Next, we wish to show that, for t sufficiently large, $\sigma \in \Lambda$. This requirement is equivalent to demonstrating that $\text{PRIMAL}(\sigma) \leq 1$ and that $\sigma \geq \mathbf{0}$. To show that $\text{PRIMAL}(\sigma) \leq 1$, note that $\text{PRIMAL}(\sigma) = \text{DUAL}(\sigma)$ and suppose that ζ is feasible for $\text{DUAL}(\sigma)$. Then,

$$\begin{aligned} \zeta^\top \sigma &= \zeta^\top [\Xi \Gamma x - t^{-1}(q' - q)] \\ &= \zeta^\top [\Xi \Gamma \rho - t^{-1} \Xi \Gamma \text{diag}(\mu)^{-1}(n' - n) - t^{-1}(q' - q)] \\ &= \zeta^\top \Xi \Gamma \rho - t^{-1} [w_\zeta(n', q') - w_\zeta(n, q)]. \end{aligned}$$

If $\zeta \in \text{CR}(\rho)$, then

$$\zeta^\top \Xi \Gamma \rho = 1, \quad \text{and} \quad w_\zeta(n', q') - w_\zeta(n, q) \geq 0,$$

thus $\zeta^\top \sigma \leq 1$. On the other hand, if $\zeta \notin \text{CR}(\rho)$, $\zeta^\top \Xi \Gamma \rho < 1$. Therefore, in any event, for t sufficiently large, $\text{DUAL}(\rho) \leq 1$.

To show that $\sigma \geq \mathbf{0}$, note that

$$\begin{aligned} \sigma &= \Xi [\Gamma x - t^{-1}(q' - q)] \\ &= \Xi \Gamma \rho - t^{-1} [\Xi \Gamma \text{diag}(\mu)^{-1}(n' - n) + \Xi(q' - q)]. \end{aligned}$$

By assumption, $\Xi \Gamma \rho > \mathbf{0}$. Therefore, for t sufficiently large enough, $\sigma \geq \mathbf{0}$. ■

8.2. Fixed Points: Characterization

Note that the optimization problem $\text{bALGD}(n, q)$ has a convex feasible set with a strictly convex and coercive objective function (see, e.g., [1]). By standard arguments from theory of convex optimization, it follows that an optimal solution exists and is unique. Hence, we can make the following definition:

Definition 6 (Lifting Map). *Given a critically scaled system, we define the **lifting map** $\Delta: \mathbb{R}_+^{\mathcal{F}} \times \mathbb{R}_+^{\mathcal{E}} \rightarrow \mathbb{R}_+^{\mathcal{F}} \times \mathbb{R}_+^{\mathcal{E}}$ to be the function that maps a state (n, q) to the unique solution of the optimization problem $\text{bALGD}(n, q)$.*

The main result of this section is to characterize the invariant states of fluid model as the fixed points of lifting map Δ .

Theorem 10. *A state $(n, q) \in \mathbb{R}_+^{\mathcal{F}} \times \mathbb{R}_+^{\mathcal{E}}$ is an invariant state of a fluid model solution under the MWUM- α policy if and only if it is a fixed point of Δ , i.e.,*

$$(n, q) = \Delta(n, q).$$

Proof. The proof follows by establishing equivalence of the following statements, for every state (n, q) :

- (i) $(n, q) = \Delta(n, q)$.
- (ii) Any fluid model solution satisfying the initial condition $(n(0), q(0)) = (n, q)$ has $(n(t), q(t)) = (n, q)$ for all t .
- (iii) There exists a fluid model solution with $(n(t), q(t)) = (n, q)$ for all t .
- (iv) (n, q) satisfy

$$(\Gamma\rho)^\top q^\alpha = \max_{\pi \in \mathcal{S}} \pi^\top (I - R)q^\alpha, \quad (47)$$

$$\rho_f q_{\iota(f)} = n_f, \quad \forall f \in \mathcal{F}. \quad (48)$$

(i) \Rightarrow (ii): If $(n, q) = \Delta(n, q)$, then it solves $b\text{ALGD}(n, q)$. Consider a fluid model solution with an initial state $(n(0), q(0)) = (n, q)$. By Lemma 3, it follows that, for all t , $L_\alpha(n(t), q(t)) \leq L_\alpha(n, q)$. From the fluid model equations (F1)–(F12), $(n(t), q(t))$ is $\text{ALGP}(n, q)$ feasible, for all t . Therefore, it follows that $(n(t), q(t))$ is an optimal solution of $\text{ALGP}(n, q)$, and, by Lemma 9, of $b\text{ALGD}(n, q)$. Since $b\text{ALGD}(n, q)$ has (n, q) as its unique solution, it follows that $(n(t), q(t)) = (n, q)$, for all t .

(ii) \Rightarrow (iii): This follows in a straightforward manner by considering the arguments in Theorem 1 with initial conditions given by (n, q) .

(iii) \Rightarrow (iv): Consider a fluid model solution that satisfies $(n(t), q(t)) = (n, q)$, for all t . Then, for any regular point t , we have $\dot{n}(t) = \mathbf{0}$ and $\dot{q}(t) = \mathbf{0}$. Using (F2)–(F4), it follows that $x(t) \triangleq \dot{\tilde{x}}(t) = \rho$. For any $f \in \mathcal{F}$, if $n_f = n_f(t) > 0$ and $x_f(t) = \rho_f < C$, then by (F9) it must be that $x_f(t) = n_f(t)/q_{\iota(f)}(t)$. Therefore, $\rho_f q_{\iota(f)} = n_f$. Similarly, if $n_f = 0$, it must be that $q_{\iota(f)} = 0$ by (F11).

Now, define $H(t) \triangleq \mathbf{1}^\top q^{1+\alpha}(t)$. Since $q(\cdot)$ is constant, applying (F5), (F6), (F12), it must be that for every regular t ,

$$0 = \dot{H}(t) = \dot{q}(t)^\top q^\alpha(t) = \left[\Gamma\rho - \left(I - R^\top \right) \Pi \dot{s}(t) \right]^\top q^\alpha(t).$$

Applying (F7) and (F10),

$$0 = (\Gamma\rho)^\top q^\alpha - \max_{\pi \in \mathcal{S}} \pi^\top (I - R)q^\alpha.$$

(iv) \Rightarrow (i): Suppose (n, q) satisfy (47)–(48). Define $(n', q') \triangleq \Delta(n, q)$. Since (n', q') solves the optimization problem $b\text{ALGD}(n, q)$, by Lemma 9, there exists (t, x, σ) so that (n', q', t, x, σ) is an optimal solution for $\text{ALGP}(n, q)$. This solution must satisfy

$$n' = n + t[\nu - \text{diag}(\mu)x], \quad q' = q + t \left[\Gamma x - \left(I - R^\top \right) \sigma \right].$$

Now consider the trajectory

$$(n(\tau), q(\tau)) \triangleq (n, q) + \frac{\tau}{t}(n' - n, q' - q), \quad \forall \tau \in [0, t].$$

Define J to be the Lyapunov function L_α evaluated along this path, i.e., $J(\tau) \triangleq L_\alpha(n(\tau), q(\tau))$. Then,

$$\begin{aligned} \frac{\dot{J}(0)}{1 + \alpha} &= \sum_{f \in \mathcal{F}} \frac{n_f^\alpha (\nu_f - \mu_f x_f)}{\mu_f \rho_f^\alpha} + (\Gamma x)^\top q^\alpha - \sigma^\top (I - R) q^\alpha \\ &= \underbrace{\left(\sum_{f \in \mathcal{F}} \frac{n_f^\alpha (\nu_f - \mu_f x_f)}{\mu_f \rho_f^\alpha} + (\Gamma \delta)^\top q^\alpha \right)}_X + \underbrace{\left((\Gamma \rho)^\top q^\alpha - \sigma^\top (I - R) q^\alpha \right)}_Y, \end{aligned}$$

where $\delta \triangleq x - \rho$.

First, consider Y . Since $\sigma \in \Lambda$, there exists some $s \in \mathbb{R}_+^S$ with $\mathbf{1}^\top s \leq 1$ and $\sigma \leq \Pi s$. From the monotonicity of \mathcal{S} , we can pick s so that $\sigma = \Pi s$. Therefore,

$$\sigma^\top (I - R) q^\alpha = s^\top \Pi^\top (I - R) q^\alpha \leq \max_{\pi \in \mathcal{S}} \pi^\top (I - R) q^\alpha.$$

Then, by (47), it follows that $Y \geq 0$. Now, consider X , and note that $X = 0$ by (48) along with the fact that

$$X = \sum_{f \in \mathcal{F}} \left(\frac{n_f^\alpha (\rho_f - x_f)}{\rho_f^\alpha} + \delta_f q_{i(f)}^\alpha \right) = \sum_{f \in \mathcal{F}} \delta_f \left(q_{i(f)}^\alpha - \frac{n_f^\alpha}{\rho_f^\alpha} \right). \quad (49)$$

Thus, we have that $\dot{J}(0) \geq 0$. Since $J(\tau)$ is a convex function, this implies that $J(0) \leq J(t)$, i.e., $L_\alpha(n, q) \leq L_\alpha(n', q')$. Due to uniqueness of the optimal solution to $b\text{ALGD}(n, q)$, it follows that $(n', q') = (n, q)$. \blacksquare

8.3. Fixed Points: Attractiveness

We will now establish the attractiveness of the space of fixed points. Specifically, we will show that starting from any initial state, the fluid trajectory converges (arbitrarily close to) space of fixed points, in finite time.

