

QUALITATIVE PROPERTIES OF α -FAIR POLICIES IN BANDWIDTH-SHARING NETWORKS

BY D. SHAH AND J. N. TSITSIKLIS AND Y. ZHONG*

*Massachusetts Institute of Technology,
University of California, Berkeley*

We consider a flow-level model of a network operating under an α -fair bandwidth sharing policy (with $\alpha > 0$) proposed by Roberts and Massoulié (2000). This is a probabilistic model that captures the long-term aspects of bandwidth sharing between users or flows in a communication network.

We study the transient properties as well as the steady-state distribution of the model. In particular, for $\alpha \geq 1$, we obtain bounds on the maximum number of flows in the network over a given time horizon, by means of a maximal inequality derived from the standard Lyapunov drift condition. As a corollary, we establish the full state space collapse property for all $\alpha \geq 1$.

For the steady-state distribution, we obtain explicit exponential tail bounds on the number of flows, for any $\alpha > 0$, by relying on a norm-like Lyapunov function. As a corollary, we establish the validity of the diffusion approximation developed by Kang et al (2009), in steady state, for the case where $\alpha = 1$ and under a local traffic condition.

1. Introduction. We consider a flow-level model of a network that operates under an α -fair bandwidth-sharing policy, and establish a variety of new results on the resulting performance. These results include tail bounds on the size of a maximal excursion during a finite time interval, finiteness of expected queue sizes, exponential tail bounds under the steady-state distribution, and the validity of the heavy-traffic diffusion approximation in steady state. We note that our results are to a great extent parallel and complementary to our work on packet-switched networks, which was reported

*April 12, 2011; revised September 2012. This work was supported by NSF project CCF-0728554. The first two authors are affiliated with the Laboratory for Information and Decision Systems as well as the Operations Research Center at MIT. The last author is affiliated with the Algorithms, Machines and People's Laboratory at the University of California, Berkeley. Emails: {devavrat,jnt}@mit.edu, zhyu4118@berkeley.edu. The authors are grateful to Kuang Xu for reading the manuscript and his comments.

AMS 2000 subject classifications: Primary 60K20, 68M12; Secondary 68M20

Keywords and phrases: Bandwidth-Sharing Policy, State Space Collapse, α -fair, Heavy Traffic

in [18].

In the remainder of this section, we put our results in perspective by comparing them with earlier work, and conclude with some more details on the nature of our contributions.

1.1. Background. The flow-level network model that we consider was introduced by Roberts and Massoulié [17] to study the dynamic behavior of Internet flows. It builds on a static version of the model that was proposed earlier by Kelly, Maulloo, and Tan [15], and subsequently generalized by Mo and Walrand [16] who introduced a class of “fair” bandwidth-sharing policies parameterized by $\alpha > 0$.

The most basic question regarding flow-level models concerns necessary and sufficient conditions for stability, that is, for the existence of a steady-state distribution for the associated Markov process. This question was answered by Bonald and Massoulié [4] for the case of α -fair policies with $\alpha > 0$, and by de Veciana et al. [6] for the case of max-min fair policies ($\alpha \rightarrow \infty$) and proportionally fair ($\alpha = 1$) policies. In all cases, the stability conditions turned out to be the natural deterministic conditions based on mean arrival and service rates.

Given these stability results, the natural next question is whether the steady-state expectation of the number of flows in the system is finite and if so, to identify some nontrivial upper bounds. When $\alpha \geq 1$, the finiteness question can be answered in the affirmative, and explicit bounds can be obtained, by exploiting the same Lyapunov drift inequality that had been used in earlier work to establish stability. However, this approach does not seem to apply to the case where $\alpha \in (0, 1)$, which remained an open problem; this is one of the problems that we settle in this paper.

A more refined analysis of the number of flows present in the system concerns exponentially decaying bounds on the tail of its steady-state distribution. We provide results of this form, together with explicit bounds for the associated exponent. While a result of this type was not previously available, we take note of related recent results by Stolyar [21] and Venkataramanan and Lin [23] who provide a precise asymptotic characterization of the exponent of the tail probability, in steady state, for the case of switched networks (as opposed to flow-level network models). (To be precise, their results concern the $(1 + \alpha)$ norm of the vector of flow counts under maximum weight or pressure policies parameterized by $\alpha > 0$.) We believe that their methods extend to the model considered here, without much difficulty. However, their approach leads to a variational characterization that appears to be difficult to evaluate (or even bound) explicitly. We also take note of work by

Subramanian [22], who establishes a large deviations principle for a class of switched network models under maximum weight or pressure policies with $\alpha = 1$.

The analysis of the steady-state distribution for underloaded networks provides only partial insights about the transient behavior of the associated Markov process. As an alternative, the heavy-traffic (or diffusion) scaling of the network can lead to parsimonious approximations for the transient behavior. A general two-stage program for developing such diffusion approximations has been put forth by Bramson [5] and Williams [24], and has been carried out in detail for certain particular classes of queueing network models. To carry out this program, one needs to: (i) provide a detailed analysis of a related fluid model when the network is critically loaded; and, (ii) identify a unique distributional limit of the associated diffusion-scaled processes by studying a related Skorohod problem. The first stage of the program was carried out by Kelly and Williams [14] who identified the invariant manifold of the associated critically loaded fluid model. This further led to the proof by Kang et al. [13] of a multiplicative state space collapse property, similar to results by Bramson [5]. We note that the above summarized results hold under α -fair policies with an arbitrary $\alpha > 0$. The second stage of the program has been carried out for the proportionally fair policy ($\alpha = 1$) by Kang et al. [13], under a technical *local traffic* condition, and more recently, by Ye and Yao [25], under a somewhat less restrictive technical condition. We note however that when $\alpha \neq 1$, a diffusion approximation has not been established. In this case, it is of interest to see at least whether properties that are stronger than multiplicative state space collapse can be derived, something that is accomplished in the present paper.

The above outlined diffusion approximation results involve rigorous statements on the finite-time behavior of the original process. Kang et al. [13] further established that for the particular setting that they consider, the resulting diffusion approximation has an elegant product-form steady-state distribution; this result gives rise to an intuitively appealing interpretation of the relation between the congestion control protocol utilized by the flows (the end-users) and the queues formed inside the network. It is natural to expect that this product-form steady-state distribution is the limit of the steady-state distributions in the original model under the diffusion scaling. Results of this type are known for certain queueing systems such as generalized Jackson networks; see the work by Gamarnik and Zeevi [10]. On the other hand, the validity of such a steady-state diffusion approximation was not known for the model considered in [13]; it will be established in the present paper.

1.2. *Our contributions.* In this paper, we advance the performance analysis of flow-level models of networks operating under an α -fair policy, in both the steady-state and the transient regimes.

For the transient regime, we obtain a probabilistic bound on the maximal (over a finite time horizon) number of flows, when operating under an α -fair policy with $\alpha \geq 1$. This result is obtained by combining a Lyapunov drift inequality with a natural extension of Doob's maximal inequality for non-negative supermartingales. Our probabilistic bound, together with prior results on multiplicative state space collapse, leads immediately to a stronger property, namely, full state space collapse, for the case where $\alpha \geq 1$.

For the steady-state regime, we obtain non-asymptotic and explicit bounds on the tail of the distribution of the number of flows, for any $\alpha > 0$. In the process, we establish that, for any $\alpha > 0$, all moments of the steady-state number of flows are finite. These results are proved by working with a *normed* version of the Lyapunov function that was used in prior work. Specifically, we establish that this normed version is also a Lyapunov function for the system (i.e., it satisfies a drift inequality). It also happens to be a Lipschitz continuous function and this helps crucially in establishing exponential tail bounds, using results of Hajek [12] and Bertsimas et al. [2].

The exponent in the exponential tail bound that we establish for the distribution of the number of flows is proportional to a suitably defined distance ("gap") from critical loading; this gap is of the same type as the familiar $1 - \rho$ term, where ρ is the usual load factor in a queueing system. This particular dependence on the load leads to the tightness of the steady-state distributions of the model under diffusion scaling. It leads to one of our main results, namely, the validity of the diffusion approximation, in steady state, when $\alpha = 1$ and a local traffic condition holds.

1.3. *Organization.* The rest of the paper is organized as follows. In Section 2, we define the notation and some of the terminology that we will employ. We also describe the flow-level network model, as well as the weighted α -fair bandwidth-sharing policies. In Section 3, we provide formal statements of our main results. The transient analysis is presented in Section 4. We start with a general lemma, and specialize it to obtain a maximal inequality under α -fair policies, when $\alpha \geq 1$. We then apply the latter inequality to prove full state space collapse when $\alpha \geq 1$.

We then proceed to the steady-state analysis. In Section 5, we establish a drift inequality for a suitable Lyapunov function, which is central to our proof of exponential upper bounds on tail probabilities. We prove the exponential upper bound on tail probabilities in Section 6. The validity of

the heavy-traffic steady-state approximation is established in Section 7. We conclude the paper with a brief discussion in Section 8.

2. Model and Notation.

2.1. Notation. We introduce here the notation that will be employed throughout the paper. We denote the real vector space of dimension M by $\mathbb{R}^M$, the set of nonnegative M -tuples by $\mathbb{R}_+^M$, and the set of positive M -tuples by $\mathbb{R}_p^M$. We write $\mathbb{R}$ for $\mathbb{R}^1$, $\mathbb{R}_+$ for $\mathbb{R}_+^1$, and $\mathbb{R}_p$ for $\mathbb{R}_p^1$. We let $\mathbb{Z}$ be the set of integers, $\mathbb{Z}_+$ the set of nonnegative integers, and $\mathbb{N}$ the set of positive integers. Throughout the paper, we reserve bold letters for vectors and plain letters for scalars.

For any vector $\mathbf{x} \in \mathbb{R}^M$, and any $\alpha > 0$, we define

$$\|\mathbf{x}\|_\alpha = \left(\sum_{i=1}^M |x_i|^\alpha \right)^{1/\alpha},$$

and we define $\|\mathbf{x}\|_\infty = \max_{i \in \{1, \dots, M\}} |x_i|$. For any two vectors $\mathbf{x} = (x_i)_{i=1}^M$ and $\mathbf{y} = (y_i)_{i=1}^M$ of the same dimensions, we let $\langle \mathbf{x}, \mathbf{y} \rangle = \sum_{i=1}^M x_i y_i$ be the inner product of $\mathbf{x}$ and $\mathbf{y}$. We let $\mathbf{e}_i$ be the i -th unit vector in $\mathbb{R}^M$, and $\mathbf{1}$ the vector of all ones. For a set $\mathcal{S}$, we denote its cardinality by $|\mathcal{S}|$, and its indicator function by $\mathbb{I}_{\mathcal{S}}$. For a matrix $\mathbf{A}$, we let $\mathbf{A}^T$ denote its transpose.

2.2. Flow-Level Network Model.

The Model. We adopt the model and notation in [14]. As explained in detail in [14], this model faithfully captures the long-term (or macro level) behavior of congestion control in the current Internet.

Let time be continuous and indexed by $t \in \mathbb{R}_+$. Consider a network with a finite set $\mathcal{J}$ of resources and a set $\mathcal{I}$ of routes, where a route is identified with a non-empty subset of the resource set $\mathcal{J}$. Let $\mathbf{A}$ be the $|\mathcal{J}| \times |\mathcal{I}|$ matrix with $A_{ji} = 1$ if resource j is used by route i , and $A_{ji} = 0$ otherwise. Assume that $\mathbf{A}$ has rank $|\mathcal{J}|$. Let $\mathbf{C} = (C_j)_{j \in \mathcal{J}}$ be a *capacity* vector, where we assume that each entry C_j is a given positive constant. Let the number of flows on route i at time t be denoted by $N_i(t)$, and define the flow vector at time t by $\mathbf{N}(t) = (N_i(t))_{i \in \mathcal{I}}$. For each route i , new flows arrive as an independent Poisson process of rate ν_i . Each arriving flow brings an amount of work (data that it wishes to transfer) which is an exponentially distributed random variable with mean $1/\mu_i$, independent of everything else. Each flow gets service from the network according to a bandwidth-sharing policy. Once a flow is served, it departs the network.

The α -Fair Bandwidth-Sharing Policy. A bandwidth sharing policy has to allocate rates to flows so that capacity constraints are satisfied at each time instance. Here we discuss the popular α -fair bandwidth-sharing policy, where $\alpha > 0$. At any time, the bandwidth allocation depends on the current number of flows $\mathbf{n} = (n_i)_{i \in \mathcal{I}}$. Let Λ_i be the *total* bandwidth allocated to route i under the α -fair policy: each flow of type i gets rate Λ_i/n_i if $n_i > 0$, and $\Lambda_i = 0$ if $n_i = 0$. Under an α -fair policy, the bandwidth vector $\mathbf{\Lambda}(\mathbf{n}) = (\Lambda_i(\mathbf{n}))_{i \in \mathcal{I}}$ is determined as follows.

If $\mathbf{n} = \mathbf{0}$, then $\mathbf{\Lambda} = \mathbf{0}$. If $\mathbf{n} \neq \mathbf{0}$, then let $\mathcal{I}_+(\mathbf{n}) = \{i \in \mathcal{I} : n_i > 0\}$. For $i \notin \mathcal{I}_+(\mathbf{n})$, set $\Lambda_i(\mathbf{n}) = 0$. Let $\mathbf{\Lambda}_+(\mathbf{n}) = (\Lambda_i(\mathbf{n}))_{i \in \mathcal{I}_+(\mathbf{n})}$. Then, $\mathbf{\Lambda}_+(\mathbf{n})$ is the *unique* maximizer in the optimization problem

$$\begin{aligned} (1) \quad & \text{maximize} && G_n(\mathbf{\Lambda}_+) && \text{over } \mathbf{\Lambda} \in \mathbb{R}_+^{|\mathcal{I}|} \\ (2) \quad & \text{subject to} && \sum_{i \in \mathcal{I}_+(\mathbf{n})} A_{ji} \Lambda_i \leq C_j, && \forall j \in \mathcal{J}, \end{aligned}$$

where

$$G_n(\mathbf{\Lambda}_+) = \begin{cases} \sum_{i \in \mathcal{I}_+(\mathbf{n})} \kappa_i n_i^\alpha \frac{\Lambda_i^{1-\alpha}}{1-\alpha}, & \text{if } \alpha \in (0, \infty) \setminus \{1\}, \\ \sum_{i \in \mathcal{I}_+(\mathbf{n})} \kappa_i n_i \log \Lambda_i, & \text{if } \alpha = 1. \end{cases}$$

Here, for each $i \in \mathcal{I}$, κ_i is a positive weight assigned to route i .

Some crucial properties of $\mathbf{\Lambda}(\mathbf{n})$ are as follows (see Appendix A of [14]):

- (i) $\Lambda_i(\mathbf{n}) > 0$ for every $i \in \mathcal{I}_+(\mathbf{n})$;
- (ii) $\mathbf{\Lambda}(r\mathbf{n}) = \mathbf{\Lambda}(\mathbf{n})$ for $r > 0$;
- (iii) For every $\mathbf{n}$ and every $i \in \mathcal{I}_+(\mathbf{n})$, the function $\Lambda_i(\cdot)$ is continuous at $\mathbf{n}$.

Flow Dynamics. The flow dynamics are described by the evolution of the flow vector $\mathbf{N}(t) = (N_i(t))_{i \in \mathcal{I}}$, a Markov process with infinitesimal transition rate matrix $\mathbf{q}$ given by

$$(3) \quad q(\mathbf{n}, \mathbf{n} + \mathbf{m}) = \begin{cases} \nu_i, & \text{if } \mathbf{m} = \mathbf{e}_i, \\ \mu_i \Lambda_i(\mathbf{n}), & \text{if } \mathbf{m} = -\mathbf{e}_i, \text{ and } n_i \geq 1, \\ 0, & \text{otherwise,} \end{cases}$$

where for each i , $\nu_i > 0$ and $\mu_i > 0$ are the arrival and service rates defined earlier, and $\mathbf{e}_i$ is the i -th unit vector.

Capacity Region. Flows of type i bring to the system an average of $\rho_i = \nu_i/\mu_i$ units of work per unit time. Therefore, in order for the Markov process $\mathbf{N}(\cdot)$ to be positive recurrent, it is necessary that

$$(4) \quad \mathbf{A}\boldsymbol{\rho} < \mathbf{C}, \text{ componentwise.}$$

We note that under the α -fair bandwidth-sharing policy, Condition (4) is also sufficient for positive recurrence of the process $\mathbf{N}(\cdot)$ [4, 6, 14].

2.3. A note on our use of constants. Our results and proofs involve various constants; some are absolute constants, some depend only on the structure of the network, and some depend (smoothly) on the traffic parameters (the arrival and service rates). It is convenient to distinguish between the different types of constants, and we define here the terminology that we will be using.

The term *absolute constant* will be used to refer to a quantity that does not depend on any of the model parameters. The term *network-dependent constant* will be used to refer to quantities that are completely determined by the structure of the underlying network and policy, namely, the incidence matrix $\mathbf{A}$, the capacity vector $\mathbf{C}$, the weight vector $\boldsymbol{\kappa}$, and the policy parameter α .

Our analysis also involves certain quantities that depend on the traffic parameters, namely, the arrival and service parameters $\boldsymbol{\mu}$ and $\boldsymbol{\nu}$. These quantities are often given by complicated expressions that would be inconvenient to carry through the various arguments. It turns out that the only property of such quantities that is relevant to our purposes is the fact they change continuously as $\boldsymbol{\mu}$ and $\boldsymbol{\nu}$ vary over the open positive orthant. (This still allows these quantities to be undefined or discontinuous on the boundary of the positive orthant.) We abstract this property by introducing, in the definition that follows, the concept of a *(positive) load-dependent constant*.