Given $\varepsilon > 0$, define

$$\mathcal{J}_\varepsilon \triangleq \{(n, q) \in \mathbb{R}_+^{\mathcal{F}} \times \mathbb{R}_+^{\mathcal{E}} : \|(n, q) - \Delta(n, q)\|_1 < \varepsilon\}.$$

In other words, \mathcal{J}_ε is the set of states (n, q) which are ε -approximate fixed points (in an ℓ_1 -norm sense) of the lifting map. Given a fluid trajectory $(n(\cdot), q(\cdot))$, define

$$h_\varepsilon(n(\cdot), q(\cdot)) \triangleq \inf \{t \geq 0 : (n(s), q(s)) \in \mathcal{J}_\varepsilon, \forall s \geq t\}.$$

In other words, $h_\varepsilon(n(\cdot), q(\cdot))$ is the amount of time required for the trajectory $(n(\cdot), q(\cdot))$ to reach and subsequently remain in the set \mathcal{J}_ε .

Theorem 11. For any $\varepsilon > 0$, there exists $H_\varepsilon > 0$ so that if $(n(\cdot), q(\cdot))$ is a fluid trajectory of the MWUM- α policy in a critically loaded system, with initial condition satisfying $\|(n(0), q(0))\|_\infty \leq 1$, then

$$h_\varepsilon(n(\cdot), q(\cdot)) \leq H_\varepsilon.$$

In order to prove Theorem 11, we require the following technical lemma:

Lemma 12. Under the MWUM- α policy, the lifting map Δ is continuous. Further, Δ is also positively homogeneous, i.e., for all $(n, q) \in \mathbb{R}_+^{\mathcal{F}} \times \mathbb{R}_+^{\mathcal{E}}$ and $\kappa > 0$,

$$\Delta(\kappa n, \kappa q) = \kappa \Delta(n, q).$$

Proof. To establish continuity, it suffices to prove that if $(n^k, q^k) \rightarrow (n, q)$, then $\Delta(n^k, q^k) \rightarrow \Delta(n, q)$. By definition of Δ , we have that $L_\alpha(\Delta(n, q)) \leq L_\alpha(n, q)$. Therefore, since the convergent sequence $\{(n^k, q^k)\}$ lies in a compact set, the sequence $\{\Delta(n^k, q^k)\}$ is contained in a compact set also. Define $x^k \triangleq \Delta(n^k, q^k)$ and $x \triangleq \Delta(n, q)$. Then, there exists a convergence subsequence of $\{x^k\}$ that converges to some \hat{x} . In what follows, that $\hat{x} = x$, and thereby complete the proof of continuity of Δ .

Suppose that $\hat{x} \neq x$. By passing to a subsequence, without loss of generality, assume that $x^k \rightarrow \hat{x}$. Since x^k is feasible for $b\text{ALGD}(n^k, q^k)$, and $(n^k, q^k) \rightarrow (n, q)$, it follows that \hat{x} is feasible for $b\text{ALGD}(n, q)$. Since x is the unique optimal solution to $b\text{ALGD}(n, q)$, we have that $L_\alpha(x) < L_\alpha(\hat{x})$.

Now, define

$$\varepsilon^k \triangleq \left(\max_{\zeta \in \text{CR}^*(\rho)} \frac{w_\zeta(n^k, q^k) - w_\zeta(n, q)}{w_\zeta(\mathbf{1})} \right)^+,$$

where $(x)^+ \triangleq \max(0, x)$. Consider $\tilde{x}^k \triangleq x + \varepsilon^k \mathbf{1}$. By the definition of ε^k , it follows that \tilde{x}^k is feasible for $b\text{ALGD}(n^k, q^k)$. Then, $L_\alpha(x^k) \leq L_\alpha(\tilde{x}^k)$. Now, as $k \rightarrow \infty$, $L_\alpha(x^k) \rightarrow L_\alpha(x)$. Further, $\varepsilon^k \rightarrow 0$, so $\tilde{x}^k \rightarrow x$ and $L_\alpha(\tilde{x}^k) \rightarrow L_\alpha(x)$. Then, $L_\alpha(\hat{x}) \leq L_\alpha(x)$. By contradiction, this establishes the continuity of Δ .

The positive homogeneity of Δ follows directly from the definition of the optimization problem $b\text{ALGD}$. ■

Proof of Theorem 11. Given $\delta > 0$, define

$$\begin{aligned} \mathcal{D} &\triangleq \{(n, q) \in \mathbb{R}_+^{\mathcal{F}} \times \mathbb{R}_+^{\mathcal{E}} : L_\alpha(n, q) \leq L_\alpha(\mathbf{1})\}, \\ \mathcal{I} &\triangleq \{(n, q) \in \mathcal{D} : (n, q) = \Delta(n, q)\}, \\ \mathcal{I}_\delta &\triangleq \{(n, q) \in \mathcal{D} : \|(n, q) - (n', q')\|_1 < \delta, (n', q') \in \mathcal{I}\}, \\ \mathcal{K}_\delta &\triangleq \{(n, q) \in \mathcal{D} : K(n, q) < K(n', q'), \forall (n', q') \in \mathcal{D} \setminus \mathcal{I}_\delta\}. \end{aligned}$$

where $K(n, q) \triangleq L_\alpha(n, q) - L_\alpha(\Delta(n, q))$. The result can be established by showing that the following hold:

- (i) $K(n(t), q(t))$ is non-increasing in t .
- (ii) For $\delta > 0$ sufficiently small, $\mathcal{I} \subset \mathcal{K}_\delta \subset \mathcal{I}_\delta \subset \mathcal{J}_\varepsilon$.
- (iii) Starting from any initial condition in \mathcal{D} (this includes all (n, q) with $\|(n, q)\|_\infty \leq 1$), the time to hit \mathcal{K}_δ is bounded uniformly.

In particular, (iii) implies that starting from any state in \mathcal{D} , the fluid trajectory hits the set \mathcal{K}_δ in finite time. By (i), once the trajectory is in set \mathcal{K}_δ , it remains in that set forever. By (ii), $\mathcal{K}_\delta \subset \mathcal{J}_\varepsilon$, and the result follows. To complete the proof, (i), (ii) and (iii) need to be justified.

(i): Since L_α is a Lyapunov function, then $L_\alpha(n(t), q(t))$ is non-increasing over time under any fluid trajectory (cf. Lemma 3). From Lemma 5, the constraints in the optimization problem $bALGD(n(t), q(t))$ can only become more restrictive over time. Therefore, as the time t increases, the cost of the optimal solution of $bALGD(n(t), q(t))$, i.e., $L_\alpha(\Delta(n(t), q(t)))$, can only be non-decreasing. Therefore, $K(n(t), q(t))$ is non-increasing over time.

(ii): First, consider claim $\mathcal{I} \subset \mathcal{K}_\delta$ for any $\delta > 0$. Δ is continuous by Lemma 12. The constraints, one for each $\zeta \in CR^*(\rho)$, in $bALGD(n, q)$ are continuous with respect to (n, q) . And $(n(t), q(t))$ continuous with over t . Therefore, both functions $L_\alpha(n(t), q(t))$ and $K(n(t), q(t))$ are continuous with respect to t . Now, \mathcal{D} is closed and bounded and \mathcal{I}_δ is open, hence $\mathcal{D} \setminus \mathcal{I}_\delta$ is closed and bounded. Therefore, the infimum of the continuous function K is achieved over this set. Since $\mathcal{I} \subset \mathcal{I}_\delta$ for any $\delta > 0$, by the definition of \mathcal{I} , this this infimum must be strictly positive. However, over \mathcal{I} value of K is 0. Therefore, it follows that $\mathcal{I} \subset \mathcal{K}_\delta$. The claim that $\mathcal{K}_\delta \subset \mathcal{I}_\delta$ is trivial since if, $(n, q) \in \mathcal{D} \setminus \mathcal{I}_\delta$ then $K(n, q)$ is greater than or equal to the infimum over that set, hence $(n, q) \notin \mathcal{K}_\delta$. Finally, to establish that $\mathcal{I}_\delta \subset \mathcal{J}_\varepsilon$, recall again that Δ is continuous and, hence, uniformly continuous over \mathcal{D} . Therefore, for any $\varepsilon > 0$ there exists $\delta > 0$ such that

$$\|(n, q) - (n', q')\|_1 < \delta \quad \Rightarrow \quad \|\Delta(n, q) - \Delta(n', q')\|_1 < \varepsilon/2.$$

Consider any $(n, q) \in \mathcal{I}_\delta$ and $(n', q') \in \mathcal{I}$ with $\|(n, q) - (n', q')\|_1 < \delta$. Then,

$$\begin{aligned} \|(n, q) - \Delta(n, q)\|_1 &\leq \|(n, q) - (n', q')\|_1 + \|(n', q') - \Delta(n', q')\|_1 + \|\Delta(n', q') - \Delta(n, q)\|_1 \\ &\leq \delta + 0 + \varepsilon/2 \\ &< \varepsilon, \end{aligned}$$

for small enough choice of δ . This completes the proof of (ii).