DEFINITION 2.1. *Consider a family of bandwidth-sharing networks with common parameters $(\mathbf{A}, \mathbf{C}, \boldsymbol{\kappa}, \alpha)$, but varying traffic parameters $(\boldsymbol{\mu}, \boldsymbol{\nu})$. A quantity K will be called a (positive) load-dependent constant if for networks in that family it is determined by a relation of the form $K = f(\boldsymbol{\mu}, \boldsymbol{\nu})$, where $f : \mathbb{R}_p^{|\mathcal{I}|} \times \mathbb{R}_p^{|\mathcal{I}|} \rightarrow \mathbb{R}_p$ is a continuous function on the open positive orthant $\mathbb{R}_p^{|\mathcal{I}|} \times \mathbb{R}_p^{|\mathcal{I}|}$.*

A key property of a load-dependent constant, which will be used in some of the subsequent proofs, is that it is by definition positive and furthermore (because of continuity), bounded above and below by positive network-dependent constants if we restrict $\boldsymbol{\mu}$ and $\boldsymbol{\nu}$ to a compact subset of the open

positive orthant. A natural example of a load-dependent constant is the load factor $\rho_i = \nu_i/\mu_i$. (Note that this quantity diverges as $\mu_i \rightarrow 0$.)

We also define the *gap* of a underloaded bandwidth-sharing network.

DEFINITION 2.2. *Consider a family of bandwidth-sharing networks with common parameters $(\mathbf{A}, \mathbf{C}, \boldsymbol{\kappa}, \alpha)$ and with varying traffic parameters $(\boldsymbol{\mu}, \boldsymbol{\nu})$ that satisfy $\mathbf{A}\boldsymbol{\rho} < \mathbf{C}$. The gap of a network with traffic parameters $(\boldsymbol{\mu}, \boldsymbol{\nu})$ in the family, denoted by $\varepsilon(\boldsymbol{\rho})$, is defined by*

$$\varepsilon(\boldsymbol{\rho}) \triangleq \sup\{\tilde{\varepsilon} > 0 : (1 + \tilde{\varepsilon})\mathbf{A}\boldsymbol{\rho} \leq \mathbf{C}\}.$$

We sometimes write ε for $\varepsilon(\boldsymbol{\rho})$ when there is no ambiguity. Note also that $\varepsilon(\boldsymbol{\rho})$ plays the same role as the term $1 - \rho$ in a queueing system with load ρ .

2.4. Uniformization. Uniformization is a well-known device which allows us to study a continuous-time Markov process by considering an associated discrete-time Markov chain with the same stationary distribution. We provide here some details and the notation that we will be using.

Recall that the Markov process $\mathbf{N}(\cdot)$ of interest has dynamics given by (3). Let $\Xi(\mathbf{n}) = \sum_{\tilde{\mathbf{n}}} q(\mathbf{n}, \tilde{\mathbf{n}})$ be the aggregate transition rate at state $\mathbf{n}$. The *embedded jump chain* of $\mathbf{N}(\cdot)$ is a discrete-time Markov chain with the same state space $\mathbb{Z}_+^{|\mathcal{I}|}$, and with transition probability matrix $\mathbf{P}$ given by

$$P(\mathbf{n}, \tilde{\mathbf{n}}) = \frac{q(\mathbf{n}, \tilde{\mathbf{n}})}{\Xi(\mathbf{n})}.$$

The so-called *uniformized Markov chain* is an alternative, more convenient, discrete-time Markov chain, denoted $\left(\tilde{\mathbf{N}}(\tau)\right)_{\tau \in \mathbb{Z}_+}$, to be defined shortly.

We first introduce some more notation. Consider the aggregate transition rates $\Xi(\mathbf{n}) = \sum_{\tilde{\mathbf{n}}} q(\mathbf{n}, \tilde{\mathbf{n}})$. Since every route uses at least one resource, we have $\Lambda_i(\mathbf{n}) \leq \max_{j \in \mathcal{J}} C_j$, for all $i \in \mathcal{I}$. Then, by (3), we have

$$\Xi(\mathbf{n}) = \sum_{\tilde{\mathbf{n}}} q(\mathbf{n}, \tilde{\mathbf{n}}) \leq \sum_{i \in \mathcal{I}} (\nu_i + \mu_i \Lambda_i(\mathbf{n})) \leq \sum_{i \in \mathcal{I}} \left(\nu_i + \mu_i \max_{j \in \mathcal{J}} C_j \right).$$

We define $\Xi \triangleq \sum_{i \in \mathcal{I}} (\nu_i + \mu_i \max_{j \in \mathcal{J}} C_j)$, and modify the rates of self-transitions (which were zero in the original model) to

$$(5) \quad q(\mathbf{n}, \mathbf{n}) := \Xi - \Xi(\mathbf{n}).$$

Note that Ξ is a positive load-dependent constant. We define a transition probability matrix $\tilde{\mathbf{P}}$ by

$$\tilde{P}(\mathbf{n}, \tilde{\mathbf{n}}) \triangleq \frac{q(\mathbf{n}, \tilde{\mathbf{n}})}{\Xi}.$$

DEFINITION 2.3. *The uniformized Markov chain $(\tilde{\mathbf{N}}(\tau))_{\tau \in \mathbb{Z}_+}$ associated with the Markov process $\mathbf{N}(\cdot)$ is a discrete-time Markov chain with the same state space $\mathbb{Z}_+^{|\mathcal{I}|}$, and with transition matrix $\tilde{\mathbf{P}}$ defined as above.*

As remarked earlier, the Markov process $\mathbf{N}(\cdot)$ that describes a bandwidth-sharing network operating under an α -fair policy is positive recurrent, as long as the system is underloaded, i.e., if $\mathbf{A}\boldsymbol{\rho} < \mathbf{C}$. It is not hard to verify that $\mathbf{N}(\cdot)$ is also irreducible. Therefore, the Markov process $\mathbf{N}(\cdot)$ has a unique stationary distribution. The chain $\tilde{\mathbf{N}}(\cdot)$ is also positive recurrent and irreducible, because $\mathbf{N}(\cdot)$ is, and by suitably increasing Ξ if necessary, it can be made aperiodic. Thus $\tilde{\mathbf{N}}(\cdot)$ has a unique stationary distribution as well. A crucial property of the uniformized chain $\tilde{\mathbf{N}}(\cdot)$ is that this unique stationary distribution is the same as that of the original Markov process $\mathbf{N}(\cdot)$; see, e.g., [9].

2.5. *A Mean Value Theorem.* We will be making extensive use of a second-order mean value theorem [1], which we state below for easy reference.

PROPOSITION 2.4. *Let $g : \mathbb{R}^M \rightarrow \mathbb{R}$ be twice continuously differentiable over an open sphere S centered at a vector $\mathbf{x}$. Then, for any $\mathbf{y}$ such that $\mathbf{x} + \mathbf{y} \in S$, there exists $\theta \in [0, 1]$ such that*

$$(6) \quad g(\mathbf{x} + \mathbf{y}) = g(\mathbf{x}) + \mathbf{y}^T \nabla g(\mathbf{x}) + \frac{1}{2} \mathbf{y}^T H(\mathbf{x} + \theta \mathbf{y}) \mathbf{y},$$

where $\nabla g(\mathbf{x}) \triangleq \left[\frac{\partial g(\mathbf{x})}{\partial x_i} \right]_{i=1}^M \in \mathbb{R}^M$ is the gradient of g at $\mathbf{x}$, and $H(\mathbf{x}) \triangleq \left[\frac{\partial^2 g(\mathbf{x})}{\partial x_i \partial x_j} \right]_{i,j=1}^M \in \mathbb{R}^{M \times M}$ is the Hessian of the function g at $\mathbf{x}$.

3. Summary of Results. In this section, we summarize our main results for both the transient and the steady-state regime. The proofs are given in subsequent sections.

3.1. *Transient Regime.* Here we provide a simple inequality on the maximal excursion of the number of flows over a finite time interval, under an α -fair policy with $\alpha \geq 1$.

THEOREM 3.1. *Consider a bandwidth-sharing network operating under an α -fair policy with $\alpha \geq 1$, and assume that $\mathbf{A}\boldsymbol{\rho} < \mathbf{C}$. Suppose that $\mathbf{N}(0) =$*

0. Let $N^*(T) = \sup_{t \in [0, T], i \in \mathcal{I}} N_i(t)$, and let ε be the gap. Then, for any $b > 0$,

$$(7) \quad \mathbb{P}(N^*(T) \geq b) \leq \frac{KT}{\varepsilon^{\alpha-1} b^{\alpha+1}},$$

for some positive load-dependent constant K .

As an important application, in Section 4.3, we will use Theorem 3.1 to prove a full state space collapse result, when $\alpha \geq 1$. (As discussed in the introduction, this property is stronger than multiplicative state space collapse.) The precise statement can be found in Theorem 4.10.

3.2. Stationary Regime. As noted earlier, the Markov process $\mathbf{N}(\cdot)$ has a unique stationary distribution, which we will denote by π . We use $\mathbb{E}_\pi$ and $\mathbb{P}_\pi$ to denote expectations and probabilities under π .

Exponential Bound on Tail Probabilities. For an α -fair policy, and for any $\alpha \in (0, \infty)$, we obtain an explicit exponential upper bound on the tail probabilities for the number of flows, in steady state. This will be used to establish an “interchange of limits” result in Section 7. See Theorem 7.6 for more details.

THEOREM 3.2. *Consider a bandwidth-sharing network operating under an α -fair policy with $\alpha > 0$, and assume that $\mathbf{A}\rho < \mathbf{C}$. Let ε be the gap. There exist positive constants B , K , and ξ such that for all $\ell \in \mathbb{Z}_+$:*

$$(8) \quad \mathbb{P}_\pi(\|\mathbf{N}\|_\infty \geq B + 2\xi\ell) \leq \left(\frac{\xi}{\xi + \varepsilon K} \right)^{\ell+1}.$$

Here ξ and K are load-dependent constants, and B takes the form K'/ε when $\alpha \geq 1$, and $K'/\min\{\varepsilon^{1/\alpha}, \varepsilon\}$ when $\alpha \in (0, 1)$, with K' being a positive load-dependent constant. In particular, all moments of $\|\mathbf{N}\|_\infty$ are finite under the stationary distribution π , i.e., $\mathbb{E}_\pi[\|\mathbf{N}\|_\infty^k] < \infty$ for every $k \in \mathbb{N}$.

Here we note that Theorem 3.2 implies the following. The system load $L(\rho)$, defined by $L(\rho) \triangleq \frac{1}{1+\varepsilon(\rho)}$, satisfies $L(\rho) \approx 1 - \varepsilon(\rho)$ when $\varepsilon = \varepsilon(\rho)$ is small, i.e., when the system approaches criticality. Then, an immediate consequence of the bound (8) is that

$$\begin{aligned} \limsup_{\gamma \rightarrow \infty} \frac{1}{\gamma} \log \mathbb{P}_\pi(\|\mathbf{N}\|_\infty \geq \gamma) &\lesssim \frac{1}{2\xi} \log \left(\frac{\xi}{\xi + \varepsilon K} \right) \\ &\approx -\frac{K\varepsilon}{2\xi^2} \approx -\frac{K}{2\xi^2} (1 - L(\rho)). \end{aligned}$$

Note that $\frac{K}{2\xi^2}$ is a load-dependent constant. Thus Theorem 3.2 shows that the large-deviations exponent of the steady-state number of flows is upper bounded by $-(1 - L(\rho))$, up to a multiplicative constant.

Interchange of Limits ($\alpha = 1$). As discussed in the introduction, when $\alpha = 1$, Theorem 3.2 leads to the tightness (Lemma 7.7) of the steady-state distributions of the model under diffusion scaling. This in turn leads to Theorem 7.6 and Corollary 7.10, on the validity of the diffusion approximation in steady state. As the statements of these results require a significant amount of preliminary notation and background (which is introduced in Section 7), we give here an informal statement.

INTERCHANGE OF LIMITS THEOREM (informal statement): *Consider a sequence of flow-level networks operating under the proportionally fair policy. Let $\mathbf{N}^r(\cdot)$ be the flow-vector Markov process associated with the r th network, let ε^r be the corresponding gap, and let $\hat{\pi}^r$ be the stationary distribution of $\varepsilon^r \mathbf{N}^r(\cdot)$. As $\varepsilon^r \rightarrow 0$, and under certain technical conditions, $\hat{\pi}^r$ converges weakly to the stationary distribution of an associated limiting process.*

4. Transient Analysis ($\alpha \geq 1$). In this section, we present a transient analysis of the α -fair policies with $\alpha \geq 1$. First we present a general maximal lemma, which we then specialize to our model. In particular, we prove a refined drift inequality for the Lyapunov function given by

$$(9) \quad F_\alpha(\mathbf{n}) = \frac{1}{\alpha + 1} \sum_{i \in I} \nu_i \kappa_i \mu_i^{\alpha-1} \left(\frac{n_i}{\nu_i} \right)^{\alpha+1}.$$

This Lyapunov function and associated drift inequalities have played an important role in establishing positive recurrence (cf. [4], [6], [14]) and multiplicative state space collapse (cf. [13]) for α -fair policies. We combine our drift inequality with the maximal lemma to obtain a maximal inequality for bandwidth-sharing networks. We then apply the maximal inequality to prove full state space collapse when $\alpha \geq 1$.

4.1. The Key Lemma. Our analysis relies on the following lemma.

LEMMA 4.1. *Let $(\mathcal{F}_n)_{n \in \mathbb{Z}_+}$ be a filtration on a probability space. Let $(X_n)_{n \in \mathbb{Z}_+}$ be a nonnegative $\mathcal{F}_n$ -adapted stochastic process that satisfies*

$$(10) \quad \mathbb{E}[X_{n+1} \mid \mathcal{F}_n] \leq X_n + B_n$$

where the B_n are nonnegative random variables (not necessarily $\mathcal{F}_n$ -adapted) with finite means. Let $X_n^ = \max\{X_0, \dots, X_n\}$ and suppose that $X_0 = 0$.*

Then, for any $a > 0$ and any $T \in \mathbb{Z}_+$,

$$\mathbb{P}(X_T^* \geq a) \leq \frac{\sum_{n=0}^{T-1} \mathbb{E}[B_n]}{a}.$$

This lemma is a simple consequence of the following standard maximal inequality for nonnegative supermartingales (see for example, Exercise 4, Section 12.4, of [11]).

THEOREM 4.2. *Let $(\mathcal{F}_n)_{n \in \mathbb{Z}_+}$ be a filtration on a probability space. Let $(Y_n)_{n \in \mathbb{Z}_+}$ be a nonnegative $\mathcal{F}_n$ -adapted supermartingale, i.e., for all n ,*

$$\mathbb{E}[Y_{n+1} \mid \mathcal{F}_n] \leq Y_n.$$

Let $Y_T^* = \max\{Y_0, \dots, Y_T\}$. Then,

$$\mathbb{P}(Y_T^* \geq a) \leq \frac{\mathbb{E}[Y_0]}{a}.$$

PROOF OF LEMMA 4.1. First note that if we take the conditional expectation of both sides of (10), given $\mathcal{F}_n$, we have

$$\mathbb{E}[X_{n+1} \mid \mathcal{F}_n] \leq \mathbb{E}[X_n \mid \mathcal{F}_n] + \mathbb{E}[B_n \mid \mathcal{F}_n] = X_n + \mathbb{E}[B_n \mid \mathcal{F}_n].$$

Fix $T \in \mathbb{Z}_+$. For any $n \leq T$, define

$$Y_n = X_n + \mathbb{E} \left[\sum_{k=n}^{T-1} B_k \mid \mathcal{F}_n \right].$$

Then

$$\begin{aligned} \mathbb{E}[Y_{n+1} \mid \mathcal{F}_n] &= \mathbb{E}[X_{n+1} \mid \mathcal{F}_n] + \mathbb{E} \left[\mathbb{E} \left[\sum_{k=n+1}^{T-1} B_k \mid \mathcal{F}_{n+1} \right] \mid \mathcal{F}_n \right] \\ &\leq X_n + \mathbb{E}[B_n \mid \mathcal{F}_n] + \mathbb{E} \left[\sum_{k=n+1}^{T-1} B_k \mid \mathcal{F}_n \right] = Y_n. \end{aligned}$$

Thus, Y_n is an $\mathcal{F}_n$ -adapted supermartingale; furthermore, by definition, Y_n is nonnegative for all n . Therefore, by Theorem 4.2,

$$\mathbb{P}(Y_T^* \geq a) \leq \frac{\mathbb{E}[Y_0]}{a} = \frac{\mathbb{E} \left[\sum_{k=0}^{T-1} B_k \right]}{a}.$$

But $Y_n \geq X_n$ for all n , since the B_k are nonnegative. Thus,

$$\mathbb{P}(X_T^* \geq a) \leq \mathbb{P}(Y_T^* \geq a) \leq \frac{\mathbb{E} \left[\sum_{k=0}^{T-1} B_k \right]}{a}.$$

□

Since we are dealing with continuous-time Markov processes, the following corollary of Lemma 4.1 will be useful for our analysis.

COROLLARY 4.3. *Let $(\mathcal{F}_t)_{t \geq 0}$ be a filtration on a probability space. Let Z_t be a nonnegative, right-continuous $\mathcal{F}_t$ -adapted stochastic process that satisfies*

$$\mathbb{E}[Z_{s+t} | \mathcal{F}_s] \leq Z_s + Bt,$$

for all $s, t \geq 0$, where B is a nonnegative constant. Assume that $Z_0 \equiv 0$. Denote $Z_T^ \triangleq \sup_{0 \leq t \leq T} Z_t$ (which can possibly be infinite). Then, for any $a > 0$, and for any $T \geq 0$,*

$$\mathbb{P}(Z_T^* \geq a) \leq \frac{BT}{a}.$$

PROOF. The proof is fairly standard. We fix $T \geq 0$ and $a > 0$. Since Z_t is right-continuous, $Z_T^* = \sup_{t \in [0, T]} Z_t = \sup_{t \in ([0, T] \cap \mathbb{Q}) \cup \{T\}} Z_t$. Consider an increasing sequence of finite sets I_n so that $\cup_{n=1}^{\infty} I_n = ([0, T] \cap \mathbb{Q}) \cup \{T\}$, and $0, T \in I_n$ for all n . Define $Z_T^{(n)} = \sup_{t \in I_n} Z_t$. Then $(Z_T^{(n)})_{n=1}^{\infty}$ is a non-decreasing sequence, and $Z_T^{(n)} \rightarrow Z_T^*$ as $n \rightarrow \infty$, almost surely. For each $Z_T^{(n)}$, we can apply Lemma 4.1, and it is immediate that for any $b > 0$,

$$(11) \quad \mathbb{P}(Z_T^{(n)} > b) \leq \frac{BT}{b},$$

since each I_n includes both 0 and T . Since $Z_T^{(n)}$ increases monotonically to Z_T^* , almost surely, we have that $\mathbb{P}(Z_T^{(n)} > b) \leq \mathbb{P}(Z_T^{(n+1)} > b)$ for all n , and $\mathbb{P}(Z_T^{(n)} > b) \rightarrow \mathbb{P}(Z_T^* > b)$ as $n \rightarrow \infty$. The right-hand side of (11) is fixed, so

$$\mathbb{P}(Z_T^* > b) \leq \frac{BT}{b}.$$

We now take an increasing sequence b_n with $\lim_{n \rightarrow \infty} b_n = a$, and obtain

$$\mathbb{P}(Z_T^* \geq a) \leq \frac{BT}{a}.$$

□

4.2. *A Maximal Inequality for Bandwidth-Sharing Networks.* We employ the Lyapunov function (9) to study α -fair policies. This is the Lyapunov function that was used in [4], [6] and [14] to establish positive recurrence of the process $\mathbf{N}(\cdot)$ under an α -fair policy. Below we fine-tune the proof in [6] to obtain a more precise bound on the Lyapunov drift. We note that a “fluid-model” version of the following lemma appeared in the proof of Theorem 1 in [4]. For notational convenience, we drop the subscript α from F_α and write F instead.

LEMMA 4.4. *Consider a bandwidth-sharing network with $\mathbf{A}\boldsymbol{\rho} < \mathbf{C}$ operating under an α -fair policy with $\alpha > 0$. Let ε be the gap. Then, for any non-zero flow vector $\mathbf{n}$,*

$$\langle \nabla F(\mathbf{n}), \boldsymbol{\nu} - \boldsymbol{\mu}\boldsymbol{\Lambda}(\mathbf{n}) \rangle \leq -\varepsilon \langle \nabla F(\mathbf{n}), \boldsymbol{\nu} \rangle,$$

where $\langle \cdot, \cdot \rangle$ denotes the standard inner product, $\nabla F(\mathbf{n})$ denotes the gradient of F , and $\boldsymbol{\mu}\boldsymbol{\Lambda}(\mathbf{n})$ is the vector $(\mu_i \Lambda_i(\mathbf{n}))_{i \in \mathcal{I}}$.

PROOF. We have

$$\begin{aligned} \langle \nabla F(\mathbf{n}), \boldsymbol{\nu} - \boldsymbol{\mu}\boldsymbol{\Lambda}(\mathbf{n}) \rangle &= \sum_{i \in \mathcal{I}} \frac{1}{\mu_i} \kappa_i \left(\frac{n_i}{\rho_i} \right)^\alpha (\nu_i - \mu_i \Lambda_i(\mathbf{n})) \\ &= \sum_{i \in \mathcal{I}} \kappa_i \left(\frac{n_i}{\rho_i} \right)^\alpha (\rho_i - \Lambda_i(\mathbf{n})) \\ &= \langle \nabla G_{\mathbf{n}}(\boldsymbol{\rho}_+), \boldsymbol{\rho}_+ - \boldsymbol{\Lambda}_+(\mathbf{n}) \rangle, \end{aligned}$$

where $\boldsymbol{\rho}_+ = (\rho_i)_{i \in \mathcal{I}_+(\mathbf{n})}$. Similarly we can get $\langle \nabla F(\mathbf{n}), \boldsymbol{\nu} \rangle = \langle \nabla G_{\mathbf{n}}(\boldsymbol{\rho}_+), \boldsymbol{\rho}_+ \rangle$.

Now consider the function $g : [0, 1] \rightarrow \mathbb{R}$ defined by

$$g(\theta) = G_{\mathbf{n}}(\theta(1 + \varepsilon)\boldsymbol{\rho}_+ + (1 - \theta)\boldsymbol{\Lambda}_+(\mathbf{n})).$$

Since $(1 + \varepsilon)\boldsymbol{\rho}_+$ satisfies the constraints in (2), and $\boldsymbol{\Lambda}_+(\mathbf{n})$ maximizes the strictly concave function $G_{\mathbf{n}}$ subject to the constraints in (2), we have

$$G_{\mathbf{n}}((1 + \varepsilon)\boldsymbol{\rho}_+) \leq G_{\mathbf{n}}(\boldsymbol{\Lambda}_+(\mathbf{n})), \quad \text{i.e., } g(1) \leq g(0).$$

Furthermore, since $G_{\mathbf{n}}$ is a concave function, g is also concave in θ . Thus,

$$g(0) \leq g(1) + (0 - 1)g'(1) \leq g(0) + (0 - 1)g'(1).$$

Hence, $g'(1) \leq 0$, i.e.,

$$(12) \quad \left. \frac{dg}{d\theta} \right|_{\theta=1} = \langle \nabla G_{\mathbf{n}}((1 + \varepsilon)\boldsymbol{\rho}_+), (1 + \varepsilon)\boldsymbol{\rho}_+ - \boldsymbol{\Lambda}_+(\mathbf{n}) \rangle \leq 0.$$

But it is easy to check that $\nabla G_{\mathbf{n}}((1+\varepsilon)\boldsymbol{\rho}_+) = (1+\varepsilon)^{-\alpha}\nabla G_{\mathbf{n}}(\boldsymbol{\rho}_+)$, so dividing (12) by $(1+\varepsilon)^{-\alpha}$, we have

$$\langle \nabla G_{\mathbf{n}}(\boldsymbol{\rho}_+), \boldsymbol{\rho}_+ - \boldsymbol{\Lambda}_+(\mathbf{n}) \rangle \leq -\varepsilon \langle \nabla G_{\mathbf{n}}(\boldsymbol{\rho}_+), \boldsymbol{\rho}_+ \rangle.$$

This is the same as

$$\langle \nabla F(\mathbf{n}), \boldsymbol{\nu} - \boldsymbol{\mu}\boldsymbol{\Lambda}(\mathbf{n}) \rangle \leq -\varepsilon \langle \nabla F(\mathbf{n}), \boldsymbol{\nu} \rangle.$$

□

Our next lemma provides a uniform upper bound on the expected change of $F(\tilde{\mathbf{N}}(\cdot))$ in one time step, where $\tilde{\mathbf{N}}(\cdot)$ is the uniformized chain associated with the Markov process $\mathbf{N}(\cdot)$ (cf. Definition 2.3).

LEMMA 4.5. *Let $\alpha \geq 1$. As above, consider a bandwidth-sharing network with $\mathbf{A}\boldsymbol{\rho} < \mathbf{C}$ operating under an α -fair policy. Let ε be the gap. Let $(\tilde{\mathbf{N}}(\tau))_{\tau \in \mathbb{Z}_+}$ be the uniformized chain associated with the Markov process $\mathbf{N}(\cdot)$. Then, there exists a positive load-dependent constant $\bar{K}$, such that for all $\tau \in \mathbb{Z}_+$,*

$$\mathbb{E} \left[F(\tilde{\mathbf{N}}(\tau+1)) - F(\mathbf{n}) \mid \tilde{\mathbf{N}}(\tau) = \mathbf{n} \right] \leq \bar{K}\varepsilon^{1-\alpha}.$$

PROOF. By the mean value theorem (cf. Proposition 2.4), for $\mathbf{n}, \mathbf{m} \in \mathbb{Z}_+^{|\mathcal{I}|}$, we have

$$(13) \quad F(\mathbf{n} + \mathbf{m}) - F(\mathbf{n}) = \langle \nabla F(\mathbf{n}), \mathbf{m} \rangle + \frac{1}{2} \mathbf{m}^T \nabla^2 F(\mathbf{n} + \theta \mathbf{m}) \mathbf{m},$$

for some $\theta \in [0, 1]$. We note that, for $\mathbf{m} = \pm \mathbf{e}_i$, we have

$$(14) \quad \begin{aligned} \frac{1}{2} \mathbf{m}^T \nabla^2 F(\mathbf{n} + \theta \mathbf{m}) \mathbf{m} &\leq \frac{\kappa_i \alpha}{2\mu_i \rho_i^\alpha} (n_i \pm \theta)^{\alpha-1} \\ &\leq \frac{\kappa_i \alpha}{2\mu_i \rho_i^\alpha} (n_i + 1)^{\alpha-1}, \end{aligned}$$

since $\alpha \geq 1$, and $\theta \in [0, 1]$.

As in [6], we define

$$\mathbf{Q}F(\mathbf{n}) \triangleq \sum_{\mathbf{m}} q(\mathbf{n}, \mathbf{n} + \mathbf{m}) [F(\mathbf{n} + \mathbf{m}) - F(\mathbf{n})],$$

so that $\mathbf{Q}$ is the generator of the Markov process $\mathbf{N}(\cdot)$. We now proceed to derive an upper bound for $\mathbf{Q}F(\mathbf{n})$. Using Equation (13), we can rewrite $\mathbf{Q}F(\mathbf{n})$ as

$$\begin{aligned}\mathbf{Q}F(\mathbf{n}) &= \sum_{\mathbf{m}} q(\mathbf{n}, \mathbf{n} + \mathbf{m}) \left[\langle \nabla F(\mathbf{n}), \mathbf{m} \rangle + \frac{1}{2} \mathbf{m}^T \nabla^2 F(\mathbf{n} + \theta_{\mathbf{m}} \mathbf{m}) \mathbf{m} \right] \\ &= \sum_{\mathbf{m}} q(\mathbf{n}, \mathbf{n} + \mathbf{m}) \langle \nabla F(\mathbf{n}), \mathbf{m} \rangle \\ &\quad + \frac{1}{2} \sum_{\mathbf{m}} q(\mathbf{n}, \mathbf{n} + \mathbf{m}) \mathbf{m}^T \nabla^2 F(\mathbf{n} + \theta_{\mathbf{m}} \mathbf{m}) \mathbf{m},\end{aligned}$$

for some scalars $\theta_{\mathbf{m}} \in [0, 1]$, one such scalar for each $\mathbf{m}$. From the definition of $\mathbf{q}$, we have

$$\begin{aligned}\sum_{\mathbf{m}} q(\mathbf{n}, \mathbf{n} + \mathbf{m}) \langle \nabla F(\mathbf{n}), \mathbf{m} \rangle &= \left\langle \nabla F(\mathbf{n}), \sum_{\mathbf{m}} q(\mathbf{n}, \mathbf{n} + \mathbf{m}) \mathbf{m} \right\rangle \\ &= \langle \nabla F(\mathbf{n}), \boldsymbol{\nu} - \boldsymbol{\mu} \boldsymbol{\Lambda}(\mathbf{n}) \rangle.\end{aligned}$$

From (14), for $\mathbf{m} = \pm \mathbf{e}_i$, we also have

$$\frac{1}{2} \mathbf{m}^T \nabla^2 F(\mathbf{n} + \theta_{\mathbf{m}} \mathbf{m}) \mathbf{m} \leq \kappa_i \alpha (n_i + 1)^{\alpha-1} / 2\mu_i \rho_i^\alpha.$$

Thus,

$$\begin{aligned}\mathbf{Q}F(\mathbf{n}) &\leq \langle \nabla F(\mathbf{n}), \boldsymbol{\nu} - \boldsymbol{\mu} \boldsymbol{\Lambda}(\mathbf{n}) \rangle + \sum_{i \in \mathcal{I}} \frac{\kappa_i \alpha}{2\mu_i \rho_i^\alpha} (n_i + 1)^{\alpha-1} (\nu_i + \mu_i \Lambda_i(\mathbf{n})) \\ &\leq -\varepsilon \sum_{i \in \mathcal{I}} \kappa_i \left(\frac{n_i}{\rho_i} \right)^\alpha \rho_i + \sum_{i \in \mathcal{I}} \frac{\kappa_i \alpha}{2\rho_i^\alpha} (n_i + 1)^{\alpha-1} (\rho_i + \Lambda_i(\mathbf{n})) \\ &\leq -m\varepsilon \sum_{i \in \mathcal{I}} n_i^\alpha + M \sum_{i \in \mathcal{I}} (n_i + 1)^{\alpha-1},\end{aligned}$$

where the second inequality follows from Lemma 4.4, and the third by defining

$$m \triangleq \min_{i \in \mathcal{I}} \kappa_i \rho_i^{1-\alpha}, \quad M \triangleq \max_{i \in \mathcal{I}} \frac{\kappa_i \alpha}{2\rho_i^\alpha} \left(\rho_i + \max_{j \in \mathcal{J}} C_j \right),$$

and noting the fact that since $\Lambda_i(\mathbf{n}) \leq \max_{j \in \mathcal{J}} C_j$ for all i , we have $M \geq \max_{i \in \mathcal{I}} \frac{\kappa_i \alpha}{2\rho_i^\alpha} (\rho_i + \Lambda_i(\mathbf{n}))$. It is then a simple calculation to see that for every $\mathbf{n} \geq \mathbf{0}$, we have

$$\mathbf{Q}F(\mathbf{n}) \leq -m\varepsilon \sum_{i \in \mathcal{I}} n_i^\alpha + M \sum_{i \in \mathcal{I}} (n_i + 1)^{\alpha-1} \leq \tilde{K} \varepsilon^{1-\alpha},$$

for some positive load-dependent constant $\tilde{K}$. Now given $\tilde{\mathbf{N}}(\tau) = \mathbf{n}$,

$$\mathbb{E} \left[F(\tilde{\mathbf{N}}(\tau + 1)) - F(\mathbf{n}) \mid \tilde{\mathbf{N}}(\tau) = \mathbf{n} \right] = \frac{\mathbf{Q}F(\mathbf{n})}{\Xi} \leq \frac{\tilde{K}\varepsilon^{1-\alpha}}{\Xi}.$$

By setting $\bar{K} = \tilde{K}/\Xi$, we have proved the lemma. $\square$

COROLLARY 4.6. *Let $\alpha \geq 1$. As before, suppose that $\mathbf{A}\boldsymbol{\rho} < \mathbf{C}$, and let ε be the associated gap. Then, under the α -fair policy, the process $\mathbf{N}(\cdot)$ satisfies*

$$\mathbb{E} [F(\mathbf{N}(s+t)) - F(\mathbf{N}(s)) \mid \mathbf{N}(s)] \leq \tilde{K}t\varepsilon^{1-\alpha}, \quad \text{for all } t \geq 0.$$

for some positive load-dependent constant $\tilde{K}$.

PROOF. The idea of the proof is to show that the expected number of state transitions of $\mathbf{N}(\cdot)$ in the time interval $[s, s+t]$ is of order $O(t)$.