(iii): Here, we shall use Theorem 10. First observe that, by the definition of K and Lemma 5, for all $0 \leq s \leq t$,

$$K(n(s), q(s)) - K(n(t), q(t)) \geq L_\alpha(n(s), q(s)) - L_\alpha(n(t), q(t)).$$

In other words, the decrease in K is at least as much as decrease in L .

Next, we wish to argue that when the fluid trajectory belongs to the set $\mathcal{D} \setminus \mathcal{K}_\delta$ (i.e., is away from the space of fixed points \mathcal{I}) then L_α is strictly decreasing at some minimal rate. To be precise, given $(n, q) \in \mathcal{D} \setminus \mathcal{K}_\delta$, suppose that $(n(\cdot), q(\cdot))$ is a fluid model solution and t a regular point such that $(n(t), q(t)) = (n, q)$. We would like to show that

$$\mathfrak{D}(n, q) \triangleq \frac{1}{1+\alpha} \frac{d}{dt} L_\alpha(n(t), q(t)) \leq -\gamma, \quad (50)$$

for some $\gamma > 0$.

To this end, for each $f \in \mathcal{F}$, define the function $x_f: \mathbb{R}_+^{\mathcal{F}} \times \mathbb{R}_+^{\mathcal{E}} \rightarrow [0, C]$ as

$$x_f(n, q) \triangleq \begin{cases} \rho_f & \text{if } n_f = 0, \\ n_f(t)/q_{i(f)}(t) & \text{if } 0 < n_f < Cq_{i(f)}, \\ C & \text{otherwise.} \end{cases}$$

Examining (F9), it is clear that this function determines the rate allocated to f at time t , i.e., $x_f(t) = x_f(n, q)$. Recalling (17)–(24), we have that

$$\begin{aligned} \mathfrak{D}(n, q) &= \sum_{f \in \mathcal{F}} \frac{n_f^\alpha}{\mu_f \rho_f^\alpha} \dot{n}_f(t) + \sum_{e \in \mathcal{E}} q_e^\alpha \dot{q}_e(t) \\ &= \sum_{f \in \mathcal{F}} \frac{n_f^\alpha}{\rho_f^\alpha} (\rho_f - x_f(n, q)) + [\Gamma x(n, q)]^\top q^\alpha - \max_{\sigma \in \mathcal{S}} \sigma^\top (I - R) q^\alpha \\ &= \sum_{f \in \mathcal{F}} \left(\frac{n_f^\alpha}{\rho_f^\alpha} (\rho_f - x_f(n, q)) + q_{i(f)}^\alpha (x_f(n, q) - \rho_f) \right) + [\Gamma \rho]^\top q^\alpha - \max_{\sigma \in \mathcal{S}} \sigma^\top (I - R) q^\alpha \\ &= \sum_{f \in \mathcal{F}} \left(\frac{n_f^\alpha}{\rho_f^\alpha} - q_{i(f)}^\alpha \right) (\rho_f - x_f(n, q)) + T(q). \end{aligned} \tag{51}$$

Here, for convenience, we define the function $T: \mathbb{R}_+^{\mathcal{E}} \rightarrow \mathbb{R}$ by

$$T(q) \triangleq [\Gamma \rho]^\top q^\alpha - \max_{\sigma \in \mathcal{S}} \sigma^\top (I - R) q^\alpha.$$

Now, recall that $\rho_f < C$ for all $f \in \mathcal{F}$. Therefore, it follows that, for all $f \in \mathcal{F}$, if $n_f > 0$,

$$x_f(n, q) < \rho_f \iff n_f / \rho_f < q_{i(f)}(t).$$

Therefore, for all $f \in \mathcal{F}$,

$$\left(\frac{n_f^\alpha}{\rho_f^\alpha} - q_{i(f)}^\alpha \right) (\rho_f - x_f(n, q)) \leq 0, \tag{52}$$

with the inequality being strict if $n_f \neq \rho_f q_{i(f)}$ and $n_f > 0$.

Since $(n, q) \in \mathcal{D} \setminus \mathcal{K}_\delta$, it can not be a fixed point. Therefore, by part (iv) of the equivalence established in Theorem 10, one of the following two conditions holds:

- $T(q) < 0$.
- $T(q) = 0$, and there exists some $f \in \mathcal{F}$ with $n_f \neq \rho_f q_{i(f)}$ with $n_f > 0$.

Note that the $n_f > 0$ requirement of the second case follows from (F11): $n_f = 0$ would imply that $q_{i(f)} = 0$, hence if $n_f \neq \rho_f q_{i(f)}$, it must be that $n_f \neq 0$. In either of the above two cases, using (52) it is easy to see that the right hand side of (51) is strictly negative. However, this does not provide a uniform, strictly negative, bound on the drift \mathfrak{D} . If \mathfrak{D} were established to be a continuous function over set $\mathcal{D} \setminus \mathcal{K}_\delta$, then such a uniform negative bound would follow as the set $\mathcal{D} \setminus \mathcal{K}_\delta$ is closed and bounded. However, closer examination (51) reveals that \mathfrak{D} depends on the function x_f , thus is not necessarily continuous at the boundary $n_f = 0$, for any $f \in \mathcal{F}$.

This difficulty is overcome as follows. We will cover $\mathcal{D} \setminus \mathcal{K}_\delta$ by a finite collection of closed and bounded sets. On each set, we will obtain a bound on the drift \mathfrak{D} that is continuous on the set

as well as strictly negative. Hence, we will conclude that \mathfrak{D} is uniformly bounded by a negative quantity on $\mathcal{D} \setminus \mathcal{K}_\delta$, i.e., that (50) holds. The details are given next.

To begin, define the function $\mathfrak{R}: \mathbb{R}_+ \times \mathbb{R}_+ \rightarrow [0, C]$ by

$$\mathfrak{R}(n_f, q_{i(f)}) \triangleq \begin{cases} n_f/q_{i(f)} & \text{if } n_f < Cq_{i(f)}, \\ C & \text{otherwise.} \end{cases}$$

For given a vector $\mathbf{b} = [\mathbf{b}_f] \in \{-1, 0, 1\}^{|\mathcal{F}|}$, define $\mathbf{S}_\mathbf{b}$ as to be the set of $(n, q) \in \mathcal{D} \setminus \mathcal{K}_\delta$ such that, for each f ,

$$\begin{aligned} \rho_f - \mathfrak{R}(n_f, q_{i(f)}) &\geq n_f, & \text{if } \mathbf{b}_f = -1, \\ |\mathfrak{R}(n_f, q_{i(f)}) - \rho_f| &\leq n_f, & \text{if } \mathbf{b}_f = 0, \\ \mathfrak{R}(n_f, q_{i(f)}) - \rho_f &\geq n_f, & \text{if } \mathbf{b}_f = 1. \end{aligned}$$

Clearly $\mathbf{S}_\mathbf{b}$ is a closed and bounded set, since \mathfrak{R} is continuous. Given $f \in \mathcal{F}$ with $\mathbf{b}_f \neq 0$, define the function $g_f: \mathbf{S}_\mathbf{b} \rightarrow \mathbb{R}$ by

$$g_f(n, q) \triangleq \begin{cases} \rho_f^\alpha - (\rho_f - n_f)^\alpha & \text{if } \mathbf{b}_f = -1, \\ (n_f + \rho_f)^\alpha - \rho_f^\alpha & \text{if } \mathbf{b}_f = 1. \end{cases}$$

It is easy to check that for all $(n, q) \in \mathbf{S}_\mathbf{b}$,

$$|\mathfrak{R}^\alpha(n_f, q_{i(f)}) - \rho_f^\alpha| \geq g_f(n, q) \geq 0. \quad (53)$$

Further, g_f is continuous and $g_f(n, q) = 0$ if and only if $n_f = 0$. Finally, define the function $F_\mathbf{b}: \mathbf{S}_\mathbf{b} \rightarrow \mathbb{R}$ as

$$\begin{aligned} F_\mathbf{b}(n, q) &\triangleq T(q) + \sum_{f: \mathbf{b}_f=0} \left(\frac{n_f^\alpha}{\rho_f^\alpha} - q_{i(f)}^\alpha \right) (\rho_f - x_f(n, q)) \\ &\quad - \sum_{f: \mathbf{b}_f \neq 0} \min \left((C - \rho_f)(\rho_f^{-\alpha} - C^{-\alpha})n_f^\alpha, \frac{n_f^{1+\alpha}g_f(n, q)}{C^\alpha \rho_f^\alpha} \right). \end{aligned}$$

We make the following claims:

- (a) $F_\mathbf{b}$ is a continuous function over $\mathbf{S}_\mathbf{b}$.
- (b) For any $(n, q) \in \mathbf{S}_\mathbf{b}$, $\mathfrak{D}(n, q) \leq F_\mathbf{b}(n, q)$.
- (c) For any $(n, q) \in \mathbf{S}_\mathbf{b}$, $F_\mathbf{b}(n, q) < 0$.

(a): To establish this claim, it is sufficient to observe that for all f with $\mathbf{b}_f = 0$, x_f is a continuous function over $\mathbf{S}_\mathbf{b}$. To see this, note that if $\mathbf{b}_f = 0$, then $|\mathfrak{R}(n_f, q_{i(f)}) - \rho_f| \leq n_f$. Now $x_f(n, q) = \mathfrak{R}(n_f, q_{i(f)})$ if $n_f > 0$ and $x_f(n, q) = \rho_f$ if $n_f = 0$. Therefore, over $\mathbf{S}_\mathbf{b}$, we have that $x_f(n, q) = \mathfrak{R}(n_f, q_{i(f)})$ for all $n_f \geq 0$. This establishes continuity of x_f for f with $\mathbf{b}_f = 0$.