Consider the uniformized Markov chain $\tilde{\mathbf{N}}(\cdot)$ associated with the process $\mathbf{N}(\cdot)$. Denote the number of state transitions in the uniformized version of the process $\mathbf{N}(\cdot)$ in the time interval $[s, s+t]$ by τ . By the Markov property, time-homogeneity, and the definition of $\tilde{\mathbf{N}}(\cdot)$, we have

$$\begin{aligned} & \mathbb{E} [F(\mathbf{N}(s+t)) - F(\mathbf{N}(s)) \mid \mathbf{N}(s) = \mathbf{n}] \\ &= \mathbb{E} [F(\tilde{\mathbf{N}}(\tau)) - F(\tilde{\mathbf{N}}(0)) \mid \tilde{\mathbf{N}}(0) = \mathbf{n}]. \end{aligned}$$

Now, by the definition of the uniformized chain, τ and $\tilde{\mathbf{N}}(\cdot)$ are independent. Thus,

$$\begin{aligned} & \mathbb{E} [F(\tilde{\mathbf{N}}(\tau)) - F(\tilde{\mathbf{N}}(0)) \mid \tilde{\mathbf{N}}(0)] \\ &= \mathbb{E} \left[\sum_{k=0}^{\tau-1} \left(F(\tilde{\mathbf{N}}(k+1)) - F(\tilde{\mathbf{N}}(k)) \right) \mid \tilde{\mathbf{N}}(0) \right] \\ &= \mathbb{E} \left[\mathbb{E} \left[\sum_{k=0}^{\tau-1} \left(F(\tilde{\mathbf{N}}(k+1)) - F(\tilde{\mathbf{N}}(k)) \right) \mid \tilde{\mathbf{N}}(0), \tau \right] \mid \tilde{\mathbf{N}}(0) \right] \\ &= \mathbb{E} \left[\sum_{k=0}^{\tau-1} \mathbb{E} [F(\tilde{\mathbf{N}}(k+1)) - F(\tilde{\mathbf{N}}(k)) \mid \tilde{\mathbf{N}}(0), \tau] \mid \tilde{\mathbf{N}}(0) \right] \\ &= \mathbb{E} \left[\sum_{k=0}^{\tau-1} \mathbb{E} [F(\tilde{\mathbf{N}}(k+1)) - F(\tilde{\mathbf{N}}(k)) \mid \tilde{\mathbf{N}}(0)] \mid \tilde{\mathbf{N}}(0) \right] \\ &\leq \mathbb{E} \left[\sum_{k=0}^{\tau-1} \bar{K}\varepsilon^{1-\alpha} \right] = \bar{K}\varepsilon^{1-\alpha}\mathbb{E}[\tau], \end{aligned}$$

for some load-dependent constant $\bar{K}$. The fourth equality follows from the independence of τ and $\tilde{\mathbf{N}}(\cdot)$, and the inequality follows from Lemma 4.5. Since the counting process of the number of state transitions in the uniformized version of the process $\mathbf{N}(\cdot)$ is a time-homogeneous Poisson process of rate Ξ , we have $\mathbb{E}[\tau] = \Xi t$. This shows that

$$\mathbb{E}[F(\mathbf{N}(s+t)) - F(\mathbf{N}(s)) \mid \mathbf{N}(s)] \leq \bar{K}\Xi t \varepsilon^{1-\alpha}.$$

The proof is concluded by setting $\tilde{K} = \bar{K}\Xi$. $\square$

Proof of Theorem 3.1. Let $b > 0$. Then

$$\begin{aligned} \mathbb{P}(N^*(T) \geq b) &= \mathbb{P}\left(\frac{1}{\alpha+1}(N^*(T))^{\alpha+1} \geq \frac{1}{\alpha+1}b^{\alpha+1}\right) \\ &\leq \mathbb{P}\left(\sup_{t \in [0, T]} F(\mathbf{N}(t)) \geq \left(\min_{i \in \mathcal{I}} \frac{1}{\alpha+1} \kappa_i \mu_i^{\alpha-1} \nu_i^{-\alpha}\right) b^{\alpha+1}\right) \\ &\leq \frac{(\alpha+1)K'T}{\left(\min_{i \in \mathcal{I}} \kappa_i \mu_i^{\alpha-1} \nu_i^{-\alpha}\right) \varepsilon^{\alpha-1} b^{\alpha+1}} = \frac{KT}{\varepsilon^{\alpha-1} b^{\alpha+1}}, \end{aligned}$$

where the second inequality follows from Corollary 4.3 and Corollary 4.6, K' is as in Corollary 4.6, and $K = \frac{(\alpha+1)K'}{\min_{i \in \mathcal{I}} \kappa_i \mu_i^{\alpha-1} \nu_i^{-\alpha}}$. $\square$

4.3. Full State Space Collapse for $\alpha \geq 1$. Throughout this section, we assume that we have fixed $\alpha \geq 1$, and correspondingly, the Lyapunov function (9). To state the full state space collapse result for $\alpha \geq 1$, we need some preliminary definitions and the statement of the multiplicative state space collapse result.

Consider a sequence of bandwidth-sharing networks indexed by r , where r is to be thought of as increasing to infinity along a sequence. Suppose that the incidence matrix $\mathbf{A}$, the capacity vector $\mathbf{C}$ and the weights $\{\kappa_i : i \in \mathcal{I}\}$ do not vary with r . Write $\mathbf{N}^r(t)$ for the flow-vector Markov process associated with the r th network. Similarly, we write $\boldsymbol{\nu}^r$, $\boldsymbol{\mu}^r$, $\boldsymbol{\rho}^r$, etc. We assume the following *heavy-traffic* condition (cf. [13]):

ASSUMPTION 4.7. *We assume that $\mathbf{A}\boldsymbol{\rho}^r < \mathbf{C}$ for all r . We also assume that there exist $\boldsymbol{\nu}, \boldsymbol{\mu} \in \mathbb{R}_+^{|\mathcal{I}|}$ and $\boldsymbol{\theta} > \mathbf{0}$, such that $\nu_i > 0$ and $\mu_i > 0$ for all $i \in \mathcal{I}$, $\boldsymbol{\nu}^r \rightarrow \boldsymbol{\nu}$ and $\boldsymbol{\mu}^r \rightarrow \boldsymbol{\mu}$ as $r \rightarrow \infty$, and $r(\mathbf{C} - \mathbf{A}\boldsymbol{\rho}^r) \rightarrow \boldsymbol{\theta}$ as $r \rightarrow \infty$.*

Note that our assumption differs from that in [13], which allows convergence to the critical load from both overload and underload, whereas here we only allow convergence to the critical load from underload.

To state the multiplicative state space collapse result, we also need to define a *workload process* $\mathbf{W}(t)$ and a *lifting map* Δ .

DEFINITION 4.8. We first define the workload $\mathbf{w} : \mathbb{R}_+^{|\mathcal{I}|} \rightarrow \mathbb{R}_+^{|\mathcal{J}|}$ associated with a flow-vector $\mathbf{n}$ by $\mathbf{w} = \mathbf{w}(\mathbf{n}) = \mathbf{A}\mathbf{M}^{-1}\mathbf{n}$, where $\mathbf{M} = \text{diag}(\boldsymbol{\mu})$ is the $|\mathcal{I}| \times |\mathcal{I}|$ diagonal matrix with $\boldsymbol{\mu}$ on its diagonal. The workload process $\mathbf{W}(t)$ is defined to be $\mathbf{W}(t) \triangleq \mathbf{A}\mathbf{M}^{-1}\mathbf{N}(t)$, for all $t \geq 0$. We also define the lifting map Δ . For each $\mathbf{w} \in \mathbb{R}_+^{|\mathcal{J}|}$, define $\Delta(\mathbf{w})$ to be the unique value of $\mathbf{n} \in \mathbb{R}_+^{|\mathcal{I}|}$ that solves the following optimization problem:

$$\begin{aligned} & \text{minimize} && F(\mathbf{n}) \\ & \text{subject to} && \sum_{i \in \mathcal{I}} A_{ji} \frac{n_i}{\mu_i} \geq w_j, \quad j \in \mathcal{J}, \\ & && n_i \geq 0, \quad i \in \mathcal{I}. \end{aligned}$$

For simplicity, suppose that all networks start with zero flows. We consider the following diffusion scaling:

$$(15) \quad \hat{\mathbf{N}}^r(t) = \frac{\mathbf{N}^r(r^2t)}{r}, \text{ and } \hat{\mathbf{W}}^r(t) = \frac{\mathbf{W}^r(r^2t)}{r},$$

where $\mathbf{W}^r(t) = \mathbf{A}(\mathbf{M}^r)^{-1}\mathbf{N}^r(t)$, and $\mathbf{M}^r = \text{diag}(\boldsymbol{\mu}^r)$.

The following multiplicative state space collapse result is known to hold.

THEOREM 4.9 (Multiplicative State Space Collapse [13, Theorem 5.1]). Fix $T > 0$ and assume that $\alpha \geq 1$. Write $\|\mathbf{x}(\cdot)\| = \sup_{t \in [0, T], i \in \mathcal{I}} |x_i(t)|$. Then, under Assumption 4.7, and for any $\delta > 0$,

$$\lim_{r \rightarrow \infty} \mathbb{P} \left(\frac{\|\hat{\mathbf{N}}^r(\cdot) - \Delta(\hat{\mathbf{W}}^r(\cdot))\|}{\|\hat{\mathbf{N}}^r(\cdot)\|} > \delta \right) = 0.$$

We can now state and prove a full state space collapse result:

THEOREM 4.10 (Full State Space Collapse). Under the same assumptions as in Theorem 4.9, and for any $\delta > 0$,

$$\lim_{r \rightarrow \infty} \mathbb{P} \left(\|\hat{\mathbf{N}}^r(\cdot) - \Delta(\hat{\mathbf{W}}^r(\cdot))\| > \delta \right) = 0.$$

PROOF. Let $\varepsilon_r = \varepsilon(\boldsymbol{\rho}^r)$ be the gap in the r th system. Then, under Assumption 4.7, $\varepsilon_r \geq D/r$ for some network-dependent constant $D > 0$, and for r sufficiently large. By Theorem 3.1, for any $b > 0$, and for sufficiently

large r ,

$$\begin{aligned}\mathbb{P}(N^{r,*}(r^2T) \geq b) &\leq \frac{K_r r^2 T}{\varepsilon_r^{\alpha-1} b^{\alpha+1}} \\ &\leq \frac{K_r r^{1+\alpha} T}{D^{\alpha-1} b^{\alpha+1}}.\end{aligned}$$

Here, K_r is a load-dependent constant associated with the r th system, as specified in the proof of Theorem 3.1. From the proof of Theorem 3.1, note also that $K_r = f(\boldsymbol{\mu}^r, \boldsymbol{\nu}^r)$, for a function f that is continuous on the open positive orthant $\mathbb{R}_p^{|Z|} \times \mathbb{R}_p^{|Z|}$. Since $\boldsymbol{\mu}^r \rightarrow \boldsymbol{\mu} > \mathbf{0}$, and $\boldsymbol{\nu}^r \rightarrow \boldsymbol{\nu} > \mathbf{0}$, $K_r \rightarrow K \triangleq f(\boldsymbol{\mu}, \boldsymbol{\nu}) \in \mathbb{R}$. In particular, the K_r are bounded, and for all sufficiently large r ,

$$\mathbb{P}(N^{r,*}(r^2T) \geq b) \leq \frac{(K+1)r^{1+\alpha}T}{D^{\alpha-1}b^{\alpha+1}}.$$

Then, with $a = b/r$ and under the scaling in (15),

$$(16) \quad \mathbb{P}(\|\hat{\mathbf{N}}^r(\cdot)\| \geq a) \leq \frac{K+1}{D^{\alpha-1}} \cdot \frac{T}{a^{\alpha+1}},$$

for any $a > 0$.

For notational convenience, we write

$$B(r) = \|\hat{\mathbf{N}}^r(\cdot) - \Delta(\mathbf{W}^r(\cdot))\|.$$

Then, for any $a > 1$, and for sufficiently large r ,

$$\begin{aligned}\mathbb{P}(B(r) > \delta) &\leq \mathbb{P}\left(\frac{B(r)}{\|\hat{\mathbf{N}}^r(\cdot)\|} > \frac{\delta}{a} \text{ or } \|\hat{\mathbf{N}}^r(\cdot)\| \geq a\right) \\ &\leq \mathbb{P}\left(\frac{B(r)}{\|\hat{\mathbf{N}}^r(\cdot)\|} > \frac{\delta}{a}\right) + \mathbb{P}(\|\hat{\mathbf{N}}^r(\cdot)\| \geq a).\end{aligned}$$

Note that by Theorem 4.9, the first term on the right-hand side goes to 0 as $r \rightarrow \infty$, for any $a > 0$. The second term on the right-hand side can be made smaller than any, arbitrarily small, constant (uniformly, for all r), by taking a sufficiently large (cf. Equation (16)). Thus, $\mathbb{P}(B(r) \geq \delta) \rightarrow 0$ as $r \rightarrow \infty$. This concludes the proof. $\square$

5. α -Fair Policies: A Useful Drift Inequality. We now shift our focus to the steady-state regime. The key to many of our results is a *drift inequality* that holds for every $\alpha > 0$ and every $\boldsymbol{\rho} > \mathbf{0}$ with $\mathbf{A}\boldsymbol{\rho} < \mathbf{C}$. In this

section, we shall state and prove this inequality. It will be used in Section 6 to prove Theorem 3.2.

We define the Lyapunov function that we will employ. For $\alpha \geq 1$, it will be simply the weighted $(\alpha + 1)$ -norm $L_\alpha(\mathbf{n}) = \sqrt[\alpha+1]{(\alpha + 1)F_\alpha(\mathbf{n})}$ of a vector $\mathbf{n}$, where F_α was defined in (9). However, when $\alpha \in (0, 1)$, this function has unbounded second derivatives as we approach the boundary of $\mathbb{R}_+^{|\mathcal{I}|}$. For this reason, our Lyapunov function will be a suitably smoothed version of $\sqrt[\alpha+1]{(\alpha + 1)F_\alpha(\cdot)}$.

DEFINITION 5.1. *Define $h_\alpha : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ to be $h_\alpha(r) = r^\alpha$, when $\alpha \geq 1$, and*

$$h_\alpha(r) = \begin{cases} r^\alpha, & \text{if } r \geq 1, \\ (\alpha - 1)r^3 + (1 - \alpha)r^2 + r, & \text{if } r < 1, \end{cases}$$

when $\alpha \in (0, 1)$. Let $H_\alpha : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ be the antiderivative of h_α , so that $H_\alpha(r) = \int_0^r h_\alpha(s) ds$. The Lyapunov function $L_\alpha : \mathbb{R}_+^{|\mathcal{I}|} \rightarrow \mathbb{R}_+$ is defined to be

$$L_\alpha(\mathbf{n}) = \left[(\alpha + 1) \sum_{i \in \mathcal{I}} \kappa_i \mu_i^{\alpha-1} \nu_i^{-\alpha} H_\alpha(n_i) \right]^{\frac{1}{\alpha+1}}.$$

For notational convenience, define

$$(17) \quad w_i = \kappa_i \mu_i^{\alpha-1} \nu_i^{-\alpha} \text{ for each } i \in \mathcal{I},$$

so that more compactly, we have

$$F_\alpha(\mathbf{n}) = \frac{1}{\alpha + 1} \sum_{i \in \mathcal{I}} w_i n_i^{\alpha+1}, \text{ and } L_\alpha(\mathbf{n}) = \left[(\alpha + 1) \sum_{i \in \mathcal{I}} w_i H_\alpha(n_i) \right]^{1/(\alpha+1)}.$$

We will make heavy use of various properties of the functions h_α , H_α , and L_α , which we summarize in the following lemma. The proof is elementary and is omitted.

LEMMA 5.2. *Let $\alpha \in (0, 1)$. The function h_α has the following properties:*

- (i) *it is continuously differentiable with $h_\alpha(0) = 0$, $h_\alpha(1) = 1$, $h'_\alpha(0) = 1$, and $h'_\alpha(1) = \alpha$;*
- (ii) *it is increasing and, in particular, $h_\alpha(r) \geq 0$ for all $r \geq 0$;*
- (iii) *we have $r^\alpha - 1 \leq h_\alpha(r) \leq r^\alpha + 1$, for all $r \in [0, 1]$;*
- (iv) *$h'_\alpha(r) \leq 2$, for all $r \geq 0$.*

Furthermore, from (iii), we also have the following property of H_α :

(iii') $r^{\alpha+1} - 2 \leq (\alpha + 1)H_\alpha(r) \leq r^{\alpha+1} + 2$ for all $r \geq 0$.

We are now ready to state the drift inequality. Here we consider the uniformized chain $\left(\tilde{\mathbf{N}}(\tau)\right)_{\tau \in \mathbb{Z}_+}$ associated with $\mathbf{N}(\cdot)$, and the corresponding drift.

THEOREM 5.3. *Consider a bandwidth-sharing network operating under an α -fair policy with $\alpha > 0$, and assume that $\mathbf{A}\boldsymbol{\rho} < \mathbf{C}$. Let ε be the gap. Then, there exists a positive constant B and a positive load-dependent constant K , such that if $L_\alpha(\tilde{\mathbf{N}}(\tau)) > B$, then*

$$(18) \quad \mathbb{E}[L_\alpha(\tilde{\mathbf{N}}(\tau+1)) - L_\alpha(\tilde{\mathbf{N}}(\tau)) \mid \tilde{\mathbf{N}}(\tau)] \leq -\varepsilon K.$$

Furthermore, B takes the form K'/ε when $\alpha \geq 1$, and $K'/\min\{\varepsilon^{1/\alpha}, \varepsilon\}$ when $\alpha \in (0, 1)$, with K' being a positive load-dependent constant.

As there is a marked difference between the form of L_α for the two cases $\alpha \geq 1$ and $\alpha \in (0, 1)$, the proof of the drift inequality is split into two parts. We first prove the drift inequality when $\alpha \geq 1$, in which case L_α takes a nicer form, and we can apply results on F_α from previous sections. The proof for the case $\alpha \in (0, 1)$ is similar but more tedious. We note that such a qualitative difference between the two cases, $\alpha < 1$ and $\alpha \geq 1$, has also been observed in other works, such as, for example, [20].

We wish to draw attention here to the main difference from related drift inequalities in the literature. The usual proof of stability involves the Lyapunov function (9); for instance, for the α -fair policy with $\alpha = 1$ (the proportionally fair policy), it involves a weighted quadratic Lyapunov function. In contrast, we use L_α , a weighted norm function (or its smoothed version), which scales linearly along radial directions. In this sense, our approach is similar in spirit to [2], which employed piecewise linear Lyapunov functions to derive drift inequalities and then moment and tail bounds. The use of normed Lyapunov functions to establish stability and performance bounds has also been considered in other works; see, for example, [23] and [7].

5.1. Proof of Theorem 5.3: $\alpha \geq 1$. We wish to decompose the drift term in (18) into the sum of a first-order term and a second-order term, and we accomplish this by using the second-order mean value theorem (cf. Proposition 2.4). Throughout this proof, we drop the subscript α from L_α and F_α , and write L and F , respectively.

Consider the function $L(\mathbf{n}) = \left(\sum_{i \in \mathcal{I}} w_i n_i^{\alpha+1}\right)^{\frac{1}{\alpha+1}} = [(\alpha+1)F(\mathbf{n})]^{\frac{1}{\alpha+1}}$. The first derivative of L with respect to $\mathbf{n}$ is $\nabla L(\mathbf{n}) = \nabla F(\mathbf{n})/L^\alpha(\mathbf{n})$ by the

chain rule and the definition of L . The second derivative is

$$\begin{aligned}\nabla^2 L(\mathbf{n}) &= \frac{\nabla^2 F(\mathbf{n})}{L^\alpha(\mathbf{n})} - \frac{\nabla F(\mathbf{n}) \nabla L^\alpha(\mathbf{n})^T}{L^{2\alpha}(\mathbf{n})} \\ &= \frac{\nabla^2 F(\mathbf{n})}{L^\alpha(\mathbf{n})} - \alpha \frac{\nabla F(\mathbf{n}) \nabla F(\mathbf{n})^T}{L^{2\alpha+1}(\mathbf{n})},\end{aligned}$$

by the quotient rule and the chain rule.

Write $\mathbf{n}$ for $\tilde{\mathbf{N}}(\tau)$ and $\mathbf{n} + \mathbf{m}$ for $\tilde{\mathbf{N}}(\tau + 1)$, so that $\mathbf{m} = \tilde{\mathbf{N}}(\tau + 1) - \tilde{\mathbf{N}}(\tau)$. By Proposition 2.4, for some $\theta \in [0, 1]$, we have

$$\begin{aligned}(19) \quad L(\mathbf{n} + \mathbf{m}) - L(\mathbf{n}) &= \mathbf{m}^T \nabla L(\mathbf{n}) + \frac{1}{2} \mathbf{m}^T \nabla^2 L(\mathbf{n} + \theta \mathbf{m}) \mathbf{m} \\ (20) \quad &= \frac{\mathbf{m}^T \nabla F(\mathbf{n})}{L^\alpha(\mathbf{n})} + \frac{1}{2} \frac{\mathbf{m}^T \nabla^2 F(\mathbf{n} + \theta \mathbf{m}) \mathbf{m}}{L^\alpha(\mathbf{n} + \theta \mathbf{m})} \\ (21) \quad &\quad - \frac{\alpha}{2} \frac{\mathbf{m}^T \nabla F(\mathbf{n} + \theta \mathbf{m}) \nabla F(\mathbf{n} + \theta \mathbf{m})^T \mathbf{m}}{L^{2\alpha+1}(\mathbf{n} + \theta \mathbf{m})} \\ (22) \quad &\leq \frac{\mathbf{m}^T \nabla F(\mathbf{n})}{L^\alpha(\mathbf{n})} + \frac{1}{2} \mathbf{m}^T \frac{\nabla^2 F(\mathbf{n} + \theta \mathbf{m})}{L^\alpha(\mathbf{n} + \theta \mathbf{m})} \mathbf{m},\end{aligned}$$

since the term $\mathbf{m}^T \nabla F(\mathbf{n} + \theta \mathbf{m}) \nabla F(\mathbf{n} + \theta \mathbf{m})^T \mathbf{m}$ is nonnegative. We now consider the two terms in (22) separately. Recall from the proof of Lemma 4.5 that

$$\mathbb{E} [\mathbf{m}^T \nabla F(\mathbf{n}) \mid \mathbf{n}] = \frac{\langle \nabla F(\mathbf{n}), \boldsymbol{\nu} - \boldsymbol{\mu} \Lambda(\mathbf{n}) \rangle}{\Xi} \leq -\varepsilon \frac{\langle \nabla F(\mathbf{n}), \boldsymbol{\nu} \rangle}{\Xi}.$$

But $\langle \nabla F(\mathbf{n}), \boldsymbol{\nu} \rangle = \sum_{i \in \mathcal{I}} w_i \nu_i n_i^\alpha$, so

$$(23) \quad \mathbb{E} [\mathbf{m}^T \nabla F(\mathbf{n}) \mid \mathbf{n}] \leq -\varepsilon \frac{\sum_{i \in \mathcal{I}} w_i \nu_i n_i^\alpha}{\Xi},$$

and so

$$\begin{aligned}
\mathbb{E} \left[\frac{\mathbf{m}^T \nabla F(\mathbf{n})}{L^\alpha(\mathbf{n})} \mid \mathbf{n} \right] &\leq -\varepsilon \frac{\sum_{i \in \mathcal{I}} w_i \nu_i n_i^\alpha}{\Xi \left(\sum_{i \in \mathcal{I}} w_i n_i^{\alpha+1} \right)^{\frac{\alpha}{\alpha+1}}} \\
&= -\varepsilon \frac{\sum_{i \in \mathcal{I}} w_i \nu_i n_i^\alpha}{\Xi \left(\sum_{i \in \mathcal{I}} \left(w_i^{\frac{1}{\alpha+1}} n_i \right)^{\alpha+1} \right)^{\frac{\alpha}{\alpha+1}}} \\
&\leq -\varepsilon \frac{\sum_{i \in \mathcal{I}} w_i \nu_i n_i^\alpha}{\Xi \cdot \sum_{i \in \mathcal{I}} w_i^{\frac{\alpha}{\alpha+1}} n_i^\alpha} \\
&\leq -\varepsilon \frac{\max_{i \in \mathcal{I}} w_i^{\frac{1}{\alpha+1}} \nu_i}{\Xi} \\
&= -\varepsilon \frac{\max_{i \in \mathcal{I}} \kappa^{\frac{1}{\alpha+1}} \mu_i^{\frac{\alpha-1}{\alpha+1}} \nu_i^{\frac{1}{\alpha+1}}}{\Xi} \\
(24) \qquad \qquad \qquad &= -\varepsilon K,
\end{aligned}$$

where

$$(25) \qquad K = K(\alpha, \kappa, \mu, \nu) \triangleq \frac{\max_{i \in \mathcal{I}} \kappa^{\frac{1}{\alpha+1}} \mu_i^{\frac{\alpha-1}{\alpha+1}} \nu_i^{\frac{1}{\alpha+1}}}{\Xi}$$

is a positive load-dependent constant. The second inequality follows from the fact that for any vector $\mathbf{x}$, and for any $\alpha > 0$, $\|\mathbf{x}\|_{\alpha+1} \leq \|\mathbf{x}\|_\alpha$. The second to last equality follows from the definition of the w_i (cf. Equation (17)).

For the second term in (22), we wish to show that if $L(\mathbf{n})$ is sufficiently large, then

$$\frac{1}{2} \mathbf{m}^T \frac{\nabla^2 F(\mathbf{n} + \theta \mathbf{m})}{L^\alpha(\mathbf{n} + \theta \mathbf{m})} \mathbf{m} \leq \frac{\varepsilon}{2} K.$$

Note that with probability 1, either $\mathbf{m} = \mathbf{0}$ or $\mathbf{m} = \pm \mathbf{e}_i$ for some $i \in \mathcal{I}$.

Thus

$$\begin{aligned}
\frac{1}{2} \mathbf{m}^T \frac{\nabla^2 F(\mathbf{n} + \theta \mathbf{m})}{L^\alpha(\mathbf{n} + \theta \mathbf{m})} \mathbf{m} &\leq \frac{1}{2} \frac{\max_{i \in \mathcal{I}} [\nabla^2 F(\mathbf{n} + \theta \mathbf{m})]_{ii}}{L^\alpha(\mathbf{n} + \theta \mathbf{m})} \\
&= \frac{\alpha}{2} \frac{\max_{i \in \mathcal{I}} w_i (n_i + \theta m_i)^{\alpha-1}}{[\sum_{i \in \mathcal{I}} w_i (n_i + \theta m_i)^{\alpha+1}]^{\frac{\alpha}{\alpha+1}}} \\
&\leq \frac{\alpha}{2} \frac{\max_{i \in \mathcal{I}} w_i (n_i + \theta m_i)^{\alpha-1}}{w_{i_0}^{\frac{\alpha}{\alpha+1}} (n_{i_0} + \theta m_{i_0})^\alpha} \\
&\leq \frac{\alpha}{2} w_{i_0}^{\frac{1}{\alpha+1}} (n_{i_0} + \theta m_{i_0})^{-1} \\
&\leq \frac{\alpha}{2} \max_{i \in \mathcal{I}} w_i^{\frac{1}{\alpha+1}} (n_{i_0} + \theta m_{i_0})^{-1},
\end{aligned}$$

where $i_0 \in \mathcal{I}$ is such that $w_{i_0} (n_{i_0} + \theta m_{i_0})^{\alpha-1} = \max_{i \in \mathcal{I}} w_i (n_i + \theta m_i)^{\alpha-1}$.

Now note that

$$\frac{\alpha}{2} \max_{i \in \mathcal{I}} w_i^{\frac{1}{\alpha+1}} (n_{i_0} + \theta m_{i_0})^{-1} \leq \frac{\varepsilon}{2} K$$

(where K is defined in (25)) if and only if

$$n_{i_0} + \theta m_{i_0} \geq \frac{\alpha \max_{i \in \mathcal{I}} w_i^{\frac{1}{\alpha+1}}}{K} \cdot \frac{1}{\varepsilon},$$

which holds if $L(\mathbf{n}) \geq K'/\varepsilon$ for some appropriately defined load-dependent constant K' . Thus, if $L(\mathbf{n}) \geq K'/\varepsilon$, then

$$(26) \quad \frac{1}{2} \mathbf{m}^T \frac{\nabla^2 F(\mathbf{n} + \theta \mathbf{m})}{L^\alpha(\mathbf{n} + \theta \mathbf{m})} \mathbf{m} \leq \frac{\varepsilon}{2} K.$$

By adding (24) and (26), we conclude that

$$\mathbb{E}[L(\mathbf{n} + \mathbf{m}) - L(\mathbf{n}) \mid \mathbf{n}] \leq -\frac{\varepsilon}{2} K,$$

when $L(\mathbf{n}) \geq K'/\varepsilon$. □

5.2. Proof of Theorem 5.3: $\alpha \in (0, 1)$. The proof in this section is similar to that for the case $\alpha \geq 1$. We invoke Proposition 2.4 to write the drift term as a sum of terms, which we bound separately. As in the previous section, we drop the subscript α from L_α , F_α , H_α , and h_α , and write instead L , F , H , and h , respectively. Note that to use Proposition 2.4, we need L to be twice continuously differentiable. Indeed, by Lemma 5.2 (i), h is continuously

differentiable, so its antiderivative H is twice continuously differentiable, and so is L . Thus, by the second order mean value theorem, we obtain an equation similar to Equation (22):

$$\begin{aligned}
 (27) \quad L(\mathbf{n} + \mathbf{m}) - L(\mathbf{n}) &= \mathbf{m}^T \nabla L(\mathbf{n}) + \frac{1}{2} \mathbf{m}^T \nabla^2 L(\mathbf{n} + \theta \mathbf{m}) \mathbf{m} \\
 (28) \quad &\leq \frac{\sum_{i \in I} m_i w_i h(n_i)}{L^\alpha(\mathbf{n})} + \frac{1}{2} \frac{\sum_{i \in \mathcal{I}} m_i^2 w_i h'(n_i + \theta m_i)}{L^\alpha(\mathbf{n} + \theta \mathbf{m})} \\
 (29) \quad &\leq \frac{\sum_{i \in I} m_i w_i h(n_i)}{L^\alpha(\mathbf{n})} + \frac{1}{2} \frac{\max_{i \in \mathcal{I}} w_i h'(n_i + \theta m_i)}{L^\alpha(\mathbf{n} + \theta \mathbf{m})}
 \end{aligned}$$

for some constant $\theta \in [0, 1]$, and where, as before, $\tilde{\mathbf{N}}(\tau) = \mathbf{n}$ and $\tilde{\mathbf{N}}(\tau + 1) = \mathbf{n} + \mathbf{m}$, and the last inequality follows from the fact that with probability 1, either $\mathbf{m} = \mathbf{0}$, or $\mathbf{m} = \pm \mathbf{e}_i$, for some $i \in \mathcal{I}$, and that h' is nonnegative.

We now bound the two terms in (29) separately. Let us first concentrate on the term

$$\frac{\sum_{i \in I} m_i w_i h(n_i)}{L^\alpha(\mathbf{n})}.$$

By Lemma 5.2 (iii),

$$\sum_{i \in I} m_i w_i h(n_i) \leq \sum_{i \in I} m_i w_i (n_i^\alpha + 1) \leq \sum_{i \in I} m_i w_i n_i^\alpha + \sum_{i \in I} m_i w_i,$$

so

$$\frac{\sum_{i \in I} m_i w_i h(n_i)}{L^\alpha(\mathbf{n})} \leq \frac{\sum_{i \in I} m_i w_i n_i^\alpha}{L^\alpha(\mathbf{n})} + \frac{\sum_{i \in I} m_i w_i}{L^\alpha(\mathbf{n})}.$$

First consider the term $\frac{\sum_{i \in I} m_i w_i n_i^\alpha}{L^\alpha(\mathbf{n})}$. Note that $\sum_{i \in I} m_i w_i n_i^\alpha = \mathbf{m}^T \nabla F(\mathbf{n})$. We also recall from the proof of Lemma 4.4 that

$$\mathbb{E} [\mathbf{m}^T \nabla F(\mathbf{n}) \mid \mathbf{n}] = \frac{\langle \nabla F(\mathbf{n}), \boldsymbol{\nu} - \boldsymbol{\mu} \Lambda(\mathbf{n}) \rangle}{\Xi} \leq -\varepsilon \frac{\langle \nabla F(\mathbf{n}), \boldsymbol{\nu} \rangle}{\Xi}.$$

We then proceed along the same lines as in the case $\alpha \geq 1$, and obtain that if $L(\mathbf{n}) \geq K_2/\varepsilon$ for some positive load-dependent constant K_2 , then

$$\begin{aligned}
 (30) \quad \mathbb{E} \left[\frac{\sum_{i \in I} m_i w_i n_i^\alpha}{L^\alpha(\mathbf{n})} \mid \mathbf{n} \right] &\leq -\frac{3}{4} \varepsilon \frac{\max_{i \in \mathcal{I}} w_i^{\frac{1}{\alpha+1}} \nu_i}{\Xi} \\
 &= -\frac{3}{4} \varepsilon \frac{\max_{i \in \mathcal{I}} \kappa^{\frac{1}{\alpha+1}} \mu_i^{\frac{\alpha-1}{\alpha+1}} \nu_i^{\frac{1}{\alpha+1}}}{\Xi} \\
 &= -\frac{3}{4} \varepsilon K,
 \end{aligned}$$

Here as in the proof for the case $\alpha \geq 1$, $K = K(\alpha, \kappa, \mu, \nu) \triangleq \frac{\max_{i \in \mathcal{I}} \kappa^{\frac{1}{\alpha+1}} \mu_i^{\frac{\alpha-1}{\alpha+1}} \nu_i^{\frac{1}{\alpha+1}}}{\sum_{i \in \mathcal{I}} \nu_i}$ is a positive load-dependent constant.

Now consider the term $\frac{\sum_{i \in \mathcal{I}} m_i w_i}{L^\alpha(\mathbf{n})}$. With probability 1, either $\mathbf{m} = \mathbf{0}$ or $\mathbf{m} = \pm \mathbf{e}_i$ for some $i \in \mathcal{I}$, and therefore $\sum_{i \in \mathcal{I}} m_i w_i \leq \max_{i \in \mathcal{I}} w_i$. Thus,

$$\mathbb{E} \left[\frac{\sum_{i \in \mathcal{I}} m_i w_i h(n_i)}{L^\alpha(\mathbf{n})} \mid \mathbf{n} \right] \leq -\frac{3}{4} \varepsilon K + \frac{\max_{i \in \mathcal{I}} w_i}{L^\alpha(\mathbf{n})}.$$

For the second term in (29), note that with $\alpha \in (0, 1)$, Lemma 5.2(iv) implies that $h' \leq 2$, and therefore,

$$\frac{1}{2} \frac{\max_{i \in \mathcal{I}} w_i h'(n_i + \theta m_i)}{L^\alpha(\mathbf{n} + \theta \mathbf{m})} \leq \frac{\max_{i \in \mathcal{I}} w_i}{L^\alpha(\mathbf{n} + \theta \mathbf{m})}.$$

Note that $L^\alpha(\mathbf{n} + \theta \mathbf{m})$ and $L^\alpha(\mathbf{n})$ differ only by a load-dependent constant, since with probability 1, either $\mathbf{m} = \mathbf{0}$ or $\mathbf{m} = \pm \mathbf{e}_i$ for some $i \in \mathcal{I}$. Thus, if $L^\alpha(\mathbf{n}) \geq K_3/\varepsilon$ for some positive load-dependent constant K_3 , then

$$(31) \quad \frac{\max_{i \in \mathcal{I}} w_i}{L^\alpha(\mathbf{n})} + \frac{\max_{i \in \mathcal{I}} w_i}{L^\alpha(\mathbf{n} + \theta \mathbf{m})} \leq \frac{1}{4} \varepsilon K.$$

Putting (30) and (31) together, we get that if $L(\mathbf{n}) \geq K'/\min\{\varepsilon^{1/\alpha}, \varepsilon\}$, where $K' = \max\{K_3^{1/\alpha}, K_2\}$, then

$$\mathbb{E} [L(\mathbf{n} + \mathbf{m}) - L(\mathbf{n}) \mid \mathbf{n}] \leq -\frac{\varepsilon}{2} K. \quad \square$$

6. Exponential Tail Bound under α -Fair Policies. In this section, we derive an exponential upper bound on the tail probability of the stationary distribution of the flow sizes, under an α -fair policy with $\alpha > 0$. We will use the following theorem, a modification of Theorem 1 from [2].