(b): Here, we need to show that for any f with $\mathbf{b}_f \neq 0$, the term in $F_\mathbf{b}$ is larger than or equal to the corresponding term on the right hand side of (51) in magnitude and preserves the sign. That is,

$$\left| \frac{n_f^\alpha}{\rho_f^\alpha} - q_{i(f)}^\alpha \right| |\rho_f - x_f(n, q)| \geq \min \left((C - \rho_f)(\rho_f^{-\alpha} - C^{-\alpha})n_f^\alpha, \frac{n_f^{1+\alpha}g_f(n, q)}{C^\alpha \rho_f^\alpha} \right), \quad (54)$$

and

$$\left(\frac{n_f^\alpha}{\rho_f^\alpha} - q_{i(f)}^\alpha \right) (\rho_f - x_f(n, q)) = 0 \iff n_f = 0. \quad (55)$$

Now, if $n_f = 0$ then $x_f(n, q) = \rho_f$ and hence the left hand side above of (55) is 0. If $n_f > 0$, since $\mathbf{b}_f \neq 0$ and thus $|\mathfrak{R}(n_f, q_{i(f)}) - \rho_f| \geq n_f > 0$, we have $x_f \neq \rho_f$. Therefore, the left hand side of (55) is not equal to 0. Thus, (55) is established.

To prove (54), we have the following cases:

- $n_f = 0$. Here, both sides of (54) are 0.
- $0 < n_f \leq Cq_{i(f)}$. Here, we have

$$\begin{aligned} \left| \frac{n_f^\alpha}{\rho_f^\alpha} - q_{i(f)}^\alpha \right| |\rho_f - x_f(n, q)| &= \frac{q_{i(f)}^\alpha}{\rho_f^\alpha} \left| \frac{n_f^\alpha}{q_{i(f)}^\alpha} - \rho_f^\alpha \right| |\rho_f - x_f(n, q)| \\ &= \frac{q_{i(f)}^\alpha}{\rho_f^\alpha} |\mathfrak{R}^\alpha(n_f, q_{i(f)}) - \rho_f^\alpha| |\rho_f - \mathfrak{R}(n_f, q_{i(f)})| \\ &\geq \frac{n_f q_{i(f)}^\alpha g_f(n, q)}{\rho_f^\alpha} \geq \frac{n_f^{1+\alpha} g_f(n, q)}{C^\alpha \rho_f^\alpha}, \end{aligned}$$

where we have used the fact that $\mathbf{b}_f \neq 0$ and (53).

- $0 \leq Cq_{i(f)} < n_f$. Here, since $\rho_f < C$, we have that

$$\left| \frac{n_f^\alpha}{\rho_f^\alpha} - q_{i(f)}^\alpha \right| |\rho_f - x_f(n, q)| = \left(\frac{n_f^\alpha}{\rho_f^\alpha} - q_{i(f)}^\alpha \right) (C - \rho_f) \geq (C - \rho_f)(\rho_f^{-\alpha} - C^{-\alpha})n_f^\alpha.$$

(c): Suppose that $(n, q) \in \mathbf{S}_b$. We wish to establish that $F_b(n, q) < 0$. Since (n, q) is not an invariant point, by part (iv) of the equivalence established in Theorem 10 and by (F11), one of the following two conditions holds:

- $T(q) < 0$. In this case, using (52), clearly $F_b(n, q) < 0$.
- $T(q) = 0$, and there exists some $f \in \mathcal{F}$ with $n_f \neq \rho_f q_{i(f)}$ and $n_f > 0$. Here, if $\mathbf{b}_f = 0$, then since the inequality in (52) must be strict, we have $F_b(n, q) < 0$. On the other hand, suppose that $\mathbf{b}_f \neq 0$. Since $n_f > 0$, we have that

$$\min \left((C - \rho_f)(\rho_f^{-\alpha} - C^{-\alpha})n_f^\alpha, \frac{n_f^{1+\alpha} g_f(n, q)}{C^\alpha \rho_f^\alpha} \right) > 0,$$

and it follows that $F_b(n, q) < 0$.

Now, given claims (a) and (c), it follows that

$$\sup_{(n, q) \in \mathbf{S}_b} F_b(n, q) \leq -\gamma_b < 0,$$

for some $\gamma_b > 0$. Using claim (b), we have that

$$\mathfrak{D}(n, q) \leq -\gamma_b < 0,$$

for all $(n, q) \in \mathbf{S}_b$. Since $\mathcal{D} \setminus \mathcal{K}_\delta = \cup_{\mathbf{b} \in \{-1, 0, 1\}^{\mathcal{F}}} \mathbf{S}_b$, we have that, for all $(n, q) \in \mathcal{D} \setminus \mathcal{K}_\delta$,

$$\mathfrak{D}(n, q) \leq -\gamma < 0,$$

where

$$\gamma \triangleq \min_{\mathbf{b} \in \{-1, 0, 1\}^{\mathcal{F}}} \gamma_{\mathbf{b}}.$$

This completes the proof of Theorem 11. ■

8.4. Fixed Points: Optimality

The following theorem characterizes the cost associated with an invariant state, relative to the effective cost. The effective cost represents the lowest cost achievable under *any* policy (cf. Theorem 6). Hence, this result implies that the invariant states of the MWUM- α policy are cost optimal, as $\alpha \rightarrow 0^+$. Combining this with the convergence result of Theorem 11, as $\alpha \rightarrow 0^+$, the state of the system under the MWUM- α policy will converge to a state that cannot be improved upon by any other policy.

Theorem 13. *Suppose (n^*, q^*) is an invariant state of a critically loaded system under the MWUM- α policy. Then,*

$$c(n^*, q^*) \leq (1 + \beta(\alpha))c^*(n^*, q^*), \quad (56)$$

where $\beta(\alpha) \rightarrow 0$ as $\alpha \rightarrow 0^+$.

Proof. Suppose (n^*, q^*) is an invariant state. Define (n', q') to be an optimal solution to the effective cost LP $c^*(n^*, q^*)$, defined by (41). Clearly

$$L_\alpha(n^*, q^*) \leq L_\alpha(n', q'),$$

since (n^*, q^*) is optimal for $b\text{ALGD}(n^*, q^*)$, and (n', q') is feasible for $b\text{ALGD}(n^*, q^*)$. Then, following the same argument as in Lemma 7,

$$c(n^*, q^*) \leq (1 + \beta(\alpha))c(n', q') = (1 + \beta(\alpha))c^*(n^*, q^*),$$

where $\beta(\alpha) \rightarrow 0$ as $\alpha \rightarrow 0^+$. ■

9. Discussion and Future Work

We have provided a model of a communications network that operates at the packet-level with the goal of achieving end-to-end performance at the flow-level. The proposed MWUM- α control policy achieves this goal by means of the maximum weight- α packet-level scheduling along with the α -fair rate allocation. We established the positive recurrence of the system by means of fluid model when the system is underloaded. For the critically loaded fluid model, we established path-wise constant factor optimality; the constant factor depends α and the *balance factor*.

There are several interesting directions for future work. To start with, by characterizing the invariant manifold of the critically loaded fluid model and establishing its attractiveness, the work

here should lead to the multiplicative state-space collapse property in a relatively straightforward manner following the method of Bramson [4]. As the next step, establishing the strong state-space collapse property would require bounding the maximal deviation in the system state over certain time-horizon. We strongly believe that under MWUM- α control policy for $\alpha \geq 1$, this should follow from a recently developed Lyapunov function based maximal inequality by Shah, Tsitsiklis and Zhong [23]. However, further obtaining a complete characterization of the diffusion (heavy traffic) approximation seems to be far more non-trivial question. Finally, the results about path-wise constant factor optimality of critically loaded fluid model seem to suggest the possibility of such constant factor optimality of MWUM- α control policy under diffusion approximation.

References

- [1] D. P. Bertsekas. *Nonlinear Programming*. Athena Scientific, 2nd edition, 1999.
- [2] P. Billingsley. *Convergence of Probability Measures*. Wiley-Interscience, New York, 2nd edition, 1999.
- [3] T. Bonald and L. Massoulié. Impact of fairness on Internet performance. In *Proceedings of the 2001 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems*, pages 82–91, Cambridge, MA, June 2001.
- [4] M. Bramson. State space collapse with application to heavy traffic limits for multiclass queueing networks. *Queueing Systems*, 30(1–2):89–148, November 1998.
- [5] J. G. Dai. On positive Harris recurrence of multiclass queueing networks: a unified approach via fluid limit models. *The Annals of Applied Probability*, 5(1):49–77, 1995.
- [6] J. G. Dai and W. Lin. Maximum pressure policies in stochastic processing networks. *Operations Research*, 53(2):197–218, March-April 2005.
- [7] J. G. Dai and B. Prabhakar. The throughput of switches with and without speed-up. In *Proceedings of IEEE Infocom*, pages 556–564, 2000.
- [8] G. de Veciana, T.-J. Lee, and T. Konstantopoulos. Stability and performance analysis of networks supporting elastic services. *IEEE/ACM Transactions on Networking*, 9(1):2–14, 2001.
- [9] A. Dembo and O. Zeitouni. *Large Deviations Techniques and Applications*. Springer, New York, 2nd edition, 2003.
- [10] A. Eryilmaz and R. Srikant. Joint congestion control, routing and MAC for stability and fairness in wireless networks. *IEEE Journal on Selected Areas in Communications*, 24(8):1514–1524, August 2006.
- [11] W. Kang, F. P. Kelly, N. H. Lee, and R. J. Williams. Fluid and brownian approximations for an internet congestion control model. In *Proceedings of the 43rd IEEE Conference on Decision and Control*, volume 4, pages 3938–3943, December 2004.
- [12] F. P. Kelly, A. K. Maulloo, and D. K. H. Tan. Rate control for communication networks: shadow prices, proportional fairness and stability. *Journal of the Operational Research Society*, 49(3):237–252, March 1998.
- [13] F. P. Kelly and R. J. Williams. Fluid model for a network operating under a fair bandwidth-sharing policy. *The Annals of Applied Probability*, 14(3):1055–1083, 2004.
- [14] I. Keslassy and N. McKeown. Analysis of scheduling algorithms that provide 100% throughput in input-queued switches. In *Proceedings of Allerton Conference on Communication, Control, and Computing*, 2001.