THEOREM 6.1. *Let $\mathbf{X}(\cdot)$ be an irreducible and aperiodic discrete-time Markov chain with a countable state space $\mathcal{X}$. Suppose that there exists a Lyapunov function $\Phi : \mathcal{X} \rightarrow \mathbb{R}_+$ with the following properties:*

- (a) Φ has **bounded increments**: there exists $\xi > 0$ such that for all τ , we have

$$|\Phi(\mathbf{X}(\tau + 1)) - \Phi(\mathbf{X}(\tau))| \leq \xi, \text{ almost surely};$$

- (b) **Negative drift:** there exist $B > 0$ and $\gamma > 0$ such that whenever $\Phi(\mathbf{X}(\tau)) > B$,

$$\mathbb{E}[\Phi(\mathbf{X}(\tau + 1)) - \Phi(\mathbf{X}(\tau)) \mid \mathbf{X}(\tau)] \leq -\gamma.$$

Then, a stationary probability distribution π exists, and we have an exponential upper bound on the tail probability of Φ under π : for any $\ell \in \mathbb{Z}_+$,

$$(32) \quad \mathbb{P}_\pi(\Phi(\mathbf{X}) > B + 2\xi\ell) \leq \left(\frac{\xi}{\xi + \gamma} \right)^{\ell+1}.$$

In particular, in steady state, all moments of Φ are finite, i.e., for every $k \in \mathbb{N}$,

$$\mathbb{E}_\pi[\Phi^k(\mathbf{X})] < \infty.$$

Theorem 6.1 is identical to Theorem 1 in [2] except that [2] imposed the additional condition $\mathbb{E}_\pi[\Phi(\mathbf{X})] < \infty$. However, the latter condition is redundant. Indeed, using Foster-Lyapunov criteria (see [8], for example), conditions (a) and (b) in Theorem 6.1 imply that the Markov chain $\mathbf{X}$ has a unique stationary distribution π . Furthermore, Theorem 2.3 in [12] establishes that under conditions (a) and (b), all moments of $\Phi(\mathbf{X})$ are finite in steady state. We note that Theorem 2.3 in [12] and Theorem 1 of [2] provide the same qualitative information (exponential tail bounds for $\Phi(\mathbf{X})$). However, [2] contains the more precise bound (32), which we will use to prove Theorem 7.6 in Section 7.

PROOF OF THEOREM 3.2. The finiteness of the moments follows immediately from the bound in (32), so we only prove the exponential bound (32). We apply Theorem 6.1 to the Lyapunov function L_α and the uniformized chain $\tilde{\mathbf{N}}(\cdot)$. Again, denote the stationary distribution of $\tilde{\mathbf{N}}(\cdot)$ by π , and note that this is also the unique stationary distribution of $\mathbf{N}(\cdot)$. The proof consists of verifying conditions (a) and (b).

(a) *Bounded Increments.* We wish to show that with probability 1, there exists ξ such that

$$|L_\alpha(\tilde{\mathbf{N}}(\tau + 1)) - L_\alpha(\tilde{\mathbf{N}}(\tau))| \leq \xi.$$

As usual, write $\mathbf{n} = \tilde{\mathbf{N}}(\tau)$ and $\mathbf{n} + \mathbf{m} = \tilde{\mathbf{N}}(\tau + 1)$, then $\mathbf{m} = \mathbf{0}$ or $\mathbf{m} = \pm \mathbf{e}_i$ for some $i \in \mathcal{I}$ with probability 1. For $\alpha \geq 1$,

$$L_\alpha(\mathbf{n}) = \left[\sum_{i \in \mathcal{I}} w_i n_i^{\alpha+1} \right]^{\frac{1}{\alpha+1}},$$

and for $\alpha \in (0, 1)$, by Lemma 5.2 (iii'), we have

$$\sum_{i \in \mathcal{I}} w_i n_i^{\alpha+1} - 2 \sum_{i \in \mathcal{I}} w_i \leq (\alpha + 1) \sum_{i \in \mathcal{I}} w_i H_\alpha(n_i) \leq \sum_{i \in \mathcal{I}} w_i n_i^{\alpha+1} + 2 \sum_{i \in \mathcal{I}} w_i.$$

In general, for $r, s \geq 0$ and $\beta \in [0, 1]$,

$$(33) \quad (r + s)^\beta \leq r^\beta + s^\beta.$$

Thus, by inequality (33),

$$\left[\sum_{i \in \mathcal{I}} w_i n_i^{\alpha+1} \right]^{\frac{1}{\alpha+1}} - \left[2 \sum_{i \in \mathcal{I}} w_i \right]^{\frac{1}{\alpha+1}} \leq L_\alpha(\mathbf{n}) \leq \left[\sum_{i \in \mathcal{I}} w_i n_i^{\alpha+1} \right]^{\frac{1}{\alpha+1}} + \left[2 \sum_{i \in \mathcal{I}} w_i \right]^{\frac{1}{\alpha+1}}.$$

Hence, for any $\alpha > 0$,

$$\begin{aligned} |L_\alpha(\mathbf{n} + \mathbf{m}) - L_\alpha(\mathbf{n})| &\leq \left| \left[\sum_{i \in \mathcal{I}} w_i (n_i + m_i)^{\alpha+1} \right]^{\frac{1}{\alpha+1}} - \left[\sum_{i \in \mathcal{I}} w_i n_i^{\alpha+1} \right]^{\frac{1}{\alpha+1}} \right| \\ &\quad + 2 \left[2 \sum_{i \in \mathcal{I}} w_i \right]^{\frac{1}{\alpha+1}} \\ &\leq \left[\sum_{i \in \mathcal{I}} w_i |m_i|^{\alpha+1} \right]^{\frac{1}{\alpha+1}} + 2 \left[2 \sum_{i \in \mathcal{I}} w_i \right]^{\frac{1}{\alpha+1}} \\ &\leq \max_{i \in \mathcal{I}} w_i^{\frac{1}{\alpha+1}} + 2 \left[2 \sum_{i \in \mathcal{I}} w_i \right]^{\frac{1}{\alpha+1}}, \end{aligned}$$

where the second last inequality follows from the triangle inequality. Thus we can take $\xi = \max_{i \in \mathcal{I}} w_i^{\frac{1}{\alpha+1}} + 2 \left[2 \sum_{i \in \mathcal{I}} w_i \right]^{\frac{1}{\alpha+1}}$, which is a load-dependent constant.

(b) *Negative Drift.* The negative drift condition is established in Theorem 5.3, with $\gamma = \varepsilon K$, for some positive load-dependent constant K .

Note that we have verified conditions (a) and (b) for the Lyapunov function L_α . To show the actual exponential probability tail bound for $\|\mathbf{N}\|_\infty$, note that $L_\alpha(\mathbf{N}) \geq K'' \|\mathbf{N}\|_\infty$, for some load-dependent constant K'' . By suitably redefining the constants B , ξ , and K , the same form of exponential probability tail bound is established for $\|\mathbf{N}\|_\infty$. $\square$

7. An Important Application: Interchange of Limits ($\alpha = 1$).

In this section, we assume throughout that $\alpha = 1$ (the proportionally-fair policy), and establish the validity of the heavy-traffic approximation for networks in steady state. We first provide the necessary preliminaries to state our main theorem, Theorem 7.6. In Section 7.2, we state and prove Theorem 7.6, which is a consequence of Lemmas 7.7 and 7.8. Further definitions and background are provided in Section 7.3, along with the proofs of Lemmas 7.7 and 7.8. All definitions and background stated in this section are taken from [14] and [13].

7.1. Preliminaries. We give a preview of the preliminaries that we will introduce before stating Theorem 7.6. The goal of this subsection is to provide just enough background to be able to state Theorem 7.5, the diffusion approximation result from [13]. To do this, we need a precise description of the process obtained in the limit, under the diffusion scaling. This limiting process is a diffusion process, called Semimartingale Reflecting Brownian Motion (SRBM) (Definition 7.3), with support on a polyhedral cone. This polyhedral cone is defined through the concept of an invariant manifold (Definition 7.2).

As in Section 4.3, we consider a sequence of networks indexed by r , where r is to be thought of as increasing to infinity along a sequence. The incidence matrix $\mathbf{A}$, the capacity vector $\mathbf{C}$, and the weight vector $\boldsymbol{\kappa}$ do not vary with r . Recall the *heavy-traffic* condition — Assumption 4.7, and the definitions of the workload $\mathbf{w}$, the workload process $\mathbf{W}$, and the lifting map Δ from Definition 4.8. We carry the notation from Section 4.3, so that $\boldsymbol{\theta} > \mathbf{0}$, and $\boldsymbol{\nu}^r \rightarrow \boldsymbol{\nu} > \mathbf{0}$, $\boldsymbol{\mu}^r \rightarrow \boldsymbol{\mu} > \mathbf{0}$ and $r(\mathbf{C} - \mathbf{A}\boldsymbol{\rho}^r) \rightarrow \boldsymbol{\theta}$ as $r \rightarrow \infty$. Recall that $\mathbf{A}\boldsymbol{\rho} = \mathbf{C}$. Let $\hat{\mathbf{N}}^r$ and $\hat{\mathbf{W}}^r$ be as in (15).

The continuity of the lifting map Δ will be useful in the sequel.

PROPOSITION 7.1 (Proposition 4.1 in [13]). *The function $\Delta : \mathbb{R}_+^{|\mathcal{J}|} \rightarrow \mathbb{R}_+^{|\mathcal{I}|}$ is continuous. Furthermore, for each $\mathbf{w} \in \mathbb{R}_+^{|\mathcal{J}|}$ and $c > 0$,*

$$(34) \quad \Delta(c\mathbf{w}) = c\Delta(\mathbf{w}).$$

DEFINITION 7.2 (Invariant manifold). *A state $\mathbf{n} \in \mathbb{R}_+^{|\mathcal{I}|}$ is called invariant if $\mathbf{n} = \Delta(\mathbf{w})$, where $\mathbf{w} = \mathbf{A}\mathbf{M}^{-1}\mathbf{n}$ is the workload, and Δ the lifting map defined in Definition 4.8. The set of all invariant states is called the invariant manifold, and we denote it by $\mathcal{M}$. We also define the workload cone $\mathcal{W}$ by $\mathcal{W} = \mathbf{A}\mathbf{M}^{-1}\mathcal{M}$, where $\mathbf{M} = \text{diag}(\boldsymbol{\mu})$ is as defined in Definition 4.8.*

The invariant manifold $\mathcal{M}$ is a polyhedral cone and admits an explicit characterization: we can write it as

$$\mathcal{M} = \left\{ \mathbf{n} \in \mathbb{R}_+^{|\mathcal{I}|} : n_i = \frac{\rho_i(\mathbf{q}^T \mathbf{A})_i}{\kappa_i} \text{ for all } i \in \mathcal{I}, \text{ for some } \mathbf{q} \in \mathbb{R}_+^{|\mathcal{J}|} \right\}.$$

Denote the j -th face of $\mathcal{M}$ by $\mathcal{M}^j$, which can be written as

$$\mathcal{M}^j \triangleq \left\{ \mathbf{n} \in \mathbb{R}_+^{|\mathcal{I}|} : n_i = \frac{\rho_i(\mathbf{q}^T \mathbf{A})_i}{\kappa_i} \text{ for all } i \in \mathcal{I}, \right. \\ \left. \text{for some } \mathbf{q} \in \mathbb{R}_+^{|\mathcal{J}|} \text{ satisfying } q_j = 0 \right\}.$$

Similarly, denote the j -th face of $\mathcal{W}$ by $\mathcal{W}^j$, which can be written as

$$\mathcal{W}^j \triangleq \mathbf{A} \mathbf{M}^{-1} \mathcal{M}^j.$$

Semimartingale Reflecting Brownian Motion (SRBM).

DEFINITION 7.3. *Define the covariance matrix*

$$\mathbf{\Gamma} = 2\mathbf{A} \mathbf{M}^{-1} \text{diag}(\boldsymbol{\nu}) \mathbf{M}^{-1} \mathbf{A}^T.$$

An SRBM that lives in the cone $\mathcal{W}$, has direction of reflection $\mathbf{e}_j$ (the j th unit vector) on the boundary $\mathcal{W}^j$ for each $j \in \mathcal{J}$, has drift $\boldsymbol{\theta}$ and covariance $\mathbf{\Gamma}$, and has initial distribution $\boldsymbol{\eta}_0$ on $\mathcal{W}$ is an adapted, $|\mathcal{J}|$ -dimensional process $\hat{\mathbf{W}}(\cdot)$ defined on some filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, \mathbb{P})$ such that:

- (i) $\mathbb{P}$ -a.s., $\hat{\mathbf{W}}(t) = \hat{\mathbf{W}}(0) + \hat{\mathbf{X}}(t) + \hat{\mathbf{U}}(t)$ for all $t \geq 0$;
- (ii) $\mathbb{P}$ -a.s., $\hat{\mathbf{W}}(\cdot)$ has continuous sample paths, $\hat{\mathbf{W}}(t) \in \mathcal{W}$ for all $t \geq 0$, and $\hat{\mathbf{W}}(0)$ has initial distribution $\boldsymbol{\eta}_0$;
- (iii) under $\mathbb{P}$, $\hat{\mathbf{X}}(\cdot)$ is a $|\mathcal{J}|$ -dimensional Brownian motion starting at the origin with drift $\boldsymbol{\theta}$ and covariance matrix $\mathbf{\Gamma}$;
- (iv) for each $j \in \mathcal{J}$, $\hat{U}_j(\cdot)$ is an adapted, one-dimensional process such that $\mathbb{P}$ -a.s.,
 - (a) $\hat{U}_j(0) = 0$;
 - (b) $\hat{U}_j$ is continuous and non-decreasing;
 - (c) $\hat{U}_j(t) = \int_0^t \mathbb{I}_{\{\hat{\mathbf{W}}(s) \in \mathcal{W}^j\}} d\hat{U}_j(s)$ for all $t \geq 0$.

The process $\hat{\mathbf{W}}(\cdot)$ is called an SRBM with the data $(\mathcal{W}, \boldsymbol{\theta}, \mathbf{\Gamma}, \{\mathbf{e}_j : j \in \mathcal{J}\}, \boldsymbol{\eta}_0)$.

Diffusion Approximation for $\alpha = 1$.

ASSUMPTION 7.4 (Local traffic). *For each $j \in \mathcal{J}$, there exists at least one $i \in \mathcal{I}$ such that $A_{ji} > 0$ and $A_{ki} = 0$ for all $k \neq j$.*

Under the local traffic condition, a diffusion approximation holds.

THEOREM 7.5 (Theorem 5.2 in [13]). *Assume that $\alpha = 1$ and that the local traffic condition, Assumption 7.4, holds. Suppose that the limit distribution of $\hat{\mathbf{W}}^r(0)$ as $r \rightarrow \infty$ is $\boldsymbol{\eta}_0$ (a probability measure on $\mathcal{W}$) and that*

$$(35) \quad \|\hat{\mathbf{N}}^r(0) - \Delta(\hat{\mathbf{W}}^r(0))\|_\infty \rightarrow 0, \quad \text{in probability, as } r \rightarrow \infty.$$

Then, the distribution of $(\hat{\mathbf{W}}^r(\cdot), \hat{\mathbf{N}}^r(\cdot))$ converges weakly (on compact time intervals) as $r \rightarrow \infty$ to a continuous process $(\hat{\mathbf{W}}(\cdot), \hat{\mathbf{N}}(\cdot))$, where $\hat{\mathbf{W}}(\cdot)$ is an SRBM with data $(\mathcal{W}, \boldsymbol{\theta}, \boldsymbol{\Gamma}, \{\mathbf{e}_j, j \in \mathcal{J}\}, \boldsymbol{\eta}_0)$ and $\hat{\mathbf{N}}(t) = \Delta(\hat{\mathbf{W}}(t))$ for all t .

7.2. *Interchange of Limits.* We now know that for $\alpha = 1$, under the local traffic condition, the diffusion approximation holds. That is, the scaled process $(\hat{\mathbf{W}}^r(\cdot), \hat{\mathbf{N}}^r(\cdot))$ converges in distribution to $(\hat{\mathbf{W}}(\cdot), \hat{\mathbf{N}}(\cdot))$, with $\hat{\mathbf{W}}(\cdot)$ being an SRBM. For any r , the scaled processes $\hat{\mathbf{N}}^r(\cdot)$ also have stationary distributions $\boldsymbol{\pi}^r$, since they are all positive recurrent. These results can be summarized in the diagram that follows.

$$\begin{array}{ccc}
 \hat{\mathbf{N}}^r(\cdot)|_{[0,T]} & \xrightarrow[\text{Theorem 7.5}]{r \rightarrow \infty} & \hat{\mathbf{N}}(\cdot)|_{[0,T]} \\
 \downarrow T \rightarrow \infty & & \downarrow T \rightarrow \infty \text{ ?} \\
 \boldsymbol{\pi}^r & \xrightarrow[\text{?}]{r \rightarrow \infty} & \hat{\boldsymbol{\pi}}
 \end{array}$$

As can be seen from the diagram, two natural questions to ask are:

1. Does the diffusion process $\hat{\mathbf{N}}(\cdot)$ have a stationary probability distribution, $\hat{\boldsymbol{\pi}}$?
2. If $\hat{\boldsymbol{\pi}}$ exists and is unique, do the distributions $\boldsymbol{\pi}^r$ converge to $\hat{\boldsymbol{\pi}}$?

Our contribution here is a positive answer to question 2. More specifically, if $\hat{\mathbf{N}}(\cdot)$ has a unique stationary probability distribution $\hat{\boldsymbol{\pi}}$, then $\boldsymbol{\pi}^r$ converges in distribution to $\hat{\boldsymbol{\pi}}$.

THEOREM 7.6. *Suppose that $\alpha = 1$ and that the local traffic condition, Assumption 7.4, holds. Suppose further that $\hat{\mathbf{N}}(\cdot)$ has a unique stationary probability distribution $\hat{\pi}$. For each r , let π^r be the unique stationary probability distribution of $\hat{\mathbf{N}}^r$. Then,*

$$\pi^r \rightarrow \hat{\pi}, \text{ in distribution, as } r \rightarrow \infty.$$

The line of proof of Theorem 7.6 is fairly standard. We first establish tightness of the set of distributions $\{\pi^r\}$ in Lemma 7.7. Letting the processes $\hat{\mathbf{N}}^r(\cdot)$ be initially distributed as $\{\pi^r\}$, we translate this tightness condition into an initial condition similar to (35), in Lemma 7.8. We then apply Theorem 7.5 to deduce the convergence of the processes $\hat{\mathbf{N}}^r(\cdot)$, which by stationarity, leads to the convergence of the distributions π^r . We state Lemmas 7.7 and 7.8 below, and defer their proofs to the next section.