- [15] X. Lin, N. B. Shroff, and R. Srikant. On the connection-level stability of congestion-controlled communication networks. *IEEE Transactions on Information Theory*, 54(5):2317–2338, May 2008.
- [16] L. Massoulié. Structural properties of proportional fairness: Stability and insensitivity. *Annals of Applied Probability*, 17(3):809–839, 2007.
- [17] L. Massoulié and J. W. Roberts. Bandwidth sharing and admission control for elastic traffic. *Telecommunication Systems*, 15(1–2):185–201, November 2000.
- [18] N. McKeown, V. Anantharam, and J. Walrand. Achieving 100% throughput in an input-queued switch. In *Proceedings of IEEE Infocom*, pages 296–302, 1996.
- [19] J. Mo and J. Walrand. Fair end-to-end window-based congestion control. *IEEE/ACM Transactions on Networking (ToN)*, 8(5):556–567, 2000.
- [20] M. J. Neely, E. Modiano, and C. Li. Fairness and optimal stochastic control for heterogeneous networks. *IEEE/ACM Transactions on Networking*, 16(2):396–409, April 2008.
- [21] P. Robert. *Stochastic Networks and Queues*. Springer, Berlin, 2003.
- [22] D. Shah. Analysis of scheduling algorithms that provide 100% throughput in input-queued switches. In *Proceedings of Allerton Conference on Communication, Control, and Computing*, 2001.
- [23] D. Shah, J. N. Tsitsiklis, and Y. Zhong. Qualitative properties of α -control policy. Under submission, November 2009.
- [24] D. Shah and D. J. Wischik. Optimal scheduling algorithm for input queued switch. In *Proceedings of IEEE INFOCOM*, 2006.
- [25] D. Shah and D. J. Wischik. The teleology of scheduling algorithms for switched networks under light load, critical load, and overload. <http://web.mit.edu/devavrat/www/shahwischik.pdf>, 2007-09.
- [26] S. Shakkottai and R. Srikant. *Network Optimization and Control*. Foundations and Trends in Networking. Now Publishers, Inc., Boston, January 2008.
- [27] A. L. Stolyar. MaxWeight scheduling in a generalized switch: State space collapse and workload minimization in heavy traffic. *Annals of Applied Probability*, 14(1):1–53, 2004.
- [28] A. L. Stolyar. Maximizing queueing network utility subject to stability: Greedy primal-dual algorithm. *Queueing Systems*, 50(4):401–457, 2005.
- [29] L. Tassiulas and P. B. Bhattacharya. Allocation of interdependent resources for maximal throughput. *Stochastic Models*, 16:27–48, 2000.
- [30] L. Tassiulas and A. Ephremides. Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks. *IEEE Transactions on Automatic Control*, 37(12):1936–1948, December 1992.
- [31] N. S. Walton. Proportional fairness and its relationship with multi-class queueing networks. To appear in *Annals of Applied Probability*, 2008.
- [32] H. Ye and D. D. Yao. Heavy-traffic optimality of a stochastic network under utility-maximizing resource control. *Operations Research*, 56(2):453–470, March-April 2008.
- [33] H.-Q. Ye, J. Ou, and X.-M. Yuan. Stability of data networks: stationary and bursty model. *Operations Research*, 53(1):107–125, January-February 2005.

Appendix A: Standard Norm Inequalities

The following lemma provides some standard norm inequalities that are used throughout the paper:

Lemma 14. *Consider a vector $x \in \mathbb{R}^d$.*

- (i) *If $0 < p \leq q \leq \infty$, then $\|x\|_q \leq \|x\|_p$.*
- (ii) *If $1 \leq p \leq \infty$ and $1/p + 1/q = 1$, then*

$$d^{-1/q} \|x\|_1 \leq \|x\|_p.$$

Appendix B: Justification of the Fluid Model

In this appendix, we will provide a proof for Theorem 1, which establishes fluid model as the formal functional law of large numbers approximation.

We begin with some technical preliminaries. Fix $T > 0$. Recall from Section 4.2 that $\mathbf{D}[0, T]$ is the space of functions from $[0, T]$ to \mathfrak{Z} , as defined in (7), that are RCLL. This space is equipped with the Skorohod metric, defined as

$$\mathbf{d}(\mathbf{x}, \mathbf{y}) \triangleq \inf_{\phi \in \Phi} \|\phi\|^o \vee \|\mathbf{x} - \phi \circ \mathbf{y}\|, \quad \text{for } \mathbf{x}, \mathbf{y} \in \mathbf{D}[0, T].$$

Here, Φ is the set of all non-decreasing functions $\phi: [0, T] \rightarrow [0, T]$ with $\phi(0) = 0$ and $\phi(T) = T$. The norm $\|\cdot\|^o$ over Φ is defined as follows: for $\phi \in \Phi$,

$$\|\phi\|^o \triangleq \sup_{0 \leq s < t \leq T} \log \left| \frac{\phi(t) - \phi(s)}{t - s} \right|.$$

By $\phi \circ \mathbf{y}$ refers to the composition $\mathbf{y}(\phi(t))$, and for any $\mathbf{x} \in \mathbf{D}[0, T]$,

$$\|\mathbf{x}\| \triangleq \sup_{t \in [0, T]} \|\mathbf{x}(t)\|_1,$$

with $\|\cdot\|_1$ being the standard ℓ_1 -norm over the product space \mathfrak{Z} .

For any $\mathbf{x} \in \mathbf{D}[0, T]$ and $0 \leq s < t \leq T$, define

$$w_{\mathbf{x}}(s, t) \triangleq \sup_{s \leq t_1, t_2 \leq t} \|\mathbf{x}(t_1) - \mathbf{x}(t_2)\|_1.$$

Further, for any $\delta > 0$, define

$$w'_{\mathbf{x}}(\delta) = \inf_{\{t_i\} \in \mathbf{S}(T, \delta)} \max_i w_{\mathbf{x}}(t_{i-1}, t_i),$$

where $\mathbf{S}(T, \delta)$ is collection of all δ -sparse decompositions $\{t_i\}$ of $[0, T]$, i.e., $0 = t_0 < t_1 < \dots < t_\ell = T$ with $t_i - t_{i-1} \geq \delta$ for all $i \geq 1$.

It can be easily checked (see [2, Chapter 3]) that the $\mathbf{D}[0, T]$ is Polish space under the metric \mathbf{d} . Let \mathcal{B}_T denote the Borel σ -algebra on $\mathbf{D}[0, T]$ with respect to the topology induced by \mathbf{d} . We will be interested in probability measures over space $(\mathbf{D}[0, T], \mathcal{B})$. We shall utilize the following well-known characterization of tightness of measures (see [2, Theorem 13.2]):

Theorem 15. *The collection of measures $\{P_\theta : \theta \in \Theta\}$ defined on $(\mathbf{D}[0, T], \mathcal{B})$ is tight if and only if the following two conditions are satisfied:*

(a)

$$\lim_{A \rightarrow \infty} \limsup_{\theta \in \Theta} P_\theta \left(\mathbf{x} \in \mathbf{D}[0, T] : \|\mathbf{x}\| \geq A \right) = 0.$$

(b) For each $\varepsilon > 0$,

$$\lim_{\delta \rightarrow 0} \limsup_{\theta \in \Theta} P_\theta \left(\mathbf{x} \in \mathbf{D}[0, T] : w'_\mathbf{x}(\delta) \geq \varepsilon \right) = 0.$$

We state the following well-known ‘concentration’ property of Poisson process that shall later be useful. It follows from the application of a standard Chernoff bound (see, for example, [9, Theorem 2.2.3]).

Proposition 16. *Consider a Poisson process of rate 1. Let $N(t)$ be the number of events of this Poisson process in time interval $[0, t]$. Then, for any $\delta \in [0, t]$,*

$$\mathbb{P} \left(|N(t) - t| \geq \delta \right) \leq 2 \exp \left(-\frac{\delta^2}{2t} \right).$$

B.1. Tightness

The first step in the proof of Theorem 1 is the following lemma, which establishes tightness of the collection of measures associated with the scaled system processes.

Lemma 17. *Under the hypotheses of Theorem 1, for each $r \geq 1$, let $\mu^{(r)}$ denote the measure of $Z^{(r)}(\cdot) \in \mathbf{D}[0, T]$. Then, the collection of measures $\{\mu^{(r)} : r \geq 1\}$ is tight.*

Proof. We will establish tightness by verifying conditions (a) and (b) of Theorem 15.