LEMMA 7.7. *Suppose that $\alpha = 1$. The set of probability distributions $\{\pi^r\}$ is tight.*

LEMMA 7.8. *Consider the stationary probability distributions π^r of $\hat{\mathbf{N}}^r(\cdot)$, and let $\{\pi^{r_k}\}$ be any convergent subsequence of $\{\pi^r\}$. Let $\hat{\mathbf{N}}^{r_\ell}(0)$ be distributed as π^{r_ℓ} for each r . Then there exists a subsequence r_ℓ of r_k such that*

$$(36) \quad \left\| \hat{\mathbf{N}}^{r_\ell}(0) - \Delta \left(\hat{\mathbf{W}}^{r_\ell}(0) \right) \right\|_\infty \rightarrow 0$$

in probability as $\ell \rightarrow \infty$, i.e., such that condition (35) holds for the subsequence $\{(\hat{\mathbf{W}}^{r_\ell}(\cdot), \hat{\mathbf{N}}^{r_\ell}(\cdot))\}$.

Proof of Theorem 7.6. Since $\{\pi^r\}$ is tight by Lemma 7.7, Prohorov's theorem implies that $\{\pi^r\}$ is relatively compact in the weak topology. Let $\{\pi^{r_k}\}$ be a convergent subsequence of the set of probability distributions $\{\pi^r\}$, and suppose that $\pi^{r_k} \rightarrow \pi$ as $k \rightarrow \infty$, in distribution.

Let $\hat{\mathbf{N}}^{r_\ell}(0)$ be distributed as π^{r_ℓ} for each r . Then by Lemma 7.8, there exists a subsequence r_ℓ of r_k such that

$$\left\| \hat{\mathbf{N}}^{r_\ell}(0) - \Delta \left(\hat{\mathbf{W}}^{r_\ell}(0) \right) \right\|_\infty \rightarrow 0$$

in probability as $\ell \rightarrow \infty$. Denote the distribution of $\hat{\mathbf{W}}^{r_\ell}(0)$ by η^{r_ℓ} . Since $\pi^{r_k} \rightarrow \pi$ as $k \rightarrow \infty$, $\pi^{r_\ell} \rightarrow \pi$ as $\ell \rightarrow \infty$ as well, and $\eta^{r_\ell} \rightarrow \eta$ as $\ell \rightarrow \infty$, for some probability distribution η .

We now wish to apply Theorem 7.5 to the sequence $\{\hat{\mathbf{N}}^{r_\ell}(\cdot)\}$. The only condition that needs to be verified is that η has support on $\mathcal{W}$. This can be

argued as follows. Let $\hat{\mathbf{N}}(0)$ have distribution $\boldsymbol{\pi}$, and let $\hat{\mathbf{W}}(0) = \mathbf{A}\mathbf{M}^{-1}\hat{\mathbf{N}}(0)$ be the corresponding workload. Then $\hat{\mathbf{W}}^{r_\ell}(0) \rightarrow \hat{\mathbf{W}}(0)$ in distribution as $r \rightarrow \infty$, and $\hat{\mathbf{W}}(0)$ has distribution $\boldsymbol{\eta}$. The lifting map Δ is continuous by Proposition 7.1, so $\Delta(\hat{\mathbf{W}}^{r_\ell}(0)) \rightarrow \Delta(\hat{\mathbf{W}}(0))$ in distribution as $r \rightarrow \infty$. This convergence, together with (36) and the fact that $\hat{\mathbf{N}}^{r_\ell}(0) \rightarrow \hat{\mathbf{N}}(0)$ in distribution, implies that $\hat{\mathbf{N}}(0)$ and $\Delta(\hat{\mathbf{W}}(0))$ are identically distributed. Now $\Delta(\hat{\mathbf{W}}(0))$ has support on $\mathcal{M}$, so $\hat{\mathbf{N}}(0)$ is supported on $\mathcal{M}$ as well, and so $\hat{\mathbf{W}}(0)$, hence $\boldsymbol{\eta}$, is supported on $\mathcal{W}$.

By Theorem 7.5, $(\hat{\mathbf{W}}^{r_\ell}(\cdot), \hat{\mathbf{N}}^{r_\ell}(\cdot))$ converges in distribution to a continuous process $(\hat{\mathbf{W}}(\cdot), \hat{\mathbf{N}}(\cdot))$. Suppose that $\hat{\mathbf{W}}(\cdot)$ and $\hat{\mathbf{N}}(\cdot)$ have unique stationary distributions $\hat{\boldsymbol{\eta}}$ and $\hat{\boldsymbol{\pi}}$, respectively. The processes $(\hat{\mathbf{W}}^{r_\ell}(\cdot), \hat{\mathbf{N}}^{r_\ell}(\cdot))$ are stationary, so $(\hat{\mathbf{W}}(\cdot), \hat{\mathbf{N}}(\cdot))$ is stationary as well. Therefore, $\hat{\mathbf{W}}(0)$ and $\hat{\mathbf{N}}(0)$ are distributed as $\hat{\boldsymbol{\eta}}$ and $\hat{\boldsymbol{\pi}}$, respectively. Since $(\hat{\mathbf{W}}^{r_\ell}(0), \hat{\mathbf{N}}^{r_\ell}(0)) \rightarrow (\hat{\mathbf{W}}(0), \hat{\mathbf{N}}(0))$ in distribution, we have that $\boldsymbol{\eta}^{r_\ell} \rightarrow \hat{\boldsymbol{\eta}}$ and $\boldsymbol{\pi}^{r_\ell} \rightarrow \hat{\boldsymbol{\pi}}$ weakly as $\ell \rightarrow \infty$. This shows that $\boldsymbol{\pi} = \hat{\boldsymbol{\pi}}$ and $\boldsymbol{\eta} = \hat{\boldsymbol{\eta}}$. Since $\{\boldsymbol{\pi}^{r_k}\}$ is an arbitrary convergent subsequence, $\hat{\boldsymbol{\pi}}$ is the unique weak limit point of $\{\boldsymbol{\pi}^r\}$, and this shows that $\boldsymbol{\pi}^r \rightarrow \hat{\boldsymbol{\pi}}$ in distribution. $\square$

For Theorem 7.6 to apply, we need to verify that $\hat{\mathbf{N}}(\cdot)$ (or equivalently, $\hat{\mathbf{W}}(\cdot)$) has a unique stationary distribution. The following theorem states that when $\kappa_i = 1$ for all $i \in \mathcal{I}$, this condition holds; more specifically, the SRBM $\hat{\mathbf{W}}(\cdot)$ has a unique stationary distribution, which turns out to have a product form.

THEOREM 7.9 (Theorem 5.3 in [13]). *Suppose that $\alpha = 1$ and $\kappa_i = 1$ for all $i \in \mathcal{I}$. Let $\hat{\boldsymbol{\eta}}$ be the measure on $\mathcal{W}$ that is absolutely continuous with respect to Lebesgue measure with density given by*

$$(37) \quad p(\mathbf{w}) = \exp(\langle \mathbf{v}, \mathbf{w} \rangle), \quad \mathbf{w} \in \mathcal{W},$$

where

$$(38) \quad \mathbf{v} = 2\boldsymbol{\Gamma}^{-1}\boldsymbol{\theta}.$$

The product measure $\hat{\boldsymbol{\eta}}$ is an invariant measure for the SRBM with state space $\mathcal{W}$, directions of reflection $\{\mathbf{e}_j, j \in \mathcal{J}\}$, drift $\boldsymbol{\theta}$, and covariance matrix $\boldsymbol{\Gamma}$. After normalization, it defines the unique stationary distribution for the SRBM.

By Theorems 7.6 and 7.9, the following corollary is immediate.

COROLLARY 7.10. *Suppose that $\alpha = 1$ and $\kappa_i = 1$ for all $i \in \mathcal{I}$. Suppose further that the local traffic condition, Assumption 7.4, holds. Let $\hat{\pi}$ be the unique stationary probability distribution of $\hat{\mathbf{N}}(\cdot)$. For each r , let π^r be the unique stationary probability distribution of $\mathbf{N}^r$. Then,*

$$\pi^r \rightarrow \hat{\pi}, \text{ in distribution, as } r \rightarrow \infty.$$

7.3. Proof of Lemmas 7.7 and 7.8.

Proof of Lemma 7.7. To establish tightness, it suffices to show that for every $y > 0$ there exists a compact set $\mathbb{K}_y \subset \mathbb{R}_+^{|\mathcal{I}|}$ such that

$$(39) \quad \limsup_{r \rightarrow \infty} \pi^r \left(\mathbb{R}_+^{|\mathcal{I}|} \setminus \mathbb{K}_y \right) \leq e^{-y}.$$

We now proceed to define the compact sets $\mathbb{K}_y$. As in the proof of Theorem 4.10, let $\varepsilon_r = \varepsilon(\rho^r)$ be the gap in the r th system. Then, under Assumption 4.7, for sufficiently large r , $\varepsilon_r \geq D/r$ for some network-dependent constant $D > 0$. Since $\alpha = 1$, Theorem 3.2 implies that for the r th system, there exist load-dependent constants $K_r > 0$ and $\xi_r > 0$ such that for every $\ell \in \mathbb{Z}_+$,

$$(40) \quad \mathbb{P}_{\pi^r} \left(\|\mathbf{N}^r\|_\infty \geq \frac{K_r}{\varepsilon_r} + 2\xi_r \ell \right) \leq \left(\frac{\xi_r}{\xi_r + \varepsilon_r} \right)^{\ell+1}.$$

By the definition of a positive load-dependent constant, there exist continuous functions f_1 and f_2 on the open positive orthant such that for all r , $K_r = f_1(\mu^r, \nu^r)$ and $\xi_r = f_2(\mu^r, \nu^r)$. Since $\mu^r \rightarrow \mu > \mathbf{0}$ and $\nu^r \rightarrow \nu > \mathbf{0}$, we have $K_r \rightarrow K \triangleq f_1(\mu, \nu) > 0$ and $\xi_r \rightarrow \xi \triangleq f_2(\mu, \nu) > 0$. Define

$$\mathbb{K}_y \triangleq \left\{ \mathbf{v} \in \mathbb{R}_+^{|\mathcal{I}|} : \|\mathbf{v}\|_\infty \leq \frac{(K+1) + 4(\xi+1)^2 \cdot y}{D} \right\}.$$

We now show that (39) holds, or equivalently, by the definition of $\mathbb{K}_y$, we show that for every $y > 0$,

$$(41) \quad \limsup_{r \rightarrow \infty} \mathbb{P}_{\pi^r} \left(\frac{1}{r} \|\mathbf{N}^r\|_\infty > \frac{(K+1) + 4(\xi+1)^2 y}{D} \right) \leq e^{-y}.$$

Let $\ell_r \triangleq \lfloor 2\xi_r y / \varepsilon_r \rfloor$, where for $z \in \mathbb{R}$, $\lfloor z \rfloor$ is the largest integer not exceeding z . By (40), we have

$$\mathbb{P}_{\pi^r} \left(\frac{1}{r} \|\mathbf{N}^r\|_\infty \geq \frac{K_r}{r\varepsilon_r} + \frac{2\xi_r \ell_r}{r} \right) \leq \left(\frac{1}{1 + \frac{\varepsilon_r}{\xi_r}} \right)^{\ell_r+1}.$$

Taking logarithms on both sides, we have

$$\log \mathbb{P}_{\pi^r} \left(\frac{1}{r} \|\mathbf{N}^r\|_\infty \geq \frac{K_r}{r\varepsilon_r} + \frac{2\xi_r \ell_r}{r} \right) \leq -(\ell_r + 1) \log \left(1 + \frac{\varepsilon_r}{\xi_r} \right).$$

Since $\varepsilon_r \rightarrow 0$ and $\xi_r \rightarrow \xi > 0$ as $r \rightarrow \infty$, $\frac{\varepsilon_r}{\xi_r} < 1$ for sufficiently large r . Since $\log(1+t) \geq t/2$ for $t \in [0, 1]$, we have

$$-(\ell_r + 1) \log \left(1 + \frac{\varepsilon_r}{\xi_r} \right) \leq -(\ell_r + 1) \frac{\varepsilon_r}{2\xi_r},$$

when r is sufficiently large. By definition, $\ell_r = \lfloor 2\xi_r y / \varepsilon_r \rfloor$, so $\ell_r + 1 \geq 2\xi_r y / \varepsilon_r$, or equivalently, $-(\ell_r + 1) \frac{\varepsilon_r}{2\xi_r} \leq -y$. Thus, when r is sufficiently large,

$$\log \mathbb{P}_{\pi^r} \left(\frac{1}{r} \|\mathbf{N}^r\|_\infty \geq \frac{K_r}{r\varepsilon_r} + \frac{2\xi_r \ell_r}{r} \right) \leq -y.$$

Consider the term $\frac{K_r}{r\varepsilon_r} + \frac{2\xi_r \ell_r}{r}$. When r is sufficiently large, $r\varepsilon_r \geq D$, $K_r \leq K + 1$, and $\xi_r \leq \xi + 1$, and so

$$\frac{K_r}{r\varepsilon_r} + \frac{2\xi_r \ell_r}{r} \leq \frac{K_r}{r\varepsilon_r} + \frac{2\xi_r (2\xi_r y)}{r\varepsilon_r} \leq \frac{K + 1}{D} + \frac{4(\xi + 1)^2 y}{D}.$$

Thus, for sufficiently large r ,

$$\begin{aligned} & \log \mathbb{P}_{\pi^r} \left(\frac{1}{r} \|\mathbf{N}^r\|_\infty > \frac{(K + 1) + 4(\xi + 1)^2 y}{D} \right) \\ & \leq \log \mathbb{P}_{\pi^r} \left(\frac{1}{r} \|\mathbf{N}^r\|_\infty \geq \frac{K_r}{r\varepsilon_r} + \frac{2\xi_r \ell_r}{r} \right) \leq -y. \end{aligned}$$

This establishes (41), and also the tightness of $\{\pi^r\}$. $\square$

Next, we prove Lemma 7.8. To this end, we need some definitions and background. In particular, we need the concept and properties of *fluid model solutions*.

DEFINITION 7.11. A fluid model solution (FMS) is an absolutely continuous function $\mathbf{n} : [0, \infty) \rightarrow \mathbb{R}_+^{|\mathcal{I}|}$ such that at each regular point ¹ $t > 0$ of $\mathbf{n}(\cdot)$, we have, for each $i \in \mathcal{I}$,

$$(42) \quad \frac{d}{dt} n_i(t) = \begin{cases} \nu_i - \mu_i \Lambda_i(\mathbf{n}(t)), & \text{if } n_i(t) > 0, \\ 0, & \text{if } n_i(t) = 0, \end{cases}$$

¹A point $t \in (0, \infty)$ is a *regular point* of an absolutely continuous function $f : [0, \infty) \rightarrow \mathbb{R}_+^{|\mathcal{I}|}$ if each component of f is differentiable at t . Since $\mathbf{n}$ is absolutely continuous, almost every time $t \in (0, \infty)$ is a regular point for $\mathbf{n}$.

and for each $j \in \mathcal{J}$,

$$(43) \quad \sum_{i \in \mathcal{I}_+(\mathbf{n}(t))} A_{ji} \Lambda_i(\mathbf{n}(t)) + \sum_{i \in \mathcal{I}_0(\mathbf{n}(t))} A_{ji} \rho_i \leq C_j,$$

where $\mathcal{I}_+(\mathbf{n}(t)) = \{i \in \mathcal{I} : n_i(t) > 0\}$ and $\mathcal{I}_0(\mathbf{n}(t)) = \{i \in \mathcal{I} : n_i(t) = 0\}$. Note that here $\mathbf{A}\boldsymbol{\rho} = \mathbf{C}$.

We now collect some properties of a FMS. The following proposition states that the invariant manifold $\mathcal{M}$ consists exactly of the stationary points of a FMS.

PROPOSITION 7.12 (Theorem 4.1 in [13]). *A vector $\mathbf{n}_0$ is an invariant state, that is, $\mathbf{n}_0 \in \mathcal{M}$, if and only if for every fluid model solution $\mathbf{n}(\cdot)$ with $\mathbf{n}(0) = \mathbf{n}_0$, we have $\mathbf{n}(t) = \mathbf{n}_0$ for all $t > 0$.*

The following theorem states that starting from any initial condition, a FMS will eventually be close to the invariant manifold $\mathcal{M}$.

THEOREM 7.13 (Theorem 5.2 in [14]). *Fix $R \in (0, \infty)$ and $\delta > 0$. There is a constant $T_{R,\delta} < \infty$ such that for every fluid model solution $\mathbf{n}(\cdot)$ satisfying $\|\mathbf{n}(0)\|_\infty \leq R$ we have*

$$d(\mathbf{n}(t), \mathcal{M}) < \delta, \quad \text{for all } t > T_{R,\delta},$$

where $d(\mathbf{n}(t), \mathcal{M}) \triangleq \inf_{\mathbf{n} \in \mathcal{M}} \|\mathbf{n} - \mathbf{n}(t)\|_\infty$ is the distance from $\mathbf{n}(t)$ to the manifold $\mathcal{M}$.

Proposition 7.14 states that the value of the Lyapunov function F_1 defined in (9) decreases along the path of any FMS.

PROPOSITION 7.14 (Corollary 6.1 in [14]). *At any regular point t of a fluid model solution $\mathbf{n}(\cdot)$, we have*

$$\frac{d}{dt} F_1(\mathbf{n}(t)) \leq 0,$$

and the inequality is strict if $\mathbf{n}(t) \notin \mathcal{M}$.

Using Proposition 7.14, and the continuity of the lifting map Δ , we can translate Theorem 7.13 into the following version, which will be used to prove Lemma 7.8.

LEMMA 7.15. *Fix $R \in (0, \infty)$ and $\delta > 0$. There is a constant $T_{R,\delta} < \infty$ such that for every fluid model solution $\mathbf{n}(\cdot)$ satisfying $\|\mathbf{n}(0)\|_\infty \leq R$ we have*

$$\|\mathbf{n}(t) - \Delta(\mathbf{w}(t))\|_\infty < \delta, \quad \text{for all } t > T_{R,\delta},$$

where $\mathbf{w}(t) = \mathbf{w}(\mathbf{n}(t))$ is the workload corresponding to $\mathbf{n}(t)$ (see Definition 4.8).

PROOF. Fix $R > 0$ and $\delta > 0$. Let $\|\mathbf{n}(0)\|_\infty \leq R$. Then,

$$F_1(\mathbf{n}(0)) = \frac{1}{2} \sum_{i \in I} \nu_i^{-1} \kappa_i n_i^2(0) \leq R',$$

where R' depends on R and the system parameters. Since $\mathbf{n}(\cdot)$ is absolutely continuous, by Proposition 7.14 and the fundamental theorem of calculus, we have that $F_1(\mathbf{n}(t)) \leq R'$ for all $t \geq 0$. Define the set

$$\mathcal{S} \triangleq \{\mathbf{n} \in \mathbb{R}_+^{|\mathcal{I}|} : F_1(\mathbf{n}) \leq R'\},$$

and its δ -fattening

$$\mathcal{S}_\delta \triangleq \{\mathbf{n} \in \mathbb{R}_+^{|\mathcal{I}|} : \|\mathbf{n} - \mathbf{n}'\| \leq \delta \text{ for some } \mathbf{n}' \in \mathcal{S}\}.$$

Note that both $\mathcal{S}$ and $\mathcal{S}_\delta$ are compact sets, and $\mathbf{n}(t) \in \mathcal{S} \subset \mathcal{S}_\delta$ for all $t \geq 0$.

Now consider the workload $\mathbf{w}$ defined in Definition 4.8. Define the set $\mathbf{w}(\mathcal{S}_\delta) = \{\mathbf{v} \in \mathbb{R}_+^{|\mathcal{J}|} : \mathbf{v} = \mathbf{w}(\mathbf{n}) \text{ for some } \mathbf{n} \in \mathcal{S}_\delta\}$. Since $\mathbf{w}$ is a linear map, there exists a load-dependent constant H such that

$$\|\mathbf{w}(\mathbf{n}) - \mathbf{w}(\mathbf{n}')\|_\infty \leq H \|\mathbf{n} - \mathbf{n}'\|_\infty,$$

for any $\mathbf{n}, \mathbf{n}' \in \mathbb{R}_+^{|\mathcal{I}|}$. Thus $\mathbf{w}(\mathcal{S}_\delta)$ is also a compact set. Since $\mathbf{n}(t) \in \mathcal{S}_\delta$ for all $t \geq 0$, $\mathbf{w}(t) \in \mathbf{w}(\mathcal{S}_\delta)$ for all $t \geq 0$. By Proposition 7.1, Δ is a continuous map, so Δ is uniformly continuous when restricted to $\mathbf{w}(\mathcal{S}_\delta)$. Therefore, there exists $\delta' > 0$ such that for any $\mathbf{w}', \mathbf{w} \in \mathbf{w}(\mathcal{S}_\delta)$ with $\|\mathbf{w}' - \mathbf{w}\|_\infty < \delta'$, $\|\Delta(\mathbf{w}') - \Delta(\mathbf{w})\|_\infty < \frac{\delta}{2}$. Thus for any $\mathbf{n}, \mathbf{n}' \in \mathcal{S}_\delta$ with $\|\mathbf{n} - \mathbf{n}'\|_\infty < \delta'/H$, we have $\|\mathbf{w}(\mathbf{n}) - \mathbf{w}(\mathbf{n}')\| \leq \delta'$, and

$$\|\Delta(\mathbf{w}(\mathbf{n})) - \Delta(\mathbf{w}(\mathbf{n}'))\|_\infty < \frac{\delta}{2}.$$

Let $\delta'' = \min\{\delta/2, \delta'/H\}$. By Theorem 7.13, there exists $T_{R,\delta''}$ such that for all $t \geq T_{R,\delta''}$,

$$d(\mathcal{M}, \mathbf{n}(t)) < \delta''.$$

In particular, there exists $\mathbf{n} \in \mathcal{M}$ (which may depend on $\mathbf{n}(t)$) such that $\|\mathbf{n} - \mathbf{n}(t)\|_\infty < \delta'' < \delta'/H$. Since $\mathbf{n}(t) \in \mathcal{S}$ and $\delta'' < \delta$, $\mathbf{n} \in \mathcal{S}_\delta$ as well. Thus

$$\|\Delta(\mathbf{w}(\mathbf{n})) - \Delta(\mathbf{w}(\mathbf{n}(t)))\|_\infty < \frac{\delta}{2}.$$

By Proposition 7.12, since $\mathbf{n} \in \mathcal{M}$, we have $\mathbf{n} = \Delta(\mathbf{w}(\mathbf{n}))$, and hence

$$\|\mathbf{n} - \Delta(\mathbf{w}(\mathbf{n}(t)))\|_\infty < \frac{\delta}{2}.$$

Thus for all $t \geq T_{R,\delta''}$,

$$\begin{aligned} \|\mathbf{n}(t) - \Delta(\mathbf{w}(t))\|_\infty &\leq \|\mathbf{n} - \mathbf{n}(t)\|_\infty + \|\mathbf{n} - \Delta(\mathbf{w}(t))\|_\infty \\ &< \delta'' + \frac{\delta}{2} \leq \frac{\delta}{2} + \frac{\delta}{2} = \delta. \end{aligned}$$

Note that δ'' depends on R , δ , and the system parameters. Thus, we can rewrite $T_{R,\delta''}$ as $T_{R,\delta}$. This concludes the proof of the lemma. $\square$

The last property of a FMS that we need is the tightness of the fluid-scaled processes $\bar{\mathbf{N}}^r$ and $\bar{\mathbf{W}}^r$, defined by

$$(44) \quad \bar{\mathbf{N}}^r(t) = \mathbf{N}^r(rt)/r, \text{ and } \bar{\mathbf{W}}^r(t) = \mathbf{W}^r(rt)/r.$$

THEOREM 7.16 (Theorem B.1 in [14]). *Suppose that $\{\bar{\mathbf{N}}^r(0)\}$ converges in distribution as $r \rightarrow \infty$ to a random variable taking values in $\mathbb{R}_+^{|\mathcal{I}|}$. Then, the sequence $\{\bar{\mathbf{N}}^r(\cdot)\}$ is C -tight², and any weak limit point $\bar{\mathbf{N}}(\cdot)$ of this sequence, almost surely satisfies the fluid model equations (42) and (43).*

Proof of Lemma 7.8. Consider the unique stationary distributions $\boldsymbol{\pi}^r$ of $\hat{\mathbf{N}}^r(\cdot)$, and $\boldsymbol{\eta}^r$ of $\hat{\mathbf{W}}^r(\cdot)$. Let $\boldsymbol{\pi}^{r_k}$ be a convergent subsequence, and suppose that $\boldsymbol{\pi}^{r_k} \rightarrow \boldsymbol{\pi}$ in distribution, as $k \rightarrow \infty$. Suppose that at time 0, $\frac{1}{r_k}\mathbf{N}^{r_k}(0)$ is distributed as $\boldsymbol{\pi}^{r_k}$. Then $\frac{1}{r_k}\mathbf{W}^{r_k}(0)$ is distributed as $\boldsymbol{\eta}_k^r$, which converges in distribution as well, say to $\boldsymbol{\eta}$.

We now use the earlier stated FMS properties to prove the lemma. Note that for all r ,

$$\frac{1}{r}\mathbf{N}^r(0) = \bar{\mathbf{N}}^r(0) = \hat{\mathbf{N}}^r(0) \quad \text{and} \quad \frac{1}{r}\mathbf{W}^r(0) = \bar{\mathbf{W}}^r(0) = \hat{\mathbf{W}}^r(0),$$

²Consider the space $\mathbf{D}^{|\mathcal{I}|}$ of functions $f : [0, \infty) \rightarrow \mathbb{R}^{|\mathcal{I}|}$ that are right-continuous on $[0, \infty)$ and have finite limits from the left on $(0, \infty)$. Let this space be endowed with the usual Skorohod topology (cf. Section 12 of [3]). The sequence $\{\bar{\mathbf{N}}^r(\cdot)\}$ is *tight* if the probability measures induced on $\mathbf{D}^{|\mathcal{I}|}$ are tight. The sequence is *C-tight* if it is tight and any weak limit point is a measure supported on the set of continuous sample paths.

and consider the fluid-scaled processes $\bar{\mathbf{N}}^{r_k}(\cdot)$ and $\bar{\mathbf{W}}^{r_k}(\cdot)$. Since $\{\bar{\mathbf{N}}^{r_k}(0)\}$ converges in distribution to $\boldsymbol{\pi}$, Theorem 7.16 implies that the sequence $\{\bar{\mathbf{N}}^{r_k}(\cdot)\}$ is C -tight, and any weak limit $\bar{\mathbf{N}}(\cdot)$ almost surely satisfies the fluid model equations. Let $\bar{\mathbf{N}}(\cdot)$ be a weak limit point of $\{\bar{\mathbf{N}}^{r_k}(\cdot)\}$, and suppose that the subsequence $\{\bar{\mathbf{N}}^{r_\ell}(\cdot)\}$ of $\{\bar{\mathbf{N}}^{r_k}(\cdot)\}$ converges weakly to $\bar{\mathbf{N}}(\cdot)$.

Let $\delta > 0$. We will show that we can find $r(\delta)$ such that for $r_\ell > r(\delta)$,

$$\mathbb{P}(\|\bar{\mathbf{N}}^{r_\ell}(0) - \Delta(\bar{\mathbf{W}}^{r_\ell}(0))\|_\infty > \delta) < \delta.$$

Since $\bar{\mathbf{N}}(0)$ is a well-defined random variable, there exists $R_\delta > 0$ such that

$$\mathbb{P}(\|\bar{\mathbf{N}}(0)\|_\infty > R_\delta) < \frac{\delta}{2}.$$

Now, for all sample paths ω such that $\|\bar{\mathbf{N}}(0)(\omega)\|_\infty \leq R_\delta$, and such that $\bar{\mathbf{N}}(\cdot)(\omega)$ satisfies the fluid model equations, Lemma 7.15 implies that there exists $T \triangleq T_{R_\delta, \delta}$ such that

$$\|\bar{\mathbf{N}}(T)(\omega) - \Delta(\bar{\mathbf{W}}(T))(\omega)\|_\infty < \delta.$$

Since $\bar{\mathbf{N}}(\cdot)$ satisfies the fluid model equations almost surely, we have

$$\mathbb{P}(\|\bar{\mathbf{N}}(T) - \Delta(\bar{\mathbf{W}}(T))\|_\infty < \delta) > 1 - \frac{\delta}{2}.$$

Now for each r , $\bar{\mathbf{N}}^r(0)$ is distributed according to the stationary distribution $\boldsymbol{\pi}^r$, so $\bar{\mathbf{N}}^r(\cdot)$ is a stationary process. Since $\bar{\mathbf{N}}^{r_\ell}(\cdot) \rightarrow \bar{\mathbf{N}}(\cdot)$ weakly as $\ell \rightarrow \infty$, $\bar{\mathbf{N}}$ is also a stationary process. Thus, $\bar{\mathbf{N}}(T)$ and $\bar{\mathbf{N}}(0)$ are both distributed according to $\boldsymbol{\pi}$. This implies that

$$\mathbb{P}(\|\bar{\mathbf{N}}(0) - \Delta(\bar{\mathbf{W}}(0))\|_\infty < \delta) > 1 - \frac{\delta}{2}.$$

Furthermore, since $\bar{\mathbf{N}}^{r_\ell}(0) \rightarrow \bar{\mathbf{N}}(0)$ in distribution,

$$\mathbb{P}(\|\bar{\mathbf{N}}^{r_\ell}(0) - \Delta(\bar{\mathbf{W}}^{r_\ell}(0))\|_\infty < \delta) \rightarrow \mathbb{P}(\|\bar{\mathbf{N}}(0) - \Delta(\bar{\mathbf{W}}(0))\|_\infty < \delta)$$

as $\ell \rightarrow \infty$. Thus there exists $r(\delta)$ such that for all $r_\ell > r(\delta)$,

$$\mathbb{P}(\|\bar{\mathbf{N}}^{r_\ell}(0) - \Delta(\bar{\mathbf{W}}^{r_\ell}(0))\|_\infty < \delta) > 1 - \delta.$$

Since $\delta > 0$ is arbitrary,

$$\|\hat{\mathbf{N}}^{r_\ell}(0) - \Delta(\hat{\mathbf{W}}^{r_\ell}(0))\|_\infty = \|\bar{\mathbf{N}}^{r_\ell}(0) - \Delta(\bar{\mathbf{W}}^{r_\ell}(0))\|_\infty \rightarrow 0,$$

in probability. □

8. Conclusion. The results in this paper can be viewed from two different perspectives. On the one hand, they provide much new information on the qualitative behavior (e.g., finiteness of the expected number of flows, bounds on steady-state tail probabilities and finite-horizon maximum excursion probabilities, etc.) of the α -fair policies for bandwidth-sharing network models. At an abstract level, our results highlight the importance of relying on a suitable Lyapunov function. Even if a network is shown to be stable by using a particular Lyapunov function, different choices and more detailed analysis may lead to more powerful bounds. At a more concrete level, we presented a generic method for deriving full state space collapse from multiplicative state space collapse, and another for deriving steady-state exponential tail bounds. The methods and results in this paper can be extended to general switched network models. Parallel results for a packet level model are detailed in [18].

REFERENCES

- [1] D. Bertsekas. *Nonlinear Programming*. Athena Scientific, 1999.
- [2] D. Bertsimas, D. Gamarnik, and J. N. Tsitsiklis. Performance of multiclass Markovian queueing networks via piecewise linear Lyapunov functions. *The Annals of Applied Probability*, 11(4):1384–1428, 2001.
- [3] P. Billingsley. *Convergence of Probability Measures*. Wiley-Interscience, 1999.
- [4] T. Bonald and L. Massoulié. Impact of fairness on Internet performance. *ACM SIGMETRICS Performance Evaluation Review*, 29(1):82–91, 2001.
- [5] M. Bramson. State space collapse with application to heavy traffic limits for multiclass queueing networks. *Queueing Systems*, 30:89–148, 1998.
- [6] G. De Veciana, T. Konstantopoulos, and T. Lee. Stability and performance analysis of networks supporting elastic services. *IEEE/ACM Transactions on Networking (TON)*, 9(1), 2001.
- [7] A. Eryilmaz and R. Srikant. Asymptotically tight steady-state queue length bounds implied by drift conditions. *Arxiv*, April 2011.
- [8] S. Foss and T. Konstantopoulos. An overview of some stochastic stability methods. *Journal of Operations Research, Society of Japan*, 47(4), 2004.
- [9] R. Gallager. *Discrete Stochastic Processes*. Kluwer Academic Publishers, 1996.
- [10] D. Gamarnik and A. Zeevi. Validity of heavy traffic steady-state approximations in generalized Jackson networks. *The Annals of Applied Probability*, 16(1):56 – 90, 2006.
- [11] G. Grimmett and D. Stirzaker. *Probability and Random Processes*. Oxford University Press, 2001.
- [12] B. Hajek. Hitting and occupation time bounds implied by drift analysis with applications. *Advances of Applied Probability*, 14:502–525, 1982.
- [13] W. Kang, F. Kelly, N. Lee, and R. Williams. State space collapse and diffusion approximation for a network operating under a fair bandwidth sharing policy. *The Annals of Applied Probability*, 19(5):1719–1780, 2009.
- [14] F. Kelly and R. Williams. Fluid model for a network operating under a fair bandwidth-sharing policy. *The Annals of Applied Probability*, 14(3):1055–1083, 2004.

- [15] 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.
- [16] J. Mo and J. Walrand. Fair end-to-end window-based congestion control. *IEEE/ACM Transactions on Networking*, 8(5):556–567, 2000.
- [17] J. Roberts and L. Massoulié. Bandwidth sharing and admission control for elastic traffic. *Telecommunication Systems*, 15:185–201, 2000.
- [18] D. Shah, J. N. Tsitsiklis, and Y. Zhong. Qualitative properties of alpha-weighted scheduling policies. In *Proceedings of the ACM SIGMETRICS*, June 2010.
- [19] D. Shah and D. J. Wischik. Switched networks with maximum weight policies: Fluid approximation and multiplicative state space collapse. *The Annals of Applied Probability*, 22(1):70–127, 2012.
- [20] R. Srikant. Models and Methods for analyzing Internet congestion control algorithms. *Advances in Communication Control Networks in the series “Lecture Notes in Control and Information Sciences,”* 308:416–419, 2005.
- [21] A. Stolyar. Large deviations of queues sharing a randomly time-varying server. *Queueing Systems*, 59(1):1–35, 2008.
- [22] V. Subramanian. LDP for max-weight scheduling over convex compact rate-regions, June 2010. *Mathematics of Operations Research*, 35(4):881–910, 2010.
- [23] V. J. Venkataramanan and X. Lin. On the large-deviations optimality of scheduling policies minimizing the drift of a Lyapunov function. In *Proceedings of the 47th Annual Allerton Conference on Communication, Control, and Computing*, September 2009.
- [24] R. Williams. Diffusion approximations for open multiclass queueing networks: sufficient conditions involving state space collapse. *Queueing Systems*, 30:27–88, 1998.
- [25] H. Ye and D. Yao. A stochastic network under proportional fair resource control - diffusion limit with multiple bottlenecks. *To appear in Operations Research*.