First, consider condition (a). It is sufficient to show that for any $\delta > 0$, there exist $K(\delta)$ and $r(\delta)$ such that, for $K \geq K(\delta)$ and $r \geq r(\delta)$,

$$\mu^{(r)} \left(\mathbf{x} \in \mathbf{D}[0, T] : \|\mathbf{x}\| \geq K \right) \leq \delta. \quad (57)$$

To establish this, fix $\delta > 0$. By definition,

$$\begin{aligned} \|\mathcal{Z}^{(r)}(\cdot)\| = \sup_{t \in [0, T]} & \left(\|Q^{(r)}(t)\|_1 + \|Z^{(r)}(t)\|_1 + \|N^{(r)}(t)\|_1 + \|S^{(r)}(t)\|_1 + \|\bar{X}^{(r)}(t)\|_1 \right. \\ & \left. + \|\mathcal{A}^{(r)}(t)\|_1 + \|D^{(r)}(t)\|_1 + \|A^{(r)}(t)\|_1 \right). \end{aligned}$$

We will bound each component of the system process $\mathcal{Z}^{(r)}(\cdot)$.

First, observe that for any $t \in [0, T]$, with probability 1,

$$\|Z^{(r)}(t)\|_1 + \|S^{(r)}(t)\|_1 + \|\bar{X}^{(r)}(t)\|_1 \leq K_1 T, \quad (58)$$

where the constant K_1 depends on system dimensions $|\mathcal{E}|$ and $|\mathcal{F}|$ and on the maximum rate allocation C . This is because at most a unit amount of scheduling can be performed per unit time, and maximal rate allocated to any flow type is at most C .

Next, consider the term $\|A^{(r)}(t)\|_1$. For each flow type $f \in \mathcal{F}$, $A_f^{(r)}(T)$ is a Poisson process with a time-varying rate that is at most C . Therefore, $A_f^{(r)}(T)$ is bounded above by $1/r$ times the total

number of events of a Poisson process of rate C in time interval $[0, rT]$. This number of events is distributionally equivalent to number of events of Poisson process of rate 1 in interval $[0, rTC]$. Therefore, using Proposition 16,

$$\mathbb{P}\left(A_f^{(r)}(T) \geq 2CT\right) \leq 2e^{-\frac{1}{2}rTC}. \quad (59)$$

It follows that for r sufficiently large, for any $f \in \mathcal{F}$,

$$\mathbb{P}\left(A_f^{(r)}(T) \geq 2CT\right) \leq \frac{\delta}{10|\mathcal{F}|}. \quad (60)$$

Then, by the union bound,

$$\mathbb{P}\left(\|A^{(r)}(T)\|_1 \geq 2|\mathcal{F}|CT\right) \leq \frac{\delta}{10}. \quad (61)$$

Next, consider term $\|Q^{(r)}(t)\|_1$. Note that

$$\sup_{t \in [0, T]} \|Q^{(r)}(t)\|_1 \leq \|Q^{(r)}(0)\|_1 + \|A^{(r)}(T)\|_1. \quad (62)$$

Now, by hypothesis, we have $Q^{(r)}(0) \rightarrow q(0)$ with probability 1 as $r \rightarrow \infty$. Therefore, the collection of vectors $\{Q^{(r)}(0)\}$ is almost surely bounded, and thus there exists a constant K_2 so that, for r sufficiently large,

$$\mathbb{P}\left(\|Q^{(r)}(0)\|_1 \geq K_2\right) \leq \frac{\delta}{10}. \quad (63)$$

It follows that for a large enough constant K_3 (dependent on $K_2, C, T, |\mathcal{E}|, |\mathcal{F}|$), and for r sufficiently large,

$$\mathbb{P}\left(\sup_{t \in [0, T]} \|Q^{(r)}(t)\|_1 \geq K_3\right) \leq \frac{\delta}{5}. \quad (64)$$

Using very similar arguments (employing Proposition 16), it is possible to bound the Poisson processes $\|D^{(r)}(t)\|_1$ and $\|\mathcal{A}^{(r)}(t)\|_1$. This, in turn, will lead to a bound on $\|N^{(r)}(t)\|_1$, since

$$\sum_{t \in [0, T]} \|N^{(r)}(t)\|_1 \leq \|N^{(r)}(0)\|_1 + \|\mathcal{A}^{(r)}(T)\|_1.$$

Therefore, there exists a constant K_4 , so that for r sufficiently large,

$$\mathbb{P}\left(\sup_{t \in [0, T]} \|D^{(r)}(t)\|_1 + \|\mathcal{A}^{(r)}(t)\|_1 + \|N^{(r)}(t)\|_1 \geq K_4\right) \leq \frac{\delta}{10}. \quad (65)$$

From the discussion above, equations (58), (61), (64), (65), and union bound, it follows that for any $\delta > 0$, there exists constants $K(\delta)$ and $r(\delta)$ such that for $K \geq K(\delta)$ and $r \geq r(\delta)$, we have that

$$\mathbb{P}\left(\|\mathcal{Z}^{(r)}(\cdot)\| \geq K\right) \leq \delta.$$

This completes the verification of (57) or equivalently, condition (a) of Theorem 15.

Next, consider condition (b) of Theorem 15. For this, it is sufficient to show that for any $\varepsilon > 0$, there exist $\delta(\varepsilon)$ and $r(\delta(\varepsilon))$ so that for $r \geq r(\delta(\varepsilon))$,

$$\mathbb{P}\left(w'_{\mathcal{Z}^{(r)}}(\delta(\varepsilon)) \geq \varepsilon\right) \leq \delta(\varepsilon),$$

with $\delta(\varepsilon) \rightarrow 0$ as $\varepsilon \rightarrow 0$.

To bound $w'_{\mathcal{Z}^{(r)}}(\delta)$, we need to find an appropriate δ -sparse decomposition $\{t_0, t_1, \dots, t_n\}$ of $[0, T]$. For this, we consider a natural decomposition: $t_i = i\delta$, for $0 \leq i < n$, and $t_n = T$, when $n = \lfloor T/\delta \rfloor$. Then, it follows that $\delta \leq t_i - t_{i-1} \leq 2\delta$ for all $0 < i \leq n$.

We wish to bound $\|\mathcal{Z}^{(r)}(s) - \mathcal{Z}^{(r)}(t)\|_1$, for $t_{i-1} \leq s, t \leq t_i$, for any $0 < i \leq n$. To this end, first note that

$$\begin{aligned} \|\mathcal{Z}^{(r)}(s) - \mathcal{Z}^{(r)}(t)\|_1 &= \|Q^{(r)}(t) - Q^{(r)}(s)\|_1 + \|Z^{(r)}(t) - Z^{(r)}(s)\|_1 + \|N^{(r)}(t) - N^{(r)}(s)\|_1 \\ &\quad + \|S^{(r)}(t) - S^{(r)}(s)\|_1 + \|\bar{X}^{(r)}(t) - \bar{X}^{(r)}(s)\|_1 + \|\mathcal{A}^{(r)}(t) - \mathcal{A}^{(r)}(s)\|_1 \\ &\quad + \|D^{(r)}(t) - D^{(r)}(s)\|_1 + \|A^{(r)}(t) - A^{(r)}(s)\|_1. \end{aligned} \quad (66)$$

As noted earlier, the terms involving $Z^{(r)}(\cdot)$, $S^{(r)}(\cdot)$ and $\bar{X}^{(r)}(\cdot)$ are collectively upper bounded by $K_1|t-s|$, with a system dependent constant K_1 , since they are all Lipschitz continuous. Therefore, for $\delta \leq \varepsilon/(10K_1)$, the sum of these terms in (66) is no more than $\varepsilon/10$ with probability 1. Next, we consider the remaining five terms in (66). As in the justification of condition (a), we will have similar argument for all these of five terms. We will focus on the term corresponding to $Q^{(r)}(\cdot)$. From (3), it follows that

$$\begin{aligned} \|Q^{(r)}(t) - Q^{(r)}(s)\|_1 &\leq \|(I - R^\top)\Pi(S^{(r)}(t) - S^{(r)}(s))\|_1 + \|(I - R^\top)(Z^{(r)}(t) - Z^{(r)}(s))\|_1 \\ &\quad + \|A^{(r)}(t) - A^{(r)}(s)\|_1. \end{aligned}$$

Now the first two terms are bounded by $K_5|t-s|$, with the constant K_5 dependent on $|\mathcal{S}|$, $|\mathcal{E}|$, and $|\mathcal{F}|$. To see this, note that both $S^{(r)}(\cdot)$ and $Z^{(r)}(\cdot)$ are Lipschitz continuous, and the matrices $(I - R^\top)\Pi$ and $(I - R^\top)$ are finite dimensional (with dimension dependent on $|\mathcal{S}|$, $|\mathcal{E}|$, and $|\mathcal{F}|$), and with each entry bounded by constants. It follows that by choosing $\delta \leq \varepsilon/(10K_5)$, we have that

$$\|(I - R^\top)\Pi(S^{(r)}(t) - S^{(r)}(s))\|_1 + \|(I - R^\top)(Z^{(r)}(t) - Z^{(r)}(s))\|_1 \leq \frac{\varepsilon}{10},$$

with probability 1. Now, to bound the contribution of term $\|A^{(r)}(t) - A^{(r)}(s)\|_1$, we can utilize arguments used for obtaining (61) with T replaced by $|t-s|$. Note that here we need to use ‘memory-less’ property of the Poisson process crucially. As a conclusion, we obtain that there exists a constant K_6 so that is $\delta \leq \varepsilon/(10K_6)$, then for sufficiently large r ,

$$\mathbb{P}\left(\|A^{(r)}(t) - A^{(r)}(s)\|_1 \leq \varepsilon/10\right) \geq 1 - \frac{\delta^2}{10T}. \quad (67)$$

From the above discussion, it follows that for any interval $[t_{i-1}, t_i]$, as long as we choose δ sufficiently small and r sufficiently large,

$$\mathbb{P}\left(\sup_{t_{i-1} \leq s, t \leq t_i} \|Q^{(r)}(t) - Q^{(r)}(s)\|_1 \leq \varepsilon/5\right) \geq 1 - \frac{\delta^2}{10T}. \quad (68)$$

In a very similar manner (using Proposition 16), we obtain the following: there exists a constant

K_7 , so that if $\delta = \varepsilon/K_7$ and r is sufficiently large, for any $0 < i \leq n$,

$$\mathbb{P} \left(\sup_{t_{i-1} \leq s, t \leq t_i} \|N^{(r)}(t) - N^{(r)}(s)\|_1 \leq \varepsilon/10 \right) \geq 1 - \frac{\delta^2}{10T}, \quad (69)$$

$$\mathbb{P} \left(\sup_{t_{i-1} \leq s, t \leq t_i} \|\mathcal{A}^{(r)}(t) - \mathcal{A}^{(r)}(s)\|_1 \leq \varepsilon/10 \right) \geq 1 - \frac{\delta^2}{10T}, \quad (70)$$

$$\mathbb{P} \left(\sup_{t_{i-1} \leq s, t \leq t_i} \|A^{(r)}(t) - A^{(r)}(s)\|_1 \leq \varepsilon/10 \right) \geq 1 - \frac{\delta^2}{10T}, \quad (71)$$

$$\mathbb{P} \left(\sup_{t_{i-1} \leq s, t \leq t_i} \|D^{(r)}(t) - D^{(r)}(s)\|_1 \leq \varepsilon/10 \right) \geq 1 - \frac{\delta^2}{10T}. \quad (72)$$

From (66)-(72) and the discussion above, it follows that there exists a constant K' so that is $\varepsilon \leq K$ and r is sufficiently large, by a union bound over at most T/δ intervals in the partition, we obtain that

$$\mathbb{P} \left(\max_i \sup_{t_{i-1} \leq s, t \leq t_i} \|\mathcal{Z}^{(r)}(t) - \mathcal{Z}^{(r)}(s)\|_1 \leq \varepsilon \right) \geq 1 - \delta.$$

The result follows. ■

B.2. Proof of Theorem 1

We are now ready to prove Theorem 1.

Given that the collection of measures $\{\mu^{(r)} : r \geq 1\}$ is tight, for any sequence $\{r_k : k \in \mathbb{N}\} \subset \mathbb{R}$ with $r_k \rightarrow \infty$ as $k \rightarrow \infty$, there exists a further subsequence $\{r_{k_\ell}\}$ and limit point $\mu^{(\infty)}$ such that $\mu^{(r_{k_\ell})}$ converges weakly to $\mu^{(\infty)}$ as $\ell \rightarrow \infty$. By restricting to this subsequence, assume that $\mu^{(r_k)} \Rightarrow \mu^{(\infty)}$. Since these measures are defined on a Polish space, by the Skorohod representation theorem, there exists a probability space over which we can define, for all k , random variables $\mathcal{Z}^{(r_k)}(\cdot)$ and $\mathfrak{z}(\cdot)$ that are distributed according to $\mu^{(r_k)}$ and $\mu^{(\infty)}$, respectively, and where the $\mathcal{Z}^{(r_k)}(\cdot)$ almost surely converges to $\mathfrak{z}(\cdot)$, in the Skorohod metric. We will use this setting to argue that the limiting random variable $\mathfrak{z}(\cdot)$ satisfies the appropriate fluid model equations with probability 1. Subsequently, we will establish that under an arbitrary control policy, $\mu^{(\infty)}(\text{FMS}(T)) = 1$, while under the MWUM- α policy, $\mu^{(\infty)}(\text{FMS}^\alpha(T)) = 1$. Since $\mu^{(r_k)} \Rightarrow \mu^{(\infty)}$, from definition of weak convergence, it follows that, under an arbitrary control policy, $\liminf_{k \rightarrow \infty} \mu^{(r_k)}(\text{FMS}_\varepsilon(T)) = 1$, and under the MWUM- α policy, $\liminf_{k \rightarrow \infty} \mu^{(r_k)}(\text{FMS}_\varepsilon^\alpha(T)) = 1$, for any $\varepsilon > 0$. This will imply the desired result.

To this end, we start by establishing that $\mu^{(\infty)}(\text{FMS}(T)) = 1$. That is, we need to show that equations (F1)-(F8) are satisfied.

(F1), (F2), (F5): We start with (F1). Among the components of $\mathcal{Z}^{(r_k)}(\cdot)$, $Z^{(r_k)}(\cdot)$, $S^{(r_k)}(\cdot)$ and $\bar{X}^{(r_k)}(\cdot)$ are Lipschitz continuous by construction over $[0, T]$. Since $\mathcal{Z}^{(r_k)}(\cdot)$ converges almost surely to $\mathfrak{z}(\cdot)$ with respect to the Skorohod metric \mathbf{d} , it follows that the corresponding components of $\mathfrak{z}(\cdot)$, $z(\cdot)$, $s(\cdot)$ and $\bar{x}(\cdot)$, are Lipschitz continuous. Equivalently, this follows by the Arzelà-Ascoli theorem.

Now, to establish Lipschitz continuity of the other components of \mathfrak{z} we will use the following result:

Lemma 18. *Given a fixed $T > 0$, consider a Poisson process \mathcal{P} with time-varying rate given by $\gamma(t)$, for $t \geq 0$. Assume there exists constant $K > 0$ such that $\gamma(t) \in [0, K]$ for all $t \geq 0$ and that $\gamma(t)$ depends only on events that happen up to time t . Consider a sequence $\{\theta_i\} \subset \mathbb{R}_+$ with $\theta_i \rightarrow \infty$ as $i \rightarrow \infty$. Define the scaled process*

$$\mathcal{P}^i(t) = \frac{1}{\theta_i} \mathcal{P}(\theta_i t),$$

for $t \in [0, T]$. Also, define the processes

$$\gamma^i(t) = \frac{1}{\theta_i} \gamma(\theta_i t), \quad \text{and} \quad \bar{\gamma}^i(t) = \int_0^t \gamma^i(s) ds,$$

for $t \in [0, T]$, where we assume $\bar{\gamma}^i(\cdot)$ is well-defined. Assume that $\mathcal{P}^i(\cdot)$ converges weakly to $\mathcal{P}^\infty(\cdot)$. Then, the sample paths (over $[0, T]$) of $\mathcal{P}^\infty(\cdot)$ are Lipschitz continuous with probability 1. Further, assume that $\bar{\gamma}^i(\cdot)$ converges (u.o.c.) to $\bar{\gamma}(\cdot)$ over $[0, T]$. Then, $\mathcal{P}^\infty(t) = \bar{\gamma}(t)$, for all $t \in [0, T]$.

Proof. This is a well-known property of Poisson processes. We describe the key steps of the proof. First, using the concentration property of Proposition 16, it can be established that the sample paths of scaled Poisson process are approximately Lipschitz. Second, using a variation of the Arzelà-Ascoli theorem (e.g., see [4, Lemma 4.2] or [33, Lemma 6.3]), it can be established that the limit points of such approximately Lipschitz sample paths are in fact Lipschitz. Finally, note the fact that number of events under a Poisson process with time-varying rate $\gamma(\cdot)$ over time interval $[0, t]$ is distributionally equivalent to number of events under a unit rate Poisson process over time interval $[0, \int_0^t \gamma(s) ds]$. As long as $\gamma(\cdot)$ is uniformly bounded by some constant, say K , the functional strong law of large numbers for scaled unit rate Poisson process over $[0, KT]$ can be used to obtain the final desired claim. An interested reader may find the details to this argument in many places in literature (e.g., [16, Appendix A]). ■

Now consider components $\mathcal{A}^{(r_k)}(\cdot)$, $A^{(r_k)}(\cdot)$, and $D^{(r_k)}(\cdot)$. By their construction, these are Poisson processes with possibly time-varying rates that are always uniformly bounded. These processes converge (over $[0, T]$) to the corresponding components $\mathbf{a}(\cdot)$, $a(\cdot)$ and $d(\cdot)$ of $\mathfrak{z}(\cdot)$. Therefore, by immediate application of Lemma 18, we obtain that $\mathbf{a}(\cdot)$, $a(\cdot)$ and $d(\cdot)$ are Lipschitz continuous.

Finally, the Lipschitz continuity of $n(\cdot)$ and $q(\cdot)$ is established if (F2) and (F5) hold — this is because all of the other components of $\mathfrak{z}(\cdot)$ are Lipschitz continuous. To this end, recall that equations (2) and (3) are satisfied by scaled system $\mathcal{Z}^{(r_k)}$ for all $t \in [0, T]$ by definition. These equations are preserved under the almost sure convergence $\mathcal{Z}^{(r_k)}(\cdot) \rightarrow \mathfrak{z}(\cdot)$. Thus, (F2) and (F5) are satisfied.

(F3), (F4), (F6): These equations follow immediately by applying the later part of Lemma 18 for Poisson process (possibly time-varying) $\mathcal{A}^{(r_k)}(\cdot)$, $A^{(r_k)}(\cdot)$ and $D^{(r_k)}(\cdot)$.

(F7), (F8): Among remaining equations, first note that (F8) follows because $Z^{(r_k)}(\cdot)$, $S^{(r_k)}(\cdot)$ and $\bar{X}^{(r_k)}(\cdot)$ are non-decreasing and this property is preserved under the almost sure convergence of $\mathcal{Z}^{(r_k)}(\cdot)$ to $\mathfrak{z}(\cdot)$. A similar argument establishes (F7).

Now, consider a system that operates under the MWUM- α control policy. We wish to establish that $\mu^{(\infty)}(\text{FMS}^\alpha(T)) = 1$. This involves further demonstrating that (F9)–(F12) are satisfied.

(F9): Consider any fixed flow type $f \in \mathcal{F}$ and any regular point $t \in (0, T)$. Now, if $n_f(t) = 0$, then it must be that $\dot{n}_f(t) = 0$. This is because of the following argument, utilizing non-negativity of $n_f(\cdot)$: suppose that either $\dot{n}_f(t) > 0$ or $\dot{n}_f(t) < 0$. Then, there exist times $t^- < t$ or $t^+ > t$ such that $n_f(t^-) < 0$ or $n_f(t^+) < 0$ — this is a contradiction. Given $\dot{n}_f(t) = 0$, by (F2) we have $\dot{\mathbf{a}}_f(t) = \dot{d}_f(t)$. By (F3) and (F4), it immediately follows that $x_f(t) = \dot{x}_f(t) = \nu_f/\mu_f$.

Now, suppose $n_f(t) > 0$. Recall that $N_f^{(r_k)}(\cdot)$ converges to $n_f(\cdot)$ as $k \rightarrow \infty$ under the Skorohod metric. Therefore, it follows that there exists a $\delta > 0$ such that, for k sufficiently large, $N_f^{(r_k)}(s) \geq \delta$ for all $s \in [t - \delta, t + \delta]$. We will consider k to be sufficiently large for this to hold. Since t is a regular point, $\bar{x}_f(t)$ is differentiable. Consider any $0 < \varepsilon < \delta$. Then, using the fact that $N_f^{(r_k)}(s) > 0$ for $s \in [t - \delta, t + \delta]$ and the radial invariance property of rate allocation policy, we obtain

$$\begin{aligned} \bar{X}_f^{(r_k)}(t + \varepsilon) - \bar{X}_f^{(r_k)}(t) &= \frac{1}{r_k} \int_{r_k t}^{r_k t + r_k \varepsilon} X_f(s) ds \\ &= \frac{1}{r_k} \int_{r_k t}^{r_k t + r_k \varepsilon} \left(\operatorname{argmax}_{x \in [0, C]} \frac{x^{1-\alpha} (N_f^{(r_k)})^\alpha}{1-\alpha} - (Q_{i(f)}^{(r_k)})^\alpha x \right) ds. \end{aligned} \quad (73)$$

Define the function $\mathfrak{R}: (0, \infty) \times \mathbb{R}_+ \rightarrow [0, C]$ by

$$\mathfrak{R}(n, q) \triangleq \operatorname{argmax}_{x \in [0, C]} \frac{x^{1-\alpha} n^\alpha}{1-\alpha} - q^\alpha x.$$

It can be easily checked that

$$\mathfrak{R}(n, q) = \begin{cases} n/q & \text{if } n < Cq, \\ C & \text{otherwise.} \end{cases}$$

Therefore, it follows that \mathfrak{R} is a continuous function. Further, $N_f^{(r_k)}(\cdot)$ and $Q_{i(f)}^{(r_k)}(\cdot)$ are continuous as functions of time. Therefore, treating $\mathfrak{R}(N_f^{(r_k)}(\cdot), Q_{i(f)}^{(r_k)}(\cdot))$ as a function of time, it is continuous and takes values in $[0, C]$. Over the bounded interval $[t, t + \varepsilon]$, it must achieve a minimum and a maximum, which we will denote by $\mathfrak{R}_{\min}(k, t, \varepsilon)$ and $\mathfrak{R}_{\max}(k, t, \varepsilon)$, respectively. From (73), it follows that

$$\mathfrak{R}_{\min}(k, t, \varepsilon) \leq \frac{\bar{X}_f^{(r_k)}(t + \varepsilon) - \bar{X}_f^{(r_k)}(t)}{\varepsilon} \leq \mathfrak{R}_{\max}(k, t, \varepsilon). \quad (74)$$

Now, since $(\bar{X}_f^{(r_k)}(\cdot), N_f^{(r_k)}(\cdot), Q_{i(f)}^{(r_k)}(\cdot))$ converges to $(\bar{x}_f(\cdot), n_f(\cdot), q_{i(f)}(\cdot))$ as $k \rightarrow \infty$, it follows (due to the appropriate continuity of \mathfrak{R}_{\min} , and \mathfrak{R}_{\max}) that

$$\mathfrak{R}_{\min}(t, \varepsilon) \leq \frac{\bar{x}_f(t + \varepsilon) - \bar{x}_f(t)}{\varepsilon} \leq \mathfrak{R}_{\max}(t, \varepsilon). \quad (75)$$

Here, $\mathfrak{R}_{\min}(t, \varepsilon)$ and $\mathfrak{R}_{\max}(t, \varepsilon)$ correspond to the minima and maxima of $\mathfrak{R}(n(\cdot), q(\cdot))$ over $[t, t + \varepsilon]$. Now, taking $\varepsilon \rightarrow 0$ in (75), invoking the continuity of $(n(\cdot), q(\cdot))$ and subsequently of \mathfrak{R} , and recalling that t is a regular point, we obtain

$$x_f(t) = \dot{x}_f(t) = \operatorname{argmax}_{x \in [0, C]} \frac{x^{1-\alpha} n_f^\alpha(t)}{1-\alpha} - q_{i(f)}^\alpha(t) x,$$

when $n_f(t) > 0$.

(F10): Suppose $t \in (0, T)$ is a regular point and consider a schedule $\pi \in \mathcal{S}$. Assume there exists a schedule $\sigma \in \mathcal{S}$ with

$$\pi^\top (I - R)q^\alpha(t) < \sigma^\top (I - R)q^\alpha(t). \quad (76)$$

We wish to establish that $\dot{s}_\pi(t) = 0$. Since $Q^{(r_k)}(\cdot)$ converges to $q(\cdot)$ as $k \rightarrow \infty$ the inequality (76) is strict. It follows that there exists $\delta > 0$ such that for all k sufficiently large and for all $s \in [t - \delta, t + \delta]$,

$$\pi^\top (I - R)(Q^{(r_k)}(s))^\alpha < \sigma^\top (I - R)(Q^{(r_k)}(s))^\alpha.$$

Then, for unscaled system, the weight of schedule π is strictly less than the weight of schedule σ throughout the time-interval $[r_k(t - \delta), r_k(t + \delta)]$. Therefore, as per the MWUM- α scheduling policy, the schedule π is never chosen in this time period. That is, for the scaled system, we have

$$S_\pi^{(r_k)}(t - \delta) = S_\pi^{(r_k)}(t + \delta).$$

Then, as $k \rightarrow \infty$

$$s_\pi(t + \delta) - s_\pi(t - \delta) = 0.$$

Thus, $\dot{s}_\pi(t) = 0$.

(F11): Consider a regular point $t \in (0, T)$ with $n_f(t) = 0$, for some $f \in \mathcal{F}$. By (F9), we have $x_f(t) = \rho_f$. Suppose that $q_{i(f)}(t) > 0$. For any $\varepsilon_1, \varepsilon_2 > 0$, there must exist $\delta > 0$ so that, if $s \in (t - \delta, t + \delta)$,

$$q_{i(f)}(s) \geq \varepsilon_1, \quad \text{and} \quad n_f(s) \leq \varepsilon_2.$$

Therefore, if k is sufficiently large, we must have

$$Q_{i(f)}^{(r_k)}(s) \geq \varepsilon_1/2, \quad \text{and} \quad N_f^{(r_k)}(s) \leq 2\varepsilon_2.$$

Equivalently for the unscaled system,

$$Q_{i(f)}(r_k s) \geq r_k \varepsilon_1/2, \quad \text{and} \quad N_f(r_k s) \leq 2r_k \varepsilon_2.$$

This that, for k sufficiently large and for all $s \in (t - \delta, t + \delta)$, the rate allocation in the unscaled system must satisfy

$$X_f(r_k s) \leq \frac{4\varepsilon_2}{\varepsilon_1}.$$

This can be made smaller than $\rho_f/2$ by the appropriate choice of ε_2 , and this contradicts the fact that $x_f(t) = \rho_f$. Therefore, it must be that $q_{i(f)}(t) = 0$.

(F12): This follows in a straightforward manner from the invariant (for the unscaled system) that $Z(\tau) = \mathbf{0}$ for all $\tau \in \mathbb{Z}_+$.